

NOS Version 2

Analysis Handbook

Reference

This product is intended for use only as described in this document. Control Data cannot be responsible for the proper functioning of undescribed features and parameters.

Manual History

Revision	System Version	PSR Level	Date
A	2.0	562	April 1982
B	2.1	580	January 1983
C	2.2	596	October 1983
D	2.3	617	October 1984
E	2.4.1	630	March 1985
F	2.4.2	642	October 1985
G	2.5.1	664	September 1986
H	2.5.1	670	December 1986
J	2.5.2	678	April 1987
K	2.5.3	688	September 1987
L	2.6.1	700	April 1988
M	2.7.1	716	December 1988
N	2.7.1	739	December 1989
P	2.7.1	750	June 1990
R	2.7.2	774	June 1991
T	2.7.4	797	June 1992
U	2.8.3	840	July 1994

Revision U of this manual, printed July 1994, reflects NOS 2.8.3 at PSR level 840.

This edition obsoletes all previous editions.

Contents

About This Manual	13
Audience	13
Organization	13
Conventions	14
Related Publications	15
Submitting Comments	20
Central Software Support Hotline	20
Controlware Utilities	1-1
Loading Controlware	1-1
Dumping Controlware	1-4
Deadstart	2-1
Modifying the CMRDECK	2-2
Modifying the EQPDECK	2-3
Modifying the APRDECKs	2-21
Modifying the IPRDECK	2-22
Loading the System	2-23
Initiating Job Processing	2-25
Preparing for Recovery Deadstart	2-27
Level 0 Initial Deadstart	2-35
Error Processing	2-35
Deadstart Decks	3-1
CMRDECK	3-2
EQPDECK	3-13
APRDECK	3-103
IPRDECK	3-107
LIBDECK	3-161
DIS Operations	4-1
DIS Job Dayfile Display (A)	4-3
DIS Job Status Display (B)	4-4
DIS Memory Displays (C, D, F, G)	4-6
DIS Exchange Package Display (X)	4-8
DIS Directory Display (Z)	4-9
Console Operation	4-10
Display Selection Commands	4-13
DIS Keyboard Entries	4-14
Memory Entry Commands	4-18
PP Call Commands	4-20
DSD Commands	5-1
Display Selection Commands	5-2
Dayfile Commands	5-3
Queued File Utility Commands	5-4

Job Processing Control Commands	5-4
Peripheral Equipment Control Commands	5-24
Subsystem Control Commands	5-50
System Control Commands	5-56
Secured System Control Commands	5-70
Memory Entry Commands	5-72
Channel Control Commands	5-74
Extended Memory Flag Register Commands	5-75
Breakpoint Package Commands	5-75
Debugging Commands	5-84
Express Deadstart Dump Interpreter (DSDI)	6-1
Calling the Express Deadstart Dump Interpreter	6-3
Input Directives	6-6
Interactive Use of DSDI	6-52
Printer Output Listing Examples	6-62
Install Command	7-1
K-Display Utilities	8-1
FLAW K Display	8-3
INITIALIZE K Display	8-8
Machine Recovery (MREC) Utility K Display	8-16
MREC Procedures	8-17
MREC Unit and Controller Reservations	8-22
Mass Storage Extended Subsystem (MSE) K Display	8-24
Network Access Method (NAM) K Display	8-26
Queue File Transfer Facility (QTF) K Display	8-40
REDEFINE K Display	8-59
Remote Batch Facility (RBF) K Display	8-70
Remote Host Facility (RHF) K Display	8-73
SCOPE 2 Station Facility (SSF) K Displays	8-82
File Transfer Limit Commands	8-84
Transaction Facility (TAF) K Displays	8-86
L-Display Utilities	9-1
FOTD L Display	9-2
LIDOU L Display	9-3
QDSPLAY L Display	9-6
SCTD L Display	9-9
SDSPLAY L Display	9-11
SUBSYST L Display	9-19
LID/RHF Configuration Files	10-1
LID Configuration File	10-1
RHF Configuration Files	10-8
QTF Configuration Requirements	10-20
Mass Storage Extended Subsystem (MSE)	11-1
Introduction	11-1

MSE Utilities	11-13
MSE Operational Procedures	11-77
Tape and Optical Disk Alternate Storage	12-1
Introduction	12-1
MAGNET Command	12-11
GENPFD Utility	12-11
Multimainframe Operations	13-1
Introduction	13-1
Linked SDM Operation	13-2
Independent SDM Operation	13-11
NAD Maintenance Utilities	14-1
Dump NAD Memory (DMPNAD)	14-1
Maintenance Host Facility (MHF)	14-2
Listing NAD Dumps	14-3
Network Operations	15-1
Network Organization	15-1
NAM Startup	15-5
NAM Shutdown	15-8
Network Control by HOP	15-9
Network Control by NOP	15-21
Network Control by DOP	15-41
Recent HISTORY Command	15-49
Report Unsolicited Status Command	15-50
Send Message Command	15-51
Status Network Element Command	15-52
Summary of Network Operation Commands	15-56
Network Utilities	16-1
NAM Initialization (NAMI)	16-1
Network Dump Collector (COLLECT)	16-5
Load File Generator (LFG)	16-8
NPU Dump Analyzer (NDA)	16-11
PIP Dump Analyzer (LISTPPM)	16-28
AIP Trace With IAF	16-29
Network Description Files	16-31
IAF Stimulator	16-33
Permanent File Utilities	17-1
Introduction	17-1
Calling the Utilities	17-36
Parameters for the Permanent File Utilities	17-45
Permanent File Utility Routines	17-62
Queued/Dayfile Utilities	18-1
Queued File Utilities	18-1

Dayfile Dumping Utilities	18-55
SYSEEDIT	19-1
Directives	19-3
Record Types	19-10
System File Initialization	20-1
ISF	20-1
Deadstart Sequencing	20-4
Tracer/Probe Utilities	21-1
Tracer Utility	21-1
PROBE Utility	21-36
Cartridge Tape Subsystem (CTS) and Automated Cartridge Subsystem (ACS)	22-1
Cartridge Tape Subsystem	22-1
Automated Cartridge Subsystem	22-6
Cross-Coupled Controllers for CTS and ACS	22-13
Requirements for CTS and ACS	22-16
Constraints for CTS and ACS	22-16
Performance Considerations for CTS and ACS	22-17
Maintenance for CTS and ACS	22-17
Character Sets	A-1
Character Set Anomalies	A-2
Character Set Tables	A-2
Glossary	B-1
Scope 2 Station Facility	C-1
Equipment Configuration	C-1
SSF File Transfers	C-1
Error Logout	C-3
Status/Control Register Simulator	D-1
Using the Simulator	D-2
Simulator Commands	D-5
580 Programmable Format Control	E-1
Building Programmable Format Control Arrays	E-2
Adding Programmable Format Control Arrays	E-3
533/536, 537, and 585 Printer Electronic Vertical Format Unit (EVFU)	F-1
EVFU Directives for 533/536 Printers	F-1
881/883 Pack Reformatting Utility	G-1
FORMAT Command	G-2

Input Formats	G-4
Output Formats	G-6
Accessing Disk Devices	G-11
NOS/VE Address Formats	H-1
Management Of Storage Media Defects	I-1
Detecting Failing Devices	I-1
Dealing With Known Media Defects	I-3
Display Disk File (DDF) Utility	J-1
Bringing Up DDF	J-1
Keyboard Input	J-4
Console Messages	J-5
DDF Commands	J-6
DDF Examples	J-12
Pack Recovery	J-16
PACKER Utility	K-1
Operational Overview	K-1
Command Format	K-4
PACKER Examples	K-6
Index	Index-1

Figures

2-1. Deadstart Process	2-1
2-2. Equipment Status Display	2-8
2-3. Mass Storage Status Display	2-12
2-4. Mass Storage Initialization Status Display	2-16
2-5. Controlware Status Display	2-18
2-6. Disk Thresholds Display	2-19
2-7. Level 3 Deadstart Left Screen Display	2-30
2-8. Level 3 Deadstart Right Screen Display	2-31
3-1. Expander Addressing Map	3-47
3-2. 844 Expander Configuration With 20 Drives	3-48
3-3. Vertically Adjacent Three-Unit Device	3-49
3-4. Horizontally Adjacent Two-Unit Device	3-49
3-5. Hardware Configured into Seven Devices	3-50
3-6. Dual-Channel Access to 5830 Drives	3-54
3-7. Example of Ranges of Service Class Priorities	3-119
4-1. DIS Job Dayfile Display (A)	4-3
4-2. DIS Job Status Display (B)	4-4
4-3. DIS Data Storage Display (F)	4-6
4-4. DIS Program Storage Display (G)	4-7
4-5. DIS Exchange Package Display (X)	4-8
4-6. DIS Directory Display (Z)	4-9

5-1.	Record of Original Values in W,M Display (Delay Values)	5-12
5-2.	Record of Original Values in L Display SDSPLAY	5-16
5-3.	Record of Original Values in L Display SDSPLAY	5-23
5-4.	PP Breakpoint V Display	5-78
5-5.	CPUMTR Breakpoint Display	5-82
6-1.	C Directive Output	6-54
6-2.	CP Directive Output	6-55
6-3.	D Directive Output	6-56
6-4.	PP Directive Output	6-57
6-5.	Q Directive Output	6-58
6-6.	Sample Alternate Directives File	6-61
6-7.	Sample PP Directive Output	6-63
6-8.	Sample AP Directive Output	6-64
6-9.	Sample CP Directive Output (Continued)	6-65
6-10.	Sample MST Directive Output	6-70
6-11.	Sample C Directive Output	6-71
6-12.	Sample FNT Directive Output	6-72
6-13.	Sample PRORF Directive Output	6-72
8-1.	FLAW K Display (Left Screen)	8-3
8-2.	FLAW K Display (Right Screen)	8-4
8-3.	FLAW Utility K Display (Left Screen)	8-7
8-4.	INITIALIZE Command K Display	8-9
8-5.	Machine Configurations	8-16
8-6.	MREC K Display (Left Screen)	8-18
8-7.	MREC K Display (Right Screen) (Page 1 of 2)	8-21
8-8.	MREC K Display (Right Screen) (Page 2 of 2)	8-21
8-9.	MSE K Display (Message Format)	8-25
8-10.	MSE K Display (Storage Module Format)	8-25
8-11.	NAM K-Display Format	8-26
8-12.	NAM K Display	8-27
8-13.	NAM Status Display	8-36
8-14.	Connection Slot Array and PID/Selection Class Matrix	8-41
8-15.	QTF HELP K Display (Right Screen)	8-49
8-16.	General Format of QTF Status K Display	8-52
8-17.	QTF STATUS,PID K Display	8-54
8-18.	QTF STATUS,SC K Display	8-56
8-19.	QTF STATUS,TRANSFER K Display	8-57
8-20.	REDEFINE K Display (Left Screen)	8-61
8-21.	REDEFINE K Display (Right Screen)	8-61
8-22.	Reconfiguration Run Output	8-65
8-23.	RBF K Display	8-70
8-24.	Sample LCN	8-75
8-25.	Application Table Display	8-77
8-26.	Network Path Table Display	8-78
8-27.	TAF Initialization K Display (Left Screen)	8-86
8-28.	TAF Status K Display (Left Screen)	8-90
8-29.	CRMTASK Commands K Display (Left Screen)	8-96
8-30.	CRM Status K Display (Left Screen)	8-98
8-31.	CRM Data Base Status K Display (Left Screen)	8-98
8-32.	CRM File Status K Display (Left Screen)	8-99
9-1.	FOTD L Display	9-2
9-2.	LIDOU L Display (LID)	9-4
9-3.	LIDOU L Display (HELP) (Page 1 of 2)	9-4

9-4.	LIDOU L Display (HELP) (Page 2 of 2)	9-5
9-5.	Initial QDSPLAY L Display (Page 1 of 2)	9-6
9-6.	Initial QDSPLAY L Display (Page 2 of 2)	9-7
9-7.	SCTD L Display (Page 1 of 2)	9-10
9-8.	SCTD L Display (Page 2 of 2)	9-10
9-9.	SDSPLAY-HELP Display (Page 1 of 3)	9-12
9-10.	SDSPLAY-HELP Display (Page 2 of 3)	9-13
9-11.	SDSPLAY-HELP Display (Page 3 of 3)	9-13
9-12.	SDSPLAY-CLASS Display (Page 1 of 3)	9-14
9-13.	SDSPLAY-CLASS Display (Page 2 of 3)	9-15
9-14.	SDSPLAY-CLASS Display (Page 3 of 3)	9-15
10-1.	Equipment Configuration	10-4
10-2.	LIDCMid Files for Example Equipment Configuration	10-5
10-3.	Sample LCN Network	10-18
10-4.	Network Configuration Statements for the RHF on Mainframe MFA	10-19
10-5.	Network Configuration Statements for the RHF on Mainframe MFD	10-20
10-6.	Procedure File Example	10-23
11-1.	SSLABEL Report File	11-23
11-2.	SSMOVE Report: List of Directives	11-37
11-3.	SSMOVE Report: Run-time Parameter Values	11-37
11-4.	SSMOVE Report: Optional List of Files	11-38
11-5.	SSMOVE Report: Device Status Report	11-39
11-6.	SSMOVE Report: Destage Abandonment	11-39
11-7.	SSVAL Validation Report	11-48
11-8.	SSUSE Basic Report <i>(Continued)</i>	11-53
11-9.	SSUSE Optional Report A	11-55
11-10.	SSUSE Optional Report B	11-56
11-11.	SSUSE Optional Report C	11-57
11-12.	SSUSE Optional Report D	11-58
11-13.	SSDEBUG: Executing SSVAL in Problem-Fixing Mode	11-64
11-14.	Full Stripe Disk File Format	11-66
11-15.	Partial Stripe Disk File Format	11-67
11-16.	SM K Display	11-75
11-17.	CU K Display	11-76
13-1.	Deadstarting a Machine in a Multimainframe Complex	13-5
15-1.	Sample 255x NPU Network	15-3
16-1.	Sample NDA Output File <i>(Continued)</i>	16-17
16-2.	Sample Session File With Tasks for Use by STIMULA	16-35
16-3.	Sample Session for Use by STIMULA <i>(Continued)</i>	16-36
16-4.	Sample Session File for Internal Stimulation	16-38
16-5.	STIMULA K Display (Left Screen)	16-38
16-6.	STIMULA Parameters (Left Screen)	16-40
16-7.	STIMULA K-Display Commands (Right Screen)	16-43
16-8.	Parameter Matrix (Left Screen)	16-45
16-9.	Mixed Mode Parameters (Right Screen)	16-46
16-10.	Output Statistics	16-47
16-11.	Postprocessing Example	16-48
17-1.	Functional Overview of the Permanent File Utility Overlays	17-1
17-2.	Example of Multifile Archive File Structure	17-3
17-3.	Track Reservation Table Word	17-23
17-4.	Bit Settings for Track Link Bytes	17-24
17-5.	Linkage Format With Example	17-25
17-6.	Example of Families in Two Configurations	17-27

17-7. Example of Transferring Family Access	17-29
17-8. Initial Permanent File Utilities K Display (Left Screen)	17-38
17-9. Permanent File Utilities K Display (Right Screen)	17-38
17-10. PFDUMP K Display (Left Screen) (Page 1 of 5)	17-40
17-11. PFDUMP K Display (Left Screen) (Page 2 of 5)	17-40
17-12. PFDUMP K Display (Left Screen) (Page 3 of 5)	17-41
17-13. PFDUMP K Display (Left Screen) (Page 4 of 5)	17-41
17-14. PFDUMP K Display (Left Screen) (Page 5 of 5)	17-42
17-15. Sample Directory Produced by PFATC,LO=T.	17-66
17-16. Sample Directory Produced by PFCAT,LO=T,DN=40,UI=4054	17-68
17-17. Mass Storage Table Report	17-68
17-18. Cumulative Statistics by User Index PFCAT,LO=S,DN=40.	17-69
17-19. Cumulative Statistics for Entire Device PFCAT,LO=S,DN=40. <i>(Continued)</i>	17-70
17-20. Catalog of Catalog Image Record	17-77
17-21. PFDUMP/PFLOAD Example (Sheet 1 of 3)	17-83
17-22. PFDUMP/PFLOAD Example (Sheet 2 of 3)	17-84
17-23. PFDUMP/PFLOAD Example (Sheet 3 of 3)	17-85
18-1. QFSP K Display (Left Screen)	18-3
18-2. QFSP K Display (Right Screen) (Page 1 of 4)	18-5
18-3. QFSP K Display (Right Screen) (Page 2 of 4)	18-5
18-4. QFSP K Display (Right Screen) (Page 3 of 4)	18-6
18-5. QFSP K Display (Right Screen) (Page 4 of 4)	18-6
18-6. DFTERM K Display (Left Screen)	18-28
18-7. DFTERM Dayfile List (Right Screen)	18-29
18-8. QALTER/QFTLIST K Display (Left Screen) (Page 1 of 3)	18-32
18-9. QALTER/QFTLIST K Display (Left Screen) (Page 2 of 3)	18-32
18-10. QALTER/QFTLIST K Display (Left Screen) (Page 3 of 3)	18-33
18-11. QALTER/QFTLIST LIST Command Display (Right Screen)	18-34
18-12. QALTER/QFTLIST LIST=qft Command Display (Right Screen)	18-34
18-13. QALTER Printer Listing/Output File	18-35
18-14. Initial QDUMP K Display (Left Screen) (Page 1 of 3)	18-37
18-15. Initial QDUMP K Display (Left Screen) (Page 2 of 3)	18-37
18-16. Initial QDUMP K Display (Left Screen) (Page 3 of 3)	18-38
18-17. Secondary QDUMP K Display (Left Screen)	18-38
18-18. QDUMP Output File	18-39
18-19. QLOAD K Display (Left Screen) (Page 1 of 3)	18-42
18-20. QLOAD K Display (Left Screen) (Page 2 of 3)	18-42
18-21. QLOAD K Display (Left Screen) (Page 3 of 3)	18-43
18-22. QLOAD Output File	18-43
18-23. Initial QMOVE K Display (Left Screen) (Page 1 of 3)	18-46
18-24. Initial QMOVE K Display (Left Screen) (Page 2 of 3)	18-46
18-25. Initial QMOVE K Display (Left Screen) (Page 3 of 3)	18-47
18-26. Secondary QMOVE K Display (Left Screen)	18-47
18-27. QMOVE Output File	18-48
18-28. Initial QREC K Display (Left Screen) (Page 1 of 3)	18-51
18-29. Initial QREC K Display (Left Screen) (Page 2 of 3)	18-52
18-30. Initial QREC K Display (Left Screen) (Page 3 of 3)	18-52
18-31. Inactive Queues K Display (Right Screen)	18-53
18-32. Secondary QREC K Display (Left Screen)	18-53
18-33. QREC Output File	18-54
21-1. Example of TRACER Output <i>(Continued)</i>	21-8
21-2. Example of PROBE Output <i>(Continued)</i>	21-38
22-1. CTS Minimum Configuration	22-3

22-2.	Dual-State Fixed Configuration for CTS	22-4
22-3.	Dual-State Flexible Configuration for CTS	22-5
22-4.	ACS Minimum Configuration	22-9
22-5.	ACS Configuration with Two LSMs	22-12
22-6.	CTS Cross-Coupled Controllers	22-14
22-7.	ACS Cross-Coupled Controllers	22-15
A-1.	Conversion Differences	A-17
D-1.	Simulator K Display (Left Screen)	D-3
D-2.	Simulator Commands K Display (Right Screen)	D-4
F-1.	EVFU Load File	F-4
G-1.	FORMAT Output, MODE=RESTORE	G-7
G-2.	FORMAT Output, MODE=ALTER	G-9
G-3.	FORMAT Output, MODE=FETCH	G-10

Tables

2-1.	EQPDECK Entries	2-4
2-2.	Levels of System Deadstart: Information Recovered	2-28
3-1.	Table of Central Memory Sizes	3-12
3-2.	Required and Optional Keywords	3-19
3-3.	Mass Storage Functions for Various Device Types	3-42
3-4.	Mass Storage Configuration for Two Systems at One Site	3-43
3-5.	PF Entry Parameter Settings (type, dm, name, device, and sm)	3-81
3-6.	Information for Setting or Clearing Flaw Areas	3-106
3-7.	Example Set of Job Control Parameters, Part I	3-116
3-8.	Example Set of Job Control Parameters, Part II	3-117
3-9.	Example Set of Job Control Parameters, Part III	3-118
3-10.	CPM Default Values for n_i	3-121
4-1.	PP Call Formats	4-20
8-1.	Flawing Information for Disks and Extended Memory	8-7
8-2.	Device Initialization Options	8-10
8-3.	Track Flawing Options	8-11
8-4.	MREC Parameters	8-20
8-5.	MREC Commands	8-20
8-6.	NAM K-Display Fields	8-26
8-7.	NAM Mode Commands	8-29
8-8.	Equipment Requirements for Reconfiguration	8-59
8-9.	Reconfiguration Parameters	8-63
8-10.	Reconfiguration Commands	8-64
10-1.	Suggested Code Conversion Parameters	10-15
15-1.	Parameter Record Feature Descriptions	15-7
15-2.	NS Activity	15-14
15-3.	Load/Dump Status Messages	15-14
17-1.	Permanent File Utility Parameters	17-45
18-1.	Queued File Utility Parameters (DFTERM through QFTLIST)	18-12
18-2.	Queued File Utility Parameters (QLIST through QREC)	18-13
20-1.	Initialize and Release Parameters of Fast-Attach Files	20-3
21-1.	Fast Loop Data Items	21-18
21-2.	Medium Loop Data Items	21-19
21-3.	Slow Loop Data Items	21-20
22-1.	CTS Hardware	22-2

22-2.	ACS Hardware	22-7
A-1.	Character Sets for Interactive Jobs	A-5
A-2.	Character Sets for Batch Jobs	A-9
A-3.	ASCII to 6/12-Bit Display Code Conversion	A-13
A-4.	Nine-Track ASCII Coded Tape Conversion	A-18
A-5.	Nine-Track EBCDIC Coded Tape Conversion	A-20
A-6.	Seven-Track Coded Tape Conversions	A-23
E-1.	Released Programmable Format Control Arrays	E-4
F-1.	Released EVFU Load Image: Print Lines and Channel Numbers	F-5

About This Manual

This manual describes the CONTROL DATA® Network Operating System (NOS) Version 2. NOS 2 operates on the following computer systems:

- CONTROL DATA CYBER 180 Computer Systems Models 810, 830, 835, 840, 845, 850, 855, 860, 870, 960, 970, 990, 994, and 995
- CONTROL DATA CYBER 170 Computer Systems Models 171, 172, 173, 174, 175, 176, 720, 730, 740, 750, 760, 815, 825, 835, 845, 855, 865, and 875
- CONTROL DATA CYBER 70 Computer Systems Models 71, 72, 73, and 74
- CONTROL DATA 6000 Computer Systems

Audience

This manual assumes you are a site analyst. It assumes you are familiar with the hardware of your computer system(s) and that you understand the functions of the various components of NOS.

Organization

This manual includes information required for the day-to-day maintenance of the operating system and for troubleshooting. Topics discussed include the mass storage subsystems, network operations, the K and L utilities, backing up and reloading files, deadstart, and DIS operations.

The appendixes include character set tables; a glossary; and descriptions of the SCOPE 2 Station Facility, the status/control register simulator, programmable format control for 580 printers, disk pack reformatting for 881/883 units, address formats for NOS/VE, management of storage media defects, and the display disk file utility.

Conventions

The following conventions are used in this manual:

examples	Examples of user entries and computer responses are shown in a font that resembles computer output.
lowercase	In a format, lowercase letters represent values you choose.
Numbers	All numbers are decimal unless otherwise noted.
UPPERCASE	In a format, uppercase letters represent reserved words defined by the system for specific purposes. You must use these words exactly as shown.
Vertical bar	A vertical bar in the margin indicates a technical change.

The CONTROL DATA 18002-2 console is available as an option for CYBER 180 Models 810 and 830 Systems using NOS 2.3, PSR level 617 or later operating systems. This product includes a CONTROL DATA 634B display terminal (also known as the 721-21 display terminal) and an AV117A cable. This console is referred to throughout the manual as the CC634B.

The CONTROL DATA 19003 console is available as an option for certain CYBER 180-class machines. This product includes a video monitor; keyboard; 40-Mbyte hard disk (Winchester) drive; 1.2-Mbyte, 5-1/2-in floppy disk drive; 640-Kbyte RAM memory; one parallel printer port; and nine RS-232-C serial ports. This console is referred to throughout this manual as the CC598B console.

Models 815, 825, 835, 845, and 855 of the CYBER 170 Computer Systems share many of the functional and architectural attributes of the CYBER 180 Computer Systems. This manual uses the term CYBER 180-class machines when describing these similar models collectively.

Extended memory for models 865 and 875 and CYBER 180-class machines is unified extended memory (UEM) and may also include either extended core storage (ECS), extended semiconductor memory (ESM), or STORNET. Extended memory for model 176 is large central memory extended (LCME) and may also include ECS, ESM, or STORNET. Extended memory for all other NOS computer systems is either ECS, ESM, or STORNET.

In this manual, ECS refers to ECS, ESM, and STORNET; and extended memory refers to all forms of extended memory unless otherwise noted. However, when referencing extended memory in the context of a linked shared device multmainframe complex or distributive data path (DDP) access, UEM and LCME are excluded. ECS, ESM, and STORNET are the only forms of extended memory that can be shared in a linked shared device multmainframe complex and can be accessed by a DDP.

(Manuals dealing with the various form of extended memory are listed under Related Publications.)

In this manual, the terms tapes and magnetic tapes refer to reel magnetic tapes and cartridge tapes for ACS and CTS. The term MSE refers to cartridges for the 7990. The term cartridge alternate storage refers to alternate storage on MSE or optical disk. The term tape alternate storage refers to alternate storage on tapes.

Related Publications

All of the following manuals are available through Control Data sales offices or through:

Control Data
Literature and Distribution Services ARHLDS
4201 Lexington Avenue N.
St. Paul, MN 55126-6198

You can also call (612) 482-3800 or 482-3801, or FAX your inquiry to (612) 482-3813. (If you are a Control Data employee, use the Controlnet number 235-3800, 235-3801, or 235-3813.)

The reader should be thoroughly familiar with the material in the following NOS publications.

Manual Title	Publication Number
NOS Version 2 Operations Handbook	60459310
NOS Version 2 Reference Set, Volume 2, Guide to System Usage	60459670
NOS Version 2 Reference Set, Volume 3, System Commands	60459680

The following lists contain manuals that provide additional information about NOS and its product set. For the reader's convenience, these are grouped according to topic: ACS manuals, CDCNET manuals, extended memory manuals, hardware manuals, NOS 2 manuals, and optional product manuals. In addition, the NOS System Information manual contains brief descriptions of all NOS operating system and NOS product set manuals. It is accessed by logging into NOS and entering the EXPLAIN command.

ACS Manuals

The following list contains manuals that describe the CONTROL DATA 5744 Automated Cartridge Subsystem (ACS) software; an ACS hardware manual is listed under the Hardware Manuals subsection.

Manual Title	Publication Number
5744 Automated Cartridge Subsystem, Automated Cartridge Exercise System, User's Guide	60000517
5744 Automated Cartridge Subsystem, UNIX-Based Cartridge Library Server Workstation, Installation and Service	60000516
5744 Automated Cartridge Subsystem, UNIX-Based Cartridge Library Server Workstation, Programmer's Guide	60000518
5744 Automated Cartridge Subsystem, UNIX-Based Cartridge Library Server Workstation, System Administrator's Guide	60000519
5744 Automated Cartridge Subsystem, TRANSF Client System Interface, Programmer's Guide	60000521

CDCNET Manuals

The following list contains manuals that describe the Control Data Distributed Communications Network (CDCNET).

Manual Title	Publication Number
CDCNET Batch Device User Guide	60463863
CDCNET Conceptual Overview	60461540
CDCNET Configuration Guide	60461550
CDCNET Hardware Installation and Troubleshooting	60000348
CDCNET Network Operations and Analysis	60461520
CDCNET Terminal Interface Usage	60463850

Extended Memory Manuals

Programming information for the various forms of extended memory can be found in the COMPASS Version 3 Reference Manual (publication number 60492600) and in the appropriate computer system hardware reference manual. Hardware descriptions of the various forms of extended memory can be found in the following manuals.

Manual Title	Publication Number
CYBER 5380-100 STORNET Subsystem (SNSS) Hardware Reference	60000188
Extended Core Storage Reference Manual	60347100
Extended Core Storage II and Distributive Data Path Reference Manual	60430000
Extended Semiconductor Memory Hardware Reference Manual	60455990

Hardware Manuals

The following list contains manuals that describe Control Data computer systems and related equipment.

Manual Title	Publication Number
CYBER 70 Model 71 Computer System Hardware Reference Manual	60453300
CYBER 70 Model 72 Computer System Hardware Reference Manual	60347000
CYBER 170 Computer Systems Models 171 through 175 (Levels A, B, C) Model 176 (Level A, B, C) Hardware Reference Manual	60420000
CYBER 170 Computer Systems Models 720, 730, 740, 750, and 760 Model 176 (Level B/C) Hardware Reference Manual	60456100
CYBER 170 Computer Systems Models 815 and 825 Hardware Reference Manual	60469350
CYBER 170 Computer Systems Models 835, 845, and 855 CYBER 180 Computer Systems Models 835, 840, 845, 850, 855, 860, 970, and 990 CYBER 990E, 994, and 995E Computer Systems CYBER 170 State Hardware Reference Manual	60469290
CYBER 170 Computer Systems Models 835, 845, and 855 CYBER 180 Computer Systems Models 835, 845, and 855 Hardware Operator's Guide	60458390
CYBER 170 Computer Systems Models 865 and 875 Hardware Reference Manual	60458920
CYBER 180 Models 810 and 830 Computer Systems Hardware Operator's Guide	60469440
CYBER 180 Models 810 and 830 Computer Systems Hardware Reference Manual	60469420
CYBER 840A, 850A, 860A, and 870A Computer Systems Hardware Reference Manual	60463560
5830 Disk Array Subsystem (DAS) Configuration Guide	60000494
5830 Disk Array Subsystem (DAS) Reference	60000552
5831 Disk Array Subsystem (DAS) Controller Reference	60000551
5870 Printer User's Reference Manual	60462720
CYBER 960 Computer Systems CYBER 170 State Hardware Reference Manual	60000127
19003 System Console CC598-A/B Operations and Maintenance Guide	60463610
380-170 Network Access Device Hardware Reference Manual	60458500
5744 Automated Cartridge Subsystem (ACS) User's Guide	60000459

NOS 2 Manuals

The following list contains NOS 2 manuals.

Manual Title	Publication Number
COMPASS Version 3 Reference Manual	60492600
CYBER Initialization Package (CIP) Reference Manual	60457180
CYBER Loader Version 1 Reference Manual	60429800
CYBER Record Manager Advanced Access Methods Version 2 Reference Manual	60499300
CYBER Record Manager Basic Access Methods Version 1.5 Reference Manual	60495700
Modify Version 1 Reference Manual	60450100
Modify Instant	60450200
NOS Online Maintenance Software Reference Manual	60454200
NOS Version 2 Administration Handbook	60459840
NOS Version 2 Applications Programmer's Instant	60459360
NOS Version 2 Diagnostic Index	60459390
NOS Version 2 Installation Handbook	60459320
NOS Version 2 Reference Set, Volume 1 Introduction to Interactive Usage	60459660
NOS Version 2 Reference Set, Volume 4 Program Interface	60459690
NOS Version 2 Security Administrator's Handbook	60460410
NOS Version 2 System Overview	60459270
NOS Version 2 Systems Programmer's Instant	60459370

Optional Product Manuals

The following list contains manuals that describe optional products.

Manual Title	Publication Number
Binary Maintenance Log (BML) Message Formats	60459940
COBOL Version 5 Reference Manual	60497100
Communications Control Program Version 3 Diagnostic Handbook	60471500
CYBER Cross System Version 1 Build Utilities Reference Manual	60471200
CYBER Supermini Operations User's Guide	60459850
FORTRAN Extended Version 4 Reference Manual	60497800
FORTRAN Version 5 Reference Manual	60481300
Message Control System Version 1 Reference Manual	60480300
MSSI Version 3 Reference Manual	60458820
Network Access Method Version 1/Communications Control Program Version 3 Terminal Interfaces Reference Manual	60480600
Network Access Method Version 1 Host Application Programming Reference Manual	60499500
Network Access Method Version 1 Network Definition Language Reference Manual	60480000
NOS Version 2 Full Screen Editor User's Guide	60460420
NOS Optical Storage Subsystem QSE32835 Reference Manual	60000351
NOS Version 2 Screen Formatting Reference Manual	60460430
NOS Version 2 Tape Management System (TMS) Site Operations Manual	60463350
NOS Version 2 Tape Management System (TMS) User Reference Manual	60463110
NOS/VE System Analyst Reference Set System Performance and Maintenance Manual	60463915
Remote Batch Facility Version 1 Reference Manual	60499600
Remote Host Facility Access Method Reference Manual	60459990
Remote Host Facility Usage	60460620
SCOPE Version 2.1 Operator's Guide	60455090
TAF Version 1 Reference Manual	60459500
TAF Version 1 User's Guide	60459520
TAF/CRM Data Manager Version 1 Reference Manual	60459510
Update Version 1 Reference Manual	60449900

Submitting Comments

If you have comments concerning this manual, mail your comments to:

Control Data
Information Services ARH219
4201 Lexington Avenue N.
St. Paul, MN 55126-6198

If you have access to SOLVER, the Control Data online facility for reporting problems, you can use it to submit comments about this manual. Use NS2 as the product identifier, and include the name and publication number of the manual.

If you have questions about the packaging and/or distribution of printed manuals, call or write Literature and Distribution Services as described in Related Publications earlier in this preface.

Central Software Support Hotline

Control Data's Central Software Support maintains a hotline to assist you if you have trouble using our products. If you need help beyond that provided in the documentation, or find the product does not perform as described, call one of the following numbers and a support analyst will work with you.

From the USA and Canada: 1-800-345-6628

From other countries: 1-612-482-3434

Controlware Utilities

1

Loading Controlware	1-1
Control Module Controlware	1-1
Disk Controlware	1-2
Network Access Device (NAD) Controlware	1-2
Dumping Controlware	1-4

This section describes the utilities used for loading and dumping controlware.

Loading Controlware

By using the LOADBC utility, you can download control module controlware, disk controlware, or network access device (NAD) controlware to the associated controller.

Control Module Controlware

You can use the LOADBC utility to load controlware into a control module for the 834 or 836 Disk Storage Subsystem. The calling job must be of system origin or you must be validated for system origin privileges, and the system must be in engineering mode (refer to the DSD ENABLE command in section 5). LOADBC will issue appropriate messages to indicate the success or failure of the attempt to load controlware.

The format of the command is:

```
LOADBC, EQ=est, F=loadfile.
```

Parameter	Description
EQ=est	est is the EST ordinal of the control module in which to load the controlware.
F=loadfile	Name of the local file from which control module controlware is to be loaded. If F=loadfile is specified, local file loadfile must contain the control module controlware in binary format and an appropriate header (refer to the NOS Version 2 Installation Handbook). If F=loadfile is omitted, controlware is read from the system library SYSTEM.

Disk Controlware

You can initiate downloading of disk controlware only from the system console. Also, you can load the disk controlware to a channel only if it is either active or down and unassigned. The calling job must be of system origin or you must be validated for system origin privileges. The system must be in engineering mode (refer to the DSD ENABLE command in section 5). LOADBC will issue appropriate messages to indicate the success or failure of the disk controlware load attempt.

The format of the command is:

```
LOADBC, C=ch, F=loadfile, D=dumpfile.
```

Parameter	Description
C=ch	ch is a 2-digit octal number of the channel to which the disk controlware is to be loaded. The controlware can be loaded only if the channel status is UP or if the channel status is DOWN and not assigned to a maintenance user.
F=loadfile	Name of the local file from which disk controlware is to be loaded. If F=loadfile is specified, local file loadfile must contain the disk controlware in binary format and an appropriate header (refer to the NOS Version 2 Installation Handbook). If F=loadfile is omitted, controlware is read from the system library SYSTEM.
D=dumpfile	Name of the local file to which 7155/7165/7255 disk controlware is to be dumped before reloading. This parameter is ignored for other types of controllers and an informative message is issued. LOADBC performs a binary comparison between the old and new controlware and writes this data to a file that can be processed by DSDI by using the DMB parameter (refer to section 6 for information on DSDI).

Network Access Device (NAD) Controlware

The LOADBC utility can be used to load NAD controlware into local NADs (380-170) and remote NADs (380-170, 380-200, 380-370, and 380-110). Since the NAD controlware is not automatically loaded at deadstart, LOADBC must be used before a local NAD can be used by the operating system. NAD controlware may be automatically loaded by the Remote Host Facility (RHF) when RHF is initiated. Refer to the RHF K display in section 8.

LOADBC can be called from the console or a batch job. When loading 380-170 controlware into a local NAD, the EST entry associated with the NAD's channel number must be OFF or the controlware-not-loaded flag must be set.

When loading a remote NAD, a local NAD that is not reserved for maintenance must be defined in the EST. The EST entry must be ON. Controlware must be loaded and running in the local NAD before loading the remote NAD.

Remote NAD loading operations can occur concurrently with RHF use of the local NAD. However, extreme care should be exercised when performing a remote NAD load to ensure that the correct remote NAD is being loaded and that the remote NAD is not being used by the mainframe to which it is connected. LOADBC will issue appropriate messages to indicate the success or failure of the NAD controlware load attempt.

The format of the command is:

LOADBC, P₁, P₂, . . . , P_n.

P _i	Description
C=ch	The 2-digit octal number of the channel to which the NAD controlware is to be loaded. This parameter is required.
F=filename	Name of the local file from which NAD controlware is to be loaded. If F=filename is specified, local file filename must contain the NAD controlware in binary format and an appropriate header (refer to the NOS Version 2 Installation Handbook). If F=filename is not specified, the NAD controlware type specified by the TY parameter is read from the system library SYSTEM.

The following parameters apply only when loading remote NAD controlware.

Parameter	Description								
AC=aaaa	The 4-digit hexadecimal number specifying the remote NAD's access code (refer to the RHF K display in section 8). The default is AC=0000.								
LT=t ₀ t ₁ t ₂ t ₃	The 4-digit binary bit pattern specifying the local trunk control units (TCUs) that are enabled. At least one TCU enable must be specified for remote NAD controlware loading. t _n =1 enables TCU _n . For example, LT=1010 indicates that the local TCUs 0 and 2 are enabled.								
ND=nn	The 2-digit hexadecimal number specifying the remote NAD's address (refer to the RHF K display in section 8). This parameter is required for remote NAD loads.								
TY=value	Type of controlware to be loaded.								
	<table border="1"> <thead> <tr> <th>value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>170</td> <td>CYBER 170 controlware (380-170)</td> </tr> <tr> <td>IBM</td> <td>IBM controlware (380-370)</td> </tr> <tr> <td>MIN</td> <td>Minicomputer controlware (380-110)</td> </tr> </tbody> </table>	value	Description	170	CYBER 170 controlware (380-170)	IBM	IBM controlware (380-370)	MIN	Minicomputer controlware (380-110)
value	Description								
170	CYBER 170 controlware (380-170)								
IBM	IBM controlware (380-370)								
MIN	Minicomputer controlware (380-110)								

The default value is TY=170.

Under certain conditions, a remote NAD loading operation will fail on the first attempt but a second loading attempt will succeed. This loading problem can be prevented by always preceding a remote NAD loading operation with a remote NAD dumping operation to ensure the remote NAD controlware is halted before loading is attempted. For example, to load a remote NAD with NAD address 7F and access code F0F0 connected to TCU 0 of the local NAD on channel 5, enter the following commands.

```
X. DMPNAD (CH=05, ND=7F, AC=F0F0, LT=1000)
X. LOADBC (C=05, ND=7F, AC=F0F0, LT=1000)
```

Dumping Controlware

The DMPCCC utility provides the capability to dynamically dump the CYBER channel coupler (CCC) memory in an online environment. The calling job must be of system origin or you must be validated for system origin privileges, and the system must be in engineering mode (refer to the DSD ENABLE command in section 5).

The format of the command is:

```
DMPCCC, C=ch, L=outfile.
```

Parameter	Description
C=ch	ch is a 1- or 2-digit octal number of the channel from which the CCC memory is to be dumped. The specified channel number must be in the range from 0 to 13 ₈ or from 20 ₈ to 33 ₈ . The default is C=0.
L=outfile	The 1- to 7-character name of the file to which the dump is to be written. The default is L=OUTPUT.

DMPCCC lists the CCC memory in line format. Each line consists of 16 CCC memory words in hexadecimal format. Repetitive lines are suppressed.

- Modifying the CMRDECK 2-2
- Modifying the EQPDECK 2-3
 - Deadstart Displays 2-8
 - Equipment Status Display 2-8
 - Mass Storage Status Display 2-12
 - Mass Storage Initialization Status Display 2-16
 - Controlware Status Display 2-18
 - Disk Thresholds Display 2-19
- Modifying the APRDECKs 2-21
- Modifying the IPRDECK 2-22
- Loading the System 2-23
 - Messages 2-23
 - Entering the Date and Time 2-24
- Initiating Job Processing 2-25
- Preparing for Recovery Deadstart 2-27
 - Level 3 Recovery Deadstart 2-30
 - Level 1 Recovery Deadstart 2-33
 - Level 2 Recovery Deadstart 2-34
- Level 0 Initial Deadstart 2-35
- Error Processing 2-35
 - Recovery Functions 2-36
 - Error Idle Recovery 2-43
 - Removable Device 2-43
 - Nonremovable Device 2-43

Deadstart is the process that makes the system operational and ready to process jobs. After performing the appropriate deadstart procedures described in the CIP User's Handbook, you can continue the deadstart process as shown in figure 2-1 and described in detail following the figure.

This section also describes what you can do to recover if you experience system problems during the deadstart process.

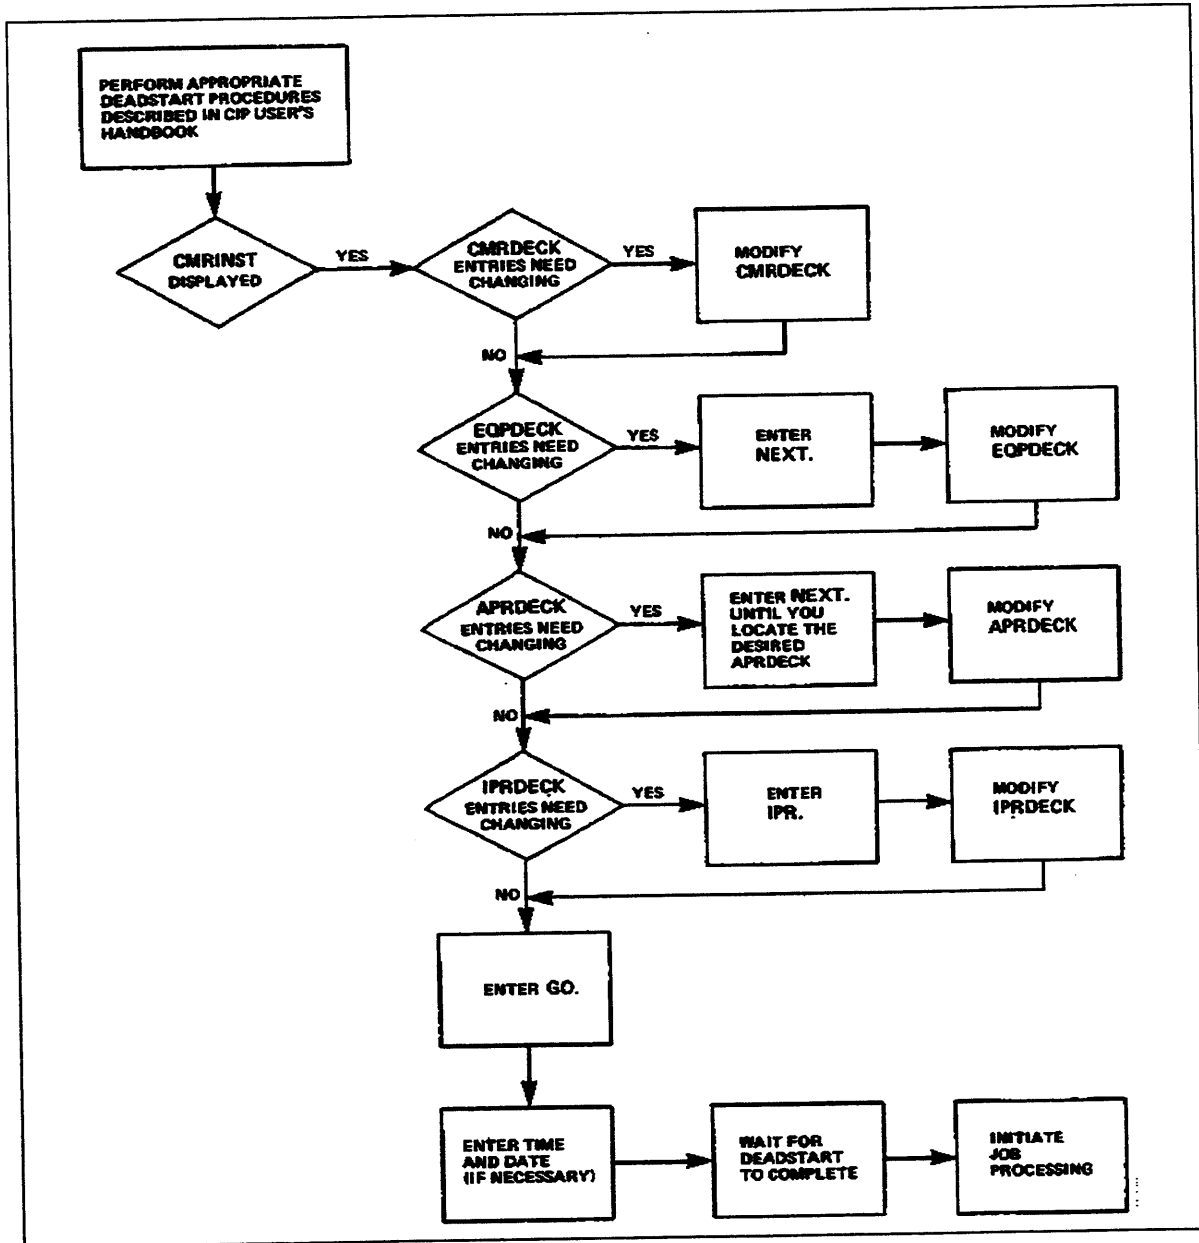


Figure 2-1. Deadstart Process

NOTE

Modifying the deadstart decks (CMRDECK, EQPDECK, APRDECKs, and IPRDECK) does not apply to a level 3 recovery deadstart. Modifications made during the last level 0, 1, or 2 deadstart are recovered during a level 3 recovery deadstart.

Modifying the CMRDECK

If bit 6 of word 13 (word 12 for warmstart on a CYBER 70 or 6000 Computer System with an active PP) is set (ppp=001), or if you select the D=Y option on the *P* display, an instruction display entitled CMRINST appears on the console screen(s) after the CTI displays on a level 0, 1, or 2 deadstart. All valid CMRDECK entries are defined in this display. Several of the entries listed are assigned system default values. These values are assumed if the entries do not appear in the CMRDECK being used. To view the contents of the CMRDECK being used, toggle from the CMRINST display to the CMRDECK display. If either the CMRDECK or CMRINST overflows two screens, the display can be paged.

Modify the CMRDECK by entering the appropriate changes or additions from the console keyboard. These entries can be made while either CMRDECK or CMRINST is being displayed. Generally, each console entry supersedes the value currently specified in the CMRDECK (or default value in CMRINST).

Refer to section 3, Deadstart Decks, for complete information on all CMRDECK entries.

NOTE

The modified CMRDECK remains in effect only until the next level 0 deadstart is performed. Changes to the CMRDECK are not recovered for the next deadstart unless a new deadstart file is created. If you want these changes to take place on the next level 0 deadstart, make the appropriate changes to the CMRDECK after NOS is up and running and use LIBEDIT to replace the record on the deadstart file.

After all CMRDECK modifications have been made and you want to modify an EQPDECK, APRDECK, or IPRDECK, enter:

NEXT.

Refer to Modifying the EQPDECK, Modifying the APRDECKs, or Modifying the IPRDECK in this section. Otherwise, to indicate that all modifications to the CMRDECK are complete and you want to begin loading the system, enter:

GO.

Modifying the EQPDECK

After completing all CMRDECK modifications, you can also modify the default EQPDECK, an APRDECK, or the IPRDECK being used. If no changes need to be made to any EQPDECK, but you do need to modify an APRDECK or the IPRDECK, refer to *Modifying the APRDECKs* or *Modifying the IPRDECK* later in this section.

To modify an EQPDECK, enter

NEXT.

while the CMRDECK or CMRINST is being displayed.

You can make changes when the EQPDECK, EQPINST, or any one of the deadstart displays is displayed at the console screen (refer to *Deadstart Displays* described later in this section).

Table 2-1 describes the entries that can be made only at the console keyboard at deadstart time and cannot be stored in the EQPDECK on the deadstart file. Refer to section 3, *Deadstart Decks*, for complete information concerning all EQPDECK entries.

After making the changes to EQPDECK you can toggle through the deadstart displays to ensure all the changes are made.

NOTE

The modified EQPDECK remains in effect only until the next level 0 deadstart is performed. Changes to the EQPDECK are not recovered for the next deadstart unless a new deadstart file is created. If you want these changes to take place on the next level 0 deadstart, make the appropriate changes to the EQPDECK after NOS is up and running and use LIBEDIT to replace the record on the deadstart file.

After all EQPDECK modifications have been made and you want to modify an APRDECK or IPRDECK, enter:

NEXT.

Refer to *Modifying the APRDECKs* or *Modifying the IPRDECK* later in this section. Otherwise, to indicate that all modifications to the EQPDECK are complete and you want to begin loading the system, enter:

GO.

Table 2-1. EQPDECK Entries

Entry	Function
AUTOLOAD.	Toggles the selection of buffer controller autoloading for all 7054/7154/7152/7155/7165/7255 control module controllers. This entry is valid only when entered from the console keyboard. That is, the AUTOLOAD entry cannot be included as part of the EQPDECK on the deadstart file. This entry, which affects all 7054/7154/7152/7155/7165/7255 control module controllers defined on the mainframe being deadstarted, is not necessary for normal system operation but is provided as an aid to hardware checkout.
GRENADE.	Toggles the selection of the grenade function. This entry is valid only when entered from the console keyboard. That is, the GRENADE entry cannot be included as part of the EQPDECK on the deadstart file. The grenade function is issued to all 7054/7154/7152/7155/7255 control module controllers, once the controlware is loaded. This function causes unit reservations to be cleared on all 834, 836, and 844 units physically connected to each controller. This entry is normally used when a unit reservation from a downed mainframe exists on a device. Use this entry with caution since it can interrupt the operation of another machine that could be accessing affected units through another controller.
INITIALIZE,op,est ₁ ,est ₂ ,...,est _n .	Creates new labels for the specified mass storage devices during a level 0 deadstart. This entry is valid only when entered from the console keyboard. That is, the INITIALIZE entry cannot be included as part of the EQPDECK on the deadstart file. Before any mass storage device defined in the EQPDECK (by an EQ entry) can be used, it must have a label that can be recognized by the system. Existing labels are normally recovered automatically during all levels of system deadstart. However, if the existing label is destroyed (for example, during maintenance operations on the device) or if a new mass storage device is added to the system, you enter the INITIALIZE command to create a new label.

(Continued on next page)

Table 2-1. EQPDECK Entries

Entry	Function		
<i>(Continued from previous page)</i>			
	<table border="1"> <thead> <tr> <th data-bbox="662 331 818 363">Parameter</th> <th data-bbox="837 331 1000 363">Description</th> </tr> </thead> </table>	Parameter	Description
Parameter	Description		
op	Level of initialization: AL Total initialization. PF Permanent files. QF Queued files. DF System dayfile. AF Account dayfile. EF Error log. FP Format pack (844 or 895). MF Binary maintenance log.		
est ₁	The 1- to 3-digit number specified in the EQ entry for the device (for example, EQ005...). This is also the EST ordinal for the device. Any number of devices can be initialized with the same level of initialization, provided the command is no more than 72 characters.		

Depending upon the levels of initialization selected, all or part of the previously existing information on the device is lost when the new label is created. Total initialization of 844 or 895 format pack (AL or FP options) destroys all information on a device. The other options selectively purge information. A separate INITIALIZE entry is required for each option selected. Selecting FP results in an automatic system selection of AL. The system deletes all existing files, including a system deadstart file, from a device initialized with the AL option. CTI, MSL, and HIVS information is not deleted when a device is initialized with the AL option. You cannot initialize the device from which you are deadstarting. You should initialize a device if you just loaded CTI, HIVS, or MSL on it.

(Continued on next page)

Table 2-1. EQPDECK Entries

Entry	Function
<i>(Continued from previous page)</i>	
	<p>No options (except AL and FP) are processed until deadstart is completed. At that time, the K display is requested and you must enter the family name (FM) and device number (DN) of the device to be initialized. This is a final check to ensure that the correct device is being initialized; the selected options are then processed.</p>
	<p>If permanent files are to reside on the device being initialized, the EQPDECK should contain a PF entry for that device. The PF entry corresponds to the EST ordinal specified in the EQ entry and indicates that permanent files can reside on the device. If the EQPDECK displayed contains a PF entry for the device being initialized, a new PF entry is not required unless the existing entry is to be altered, or the associated EQ entry is altered. Redefining the existing EQ entry clears all associated attributes. In addition, if PF entries do not exist in the EQPDECK, initializing the device causes a default family name and device number to be assigned. Thus, it is necessary to reestablish the PF entry via the console keyboard if the device is to remain a permanent file device. For this reason, it is recommended that the PF entry for all mass storage devices used for permanent files reside in the EQPDECK. Although this is recommended, it is not required.</p>
	<p>If the EQ entry in the EQPDECK displayed indicates that the status of a particular mass storage device is off when the INITIALIZE entry is made, initialize status is maintained and occurs automatically when the DSD command ON is entered for that device during normal system operation.</p>
	<p>Note that initialization of mass storage devices can also be accomplished during normal system operation via the DSD command INITIALIZE.</p>

(Continued on next page)

Table 2-1. EQPDECK Entries

Entry	Function				
<i>(Continued from previous page)</i>					
PRESET.	<p>Initializes tables (MST, TRT, MRT, DAT, BAT, BRT, and BDT) on the link device that are required for the management of shared multimainframe mass storage devices. The entry is valid only for level 0 deadstarts by the first mainframe in the multimainframe complex to deadstart. This entry is valid only when entered from the console keyboard. That is, the PRESET entry cannot be included as part of the EQPDECK on the deadstart file.</p> <p>The amount of table space reserved on the link device is determined by the number of shared devices, as specified by the MAXSHARE entry in the EQPDECK. If the MAXSHARE entry is not specified, the amount of table space reserved on the link device is determined by the number of shared device entries in the EQPDECK.</p>				
PRESET=est ₁ ,est ₂ ,...,est _n .	<p>Presets independent shared devices in a multimainframe complex. It is valid only on a level 0 deadstart by the first mainframe in the multimainframe complex. This entry is valid only when entered from the console keyboard. That is, the PRESET entry cannot be included as part of the EQPDECK on the deadstart file.</p> <table border="1" data-bbox="675 1119 1341 1188"> <thead> <tr> <th data-bbox="675 1119 727 1148">est_i</th> <th data-bbox="797 1119 959 1148">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="675 1157 727 1186">est_i</td> <td data-bbox="797 1157 1341 1186">EST ordinal of independent shared device.</td> </tr> </tbody> </table>	est_i	Description	est _i	EST ordinal of independent shared device.
est_i	Description				
est _i	EST ordinal of independent shared device.				
RESET=est ₁ ,est ₂ ,...,est _n .	<p>Rescinds all device-related attributes except those specified on the EQ entry. It restores the values specified with the last encountered EQest entry. If the DOWN command was specified, EQest_n will be up, but off. Ranges of ordinals are not allowed; each ordinal must be entered individually. This entry is valid only when entered from the console keyboard. That is, the RESET entry cannot be included as part of the EQPDECK on the deadstart file.</p>				

Deadstart Displays

When processing the EQPDECK, the following deadstart displays are available in addition to the EQPDECK and EQPINST displays. You can page through the deadstart displays in a circular manner.

Display	Description
Equipment Status	Shows the hardware configurations (refer to figure 2-2).
Mass Storage Status	Shows how the system allocates files on a mass storage device (refer to figure 2-3).
Mass Storage Initialization Status	Shows the initialization status of the devices (refer to figure 2-4).
Controlware Status	Shows the status of the controlware loaded (refer to figure 2-5).
Disk Thresholds	Shows the disk threshold values at deadstart time (refer to figure 2-6).

Equipment Status Display

The equipment status display lists the status of all the devices.

Figure 2-2 illustrates the equipment status display.

EQUIPMENT STATUS						
EST	TYPE	ST	EQ	UN	CHANNELS	
0.	RD	ON	0.	00.	00.	.
1.	DS	ON	7.	00.	10.	.
2.	NE	ON	0.	00.	00.	.
3.	TE	ON	0.	00.	00.	.

Figure 2-2. Equipment Status Display

Each entry in the display appears in the following format:

est type st eq un channels

Header	Description																																												
est	EST ordinal.																																												
type	Device type. The following device types can appear in the equipment status display.																																												
	<table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AT</td> <td>ACS Cartridge Tape Unit; a 5680 Cartridge Tape Unit attached to a 5744 Automated Cartridge Subsystem (ACS).</td> </tr> <tr> <td>CC</td> <td>Satellite Coupler.</td> </tr> <tr> <td>CM</td> <td>Control module for an 834 Disk Storage Subsystem.</td> </tr> <tr> <td>CP</td> <td>415 Card Punch.</td> </tr> <tr> <td>CR</td> <td>405 Card Reader.</td> </tr> <tr> <td>CT</td> <td>CTS Cartridge Tape Unit; a standalone 5680 Cartridge Tape Unit.</td> </tr> <tr> <td>DB-i</td> <td>885-42 Disk Storage Subsystem ($1 \leq i \leq 3$).</td> </tr> <tr> <td>DC-i</td> <td>895 Disk Storage Subsystem ($1 \leq i \leq 2$).</td> </tr> <tr> <td>DD-i</td> <td>834 Disk Storage Subsystem ($1 \leq i \leq 8$).</td> </tr> <tr> <td>DE</td> <td>Extended memory.</td> </tr> <tr> <td>DF-i</td> <td>887 Disk Storage Subsystem (4K sector; $1 \leq i \leq 3$).</td> </tr> <tr> <td>DG-i</td> <td>836 Disk Storage Subsystem ($1 \leq i \leq 3$).</td> </tr> <tr> <td>DH-i</td> <td>887 Disk Storage Subsystem (16K sector; $1 \leq i \leq 2$).</td> </tr> <tr> <td>DI-i</td> <td>844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).</td> </tr> <tr> <td>DJ-i</td> <td>844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).</td> </tr> <tr> <td>DK-i</td> <td>844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).</td> </tr> <tr> <td>DL-i</td> <td>844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).</td> </tr> <tr> <td>DM-i</td> <td>885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).</td> </tr> <tr> <td>DN</td> <td>9853 Disk Storage Subsystem.</td> </tr> <tr> <td>DP</td> <td>Distributive data path to extended memory.</td> </tr> <tr> <td>DQ-i</td> <td>885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).</td> </tr> </tbody> </table>	type	Description	AT	ACS Cartridge Tape Unit; a 5680 Cartridge Tape Unit attached to a 5744 Automated Cartridge Subsystem (ACS).	CC	Satellite Coupler.	CM	Control module for an 834 Disk Storage Subsystem.	CP	415 Card Punch.	CR	405 Card Reader.	CT	CTS Cartridge Tape Unit; a standalone 5680 Cartridge Tape Unit.	DB-i	885-42 Disk Storage Subsystem ($1 \leq i \leq 3$).	DC-i	895 Disk Storage Subsystem ($1 \leq i \leq 2$).	DD-i	834 Disk Storage Subsystem ($1 \leq i \leq 8$).	DE	Extended memory.	DF-i	887 Disk Storage Subsystem (4K sector; $1 \leq i \leq 3$).	DG-i	836 Disk Storage Subsystem ($1 \leq i \leq 3$).	DH-i	887 Disk Storage Subsystem (16K sector; $1 \leq i \leq 2$).	DI-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).	DJ-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).	DK-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).	DL-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).	DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).	DN	9853 Disk Storage Subsystem.	DP	Distributive data path to extended memory.	DQ-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).
type	Description																																												
AT	ACS Cartridge Tape Unit; a 5680 Cartridge Tape Unit attached to a 5744 Automated Cartridge Subsystem (ACS).																																												
CC	Satellite Coupler.																																												
CM	Control module for an 834 Disk Storage Subsystem.																																												
CP	415 Card Punch.																																												
CR	405 Card Reader.																																												
CT	CTS Cartridge Tape Unit; a standalone 5680 Cartridge Tape Unit.																																												
DB-i	885-42 Disk Storage Subsystem ($1 \leq i \leq 3$).																																												
DC-i	895 Disk Storage Subsystem ($1 \leq i \leq 2$).																																												
DD-i	834 Disk Storage Subsystem ($1 \leq i \leq 8$).																																												
DE	Extended memory.																																												
DF-i	887 Disk Storage Subsystem (4K sector; $1 \leq i \leq 3$).																																												
DG-i	836 Disk Storage Subsystem ($1 \leq i \leq 3$).																																												
DH-i	887 Disk Storage Subsystem (16K sector; $1 \leq i \leq 2$).																																												
DI-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).																																												
DJ-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).																																												
DK-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).																																												
DL-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).																																												
DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).																																												
DN	9853 Disk Storage Subsystem.																																												
DP	Distributive data path to extended memory.																																												
DQ-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).																																												

Header	Description
type	<i>(Continued)</i>
type	Description
DV	819 Disk Storage Subsystem (single density).
DW	819 Disk Storage Subsystem (double density).
EA-i	5830 Disk Array Subsystem; one 5832 SSD drive run in serial mode ($1 \leq i \leq 8$).
EB-i	5830 Disk Array Subsystem; two 5832 SSD drives run in parallel mode ($1 \leq i \leq 6$).
EC-i	5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).
ED-i	5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).
EE	5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.
EF	5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.
EG	5830 Disk Array Subsystem; one 5838 Elite drive.
EH	5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.
LQ	Any line printer.
LR	580-12 Line Printer.
LS	580-16 Line Printer.
LT	580-20 Line Printer.
LX	5870 Printer.
MT	Magnetic Tape Drive (7 track).
NC	380-170 Network Access Device.
ND	CDCNET Device Interface (MDI or MTI).
NP	255x Network Processing Unit.
NT	Magnetic Tape Drive (9 track).
OD	Optical Disk Drive.
RM	Two-port multiplexer (models 865, 875, and CYBER 180-class machines).
SS	MSE Controller.
TT	Internal stimulation device.

Header	Description												
type	<p><i>(Continued)</i></p> <p>The system creates the following device types at deadstart for internal use. Except for the display console, DS, physical hardware does not exist for these device types. The device types appear in the second column of the equipment status display along with the real device types.</p> <table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DS</td> <td>Display console (EST ordinal 1).</td> </tr> <tr> <td>NE</td> <td>Null equipment (EST ordinal 2).</td> </tr> <tr> <td>RD</td> <td>Used for online reconfiguration of mass storage (EST ordinal 0).</td> </tr> <tr> <td>TE</td> <td>Tape equipment (EST ordinal 3).</td> </tr> <tr> <td>TT</td> <td>Used for assignment of terminal files (EST ordinal 4).</td> </tr> </tbody> </table>	type	Description	DS	Display console (EST ordinal 1).	NE	Null equipment (EST ordinal 2).	RD	Used for online reconfiguration of mass storage (EST ordinal 0).	TE	Tape equipment (EST ordinal 3).	TT	Used for assignment of terminal files (EST ordinal 4).
type	Description												
DS	Display console (EST ordinal 1).												
NE	Null equipment (EST ordinal 2).												
RD	Used for online reconfiguration of mass storage (EST ordinal 0).												
TE	Tape equipment (EST ordinal 3).												
TT	Used for assignment of terminal files (EST ordinal 4).												
st	Equipment status (ON, OFF, IDLE, or DOWN).												
eq	Equipment number that corresponds to switch settings on the controller that connects the equipment to the channel(s). The site engineer is responsible for setting these switches. Ask your customer engineer for the correct switch settings if you are unsure.												
un	<p>Unit number (serves as ID code for unit record devices). The ID code provides a method of grouping unit record devices when a site has several units. Output from a job read in through a card reader with identifier un can only be directed to a device with the same identifier. Changing the identifier code via the ROUTE command can direct program output to a special printer.</p> <p>Most equipment has the unit number labeled on the equipment. Some equipment has the unit number as part of a switch or button. Ask your customer engineer for the correct unit numbers if you are unsure.</p>												
channels	Channel(s) on which equipment is available. An asterisk (*) instead of a period (.) following the channel number entry indicates that the channel is down.												

Mass Storage Status Display

The mass storage status display provides detailed information about all mass storage devices.

Figure 2-3 illustrates the mass storage status display.

MASS STORAGE STATUS			
EST	TYPE	STATUS	FILES
5.	DE	----AI-----	S-----
6.	DB	S----I----B--	-----R--T
7.	DB	S----I-----	-----R--T
10.	DB	S----I-----	-----R--T
11.	DB	S----I-----	-----R--T
14.	DQ	-----	-----
15.	DQ	-----	-----
16.	DQ	-----	-----
17.	DQ	-----F-	-----
20.	DQ	-----	-----
21.	DQ	-----	-----
22.	DL	--R-----	-----
23.	DL	--R-----	-----
24.	DL	--R-----	-----
25.	DL	--R-----	-----

Figure 2-3. Mass Storage Status Display

Each entry in the display appears in the following format:

Header	Description
est	EST ordinal.
type	Device type:
type	Description
DB-i	885-42 Disk Storage Subsystem ($1 \leq i \leq 3$).
DC-i	895 Disk Storage Subsystem ($1 \leq i \leq 2$).
DD-i	834 Disk Storage Subsystem ($1 \leq i \leq 8$).
DE	Extended memory.
DF-i	887 Disk Storage Subsystem (4K sector; $1 \leq i \leq 3$).
DG-i	836 Disk Storage Subsystem ($1 \leq i \leq 3$).
DH-i	887 Disk Storage Subsystem (16K sector; $1 \leq i \leq 2$).
DI-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).
DJ-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).
DK-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).
DL-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).
DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).
DN	9853 Disk Storage Subsystem.
DP	Distributive data path to extended memory.
DQ-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).
DV	819 Disk Storage Subsystem (single density).
DW	819 Disk Storage Subsystem (double density).
EA-i	5830 Disk Array Subsystem; one 5832 SSD drive run in serial mode ($1 \leq i \leq 8$).
EB-i	5830 Disk Array Subsystem; two 5832 SSD drives run in parallel mode ($1 \leq i \leq 6$).

Header	Description																												
type (Continued)	Device type.																												
	<table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>EC-i</td> <td>5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).</td> </tr> <tr> <td>ED-i</td> <td>5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).</td> </tr> <tr> <td>EE</td> <td>5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.</td> </tr> <tr> <td>EF</td> <td>5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.</td> </tr> <tr> <td>EG</td> <td>5830 Disk Array Subsystem; one 5838 Elite drive.</td> </tr> <tr> <td>EH</td> <td>5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.</td> </tr> </tbody> </table>	type	Description	EC-i	5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).	ED-i	5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).	EE	5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.	EF	5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.	EG	5830 Disk Array Subsystem; one 5838 Elite drive.	EH	5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.														
type	Description																												
EC-i	5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).																												
ED-i	5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).																												
EE	5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.																												
EF	5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.																												
EG	5830 Disk Array Subsystem; one 5838 Elite drive.																												
EH	5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.																												
status	Status conditions. Any combination of conditions can exist. The following codes are listed in the order in which they appear on the display.																												
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>S</td> <td>A copy of the system resides on this device.</td> </tr> <tr> <td>M</td> <td>Device is shared by more than one mainframe.</td> </tr> <tr> <td>R</td> <td>Device is removable.</td> </tr> <tr> <td>N</td> <td>Independent shared device.</td> </tr> <tr> <td>A</td> <td>Alternate system device.</td> </tr> <tr> <td>I</td> <td>Initialization requested or format is pending.</td> </tr> <tr> <td>P</td> <td>Preset of the device requested.</td> </tr> <tr> <td>D</td> <td>System dayfile on this device.</td> </tr> <tr> <td>C</td> <td>Account dayfile on this device.</td> </tr> <tr> <td>E</td> <td>Error log on this device.</td> </tr> <tr> <td>B</td> <td>Binary maintenance log on this device.</td> </tr> <tr> <td>F</td> <td>System default family on this device.</td> </tr> <tr> <td>G</td> <td>Checkpoint device. A copy of the checkpoint file resides on this device.</td> </tr> </tbody> </table>	status	Description	S	A copy of the system resides on this device.	M	Device is shared by more than one mainframe.	R	Device is removable.	N	Independent shared device.	A	Alternate system device.	I	Initialization requested or format is pending.	P	Preset of the device requested.	D	System dayfile on this device.	C	Account dayfile on this device.	E	Error log on this device.	B	Binary maintenance log on this device.	F	System default family on this device.	G	Checkpoint device. A copy of the checkpoint file resides on this device.
status	Description																												
S	A copy of the system resides on this device.																												
M	Device is shared by more than one mainframe.																												
R	Device is removable.																												
N	Independent shared device.																												
A	Alternate system device.																												
I	Initialization requested or format is pending.																												
P	Preset of the device requested.																												
D	System dayfile on this device.																												
C	Account dayfile on this device.																												
E	Error log on this device.																												
B	Binary maintenance log on this device.																												
F	System default family on this device.																												
G	Checkpoint device. A copy of the checkpoint file resides on this device.																												

Header	Description																				
files	Types of files that are allowed on this device. Any combination of types can exist. The following codes are listed in the order in which they appear on the display. Refer to the MSAL EQPDECK entry in section 3, Deadstart Decks.																				
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>S</td> <td>Secondary rollout.</td> </tr> <tr> <td>B</td> <td>LGO.</td> </tr> <tr> <td>L</td> <td>Local.</td> </tr> <tr> <td>P</td> <td>Primary.</td> </tr> <tr> <td>D</td> <td>Job dayfile.</td> </tr> <tr> <td>R</td> <td>Rollout.</td> </tr> <tr> <td>O</td> <td>Output.</td> </tr> <tr> <td>I</td> <td>Input.</td> </tr> <tr> <td>T</td> <td>Temporary.</td> </tr> </tbody> </table>	status	Description	S	Secondary rollout.	B	LGO.	L	Local.	P	Primary.	D	Job dayfile.	R	Rollout.	O	Output.	I	Input.	T	Temporary.
status	Description																				
S	Secondary rollout.																				
B	LGO.																				
L	Local.																				
P	Primary.																				
D	Job dayfile.																				
R	Rollout.																				
O	Output.																				
I	Input.																				
T	Temporary.																				

Mass Storage Initialization Status Display

The mass storage initialization status display shows the initialization status of mass storage devices on the system.

Figure 2-4 illustrates the mass storage initialization status display.

MASS STORAGE INITIALIZATION STATUS								
EST	TY	OPTIONS	T	IAM	DAM	FM/PN	DN	NC
5.	DE	---QSA--						
6.	DB	-----E-						
7.	DB	--PQSAED						
10.	DB	FT-----						
11.	DB							
14.	DQ		F	377.	377.	FEATURE	40.	000.
15.	DQ		X	000.	377.	R4IAE		000.
16.	DQ							
17.	DQ							
20.	DQ							
21.	DQ							

Figure 2-4. Mass Storage Initialization Status Display

Each line in the display appears in this format:

```
est      type      options    t      iam      dam      fm/pn     dn      nc
```

Header

Description

est EST ordinal.
type Device type.

type Description

DB-i 885-42 Disk Storage Subsystem (1 ≤ i ≤ 3).
DC-i 895 Disk Storage Subsystem (1 ≤ i ≤ 2).
DD-i 834 Disk Storage Subsystem (1 ≤ i ≤ 8).
DE Extended memory.
DF-i 887 Disk Storage Subsystem (4K sector; 1 ≤ i ≤ 3).
DG-i 836 Disk Storage Subsystem (1 ≤ i ≤ 3).
DH-i 887 Disk Storage Subsystem (16K sector; 1 ≤ i ≤ 2).
DI-i 844-21 Disk Storage Subsystem (1 ≤ i ≤ 8; half track).

Header	Description
--------	-------------

type (Continued)

type	Description
DJ-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; half track).
DK-i	844-21 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).
DL-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).
DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).
DN	9853 Disk Storage Subsystem.
DP	Distributive data path to extended memory.
DQ-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).
DV	819 Disk Storage Subsystem (single density).
DW	819 Disk Storage Subsystem (double density).
EA-i	5830 Disk Array Subsystem; one 5832 SSD drive run in serial mode ($1 \leq i \leq 8$).
EB-i	5830 Disk Array Subsystem; two 5832 SSD drives run in parallel mode ($1 \leq i \leq 6$).
EC-i	5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).
ED-i	5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).
EE	5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.
EF	5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.
EG	5830 Disk Array Subsystem; one 5838 Elite drive.
EH	5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.

options Initialize options. Maximum of eight options are displayed.

options	Description
F	Format pack.
T	Total system.
P	Permanent files.
Q	Queue files.
S	System dayfile.
A	Account dayfile.
E	Error log.
B	Binary maintenance log.

Header	Description						
t	Type of unit device.						
	<table border="1"> <thead> <tr> <th>t</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>F</td> <td>Family device.</td> </tr> <tr> <td>X</td> <td>Auxiliary device.</td> </tr> </tbody> </table>	t	Description	F	Family device.	X	Auxiliary device.
t	Description						
F	Family device.						
X	Auxiliary device.						
iam	Indirect access mask.						
dam	Direct access mask.						
fm/pn	Family name/pack name.						
	<table border="1"> <thead> <tr> <th>fm/pn</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>F</td> <td>Family name.</td> </tr> <tr> <td>X</td> <td>Pack name.</td> </tr> </tbody> </table>	fm/pn	Description	F	Family name.	X	Pack name.
fm/pn	Description						
F	Family name.						
X	Pack name.						
dn	Device number if the device unit is the family device.						
nc	Number of permanent file catalog tracks.						

Controlware Status Display

The controlware status display shows what type of controlware is loaded to which channels. Figure 2-5 illustrates the controlware status display.

CONTROLWARE CHANNEL	STATUS CONTROLWARE
00.	
01.	
02.	
03.	NN
04.	
05.	
07.	FM
10.	
11.	NN
12.	
13.	

Figure 2-5. Controlware Status Display

Each line in the display appears in this format:

channel controlware

Header	Description
channel	Channel on which equipment is available.
controlware	Name of the controlware loaded. Refer to the LBC EQPDECK entry in section 3, Deadstart Decks, for a list of controlware types.

Disk Thresholds Display

The disk thresholds display shows the disk threshold values at deadstart time. The system uses these threshold values as limits when monitoring disk verification failures, available disk space, and disk error processing. If a threshold value is exceeded, the system performs a corrective action, such as notifying the operator or restricting activity on the affected disk.

Figure 2-6 illustrates the disk thresholds display.

EST	TYPE	SIZE	VF	RA	LS	RE	UE
6	DL-3	3150	100	315	146	100	0

Figure 2-6. Disk Thresholds Display

Each line in the display appears in the following format:

```
est      type      size      vf      ra      ls      re      ue
```

Header	Description
est	EST ordinal.
type	Device type.
type	Description
DB-i	885-42 Disk Storage Subsystem (1 ≤ i ≤ 3).
DC-i	895 Disk Storage Subsystem (1 ≤ i ≤ 2).
DD-i	834 Disk Storage Subsystem (1 ≤ i ≤ 8).
DE	Extended memory.
DF-i	887 Disk Storage Subsystem (4K sector; 1 ≤ i ≤ 3).
DG-i	836 Disk Storage Subsystem (1 ≤ i ≤ 3).
DH-i	887 Disk Storage Subsystem (16K sector; 1 ≤ i ≤ 2).
DI-i	844-21 Disk Storage Subsystem (1 ≤ i ≤ 8; half track).
DJ-i	844-41/44 Disk Storage Subsystem (1 ≤ i ≤ 8; half track).
DK-i	844-21 Disk Storage Subsystem (1 ≤ i ≤ 8; full track).

Header	Description																																
type	<i>(Continued)</i>																																
	<table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DL-i</td> <td>844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).</td> </tr> <tr> <td>DM-i</td> <td>885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).</td> </tr> <tr> <td>DN</td> <td>9853 Disk Storage Subsystem.</td> </tr> <tr> <td>DP</td> <td>Distributive data path to extended memory.</td> </tr> <tr> <td>DQ-i</td> <td>885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).</td> </tr> <tr> <td>DV</td> <td>819 Disk Storage Subsystem (single density).</td> </tr> <tr> <td>DW</td> <td>819 Disk Storage Subsystem (double density).</td> </tr> <tr> <td>EA-i</td> <td>5830 Disk Array Subsystem; one 5832 SSD drive run in serial mode ($1 \leq i \leq 8$).</td> </tr> <tr> <td>EB-i</td> <td>5830 Disk Array Subsystem; two 5832 SSD drives run in parallel mode ($1 \leq i \leq 6$).</td> </tr> <tr> <td>EC-i</td> <td>5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).</td> </tr> <tr> <td>ED-i</td> <td>5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).</td> </tr> <tr> <td>EE</td> <td>5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.</td> </tr> <tr> <td>EF</td> <td>5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.</td> </tr> <tr> <td>EG</td> <td>5830 Disk Array Subsystem; one 5838 Elite drive.</td> </tr> <tr> <td>EH</td> <td>5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.</td> </tr> </tbody> </table>	type	Description	DL-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).	DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).	DN	9853 Disk Storage Subsystem.	DP	Distributive data path to extended memory.	DQ-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).	DV	819 Disk Storage Subsystem (single density).	DW	819 Disk Storage Subsystem (double density).	EA-i	5830 Disk Array Subsystem; one 5832 SSD drive run in serial mode ($1 \leq i \leq 8$).	EB-i	5830 Disk Array Subsystem; two 5832 SSD drives run in parallel mode ($1 \leq i \leq 6$).	EC-i	5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).	ED-i	5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).	EE	5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.	EF	5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.	EG	5830 Disk Array Subsystem; one 5838 Elite drive.	EH	5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.
type	Description																																
DL-i	844-41/44 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).																																
DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).																																
DN	9853 Disk Storage Subsystem.																																
DP	Distributive data path to extended memory.																																
DQ-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).																																
DV	819 Disk Storage Subsystem (single density).																																
DW	819 Disk Storage Subsystem (double density).																																
EA-i	5830 Disk Array Subsystem; one 5832 SSD drive run in serial mode ($1 \leq i \leq 8$).																																
EB-i	5830 Disk Array Subsystem; two 5832 SSD drives run in parallel mode ($1 \leq i \leq 6$).																																
EC-i	5830 Disk Array Subsystem; one 5833 Sabre drive run in serial mode ($1 \leq i \leq 2$).																																
ED-i	5830 Disk Array Subsystem; two 5833 Sabre drives: one for data and the other for parity ($1 \leq i \leq 2$).																																
EE	5830 Disk Array Subsystem; two 5833 Sabre drives run in parallel mode.																																
EF	5830 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.																																
EG	5830 Disk Array Subsystem; one 5838 Elite drive.																																
EH	5830 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.																																
size	Total number of logical tracks on the corresponding device type.																																
vf	Verification failure threshold.																																
ra	Restricted activity threshold.																																
ls	Low space threshold.																																
re	Recovered error threshold.																																
ue	Unrecovered error threshold.																																

Modifying the APRDECKs

After completing all EQPDECK modifications, you can also modify the default APRDECK, the APRDECK for a specific equipment, or the IPRDECK being used. You can modify an APRDECK only when you are initializing the corresponding equipment. If no changes need to be made to any APRDECK, but you do need to modify the IPRDECK, refer to Modifying the IPRDECK later in this section.

The APRDECK contains entries identifying areas of mass storage that are not usable (flaws). The APRDECK used can vary from equipment to equipment. One of the parameters specified when an equipment is defined in the EQPDECK is the APRDECK number that applies to that equipment. The default (APRD00) is selected if this parameter is not specified.

To modify an APRDECK, enter

NEXT.

while the EQPDECK or EQPINST is being displayed. The APRINST display describes the valid entries. You can toggle between the APRDECK and APRINST. Enter the changes or additions to the APRDECK from the console keyboard (refer to section 3, Deadstart Decks, for a description of the entries).

If there are no changes to the APRDECK displayed, enter

NEXT.

to go to the next APRDECK. Repeat this process until the appropriate APRDECK is displayed or until you have changed all APRDECKs needing changes.

After all APRDECK modifications are complete, you can skip to the IPRDECK by entering:

IPR.

Refer to Modifying the IPRDECK, next, for more information. Otherwise, to indicate that all modifications to the APRDECKs are complete and you want to begin loading the system, enter:

GO.

Modifying the IPRDECK

The IPRDECK contains installation parameters that describe the mode of system operation. IPRDECK modification is seldom required during deadstart since nearly all IPRDECK commands are also valid DSD commands that make the same changes during normal system operation. Generally, installation parameters changed during normal operations (with DSD commands or by modifying the IPRDECK) are retained only across a level 3 recovery deadstart.

After entering

IPR.

when all the CMRDECK, EQPDECK, or APRDECK modifications are complete or after repeatedly entering

NEXT.

to step through all the APRDECKs, the instruction display entitled IPRINST appears on the console screen(s). This display defines all valid IPRDECK entries. Most of these entries are also valid DSD commands.

To view the contents of the IPRDECK being used, toggle from the IPRINST display to the IPRDECK display. If either the IPRDECK or IPRINST overflows two screens, you can page the display.

Enter the appropriate changes or additions from the console keyboard. These entries can be made while either IPRINST or IPRDECK is being displayed. A console entry supersedes the value currently specified in the IPRDECK.

NOTE

The modified IPRDECK remains in effect only until the next level 0, 1, or 2 deadstart is performed. Changes to the IPRDECK are retained if a level 3 recovery deadstart is performed. If you want these changes to take place on the next level 0 deadstart, make the appropriate changes to the IPRDECK after NOS is up and running and use LIBEDIT to replace the record on the deadstart file.

For complete information concerning IPRDECK entries, refer to section 3, Deadstart Decks, and to section 5, DSD Commands.

To indicate that changes to the CMRDECK, EQPDECK, APRDECK, and/or IPRDECK are completed and you want to begin loading the system, enter:

GO.

Loading the System

If you are performing a level 0 or level 2 deadstart, the system library is automatically loaded from the deadstart file to each mass storage device specified in the EQPDECK as a system device. If no system device is specified, the system is loaded on the first nonremovable mass storage device in the equipment status table.

Mass storage labels are validated for all levels of deadstart. This ensures that the configuration matches the one specified in the EQPDECK. Normally, the device label validation display appears for only a few seconds. However, this display remains longer (from 30 seconds to 3 minutes) if 834, 836, 887, or 9853 disk units must first be automatically spun up in order to validate their labels.

If you specify a level 1 or level 3 recovery deadstart, the system library is not reloaded. In this case, the deadstart file is rewound and is not accessed again until another deadstart operation is performed. The system library is recovered from checkpoint information on mass storage. Central memory tables such as the system file name table (FNT), executing job table (EJT), queued file table (QFT), equipment status table (EST), and track reservation table (TRT) are either recovered from checkpoint information for level 1 or from central memory (and the link device, extended memory, if in multimainframe mode) for level 3.

Messages

If a deadstart error occurs, a message appears on the right console screen and, depending upon the nature of the error, deadstart processing may halt. Refer to Error Processing later in this section for complete information and corrective action. If the system is being loaded (level 0 or 2 only), the name of each system library program is also displayed on the right console screen as it is being loaded. This allows you to monitor deadstart progress.

The left console screen may display the message ENTER DATE YY/MM/DD. You can enter the date while the system is being loaded.

Entering the Date and Time

Each time a system deadstart function is performed, it is necessary to enter the current date and time (except for CYBER 180-class machines; they get the date and time from CTI during deadstart). The system uses the date and time (updated every second) for dayfile messages and for permanent file catalogs and directories for files being accessed. This includes the creation, last modification, and last access date and time for each permanent file in the system. It is important to enter the correct date and time in order to accurately maintain these system records. If you are performing a level 3 recovery deadstart, it is possible to recover the date and time from the previous system deadstart. However, this is not recommended since the new date and time recorded for system records would no longer be accurate.

When the system loading (or recovery) phase of deadstart is about to begin, the system checks for the correct date and time. If CTI is unable to determine the correct date and time or the mainframe is other than a CYBER 180-class, the following message appears on the left console screen requesting entry of the current date.

```
ENTER DATE YY/MM/DD.
```

Enter the current date, followed by CR, in this format:

```
yy/mm/dd.
```

Entry	Description
yy	Year; 70 through 99 (corresponds to the years 1970 through 1999) or 00 through 33 (corresponds to the years 2000 through 2033).
mm	Month; 01 through 12.
dd	Day; 01 through nn (nn is the number of days in the month).

For deadstart levels 0, 1, and 2, pressing CR without first entering the date causes the system to assume the date that the deadstart file was created. For a level 3 recovery deadstart, pressing CR alone recovers both the previous date and time (time entered during the last deadstart plus time accumulated until this deadstart).

When the system accepts the date entry, it displays the following request for entry of the current time.

```
ENTER TIME HH.MM.SS.
```

Enter the current time in this format:

```
hh.mm.ss.
```

Entry	Description
hh	Hour; 00 through 23.
mm	Minute; 00 through 59.
ss	Second; 00 through 59.

For deadstart levels 0, 1, and 2, pressing CR without first entering the time causes the system to set the time to 00.00.00. If you enter CR alone following the date prompt on a level 3 recovery, this prompt to enter time does not appear (the previous time is assumed).

Initiating Job Processing

DSD commands specified in the IPRDECK automatically initiate normal job processing except on a level 3 recovery where you must enter AUTO. If a level 1 or level 3 recovery deadstart is being performed, the system recovers all jobs and active files and resumes normal operation immediately. However, if an initial deadstart (level 0) or level 2 recovery deadstart is being performed, job processing may not be initiated immediately. The time of initiation depends upon the time it takes to load the system from the deadstart file (you can monitor progress on the right console screen). If file loading is not completed when the time entry is made, the DSD commands specified in the IPRDECK are displayed on the left screen. Until file loading completes, you can clear one or more of the DSD commands.

Clearing a command prevents it from being executed when file loading completes. In this case, you must manually enter the commands necessary to initiate job processing from the console keyboard.

To initiate job processing, enter:

AUTO.

or

MAINTENANCE.

or

SCHEDULE.

Following entry of the AUTO, MAINTENANCE, or SCHEDULE command during an initial (level 0) deadstart, the deadstart sequencing process begins. Deadstart sequencing causes job processing to be suspended until all system files in the default family are initiated. To initiate a family other than the default, enter the command:

X.ISF(FM=family)

Parameter	Description
family	Alternate family of devices.

For additional information concerning the ISF command and deadstart sequencing, refer to section 20, System File Initialization.

Normal job processing begins after the deadstart sequencing job completes. If the AUTO command is entered, the subsystems enabled in the IPRDECK are automatically assigned to specified control points. Assuming that all standard subsystems are set to be enabled, the system calls them to specific control points as shown in the following example:

Control Point Number	Job Sequence Name	Activity
1	IAF	Interactive Facility.
2	NAM	Network Access Method.
3	RHF	Remote Host Facility.
4	SMF	Screen Management Facility.
n-2 ¹	MAG	Magnetic tape subsystem executive routine.
n-1 ¹	BIO	Central site batch input/output.
n ¹	RBF	Remote Batch Facility.
n+1 ¹	SYS	System.

The MAINTENANCE command performs the same function as the AUTO command. Additionally, it assigns several maintenance routines, according to mainframe type, to available control points and runs them as normal jobs with minimum queue and CPU priorities. These are CPU or central memory test routines designed to detect hardware errors. The routines display error messages either at the control point on the B,O display or in the system error log. To display the error log, enter:

A, ERROR LOG.

You should monitor these routines from time to time. If a maintenance routine displays an error message indicating a hardware malfunction occurred, call a customer engineer. It is recommended that these programs be run at all times. The maintenance programs use little memory, are run at minimum CPU and queue priority, and are automatically rolled out if necessary; thus, system performance is not severely affected. Descriptions of the maintenance routines are in the NOS Online Maintenance Software Reference Manual.

You can use the SCHEDULE command to schedule jobs, but it does not initiate subsystems or maintenance jobs. Refer to the NOS Version 2 Operations Handbook for a description of the SCHEDULE command.

1. n is the number of control points defined by the NCP entry in the CMRDECK.

Preparing for Recovery Deadstart

Sometimes during system operation an uncorrectable error occurs that prevents further system activity. Often the situation can be corrected by deadstarting the system and recovering prior activity. The success of such a recovery depends upon the severity of the problem and the extent to which system information is destroyed. There are three levels of recovery deadstart available (levels 1, 2, and 3). Table 2-2 lists each deadstart level, including level 0 (initial deadstart), and describes the extent of recovery possible.

If the MS VALIDATION installation option is enabled during a level 1 or 2 recovery, or if both the MS VALIDATION and PF VALIDATION options are enabled during a level 3 recovery, the system:

- Verifies selected mass storage files.
- Checks files identified in the queued file table (QFT) to ensure that all tracks in the chain are reserved and that no circular linkage exists.
- Depending on file type, checks the track reservation table (TRT) to ensure that the file is preserved.
- Checks all preserved files for proper length.

If the system encounters a verification failure, it clears the queued file table entry but does not release disk space assigned to the file. If a length error is detected, the system sets error idle status and terminates recovery operations on the device. To recover from an error idle condition, refer to Error Idle Recovery under Error Processing later in this section.

During level 0 deadstart, the system verifies the length of preserved files regardless of the setting of the MS VALIDATION option. If a length error is detected, the system reads the disk chain to determine the correct length of the file, issues a message to the B,O display, and stops recovery of the device. To alter the EOI for the file and proceed with recovery, enter:

```
GO,SYS. or GO,.
```

To terminate recovery of the device, enter:

```
PAUSE,SYS.
```

Table 2-2. Levels of System Deadstart: Information Recovered

Dead-start Level	Jobs	Queued Files	Perma- nent Files	System	System Dayfiles
0	Batch jobs are rerun. ¹	Recovered as inactive.	Yes.	No.	Yes. ²
1	Recovered from last checkpoint.	Active queues recovered from last checkpoint.	Yes.	Recovered from last checkpoint.	Yes.
2	Recovered from last checkpoint.	Active queues recovered from last checkpoint	Yes.	No.	Yes.
3	Recovered from CM copy of EJT. ³	Active queues recovered from CM copy of QFT.	Yes.	Yes.	Yes.

1. The input and output queues are recovered. Rolled out and executing batch jobs are also rerun. The input files associated with these jobs are returned to the input queues, unless the device they reside on is initialized. Permanent files are recovered unless the device is initialized.

2. Dayfiles are recovered unless initialized by an INITIALIZE entry in the EQPDECK.

3. Jobs that are rolled out continue. Jobs that are in CM are aborted with EXIT processing and then rerun if possible.

CAUTION

Before attempting any level of recovery deadstart, whether it is a level 0, 1, 2, or 3 deadstart, enter a **CHECKPOINT SYSTEM** command, and wait for the **CHECKPOINT COMPLETE** message to appear on the system console. The **CHECKPOINT SYSTEM** performs the following functions:

1. All currently executing jobs are rolled out. (Ensures that these jobs can roll back and continue executing after a level 3 deadstart.)
2. All mass storage tables and track reservation tables are copied to disk, and all pending disk I/O operations are completed. (Ensures the integrity of the disks after a level 0, 1, or 2 deadstart.)
3. A system checkpoint file (which includes a copy of all the system tables in central memory) is written to disk. (Provides a system checkpoint file for a level 1 or 2 deadstart.)

If, due to a system hang or other problem, it is not possible to enter a **CHECKPOINT SYSTEM**, or if the **CHECKPOINT COMPLETE** message does not appear, the disk updates may not have been performed. If a level 0 deadstart is performed at this time, permanent file information may be lost. To prevent such a loss, always perform a level 3 deadstart after a system hang. Select the **ABORT** option on the level 3 deadstart display to abandon the deadstart; the **ABORT** option performs the disk updates and then abandons the deadstart.

The following topics provide general information concerning each level of system deadstart and recommended steps of preparation.

Level 3 Recovery Deadstart

Usually you perform a level 3 recovery deadstart following an equipment or system malfunction (for example, channel or PP hung), providing the system remains intact. Basically, the system FNT, QFT, EJT, TRT, EST, and control-point areas of central memory must be intact in order to successfully perform a level 3 recovery deadstart. However, unless you can determine that central memory is no longer intact, attempt a level 3 recovery deadstart before a level 0 deadstart. This is recommended because system activity, as it existed at the time of the malfunction, can best be recovered by performing a level 3 recovery deadstart. Only PP memory confidence testing occurs during a level 3 recovery deadstart; central memory is not affected. User jobs running at control points are either rerun or aborted by a level 3 deadstart; rolled out user jobs are rolled back in and continue to run.

Requests for device checkpoint are retained over a level 3 recovery. Therefore, if a system malfunction prevents a device checkpoint from being done, the checkpoint is processed after level 3 recovery is successfully completed. If a level 3 recovery fails, always do another level 3 recovery with the ABORT option selected before doing a level 0 deadstart. This ensures that all checkpoint processing is done correctly.

On a level 3 deadstart the CMRDECK, the EQPDECK, the APRDECKs, and the IPRDECK cannot be viewed or changed. The CMRDECK, the EQPDECK, the APRDECKs, and the IPRDECK specified during the last level 0 deadstart remain in effect. If you set the CMRDECK switch (bit 6 in word 13 of the deadstart panel) or select the D=YES option on the CTI *P* display; the system halts and displays level 3 deadstart options. The options and their default values are displayed on the left screen (figure 2-7) and instructions on the right screen (figure 2-8).

LEVEL 3 OPTIONS	
ABORT .	NO
ABORT , B .	NO
AUTOLOAD .	YES
GRENAD .	NO
AUTO .	YES

Figure 2-7. Level 3 Deadstart Left Screen Display

```

INSTRUCTIONS FOR SELECTING LEVEL 3 RECOVERY OPTIONS.
  ENTER COMMAND TO TOGGLE SELECTION.
  ENTER GO. TO CONTINUE RECOVERY.
ABORT.
  CHECKPOINT ALL DEVICES AND ABORT LEVEL 3 RECOVERY.
  SELECTING *ABORT.* DESELECTS *ABORT,B.*

ABORT,B.
  CHECKPOINT ALL NONBUFFERED DEVICES AND ABORT LEVEL
  3 RECOVERY.  SELECTING *ABORT,B.* DESELECTS *ABORT.*.

AUTOLOAD.
  TOGGLE THE SELECTION OF BUFFER CONTROLLER AUTOLOADING.

GRENADE.
  TOGGLE THE SELECTION OF THE GRENADE FUNCTION.  THE
  GRENADE FUNCTION IS ISSUED ONCE THE CONTROLWARE IS
  LOADED, CAUSING UNIT RESERVATIONS TO BE CLEARED ON
  ALL 844 UNITS PHYSICALLY CONNECTED TO EACH CONTROLLER.

AUTO.
  TOGGLE THE SELECTION OF THE DSD AUTO COMMAND.

```

Figure 2-8. Level 3 Deadstart Right Screen Display

Entering a command will toggle the level 3 deadstart selections. The ABORT command checkpoints all the devices and aborts level 3 recovery. Selecting the ABORT command automatically deselects the ABORT,B command. The ABORT,B command checkpoints all the nonbuffered devices and aborts level 3 recovery. Selecting the ABORT,B command automatically deselects the ABORT command. The AUTOLOAD command toggles the selection of buffer controller autoloading. The GRENADE command toggles the selection of the grenade function. The AUTO command toggles the selection of the DSD AUTO command. Refer to table 2-1 for more information about the AUTOLOAD and GRENADE commands. When you are done making changes, enter

GO.

to continue the deadstart recovery. If you choose not to display the level 3 options, the default values are used.

A level 3 recovery deadstart is impossible after:

- An attempted checkpoint recovery (level 1).
- An aborted level 0 (initial) deadstart.
- The MREC utility (refer section 8, K-Display Utilities) has been run for the machine to be deadstarted while in multimainframe mode.

It is recommended that you stop system activity prior to beginning the system deadstart procedure. To accomplish this, enter the following DSD commands:

Command	Description
E,M.	Displays the disk status display (E,M).
UNLOCK.	Unlocks the system console. This command is necessary only if the system console is locked.
CHECK POINT SYSTEM.	Provides for termination of job processing and for writing the contents of central memory tables to mass storage. For a complete description of this process, refer to the CHECK POINT SYSTEM command in section 5, DSD Commands.
STEP.	Prevents the system from processing PP requests. This stops all central memory input/output (I/O) operations. You should enter the STEP command after all device checkpoints are completed. Determine checkpoint status from the disk status display (E,M) (refer to the NOS Version 2 Operations Handbook).

NOTE

To recover interactive users after a level 3 recovery, sense switch 1 must be set in the IAF procedure. Sense switch 1 is set by default in the released IAF procedure.

Level 1 Recovery Deadstart

Usually you perform a level 1 recovery deadstart to resume normal processing following maintenance procedures. The system, all jobs, and all active files are recovered from checkpoint information on mass storage.

NOTE

A level 1 recovery deadstart is not intended to be a recovery process after a system/equipment malfunction. You should never attempt it after a level 3 recovery deadstart fails.

Level 1 recovery is also useful in system test situations. If two systems are being alternated, separate mass storage devices and tapes² must be available for both systems. Tapes are not repositioned after a level 1 deadstart. Thus, if a job was previously assigned to the tape unit that has been used for deadstarting, the job cannot be recovered. The tape unit should be left unloaded after recovery until it is no longer assigned to the job (job aborted).

The following rules apply when performing a level 1 recovery deadstart.

- The DSD command CHECK POINT SYSTEM must have been successfully completed immediately before the end of the last NOS operating period.
- The contents of the extended memory must not be destroyed from the time of the CHECK POINT SYSTEM command.
- Memory dumps must be completed before level 1 recovery deadstart begins since memory confidence testing destroys the contents of both central memory (except on CYBER 180-class machines) and PPs.
- The mass storage equipment configuration must be the same as specified during the most recent level 0 deadstart; that is, the same EQPDECK must be used.
- The system devices (mass storage devices on which the system library resides) must be the same as or fewer than those specified during the most recent level 0 deadstart.

It is recommended that you stop system activity before beginning the system deadstart procedure. To accomplish this, enter the following DSD commands:

Command	Description
UNLOCK.	Necessary only if console is currently locked.
CHECK POINT SYSTEM.	Provides for termination of job processing and for writing the contents of central memory tables to mass storage.
STEP.	Prevents the system from processing PP requests. This stops all central memory I/O operations.

2. Separate tapes are necessary only if tape jobs are being checkpointed.

Level 2 Recovery Deadstart

Usually you perform level 2 recovery deadstart in system test situations; it is not recommended for the normal production environment. If you select level 2 recovery, all jobs and active files are recovered from checkpoint information on mass storage as in level 1 recovery. However, no attempt is made to recover the system. Instead, the system is loaded from the deadstart file as in level 0 deadstart. In all other respects, level 2 recovery is identical to that described for level 1, and all level 1 rules apply.

Level 0 Initial Deadstart

Use level 0 or initial deadstart in cases where a recovery deadstart is not possible. This is a complete or initial load from the deadstart file. Only preserved files, which include permanent files, queued files, and system dayfiles, are recovered (preserved files are recovered on all levels of system deadstart). Because memory confidence testing destroys the contents of central memory (except on CYBER 180-class machines) and PPs, all memory dumps must be completed before deadstart begins.

NOTE

If the machine is the first machine being deadstarted in a shared device multimainframe environment, you must enter a PRESET command (refer to section 13, Multimainframe Operations).

Error Processing

If no CTI display appears when you initiate a deadstart, perform the following steps as needed. After each step, reinitiate the deadstart to see if the problem has been eliminated.

Deadstart from tape:³

1. If the unit select switch on the deadstart tape unit is not on (tape does not move), check the channel, controller, and unit selections on the deadstart panel or display to ensure they are set correctly.
2. If the unit select switch is on, the correct unit was selected; however, check word 11 of the deadstart panel or display to ensure it is set correctly.
3. Ensure that a 7-track tape is not mounted on a 9-track drive or vice versa. Also, ensure that a deadstart tape is not mounted on a tape unit that does not support the density of the deadstart tape.
4. Ensure that the deadstart tape is an I-format unlabeled tape.
5. Ensure that the card reader and tape unit (667 or 669 only) are not on the same channel and that the card reader is not on a channel with a PP. Also, ensure that two or more units do not have the same physical unit number.
6. If still no display appears after initiating the deadstart, there might be a parity error on one of the first records of the deadstart tape or the magnetic tape controller might have detected a channel parity error on a CYBER 170 Computer System.

3. When deadstarting a CYBER 180-class machine, the disk containing the CIP module must be used.

Deadstart from disk:

1. Ensure that disk power is on and the disk is ready for operation.
2. Ensure that the disk has the CIP module loaded.
3. Ensure that the deadstart panel or display is set correctly.
4. Select an alternate channel.
5. If still no display appears after initiating the deadstart, there might be a parity error on one of the first records of the deadstart file or the disk controller might have detected a channel parity error on a CYBER 170 Computer System.

For a proper understanding of the problems that can occur during deadstart, you should be familiar with several basic concepts. For example, because most errors that occur involve mass storage devices, you should be familiar with their use in the system. Each mass storage device has a label that contains descriptive information about its contents. For certain levels of recovery deadstart, this information must be consistent with corresponding information either contained in central memory or provided through deadstart procedures. Conflicts can result in the system issuing deadstart error messages. An attempt is made to recover all mass storage devices defined in the EST during all levels of system deadstart.

Recovery Functions

The specific recovery function performed depends upon the level of deadstart selected. The table on the following pages describes the recovery function performed for each deadstart level and the types of errors you can encounter. The system response to errors and the recommended action are also listed.

Refer to the NOS Version 2 Operations Handbook for information concerning all deadstart messages.

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
Read and verify the label on each mass storage device against that specified by the EQ definition for the device in BOPDECK.	0	<p>Device being recovered is not ready and, therefore, cannot be read.</p> <p>Read error occurred when attempting to read label (parity error, surface error, and so on).</p> <p>Label verification error. The label was read but could not be verified. For example, label indicates device is first unit of a three-unit multispindle pack (DK-3), but not all elements of the configuration are present.</p> <p>Bad label. Information read was not recognized as a label sector.</p>	<p>Removable.</p> <p>Nonremovable.</p> <p>All.</p> <p>Nonsystem device.</p> <p>System device (nonremovable with system residence).</p> <p>Nonsystem device.</p> <p>System device (nonremovable with system residence).</p>	<p>Error code NR and status U set in MST; device is left unavailable.</p> <p>Wait and retry. Error message appears on the right screen display.</p> <p>Advance to next track and try to read. A predetermined number of tracks are searched if error persists. If this number is exceeded, the device is treated as if a bad label existed. Refer to the description of bad label error, below.</p> <p>Error code CE and status U set in MST; device is left unavailable.</p> <ul style="list-style-type: none"> - If permanent files do not reside on the device, it is initialized automatically using parameters in existing label. - If permanent files reside on the device, the following message appears on the right console display and deadstart processing halts. CONTINUING DESTROYS PFS RECOVERY OF DEVICE IMPOSSIBLE. <p>Error code LE and status U set in MST; device is left unavailable.</p> <p>Device is initialized automatically. Any permanent files that reside on device are lost.</p>	<p>None.</p> <p>Check to ensure that all nonremovable devices are ready.</p> <p>None.</p> <p>Examine E.M display after deadstart is complete to determine corrective action.</p> <p>None.</p> <p>Type GO, or PAUSE, to proceed. If GO, is entered, the device is initialized using parameters in existing label (permanent files will be lost), or redeadstart without the system on this device.</p> <p>Device must be initialized after deadstart if it is to be used (refer to DSD command INITIALIZE in section 5).</p> <p>None.</p>

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
Chain removable devices with consecutive physical unit numbers (applies only to 844-21 (DI or DK) and 844-61/44 (DJ or DL). This will allow a multispindle pack (for example, DK-3) to be mounted on devices defined in EST as single spindle devices (for example, DK-1). Definition for devices in EST is changed to reflect a multispindle device.	1, 2, 3	Label verification error. The label was read but could not be verified. For example, label indicates device is first unit of a three-unit multispindle pack (DK-3), but units 2 of 3 and/or 3 of 3 are not present.	Removable (no active direct access files). Nonremovable or removable with active direct access files.	Error code indicates cause of error and status U set in MST; device is left unavailable. The following message appears on the right console display and deadstart processing halts. ERROR ON DEVICE WITH ACTIVE FILES. NOTE Active direct access files are attached when checkpoint is taken (level 1 or 2) or when deadstart is initiated. It is assumed that all removable devices have active direct access files.	Examine E,M display after deadstart is complete to determine corrective action. Perform level 0 deadstart or redefine configuration to match that of system being recovered.
	0, 1, 2	Label verification error. Label on remaining spindles of multispindle pack could not be verified.	Removable.	All units are returned to their original status in EST and device with label specifying multispindle pack is left unavailable (U status set in MST). The labels on remaining devices that were chained are then checked to determine if they are valid.	Examine the E,M display after deadstart is complete to determine corrective action.

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
Read TRK into central memory.	0	TRK could not be read successfully.	All.	The following message appears on the right console display and deadstart processing halts. LENGTH OF DEVICES TRK BAD - RECOVERY OF DEVICE IMPOSSIBLE.	Type GO,SYS. or PAUSE,SYS. to proceed. If GO. is entered, the device is initialized using parameters in existing label; permanent files will be lost.
	1, 2	TRK could not be read successfully.	Removable (no active direct access files).	Error code TI and status U set in MST; device is left unavailable.	Examine the E,M display after deadstart is complete to determine corrective action.
Edit TRK to remove nonpermanent file information.	0	Edit was not successful. This could be caused by specifying a new SPP/SLF entry (flawed track) APRDECK.	Nonremovable or removable with active direct access files. All.	The following message appears on the right console display and deadstart processing halts. LENGTH OF DEVICES TRK BAD. The following message appears on the right console display and deadstart processing halts. PERM. FILE LINKAGE ERROR RECOVERY OF DEVICE IMPOSSIBLE.	Perform level 0 deadstart and initialize the device. Type GO,SYS. or PAUSE,SYS. to proceed. If GO. is entered, the device is initialized using parameters in existing label; permanent files will be lost. If SPP/SLF entry was specified, redeadstart and omit SPP/SLF entry.
Initiate deadstart sequencing job.	0, 1, 2	Deadstart sequencing job could not be initiated because I/O errors were encountered on a mass storage device.	Nonremovable.	The following message appears on the right console display and deadstart processing halts. DEADSTART SEQUENCING FAILED.	Examine the A,E; E,E; and E,H displays to determine the failing device(s). Attempt an express deadstart dump for later analysis. Turn off the failing device(s) and redeadstart. Permanent files on the failing device(s) are lost.

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
	3	Deadstart sequencing job could not be initiated because I/O errors were encountered on a mass storage device.	Nonremovable.	The following message appears on the right console display and deadstart processing halts. DEADSTART SEQUENCING FAILED.	Examine the A,E; E,E; and E,M displays to determine the failing device(s). Attempt an express deadstart dump for later analysis. Restart the level 3 recovery deadstart and enter ABORT or ABORT, E to checkpoint the device and abort. Turn off the failing device(s) and redeadstart. Permanent files on the failing device(s) are lost.
Verify all mass storage devices in system to ensure permanent file integrity.	0	Two devices in same family have the same device number or the same bits set in the device mask, or two auxiliary devices have the same packname.	System device (nonremovable with system residence).	The following message appears on the right console display and deadstart processing halts. EQuest ₁ EQuest ₂ DN. PN. CONFLICTING	Perform one of the following: - Remove one of the specified devices and redeadstart. - Redeadstart and logically turn off one of the specified devices (via EQPDECK entry).
			Non-system device.	Error code DN, DM, or PN and status U set in NST; both of the devices are left unavailable.	Examine the E,M display after deadstart is complete to determine the type of error. Either remove, or logically turn off one of the devices; the remaining device is recovered automatically.
Clear interlock data for direct access files.	0	Failure to read legal system sector. The TRT indicates track is beginning of direct access chain but first sector was not a system sector.	All.	Increment count of direct access files in error. The message appears in the dayfile in the following format. EQuest files PRESERVED FILE ERRORS. est EST ordinal of files files Number of files in error.	Examine dayfile after deadstart completes. Message described in preceding column indicates the number of direct access files in error.

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
Recovering queued files.	0	<p>Interlock data specified file was in WRITE mode and last sector for file in TRT was not an EOI sector.</p> <p>Last sector of file specified by TRT is not an EOI sector.</p>	All.	<p>Clear interlocks for file and rewrite system sector. Issue the following message to error log and system dayfile:</p> <p>LENGTH ERROR filename uindex. filename Name of file on which error was encountered. uindex User index of file owner.</p> <p>Also, increment count of direct access file errors. The following message appears in the system dayfile.</p> <p>EQest count DIRECT ACCESS FILE ERRORS.</p> <p>est EST ordinal of device.</p> <p>count Total number of length errors.</p> <p>The following message is issued to the dayfile and error log:</p> <p>QF LENGTH ERROR filename. filename Name of file being recovered.</p> <p>The following message is issued to the dayfile at the end of recovery:</p> <p>EQest count QUEUED FILE ERROR COUNT.</p> <p>est EST ordinal of device.</p> <p>count File count.</p>	<p>Examine system and error log dayfiles after deadstart completes. The number of LENGTH ERROR messages issued and the count of DIRECT ACCESS FILE ERRORS should match.</p> <p>To recover files in error:</p> <ol style="list-style-type: none"> 1. ATTACH,oldfilename. 2. PURGE,oldfilename. 3. DEFINE,newfilename with oldfilename. 4. COPY,oldfilename to newfilename. or CHANGE,oldfilename/CV. <p>The utility QREC issues messages to the system dayfile indicating the location and job name of all files in error during the requeue operation.</p> <p>The message format is:</p> <p>filename Txxxx LENGTH ERROR.</p>

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
Recovery of multi-mainframe link device.	0	Extended memory resident table containing machine IDs of active mainframes is full.	Extended memory.	The following message appears on the right console display and deadstart processing halts. MAXIMUM NUMBER MIDS ACTIVE	1. Verify that the correct machine ID is specified in the CHRDECK. 2. Run MREC on active machine for specified machine ID if not initial deadstart.
	0, 1, 2, 3	A valid label track for extended memory was not found within predetermined limits.	Extended memory.	The following message appears on the right console display and deadstart processing halts. EXTENDED MEMORY LABEL TRACK NOT FOUND	Perform a level 0 deadstart with INITIALIZE and PRESET selected if no other machines are active.
	0	Machine ID selected if CHRDECK is in use.	Extended memory.	The following message appears on the right console display and deadstart processing halts. MID CURRENTLY ACTIVE	Either rereadstart specifying a different machine ID or, if the correct machine ID was entered, run MREC on another machine.
	1, 2, 3	The machine ID as maintained in CHR was not found in the extended memory resident table. This indicates that extended memory has been changed, either abnormally or by the execution of MREC on another machine.	Extended memory.	The following message appears on the right console display and deadstart processing halts. MID UNDEFINED IN EXTENDED MEMORY.	Perform a level 0 deadstart.
	0, 1, 2, 3	Error exit was taken while attempting to execute instruction accessing extended memory.	Extended memory.	The following message appears on the right console display and deadstart processing halts. EXTENDED MEMORY READ/WRITE PARITY ERRORS.	Inform a customer engineer. Recovery is impossible.

Error Idle Recovery

Setting the device error idle status helps to prevent error propagation by inhibiting new file assignment on the device and by inhibiting any new PFM, PFLOAD, or QLOAD activity from occurring on the device. The error idle status is indicated by the EI error code on the E,M display. Error idle status is set when:

- Mass storage read or write errors occur in the catalog or permit chains.
- There is an incorrect user index in the permit sector.
- Direct or indirect access files are too long or too short.

Correcting the error idle condition requires total initialization of the affected device. This can be accomplished online for a removable device, but must be performed during a level 0 deadstart for a nonremovable device. Suggested recovery procedures follow.

Removable Device

1. Dump inactive queued files, if present on the device (QDUMP).
2. Dump permanent files (PFDUMP).
3. Perform a total initialization of the device (INITIALIZE,AL,est.). If the error idle status was caused by a mass storage error, the appropriate track should be flawed.
4. Reload inactive queued files, if applicable (QLOAD).
5. Reload permanent files (PFLOAD).

Nonremovable Device

1. Dump queued files (QDUMP).
2. Terminate any dayfiles that must be saved (DFTERM).
3. Dump permanent files (PFDUMP).
4. Idle the system and perform a level 0 deadstart.
5. Perform a total initialization of the device affected by the error idle condition (INITIALIZE,AL,est.). If the error idle status was caused by a mass storage error, the appropriate track should be flawed.
6. Reload queued files (QLOAD).
7. Reload permanent files (PFLOAD).

CMRDECK	3-2
Central Memory Descriptions	3-3
EQPDECK	3-13
Dayfile Descriptions	3-15
General Description of Equipment Assignment Entries	3-18
Clear EST Assignment Entry	3-18
Equipment Assignments: Nonmass Storage	3-21
Nonstandard Equipment EST Entry	3-21
Pseudoequipments	3-22
Unit Record Equipment EST Entry	3-23
Magnetic Tape Equipment EST Entry	3-27
Mass Storage Extended Subsystem EST Entry	3-30
Stimulator Equipment EST Entry	3-31
Network Processing Unit EST Entry	3-32
CDCNET Device Interface EST Entry	3-34
Network Access Device EST Entry	3-35
Two-Port Multiplexer EST Entry	3-36
MAP III or IV Equipment EST Entry	3-37
6683 Satellite Coupler EST Entry	3-37
CYBERPLUS Ring Port EST Entry	3-38
Optical Disk EST Entry	3-39
Equipment Assignments: Mass Storage	3-40
NOS Mass Storage Concepts	3-41
Alternate System Device	3-41
Alternate Permanent File Family	3-43
Auxiliary Device	3-44
Family Device	3-45
Link Device	3-45
Master Device	3-45
Multispindle Device	3-46
844 Expander	3-46
Nonremovable Device	3-51
Removable Device	3-51
Shared Device	3-51
System Device	3-51
Temporary File Device	3-51
Buffered Disks	3-52
Disk Array Subsystem (DAS)	3-52
Dual-Channel Access for 5830 and 9853 Subsystems	3-53
Extended Memory Overview	3-55
Physical Configuration	3-55
Logical Configuration	3-56
CM and UEM Size Specification and Determination	3-57
Simultaneous Use of Two Types of Extended Memory	3-57
Specifying Ranges of EST Ordinals	3-58
EQ — Disk Equipment EST Entry	3-58
EQ — Control Module EST Entry	3-65
EQ — Extended Memory EST Entry	3-66
Examples of EQPDECK Entries for Extended Memory	3-70

Single Mainframe Examples	3-70
Multimainframe Examples	3-73
MSAL — Mass Storage Allocation Control Entry	3-77
DOWN — Down Channel Entry	3-79
UP — Up Channel Entry	3-79
PF — Permanent Files Device Assignment Entry	3-80
SYSTEM — System Library Device Assignment Entry	3-84
ASR — Alternate System Library Device Assignment Entry	3-85
SCKP — System Checkpoint File Device Assignment Entry	3-85
FAMILY — Default Family Name Assignment Entry	3-86
REMOVE — Removable Device Assignment Entry	3-86
SHARE — Shared Device Entry	3-87
MAXSHARE — Maximum Shared Device Entry	3-88
ISHARE — Independent Shared Device Entry	3-89
LBC — Load Buffer Controllers Entry	3-90
XM — Declare Extended Memory Space for Buffers or User Access	3-92
UEMIN — UEM Equipment Initialization	3-93
ACCESS — Set Access Level Limits	3-94
THRESHOLD — Set Disk Thresholds	3-95
EQPDECK Entries Made Only During Deadstart	3-96
AUTOLOAD — Toggle Autoloading	3-96
GRENADE — Clear Unit Reservations	3-96
INITIALIZE — Initialization Entry	3-96
PRESET — Preset the Link Device Entry	3-101
PRESET — Preset the Independent Shared Device Entry	3-102
RESET — Reset Device Attributes	3-102
APRDECK	3-103
APRDECK Format	3-103
CAF — Clear All Flaw Reservations	3-104
SLF — Set Logical Flaws on Any Mass Storage Device	3-105
CLF — Clear Logical Flaws on Any Mass Storage Device	3-105
SPF — Set Physical Extended Memory Track or Disk Area Flaws	3-105
CPF — Clear Physical Extended Memory Track or Disk Area Flaws	3-106
IPRDECK	3-107
Job Control Information	3-108
Job Scheduling	3-108
Control Points	3-110
Pseudo-control Points	3-111
Number of Control Points and Pseudo-control Points Available	3-111
CPU Scheduling	3-112
Memory Control	3-113
Example of Job Control Parameters	3-114
IPRDECK Entries Made Only During Deadstart	3-120
COMLIB	3-120
CPM	3-120
CSM	3-123
DISK VALIDATION	3-123
DSD	3-124
EI	3-125
EXTENDED STACK PURGING	3-125
HARDWARE FAULT INJECTION	3-125
KEYPM	3-126

MEMORY CLEARING	3-126
MICRO	3-127
NAMIAF	3-127
PROBE	3-127
SCP	3-128
SCRSIM	3-128
SECCATS	3-129
SPC	3-129
SPD	3-130
SPL	3-130
SPW	3-130
SUBCP	3-131
TCVM	3-131
TDEN	3-132
TDTY	3-132
TMS	3-133
TMSTO	3-133
TRACE	3-134
IPRDECK Entries	3-135
Subsystems	3-135
ACS TAPE PF STAGING	3-137
AUTORESTART	3-138
CARTRIDGE PF STAGING	3-138
CLASS	3-139
CPTT	3-140
DEBUG	3-140
DELAY	3-141
DDP ROLLOUT PATH	3-142
DFPT	3-142
ENGR	3-143
FLEXIBLE PARTITIONS	3-143
FORCED FORMATTING	3-144
LOCK	3-145
LOGGING	3-145
MASTER MSE	3-146
MS VALIDATION	3-146
OQSH	3-147
PCLASS	3-147
PF VALIDATION	3-148
PRIVILEGED ANALYST MODE	3-148
PRIVILEGED RDF	3-149
QUEUE	3-149
REMOVABLE PACKS	3-151
RESIDENT RDF	3-151
SECONDARY USER COMMANDS	3-151
SECURES	3-152
SERVICE	3-153
SPINDOWN	3-158
SRST	3-158
SYSTEM DEBUG	3-159
TAPE PF STAGING	3-160
UNLOCK	3-160
USER EXTENDED MEMORY	3-160

Deadstart decks are text records that reside on the deadstart file. This section describes the following deadstart decks.

Deck	Description
CMRDECK	Central memory resident deck.
EQPDECK	Equipment deck.
APRDECK	Auxiliary mass storage parameter deck.
IPRDECK	System installation parameter deck.
LIBDECK	System library deck.

This section contains detailed information about how you can modify the released versions of the deadstart decks.

CMRDECK

The central memory resident deck (CMRDECK) resides on the deadstart file as a text record that is processed during system initialization. It contains entries defining the following types of information.

- Central memory.
- Table sizes.
- Configuration information not oriented to equipment.

The deadstart file can contain up to 100₈ CMRDECKs. Having several CMRDECKs on the same deadstart file is advantageous because one file can deadstart several configurations. You can obtain a listing of all CMRDECKs by accessing the system file SYSTEM with an ASSIGN or COMMON command, then using the T parameter on the CATALOG command (refer to the NOS Version 2 Reference Set, Volume 3 for more information concerning these commands). CMRDECKs are named CMRDnn, where nn is from 00 to 77₈.

You can modify the released settings of the CMRDECK in two ways: type a new entry when the CMRDECK is displayed during deadstart, or create a new deadstart file. The usual method of modifying a CMRDECK follows.

1. Deadstart, using the released deadstart file, and select the CMRINST and the CMRDECK to be displayed (refer to section 2, Deadstart).

CMRINST lists all valid CMRDECK entries. Default values, described in this section, are assumed if the entries do not appear in the CMRDECK being used. If either CMRDECK or CMRINST overflows two screens, you can page the display.

2. Modify the released version of CMRDECK by entering the changes or additions from the system console while either the CMRDECK or CMRINST is displayed. Each console entry supersedes the value currently specified in the CMRDECK (or the default value).

NOTE

The modified CMRDECK remains in effect only until the next deadstart is performed, except for a level 3 deadstart. That is, changes to the CMRDECK are not recovered across level 0, 1, and 2 deadstarts unless a new deadstart file is created to reflect them.

3. To expedite subsequent deadstarts, modify the CMRDECK on the deadstart file using SYSGEN (refer to the NOS Version 2 Installation Handbook).

When constructing or modifying a CMRDECK, separate parameters with commas and terminate each entry with a period. When an error exists in an entry in a CMRDECK on the deadstart file, the CMRDECK is displayed and the error is indicated. This occurs even if you do not select the display CMRDECK option.

Central Memory Descriptions

The general function of central memory description entries is to assign the amount of central memory to be used for central memory resident (CMR) and the amount to be used for job processing. The simplified relationship is: the more central memory that is assigned to tables in CMR, the less is available for job field lengths.

If you will run only a few batch jobs, fewer control points may be required. Thus, you could decrease the control point area in CMR, which requires 200₈ words per control point.

The following entries are specified in the SET program with the released default values listed.

Entry Format	Released Default Value	Significance
BSP=option.	Y	Specifies whether NOS should build the SCI parameter table. If the table is built by NOS, then one copy of SCI services MDD mode and the NOS/VE console in a dual-state environment.

NOTE

For compatibility, if NOS 2.6.1 level 700 is run in a dual-state environment with a NOS/VE system which is pre-1.3.1 level 700, then you must enter the BSP entry in order to use MDD.

option	Description
Y	If Y is specified, the SCI parameter table is built in the NOS CMR.
N	If N is specified, the SCI parameter table is not built in the NOS CMR unless the NOS console is a CC598B or CC634B. In this case, the table is always built.

CLT=number.	0	Specifies the octal number of common library table (CLT) entries allowed. The number can range from 0 to 7777 ₈ .
-------------	---	--

The first entry in the table is reserved for the system; so you should equate CLT with the number of entries you want, plus 1. For example, if you want three entries, enter CLT=4.

Common library table entries are described in the IPRDECK entry COMLIB.

Entry Format	Released Default Value	Significance
DFT=mode.	Depends on mainframe model	<p>The dedicated fault tolerance (DFT) entry specifies whether to run DFT in dedicated mode or in nondedicated mode. In dedicated mode, the DFT PP is loaded during the NOS deadstart process and DFT resides in the PP for the life of the system. Dedicated mode allows the halt-on-error hardware feature of the models 960, 990, 994, and 995 to be used. It also results in capturing the relevant data about hardware errors closer to the actual point of failure. In nondedicated mode, DFT is called to process mainframe errors when necessary. The default for models 960, 990, 994, and 995 is DFT=Y; the default for all other models is DFT=N.</p> <p>mode Description</p> <p>D If D is specified, DFT runs in dedicated mode. DFT also runs in dedicated mode if the VE CMRDECK entry is specified (other than VE=*).</p> <p>N If N is specified, DFT runs in nondedicated mode (unless the VE CMRDECK entry is specified).</p>
EJT=number.	620 ₈	<p>Specifies an octal number of entries for the executing job table (EJT). The system uses the EJT entries to keep track of executing jobs.</p> <p>The maximum value for number is 7777₈; the minimum value is 3.</p>
EMB=number.	200 ₈	<p>Specifies the number of CM words for the error message buffer (EMB). Error messages are placed in the EMB before being written to one of the system dayfiles.</p> <p>The maximum value for number is 7777₈; the minimum value is 200₈.</p>
EQP=number.	Number of CMRDECK	<p>Specifies the number of the EQPDECK to use at deadstart. The EQPDECK contains equipment assignments. Up to 100₈ EQPDECKs can exist on a deadstart file; the number can range from 00 to 77₈.</p> <p>If an EQP entry is not included in the CMRDECK, the system uses the EQPDECK with the same number as the CMRDECK being used.</p>

Entry Format	Released Default Value	Significance														
FNT=number.	23 ₈	<p>Sets the octal number of entries allowed in the system file name table (FNT). The system FNT contains the system file and all fast-attach files.</p> <p>Determine the necessary number of FNT entries by using the formula:</p> $\text{number} = \text{SY} + \text{RS} + (\text{VL} + \text{PR}) * \text{FM} + \text{SFA}$ <p>Variable Description</p> <table border="1"> <thead> <tr> <th>Variable</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>SY</td> <td>Number of system files.</td> </tr> <tr> <td>RS</td> <td>Number of resource files.</td> </tr> <tr> <td>VL</td> <td>Number of VALIDUs files per family.</td> </tr> <tr> <td>PR</td> <td>Number of PROFILa files per family.</td> </tr> <tr> <td>FM</td> <td>Number of families that can be active at any one time.</td> </tr> <tr> <td>SFA</td> <td>Site defined fast-attach files.</td> </tr> </tbody> </table> <p>For a system installed with the released defaults and with no additional files added to the system FNT by local code, the maximum number of FNT entries necessary is 201₈.</p> <p>The maximum value for number is 7777₈; the minimum value is 3.</p>	Variable	Description	SY	Number of system files.	RS	Number of resource files.	VL	Number of VALIDUs files per family.	PR	Number of PROFILa files per family.	FM	Number of families that can be active at any one time.	SFA	Site defined fast-attach files.
Variable	Description															
SY	Number of system files.															
RS	Number of resource files.															
VL	Number of VALIDUs files per family.															
PR	Number of PROFILa files per family.															
FM	Number of families that can be active at any one time.															
SFA	Site defined fast-attach files.															
FOT=number.	10 ₈	<p>Sets the number of entries in the family ordinal table (FOT). Each family is allowed one entry in the table. The first entry in the FOT is reserved for system use. The system uses family name and user index for job ownership and file routing. The family ordinal is a 6-bit number that corresponds to a particular family name. The FOT maintains the family name to family ordinal relationship. If a family is unloaded and later reloaded, it continues to use the same family ordinal.</p> <p>The maximum value for number is 100₈; the minimum value is 3. The size of the FOT need not be the same for each mainframe in a multimainframe environment.</p>														

Entry Format	Released Default Value	Significance						
INB=number.	0	Specifies a block of memory for site use. The size of the installation block can range from 0 to 7777 ₈ central memory words.						
IPD=number.	First IPRDECK on deadstart file	<p>Specifies the number of the IPRDECK to use at deadstart. The IPRDECK contains installation parameters. Up to 100₈ IPRDECKs can exist on a deadstart file; the number can range from 00 to 77₈.</p> <p>If an IPD entry is not included in the CMRDECK, the first IPRDECK on the deadstart tape is processed and is not displayed.</p>						
LDT=number.	0	<p>Specifies the number of central memory words allocated to the LID table. The minimum value is 0; the maximum value is 16610₈. The value to be specified can be determined by using this formula:</p> $(3+lid) * pid$						
<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>lid</td> <td>Maximum number of LIDs allowed per PID.</td> </tr> <tr> <td>pid</td> <td>Total number of PIDs in all networks defined on this mainframe.</td> </tr> </tbody> </table>			Parameter	Description	lid	Maximum number of LIDs allowed per PID.	pid	Total number of PIDs in all networks defined on this mainframe.
Parameter	Description							
lid	Maximum number of LIDs allowed per PID.							
pid	Total number of PIDs in all networks defined on this mainframe.							
<p>One central memory word is always allocated for the LID table header.</p> <p>As an example, assume lid=5 and pid=6. By using the formula, you come up with (3+5)*6=48. Adding one word for the LID table header makes 49, or 61₈, central memory words that should be allocated to the LID table.</p>								
LIB=number.	0	Specifies the number of the LIBDECK to use at deadstart; the number can range from 00 to 77 ₈ . LIBDECK is a directive record used by SYSEDT.						
MID=id.	AA	Specifies the 2-character machine identification (id) that is associated with the mainframe. The id characters must be alphanumeric.						

Entry Format	Released Default Value	Significance
MINCM=size.	1400 ₈	Reserves an amount of central memory to be used by NOS when running in dual state with NOS/VE. The VEMEM utility can help you determine how NOS and NOS/VE will share central memory (refer to the DUAL installation procedure in the Special Product Installation Information chapter in the NOS Version 2 Installation Handbook for more information about VEMEM). (UEM is defined in the extended memory equipment EST entry). The remaining central memory is available for use by NOS/VE. Use the VE CMRDECK entry to reserve central memory for NOS/VE if UEM is defined. The minimum and default value for size is 1400 ₈ .

Parameter	Description
size	The size of central memory to reserve in words/100 ₈ . Size is a 1- to 4-digit octal value (1- to 5-digit octal value for a model 865, 875, or CYBER 180-class machine) that restricts the actual central memory size. The value cannot be 0 and cannot exceed the total number of words of memory present in the machine. Refer to table 3-1.

Size (Octal)	CM Words (Octal)	CM Words (Decimal)
1400	140000	49152
2000	200000	65536
3000	300000	98304
4000	400000	131072
6000	600000	196608
10000	1000000	262144
14000	1400000	393216
20000	2000000	524288
24000	2400000	655360
30000	3000000	786432
34000	3400000	917504
40000	4000000	1048576
44000	4400000	1179648
50000	5000000	1310720
54000	5400000	1441792
60000	6000000	1572864
64000	6400000	1703936
70000	7000000	1835008
74000	7400000	1966080
100000	10000000	2097152

Entry Format	Released Default Value	Significance				
NAME=date line.	NETWORK OPERATING SYSTEM	Specifies the system date line, which is displayed on the system console display and on the terminal when an interactive user logs in to the system.				
<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>date line</td> <td>Alphanumeric string; must be fewer than 39 characters. If the date line is not terminated with a period, one will be appended. If the date line is 38 characters, the last character must be a period.</td> </tr> </tbody> </table>			Parameter	Description	date line	Alphanumeric string; must be fewer than 39 characters. If the date line is not terminated with a period, one will be appended. If the date line is 38 characters, the last character must be a period.
Parameter	Description					
date line	Alphanumeric string; must be fewer than 39 characters. If the date line is not terminated with a period, one will be appended. If the date line is 38 characters, the last character must be a period.					
NCP=number.	12 ₈	Sets the number of control points available for job processing. Refer to Job Control Information later in this section for a discussion of the proper number of control points to select.				
<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>number</td> <td>Number of control points available in central memory; number can be from 2 to 34₈.</td> </tr> </tbody> </table>			Parameter	Description	number	Number of control points available in central memory; number can be from 2 to 34 ₈ .
Parameter	Description					
number	Number of control points available in central memory; number can be from 2 to 34 ₈ .					

Entry Format	Released Default Value	Significance										
OPSECM=n.	0	Specifies the operating system security mode. The site security administrator supplies the value for this entry. Refer to the NOS Version 2 Security Administrator's Handbook for information on this entry.										
<table border="1"> <thead> <tr> <th>n</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0 or omitted</td> <td>Sets the system to unsecured mode.</td> </tr> <tr> <td>1</td> <td>Enables multilevel security. The values of the system access level limits can be set by the SECURES IPRDECK entry or by the console command. The SECURES console command can be used to raise or lower system access level limits.</td> </tr> <tr> <td>2</td> <td>Enables multilevel security. The values of the system access level limits are set by the SECURES IPRDECK entry. The SECURES console command can be used to raise but not to lower system access level limits.</td> </tr> <tr> <td>3</td> <td>Enables multilevel security. The values of the system access level limits are set by the SECURES IPRDECK entry only; the SECURES console command is not allowed.</td> </tr> </tbody> </table>			n	Description	0 or omitted	Sets the system to unsecured mode.	1	Enables multilevel security. The values of the system access level limits can be set by the SECURES IPRDECK entry or by the console command. The SECURES console command can be used to raise or lower system access level limits.	2	Enables multilevel security. The values of the system access level limits are set by the SECURES IPRDECK entry. The SECURES console command can be used to raise but not to lower system access level limits.	3	Enables multilevel security. The values of the system access level limits are set by the SECURES IPRDECK entry only; the SECURES console command is not allowed.
n	Description											
0 or omitted	Sets the system to unsecured mode.											
1	Enables multilevel security. The values of the system access level limits can be set by the SECURES IPRDECK entry or by the console command. The SECURES console command can be used to raise or lower system access level limits.											
2	Enables multilevel security. The values of the system access level limits are set by the SECURES IPRDECK entry. The SECURES console command can be used to raise but not to lower system access level limits.											
3	Enables multilevel security. The values of the system access level limits are set by the SECURES IPRDECK entry only; the SECURES console command is not allowed.											
PCP=number.	0	Sets the number of pseudo-control points available for job processing. Refer to Job Control Information later in this section for a discussion of the proper number of pseudo-control points to select.										
<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>number</td> <td>Number of pseudo-control points available in central memory; number can be from 2 to 34₈.</td> </tr> </tbody> </table>			Parameter	Description	number	Number of pseudo-control points available in central memory; number can be from 2 to 34 ₈ .						
Parameter	Description											
number	Number of pseudo-control points available in central memory; number can be from 2 to 34 ₈ .											

Entry Format	Released Default Value	Significance
PPU=pp ₁ , pp ₂ ,...,pp _n .	All available PPs are active	Toggles the active status of the physically available PPs ¹ except for PPs 1, 2, and 10, which are always active. Active means the PP is available for system use; inactive means it is not available for system use. The PPU entry is not in the released CMRDECK; therefore, all available PPs are active. PPU is a toggle entry: each entry changes the active status of the PPs.
Parameter Description		
pp _i	Number of the PP; from 0 to 11 ₈ and 20 ₈ to 31 ₈ . Specifying an asterisk (*) on the entry instead of pp _i toggles the status of PPs 20 ₈ to 31 ₈ .	<p>This entry may be useful if PP memory is failing or if a channel is causing problems on its associated PP. For example, the following entry deactivates PP 3 and PP 4 (assuming no other entries have been made for PP 3 and PP 4).</p> <p style="text-align: center;">PPU=3, 4.</p>
QFT=number.	620 ₈	Sets the number of entries allowed for the queued file table (QFT). The system uses the QFT to manage all files in the input and output queues. number is a 1- to 4-digit number from 3 to 7777 ₈ .

1. A PP that has been turned off by CTI is physically unavailable and cannot be turned on by the PPU entry.

Entry Format	Released Default Value	Significance
VE=size.	*	<p>For NOS and NOS/VE dual state only. Sets the size of central memory reserved for NOS/VE, in octal words divided by 1000₈. You must specify this value if NOS UEM is defined by either the XM or EQ EQPDECK entry. Otherwise, you can set VE to zero and specify the amount of memory reserved for NOS with the MINCM entry. For a list of memory sizes, refer to table 3-1.</p> <p>NOS/VE requires at least 6.5 megabytes of central memory. On a machine with more than 128MB of central memory, NOS can use only the first 128MB for CM and UEM. Memory in excess of 128MB can be used only by NOS/VE.</p> <p>The VEMEM utility can help you determine how NOS/VE and NOS will share central memory (refer to the DUAL installation procedure in the NOS 2 Installation Handbook for more information about VEMEM).</p> <p>If you do not want to run NOS/VE in dual-state mode, use VE=* or omit the VE entry.</p> <p>On a dual-CPU CYBER 180-class machine, NOS cannot use the second CPU if dual-state mode is selected by the VE entry. This is because NOS/VE does not support dual-state operation in the second processor.</p> <p>To cause NOS to use both CPUs on a dual-CPU CYBER 830 or 960-32 mainframe, omit the VE entry (or specify VE=*) in the CMRDECK.</p> <p>To cause NOS to use both CPUs on a dual-CPU CYBER 860- or 990-class mainframe, select an OS LOAD STATE of NOS-NOS/BE DUAL CPU on the Operator Intervention menu under CIP. When this OS LOAD STATE is selected, the VE entry is not accepted from the CMRDECK.</p> <p>For all other CYBER 180-class machines, NOS cannot use the second CPU.</p>

Entry Format	Released Default Value	Significance
VERSION=name.	NOS 2.7.2 774	<p>Specifies the system version that is displayed on the system console display. Specify an alphanumeric string fewer than 18 characters. If name is not terminated with a period, one will be appended. If name is 17 characters, the last character must be a period.</p> <p>The operating system version number should be in character positions 5 through 7 of the name parameter, since the value in character positions 5 through 7 is available to users through the CCL symbolic name VER.</p>

Table 3-1. Table of Central Memory Sizes

Size¹	CM Words (Octal)	CM Words (Decimal)	CM Words (Hexadecimal)	CM Megabytes²
400	400000	131072	20000	1 MB
1000	1000000	262144	40000	2 MB
2000	2000000	524288	80000	4 MB
4000	4000000	1048576	100000	8 MB
10000	10000000	2097152	200000	16 MB
20000	20000000	4194304	400000	32 MB
40000	40000000	8388608	1000000	64 MB
100000	100000000	16777216	2000000	128 MB
200000	200000000	33554432	4000000	256 MB
400000	400000000	67108864	8000000	512 MB

1. Value specified for size on the VE entry.

2. MB represents megabytes. One megabyte is 1,048,576 8-bit bytes, which is equivalent to 100,000 hexadecimal 8-bit bytes or 131,072 60-bit words.

EQPDECK

The equipment deck (EQPDECK) resides on the deadstart file as a text record that is processed during system initialization. It contains entries defining the following types of information.

- Dayfile buffer size.
- Hardware configuration.
- File residence on mass storage devices.
- Permanent file family residence.
- Multimainframe configurations.
- Software attributes associated with hardware (such as forms code).

An EQPDECK entry is described under the type of information it defines.

The deadstart file can contain up to 100₈ EQPDECKs. Having several EQPDECKs on the same deadstart file is advantageous because one file can deadstart several configurations. You can obtain a listing of all EQPDECKs by accessing the system file SYSTEM with a COMMON command, then using the T parameter on the CATALOG command (refer to the NOS Version 2 Reference Set, Volume 3 for more information concerning these commands). EQPDECKs are named EQPDnn, where nn is from 00 to 77₈.

You can modify the released settings of the EQPDECK in two ways: type a new entry when the EQPDECK is displayed during deadstart, or create a new deadstart file. The usual method of modifying an EQPDECK follows.

1. Deadstart, using the released deadstart file, and select the EQPINST and the EQPDECK to be displayed (refer to section 2, Deadstart).

EQPINST lists all valid EQPDECK entries. Default values, described in this section, are assumed if the entries do not appear in the EQPDECK being used. If either EQPDECK or EQPINST overflows two screens, you can page the display.
2. Modify the released version of EQPDECK by entering the changes or additions from the system console while either the EQPDECK or EQPINST is displayed. Each console entry supersedes the value currently specified in the EQPDECK (or the default value).

NOTE

The modified EQPDECK remains in effect only until the next level 0, 1, or 2 deadstart is performed. That is, changes to the EQPDECK are retained across a level 3 deadstart, but not across a level 0, 1, or 2 deadstart (unless a new deadstart file is created to reflect them).

3. To expedite subsequent deadstarts, modify the EQPDECK on the deadstart file using SYSGEN (refer to the NOS Version 2 Installation Handbook).

Several equipment status displays are available during deadstart (refer to section 2, Deadstart).

When constructing or modifying an EQPDECK, the following restrictions apply:

- The equipment assignment entry (EQ) must precede any other assignments for a device (such as assigning it for permanent file, system, or temporary file use). If you redefine the EQ entry (by clearing the entry and retyping a new entry or by changing the device type in the entry), reenter all other assignments for that equipment. If you need to modify an EQ entry, you need only enter the changes rather than reentering the entire entry. You do not need to reenter other assignments for that equipment. This is especially useful for changing channels or units and for reconfiguring around hardware malfunctions.

Example:

The following mass storage assignments are in the EQPDECK.

```
EQ7=DI,ST=ON,UN=60,CH=1.
PF=7,F,125,125,SYSTEM,3,200.
```

If you want to specify unit 70 rather than 60, you need only enter:

```
EQ7=DI,UN=70.
```

You do not need to modify the PF entry.

- EST ordinals range from 1 to 777₈. This range depends on the value of ESMX (refer to PPCOM parameters in the NOS Version 2 Installation Handbook). Ordinals 1 through 4 are reserved for pseudoequipment EST entries that are automatically defined by the system and cannot be changed. Ordinal 5 is reserved for extended memory if it is used as a link device in a multimainframe environment (refer to EQ – Extended Memory EST Entry later in this section). Otherwise, EST ordinal 5 can be used for any equipment type.
- The device from which you are deadstarting must be defined.
- Commas must separate parameters.
- A period must terminate each entry.
- Except where explicitly specified that controllers and/or equipment can be shared between mainframes, assume they cannot be shared. For example, NOS does not support sharing a two-channel tape controller between mainframes; nor does it support sharing mass storage controllers, except as specified in section 13, Multimainframe Operations.

When an error exists in an entry in an EQPDECK on the deadstart file, the EQPDECK is displayed and the error is indicated. This occurs even if you do not select the display EQPDECK option or if you enter GO after displaying the EQPDECK.

Dayfile Descriptions

One of the functions of dayfile description entries is to assign the amount of central memory to be used for buffers. The more central memory assigned to dayfile buffers in CMR, the less central memory available for job field lengths.

If, for example, you need a large portion of central memory to run a job, it might be advisable to decrease the size of the dayfile buffers area in CMR to accommodate that job. However, when dayfile buffers are smaller, the information stored in them is written to mass storage more often, which requires more system overhead.

The following entries are specified in the SET program with the released default values listed.

Entry Format	Released Default Value	Significance						
ACCOUNT=est,length.	400 _g	<p>Sets the residence of the account dayfile and the length of the account dayfile buffer.</p> <p>The account dayfile is an accounting record containing information such as type and quantity of resources used, and job and execution times. This account information is written to the central memory account dayfile buffer during job processing; the central memory buffer is written to mass storage when it is full. The account dayfile buffer resides in CMR in the dayfile buffer area.</p>						
		<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>est</td> <td>The 1- to 3-digit octal equipment status table (EST) ordinal of the equipment on which the account dayfile is to reside. A null value means the system uses the first system device it can find. If the existing account dayfile is recovered, the account dayfile buffer resides on that equipment and the est parameter is ignored.</td> </tr> <tr> <td>length</td> <td>The 3- or 4-digit octal length of the account dayfile buffer in CMR; must be a multiple of 100_g and less than or equal to 1000_g.</td> </tr> </tbody> </table>	Parameter	Description	est	The 1- to 3-digit octal equipment status table (EST) ordinal of the equipment on which the account dayfile is to reside. A null value means the system uses the first system device it can find. If the existing account dayfile is recovered, the account dayfile buffer resides on that equipment and the est parameter is ignored.	length	The 3- or 4-digit octal length of the account dayfile buffer in CMR; must be a multiple of 100 _g and less than or equal to 1000 _g .
Parameter	Description							
est	The 1- to 3-digit octal equipment status table (EST) ordinal of the equipment on which the account dayfile is to reside. A null value means the system uses the first system device it can find. If the existing account dayfile is recovered, the account dayfile buffer resides on that equipment and the est parameter is ignored.							
length	The 3- or 4-digit octal length of the account dayfile buffer in CMR; must be a multiple of 100 _g and less than or equal to 1000 _g .							

Entry Format	Released Default Value	Significance						
DAYFILE=est,length.	400 ₈	<p>Sets the residence of the system dayfile and the length of the system dayfile buffer.</p> <p>The dayfile buffer contains the dayfile information, which is maintained in the same way as the account dayfile buffer. It resides in CMR in the dayfile buffer area.</p> <table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>est</td> <td>The 1- to 3-digit octal EST ordinal of equipment on which the dayfile is to reside. A null value means the system uses the first system device it can find. The residence of this dayfile is normally determined by the recovery of the existing dayfile. Use this parameter if no system dayfiles are recovered.</td> </tr> <tr> <td>length</td> <td>The 3- or 4-digit octal length of the system dayfile buffer in CMR; must be a multiple of 100₈ and less than or equal to 1000₈.</td> </tr> </tbody> </table>	Parameter	Description	est	The 1- to 3-digit octal EST ordinal of equipment on which the dayfile is to reside. A null value means the system uses the first system device it can find. The residence of this dayfile is normally determined by the recovery of the existing dayfile. Use this parameter if no system dayfiles are recovered.	length	The 3- or 4-digit octal length of the system dayfile buffer in CMR; must be a multiple of 100 ₈ and less than or equal to 1000 ₈ .
Parameter	Description							
est	The 1- to 3-digit octal EST ordinal of equipment on which the dayfile is to reside. A null value means the system uses the first system device it can find. The residence of this dayfile is normally determined by the recovery of the existing dayfile. Use this parameter if no system dayfiles are recovered.							
length	The 3- or 4-digit octal length of the system dayfile buffer in CMR; must be a multiple of 100 ₈ and less than or equal to 1000 ₈ .							
ERRLOG=est,length.	100 ₈	<p>Sets the residence of the error log and the length of the error log buffer.</p> <p>The error log is a record of hardware error and status messages. This information is maintained in the same way as the account dayfile buffer.</p> <table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>est</td> <td>The 1- to 3-digit octal EST ordinal of equipment on which the error log is to reside. A null value means the system uses the first system device it can find. If the existing error log is recovered, the error log buffer resides on that equipment and the est parameter is ignored.</td> </tr> </tbody> </table>	Parameter	Description	est	The 1- to 3-digit octal EST ordinal of equipment on which the error log is to reside. A null value means the system uses the first system device it can find. If the existing error log is recovered, the error log buffer resides on that equipment and the est parameter is ignored.		
Parameter	Description							
est	The 1- to 3-digit octal EST ordinal of equipment on which the error log is to reside. A null value means the system uses the first system device it can find. If the existing error log is recovered, the error log buffer resides on that equipment and the est parameter is ignored.							

Entry Format	Released Default Value	Significance						
ERRLOG=est,length. (Continued)		<table border="1"> <thead> <tr> <th data-bbox="761 333 911 363">Parameter</th> <th data-bbox="946 333 1109 363">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="761 373 841 403">length</td> <td data-bbox="946 373 1341 495">The 3- or 4-digit octal length of the error log buffer in CMR; must be a multiple of 100₈ and less than or equal to 1000₈.</td> </tr> </tbody> </table>	Parameter	Description	length	The 3- or 4-digit octal length of the error log buffer in CMR; must be a multiple of 100 ₈ and less than or equal to 1000 ₈ .		
Parameter	Description							
length	The 3- or 4-digit octal length of the error log buffer in CMR; must be a multiple of 100 ₈ and less than or equal to 1000 ₈ .							
MAINLOG=est,length.	400 ₈	Sets the residence of the binary maintenance log and the length of the binary maintenance log buffer. The binary maintenance log is a record of hardware diagnostic information.						
		<table border="1"> <thead> <tr> <th data-bbox="761 688 911 718">Parameter</th> <th data-bbox="946 688 1109 718">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="761 728 802 758">est</td> <td data-bbox="946 728 1341 1073">The 1- to 3-digit octal EST ordinal of equipment on which the binary maintenance log is to reside. A null value means the system uses the first system device it can find. If the existing binary maintenance log is recovered, the binary maintenance log buffer resides on that equipment and the est parameter is ignored.</td> </tr> <tr> <td data-bbox="761 1089 846 1119">length</td> <td data-bbox="946 1089 1341 1243">The 3- or 4-digit octal length of the binary maintenance log buffer in CMR; must be a multiple of 100₈ and less than or equal to 1000₈.</td> </tr> </tbody> </table>	Parameter	Description	est	The 1- to 3-digit octal EST ordinal of equipment on which the binary maintenance log is to reside. A null value means the system uses the first system device it can find. If the existing binary maintenance log is recovered, the binary maintenance log buffer resides on that equipment and the est parameter is ignored.	length	The 3- or 4-digit octal length of the binary maintenance log buffer in CMR; must be a multiple of 100 ₈ and less than or equal to 1000 ₈ .
Parameter	Description							
est	The 1- to 3-digit octal EST ordinal of equipment on which the binary maintenance log is to reside. A null value means the system uses the first system device it can find. If the existing binary maintenance log is recovered, the binary maintenance log buffer resides on that equipment and the est parameter is ignored.							
length	The 3- or 4-digit octal length of the binary maintenance log buffer in CMR; must be a multiple of 100 ₈ and less than or equal to 1000 ₈ .							

General Description of Equipment Assignment Entries

The general format of the EQ entry is:

```
EQest=type,keyword1=value1,...,keywordn=valuen.
```

Keyword=value entries are order independent.

Depending on the type of equipment being defined, certain keywords are required, while others are optional, as shown in table 3-2.

You can redefine an EST entry by entering an EQ entry for the EST ordinal with a different equipment type. This action clears the original entry before creating the new entry. You can also modify an EST entry by entering an EQ entry with the same equipment type and only the keyword(s) you want changed.

Clear EST Assignment Entry

Use the following entry to clear an assignment that currently exists for an EST ordinal. Clearing the assignment does not clear flaw entries for that equipment. The EQest=0 entry is not required when you are assigning the EST entry to a different type of equipment.

```
EQest=0.
```

or

```
EQest=.
```

Parameter Description

est	EST ordinal of the equipment; est can be from 5 to 777 ₈ .
-----	---

Table 3-2. Required and Optional Keywords

Equipment Type	Required Keywords	Optional Keywords
AT	CH,UN,LS,PA,DR	ST
CC	EQ,CH	ST,IB
CM	EQ,CH	CW,IB
CP	EQ,CH	ST,ID,FC,IB
CR	EQ,CH	ST,ID,IB
CT	CH,UN	ST
DB	UN,CH	ST,AP,IB
DC	UN,CH	ST,AP,IB
DD	UN	ST,AP,IB
DE	SZ	ST,MA,ET,IB
DF	UN,CH	ST,AP,IB,AP,MC,IB
DG	UN	ST,AP,IB
DH	UN,CH	ST,AP,IB
DI	UN,CH	ST,AP,IB
DJ	UN,CH	ST,AP,IB
DK	UN,CH	ST,AP,IB
DL	UN,CH	ST,AP,IB
DM	UN,CH	ST,AP,IB
DN	EQ,UN,CH	ST,AP,IB
DP	SZ,CH	ST,MA,ET,AP,MC,IB
DQ	UN,CH	ST,AP,IB

(Continued on next page)

Table 3-2. Required and Optional Keywords

Equipment Type	Required Keywords	Optional Keywords
<i>(Continued from previous page)</i>		
DV	UN,CH	ST,AP,IB
DW	UN,CH	ST,AP,IB
EA	EQ,UN,CH	ST,AP,IB
EB	EQ,UN,CH	ST,AP,IB
EC	EQ,UN,CH	ST,AP,IB
ED	EQ,UN,CH	ST,AP,IB
EE	EQ,UN,CH	ST,AP,IB
EF	EQ,UN,CH	ST,AP,IB
EG	EQ,UN,CH	ST,AP,IB
EH	EQ,UN,CH	ST,AP,IB
LR	EQ,CH	ST,ID,FC,TN,PS,IB
LS	EQ,CH	ST,ID,FC,TN,PS,IB
LT	EQ,CH	ST,ID,FC,TN,PS,IB
LX	EQ,CH	ST,ID,FC,IB
LY	EQ,CH	ST,ID,FC,IB
MP	CH	ST,IB
MT	EQ,UN,CH,TF	ST,IB
NC	CH	ST,IB
ND	EQ,CH,PI,ND,NT	ST,IB
NP	EQ,CH,PI,ND,SA	ST,IB
NT	EQ,UN,CH,TF	ST,IB
OD	EQ,UN,CH	ST,IB
RM	CH,PT	ST,IB
RP	CH	ST,IB
SS	EQ,UN,CH	ST,IB
TT	EQ,CH	ST,PT,IB
**	WO	W1

Equipment Assignments: Nonmass Storage

The following EST entries are described in this subsection.

- Nonstandard equipment
- Pseudoequipments
- Unit record equipment
- Magnetic tape equipment
- Mass Storage Extended Subsystem equipment
- Stimulator equipment
- Network processing unit
- CDCNET device interface
- Network access device
- Two-port multiplexer
- MAP III or IV equipment
- 6683 Satellite Coupler
- CYBERPLUS ring port equipment

Nonstandard Equipment EST Entry

The nonstandard equipment EST entry allows you to define nonstandard equipment or to add site debugging modifications.

Use the following format to enter the actual octal value that is to reside at that EST ordinal.

```
EQest=**,W0=value0,W1=value1
```

Parameter Description

est	EST ordinal of the equipment; est can be from 5 to 777 ₈ .
value _i	The 1- to 20-digit octal value; this value is entered in the EST word i for the specified ordinal. The word is right-justified and zero-filled if value _i has fewer than 20 digits.

Pseudoequipments

The system automatically defines EST ordinals 0 through 4 as pseudoequipment EST entries; they cannot be used for other equipment definitions. You cannot declare, change, or remove pseudoequipments.

EST Ordinal	Description
Ordinal 0 (RD)	This ordinal reserves an EST entry and MST entry to be used for device reconfiguration.
Ordinal 1 (DS)	This ordinal represents the system display console.
Ordinal 2 (NE)	<p>The system uses 2 internally to signify that a file is assigned, but that no space exists on the device. If a read is tried, end-of-information (EOI) status occurs. If a write is attempted, the data is discarded.</p> <p>For example, you can use ordinal 2 with the permanent file utility (PFDUMP) to validate the integrity of a permanent file device, without taking the time to actually create a dump file on tape. In this case, enter:</p> <pre>X.DIS. ASSIGN, NE, TAPE. PFDUMP.</pre> <p>This causes all dump data to be discarded, even though the permanent file device is read and informative messages about the permanent file device are issued to the system console. These messages are described in the NOS Version 2 Operations Handbook.</p>
Ordinal 3 (TE)	If an association is established between file name and volume serial number with an ASSIGN, LABEL, REQUEST, or VSN command, the system automatically enters EQ3 in the file's FNT/FST entry. When a tape having the desired volume serial number is assigned to the file, the system replaces EQ3 in the file's FNT/FST entry with the EST ordinal of the tape unit on which the tape is mounted. If a file that has had the file name and volume serial number association established by a VSN command is returned prior to attempting to assign the tape equipment to the file, the FNT/FST entry is canceled.
Ordinal 4 (TT)	The system assigns to this equipment a file used for either input from or output to an interactive terminal. This allows the system to determine whether a file requires the special handling needed to accomplish terminal input/output. Byte 4, set by the NAMIAF IPRDECK entry, contains the number of network terminals.

Unit Record Equipment EST Entry

A unit record equipment EST entry defines channel connected card readers, card punches, and line printers.

Unit record equipment connected to RHF or PSU through communications ports are defined in the network configuration files rather than with individual EST entries.

```
EQest=type,ST=status,TN=tn,EQ=eq,CH=ch,ID=id,FC=fc,PS=ps,IB=ib.
```

or

```
EQest=type-P,ST=status,TN=tn,EQ=eq,CH=ch,ID=id,FC=fc,PS=ps,IB=ib.
```

Parameter Description

Parameter	Description
est	EST ordinal of the unit record equipment; from 5 to 777.
type	Unit record equipment type; NOS supports the following unit record equipments:
	type Equipment
	CR Card reader 405-3447/3649
	CP Card punch 415-3446/3644
	CP Card punch 415-30
	LR Line printer 580-12
	LS
	Line printer 580-16
	LT
	Line printer 580-20
	LX
	Line printer 5870
	LY
	Line printer 5970
P	Specifies that a 580 printer is equipped with a programmable format controller.

Parameter Description

TN=tn Print train for local batch line printer; from 1 to 7. NOS supports the following print trains.

tn	Print Train	Description
0	595-1/596-1	CDC graphic 63/64-character set.
1	595-1/596-1	CDC graphic 63/64-character set.
4	595-6/596-6	ASCII graphic 95-character set or ASCII graphic 63/64-character set.
5	595-5/596-5	ASCII graphic 63/64-character set.
6	595-6/596-6	ASCII graphic 95-character set.
7	595-6/596-6	ASCII graphic 95-character set or ASCII graphic 63/64-character set.

If you set a nonsupported print train value, tn defaults to a supported value. If you omit tn or specify 2 or 3, the actual value of tn is 1.

If an invalid external characteristic (EC)² is specified, the queued file processor cannot output the file. The following shows which files will print and which files will not print for a given print train selection.

Print Train Selected	Will Print File With Specified EC	Will Not Print File With Specified EC
0	None, B4, B6	A4, A6, A9
1	None, B4, B6	A4, A6, A9
4	None, A4, A6, A9	B4, B6
5	None, A4, A6	B4, B6, A9
6	A9	None, B4, B6, A4, A6
7	None, B4, B6, A4, A6, A9	

ST=status Specifies whether unit record equipment is available for system use; enter one of these values:

status	Description
ON	Unit record equipment is available.
OFF	Unit record equipment is ignored during system operation.

2. Refer to the NOS Version 2 Reference Set, Volume 3 for a discussion of the ROUTE command EC parameter.

Parameter Description

EQ=eq Controller number for equipment; from 0 to 7. To determine the number for most unit record equipment, look at the dial switch on the controller. To determine the equipment number for an LX printer, use the following equation:

$$\text{eq} * 2 + 1 = \text{addr}$$

where addr is the 5870 host address obtained from the 5870 installer.

CH=ch Number of channel to which unit record equipment is connected; from 0 to 13₈ and from 20₈ to 33₈.

NOTE

If you want to use a card reader to perform a coldstart on a 66x tape controller, the card reader must be available on channel 12₈, 13₈, 32₈, or 33₈.

To ensure that all printers are restored to their original states (such as 8 lines per inch and auto page eject) after a master clear has been issued, all unit record equipment should be available on dedicated channels. If it is not, printers revert to 6 lines per inch, and no auto page eject status after a master clear is issued.

ID=id The 1- or 2-digit octal numeric identifier assigned to the device; from 0 to 67₈. This id is assigned to any output created by a job. For card readers, all jobs loaded from this card reader are assigned the identifier id.

FC=fc The 2-character optional forms code assigned to a line printer or card punch. If the forms code is not present, the forms code field is cleared. The forms code must either be null (not specified) or in the range from AA to 99.

NOTE

The forms code cannot be assigned to a card reader.

Parameter	Description						
PS=ps	Paper size, L or S, for local batch line printer. Default is L.						
<table border="1"> <thead> <tr> <th>ps</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>L</td> <td>Long (11 inch) paper will be mounted.</td> </tr> <tr> <td>S</td> <td>Short (8 1/2 inch) paper will be mounted.</td> </tr> </tbody> </table>		ps	Description	L	Long (11 inch) paper will be mounted.	S	Short (8 1/2 inch) paper will be mounted.
ps	Description						
L	Long (11 inch) paper will be mounted.						
S	Short (8 1/2 inch) paper will be mounted.						
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.						

Examples:

EQ11=CR, ST=ON, EQ=4, CH=12.

EQ11=CP, EQ=5, CH=12, ST=ON.

EQ20=LR, ST=ON, CH=12, EQ=6, ID=15, FC=AA.

EQ21=CR, ID=15, ST=ON, EQ=7, CH=12.

EQ22=LT-P, TN=6, ST=ON, EQ=2, CH=12.

Magnetic Tape Equipment EST Entry

The released tape subsystem supports a maximum of 16 magnetic tape units. You can change this value with the MUNIT parameter of COMSMTX; refer to the NOS 2 Installation Handbook. As a minimum, NOS requires one 639, 667, 669, 677, 679, or 698 magnetic tape unit. The format of the entry is:

EQest=AT,ST=status,CH=ch₁/ch₂,UN=un,LS=lsm,PA=panel,DR=dr.

or

EQest=CT-n,ST=status,CH=ch₁/ch₂,UN=un.

or

EQest=MT-n,ST=status,EQ=eq,UN=un,CH=ch₁/ch₂,TF=tf,IB=ib.

or

EQest=NT-n,ST=status,EQ=eq,UN=un,CH=ch₁/ch₂,TF=tf,IB=ib.

Parameter Description

est EST ordinal of the tape unit; from 5 to 777₈. Refer to the AT, CT-n, MT-n, or NT-n parameter.

AT Equipment type:

CT-n

MT-n

or

NT-n

- AT specifies Automated Cartridge Subsystem (ACS) tape units (5682 tape units attached to a 5477 Automated Cartridge Subsystem).
- CT specifies Cartridge Tape Subsystem (CTS) tape units (standalone 5682 tape units).
- MT specifies 7-track reel tape units.
- NT specifies 9-track reel tape units.

n is the total number of tape units connected to the controller. For CTS tape units, n can be from 1 to 20₈. For 677, 679, and 698 units connected to a 7021-31/32 controller, n can be from 1 to 20₈. For 667 and 669 units connected to a 7021-21/22 controller, n can be from 1 to 10₈.

The system automatically generates n EST entries with consecutive EST ordinals beginning with the ordinal specified in the est parameter. The n units begin with the unit number specified in the UN=un parameter.

For either ACS tape units or 639 units, n should not be specified. If you specify n for 639 units, specify a value of 0 (zero).

For more information on configuring CTS and ACS tape units, refer to section 22.

Parameter Description

NOTE

To clear a CT-n, MT-n, or NT-n assignment, enter an EQest=0 entry for each unit. For example, to clear EQ50=MT-4,ST=ON,..., enter:

```
EQ50=0.
EQ51=0.
EQ52=0.
EQ53=0.
```

ST=status Indicates whether the tape unit is available for access; enter one of these values:

status Description

ON	Magnetic tape unit is available for access.
OFF	Magnetic tape unit is ignored during system operation.
DOWN	All access to the magnetic tape unit is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable, or if the unit is to be used by NOS/VE in a dual-state environment (ACS, CTS, and 9-track units only). ³

CH=ch_i Number of the channel(s) to which the tape unit is connected, from 0 to 13₈ and from 20₈ to 33₈.

A controller can be connected to one or two channels, depending on the controller model. However, a maximum of two channels can be handled regardless of the number of controllers.

UN=un Number of the lowest numbered magnetic tape unit to be processed; units must have consecutive physical unit numbers; from 0 to 7 for 667 and 669 units, or from 0 to 17₈ for 677, 679, 698, and 5682 units. The unit number is on the front of the tape unit, either as a switch or a label.

LS=lsm Index of the library storage module (LSM) containing the tape unit; from 0 to 17₈. This parameter is specified only for ACS tape units.

PA=panel Index of the LSM panel containing the tape unit; from 0 to 13₈. This parameter is specified only for ACS tape units.

3. If a magnetic tape unit is accessible from more than one mainframe, concurrent use must be disabled. For reel tapes, use one of these methods: set access switches on the tape controller so that only one mainframe can access the unit at any time or make sure that the EST entry for each shared unit is defined as ON in only one of the mainframes. For cartridge tapes, make sure that the EST entry for each shared unit is defined as ON in only one of the mainframes.

Parameter Description

DR=dr	Index of the tape unit (tape drive) in the LSM panel; from 0 to 3. The index indicates the position of a cartridge tape unit in a cartridge tape cabinet. If the panel contains only two tape units, their indexes must be 0 and 2. This parameter is specified only for ACS tape units.
EQ=eq	Controller number for the tape unit; from 0 to 7. This number is usually 0 unless the controller is switched to a different number. Ask your customer engineer for the correct number if you are unsure. This parameter is omitted from the EST entry for ACS and CTS tape units.
TF=tf	Hardware features available. This parameter specifies the following hardware characteristics of the tape unit(s) being defined; select one option. This parameter is omitted for both ACS and CTS tape units.

The following options apply to 677/679/698 units:

tf	Description
ATS	The unit(s) being defined cannot process 6250-cpi group-encoded (GE) tapes. If the system detects a unit with this capability, it automatically changes this value to ATSGE or CMTGE, indicating the availability of the GE feature.
ATSGE or CMTGE	The unit(s) being defined can process 6250-cpi GE tapes. Specify this value (ATSGE for 679; CMTGE for 698) only if a unit being defined has the capability of processing GE tapes but is down. In such a case, the system would not be able to connect to the unit to determine the availability of the GE feature.

The following options apply to 667/669 units:

tf	Description
MTS	The FCOs needed to implement the block identification feature have been installed in the 7021 controller for the unit(s) being defined. If the controller is a two-channel model, the block identification feature must have been implemented on both channels.

The following options apply to 639 units:

tf	Description
IST	Specify TF=IST.

IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.
-------	--

Mass Storage Extended Subsystem EST Entry

The Mass Storage Extended Subsystem (MSE) hardware consists of a 7991 storage module and a 7990 controller.

A mainframe that runs SSSLV in slave mode must have an MSE EST entry, even though no MSE hardware is actually configured for this mainframe. A mainframe that runs SSEXEC in master mode must have an LBC EQPDECK entry with a controlware type of N8 or M8 specified.

The MSE EST entry format is:

EQ=est=SS,ST=status,EQ=eq,UN=un,CH=ch,IB=ib.

Parameter Description

est	EST ordinal of the 7990 controller; from 10 ₈ to 777 ₈ .
SS	Indicates a 7990 controller.
ST=status	Specifies whether the 7990 controller is available for use; enter one of these values:
	status Description
	ON 7990 controller is available.
	OFF 7990 is ignored during system operation.
EQ=eq	Logical controller number.
	This number sets the controller location in the unit device table within SSEXEC's memory. Enter one of these values:
	eq Description
	0 First 7990 controller.
	2 Second 7990 controller.
	4 Third 7990 controller.
	6 Fourth 7990 controller.
UN=un	7991 storage module unit number; must be 0 (zero).
CH=ch	Number of the channel to which the 7990 controller is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ . You can specify up to two channels.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

Stimulator Equipment EST Entry

NOS IAF supports a stimulator called STIMULA. However, there are other stimulators such as PACER that can use this EST entry. (PACER is a product that can be ordered from the Special Products Library.)

The format of the EST entry used by the interactive subsystem during stimulation is:

```
EQest=TT, ST=status, EQ=eq, CH=ch, PT=lines, IB=ib.
```

Parameter Description

est	EST ordinal of the interactive stimulator; from 5 to 777 ₈ .
TT	Indicates an interactive stimulator.
ST=status	Specifies whether the interactive stimulator is available for use; enter one of these values:
	status Description
	ON Stimulator is available.
	OFF Stimulator is ignored during system operation.
EQ=eq	Number of the controller; from 0 to 7. Refer to the description of the channel parameter.
CH=ch	Number of the channel; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ . The channel/controller combination must not have any equipment attached to it.
PT=lines	Number of lines to stimulate; from 1 to 1000 ₈ . If this parameter is omitted, 100 ₈ is used.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

Network Processing Unit EST Entry

A total of 16 NPU, MDI, and MTI ESTs can be available (ST=ON) at the same time.

The format of the EST entry is:

EQest=NP, ST=status, EQ=eq, PI=pip, CH=ch, ND=node, SA=sa, IB=ib.

Parameter Description

est	EST ordinal of the NPU; from 5 to 777 ₈ .
NP	Indicates a 255x NPU.
ST=status	Specifies whether the NPU is available for use; enter one of these values:
	status Description
	ON NPU is available.
	OFF NPU is ignored during system operation.
EQ=eq	Number of the controller for the NPU; from 0 to 7.
PI=pip	Peripheral interface program index, which determines which copy of the PP driver drives this NPU; from 1 to 4. Up to four EST entries can have the same PIP index (that is, one PP can drive a total of four front ends consisting of NPUs, MDIs, or MTIs).
CH=ch	Number of the channel to which the NPU is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .
ND=node	Node number of the coupler associated with the NPU being defined; from 1 to 377 ₈ . This value is the same as the NODE parameter on the COUPLER statement in the network configuration file definition. For the procedure to assign this value, refer to the Network Definition Language Reference Manual.

NOTE

The node parameter is specified as an octal value on the NPU entry. The NODE parameter is specified on NDL statements as a decimal value.

Parameter Description

SA=sa System autostart module (SAM) flag (refer to CCP in the NAM Host Application Programming Reference Manual). This parameter is required.

sa	Description
----	-------------

ON	SAM is present on the NPU.
----	----------------------------

OFF	SAM is absent. The host attempts to load the system autostart module program (SAM-P) to the NPU if the NPU is down.
-----	---

If an NPU is channel-connected to a single host in a single-host network, the NPU does not need a SAM. However, if an NPU is channel-connected to two hosts, the NPU should have a SAM; otherwise, both hosts might attempt to reload SAM-P if the NPU is down. This requires operator intervention to resolve the problem. The NPU EST entries for both hosts should specify SA=ON.

IB=ib The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

Example:

Assume that three NPUs exist on channels 4, 5, and 6, all with controller 7. The NPUs are connected to coupler nodes 2, 8, and 11, respectively. The NPU on channel 5 has a system autostart module; the others do not. The first two NPUs are to be driven by the same PP. The EST entries for these NPUs follow.

```
EQ70=NP,ST=ON,CH=4,EQ=7,PI=1,ND=2,SA=OFF.
```

```
EQ71=NP,ST=ON,CH=5,EQ=7,PI=1,ND=10,SA=ON.
```

```
EQ72=NP,ST=ON,CH=6,EQ=7,PI=2,ND=13,SA=OFF.
```

The NDLP input for the network configuration would include the following statements (refer to the Network Definition Language Reference Manual for a complete description of these statements).

```
CPL1: COUPLERNODE=2,HNAME=HOST1.
```

```
CPL2: COUPLERNODE=8,HNAME=HOST1.
```

```
CPL3: COUPLERNODE=11,HNAME=HOST1.
```

The node parameter of the EST entry and NODE on the COUPLER statement have the same numeric values, in this case 2, 8 (10₈), and 11 (13₈).

CDCNET Device Interface EST Entry

Each mainframe device interface (MDI) or mainframe terminal interface (MTI) in a CDCNET network that is connected to a CYBER mainframe must have an EST entry. A total of 16 MDI, MTI, and NPU ESTs can be available (ST=ON) at the same time. The format of the EST entry is:

EQest=ND, ST=status, EQ=eq, PI=pip, CH=ch, ND=node, NT=node, IB=ib.

Parameter Description

est	EST ordinal of the MDI or MTI; from 5 to 777 ₈ .
ND	Indicates an MDI or MTI.
ST=status	Specifies whether the MDI or MTI is available for use; enter one of these values:
	status Description
	ON MDI or MTI is available.
	OFF MDI or MTI is ignored during system operation.
EQ=eq	Number of the controller for the MDI or MTI; from 0 to 7.
PI=pip	Peripheral interface program (PIP) index, which determines which copy of the PP driver drives this MDI or MTI; from 1 to 4. Up to four EST entries can have the same PIP index (that is, one PP can drive a total of four front ends consisting of MDIs, MTIs, or NPUs).
CH=ch	Number of the channel to which the MDI or MTI is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .
ND=node ⁴	Node number of the coupler associated with the MDI or MTI being defined; from 1 to 377 ₈ .
NT=node ⁴	Node number of the MDI or MTI; from 1 to 377 ₈ . This node number identifies the Control Data Network Architecture (CDNA) transport function.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

4. The procedure for assigning node numbers and creating network configuration file definitions is described in the Network Definition Language Reference Manual.

The node numbers assigned to an MDI or MTI must be consistent with those assigned to a 255x NPU network configuration defined through NDL. To minimize the impact on existing 255x NPU networks, assign node numbers as follows:

1. Assign the lowest node numbers to NPU coupler nodes.
2. Reserve node numbers subsequent to those assigned in step 1 for NPU nodes.
3. Assign the remaining node numbers to MDI or MTI coupler nodes.
4. Reserve one node number to identify the CDNA transport function.

Network Access Device EST Entry

The format for the network access device (NAD) EST entry is:

EQest=NC, ST=status, CH=ch, IB=ib.

Parameter	Description						
est	EST ordinal of the NAD; from 5 to 777 ₈ .						
NC	Indicates a NAD.						
ST=status	Specifies whether the NAD is available for use; enter one of these values: <table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>NAD is available.</td> </tr> <tr> <td>OFF</td> <td>NAD is ignored during system operation.</td> </tr> </tbody> </table>	status	Description	ON	NAD is available.	OFF	NAD is ignored during system operation.
status	Description						
ON	NAD is available.						
OFF	NAD is ignored during system operation.						
CH=ch	Number of the channel to which the NAD is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .						
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.						

Two-Port Multiplexer EST Entry

The two-port multiplexer (TPM) must be described in the EST to allow use of the Remote Diagnostic Facility (RDF) on models 865 and 875 and CYBER 180-class machines.

The format of the TPM EST entry is:

EQest=RM,ST=status,PT=pt,CH=ch,IB=ib.

Parameter Description

Parameter	Description
est	The 1- to 3-digit octal EST ordinal of the TPM; 5 to 777 ₈ .
RM	TPM equipment type.
ST=status	Specifies whether the TPM is available for use; enter one of these values:
	status Description
	ON TPM is available.
	OFF TPM is ignored during system operation.
PT=pt	Port number to be used by RDF; 0 or 1. RDF normally uses port 1. RDF and the system console cannot use the same port number.
CH=ch	Channel number on models 865 and 875 and CYBER 180-class machines; channel 15 is required.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

MAP III or IV Equipment EST Entry

The MAP III or IV equipment EST entry allows use of the matrix algorithm processor (MAP).

The format of the MAP III or IV EST entry is:

EQest=MP,ST=status,CH=ch,IB=ib.

Parameter Description

Parameter	Description
est	EST ordinal of MAP; from 5 to 777 ₈ .
MP	Indicates MAP equipment.
ST=status	Specifies whether the MAP is available for use; enter one of these values:
	status Description
	ON MAP is available.
	OFF MAP is ignored during system operation.
CH=ch	Number of the channel to which the MAP is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

6683 Satellite Coupler EST Entry

The 6683 satellite coupler EST entry allows use of the NOS-SCOPE 2 Station Facility. The format for the entry is:

EQest=CC,ST=status,EQ=eq,CH=ch,IB=ib.

Parameter Description

Parameter	Description
est	EST ordinal of 6683 coupler; from 5 to 777 ₈ .
CC	Indicates 6683 coupler.
ST=status	Specifies whether the 6683 coupler is available for use; enter one of these values:
	status Description
	ON 6683 coupler is available.
	OFF 6683 coupler is ignored during system operation.
EQ=eq	Controller number for equipment; from 0 to 7.
CH=ch	Number of channel to which coupler is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

CYBERPLUS Ring Port EST Entry

The CYBERPLUS ring port equipment EST entry allows use of the CYBERPLUS Subsystem. The format of the entry is:

EQest=RP,ST=status,CH=ch,IB=ib.

Parameter	Description
est	EST ordinal of the ring port; from 5 to 777 ₈ .
RP	Indicates ring port.
ST=status	Specifies whether the ring port is available for use; enter one of these values:
	status Description
	ON Ring port is available.
	OFF Ring port is ignored during system operation.
CH=ch	Number of the channel to which ring port is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ . This parameter is required.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

Optical Disk EST Entry

The optical disk EST entry allows use of optical disk units from Large Storage Configurations, Incorporated. You can have both standalone optical disk units and multiple units housed in a jukebox. The format of the entry is:

EQest=OD,ST=status,CH=ch,EQ=eq,UN=un,IB=ib.

Parameter	Description								
est	EST ordinal of the optical disk unit; from 5 to 777 ₈ .								
OD	Indicates optical disk unit.								
ST=status	Specifies whether the optical disk unit is available for use; enter one of these values:								
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>Optical disk unit is available.</td> </tr> <tr> <td>OFF</td> <td>Optical disk unit is ignored during system operation.</td> </tr> <tr> <td>DOWN</td> <td>All access to the optical disk unit is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable.</td> </tr> </tbody> </table>	status	Description	ON	Optical disk unit is available.	OFF	Optical disk unit is ignored during system operation.	DOWN	All access to the optical disk unit is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable.
status	Description								
ON	Optical disk unit is available.								
OFF	Optical disk unit is ignored during system operation.								
DOWN	All access to the optical disk unit is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable.								
CH=ch	Number of the channel to which the optical disk unit is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .								
EQ=eq	Optical disk formatter (equipment) number for this unit; from 0 to 7.								
UN=un	Optical disk control module (unit) number; from 0 to 17 ₈ . Use unit numbers 0 to 7 for standalone optical disk units. Use unit numbers 10 ₈ to 17 ₈ for optical disk units housed in a jukebox.								
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST formats in the NOS Version 2 Systems Programmer's Instant.								

Equipment Assignments: Mass Storage

This subsection describes the following EST entries and mass storage equipment assignments.

Equipment	Entry
Disk equipment	EQ
Control module	EQ
Extended memory	EQ
Mass storage allocation control	MSAL
Down channel entry	DOWN
UP channel entry	UP
Permanent files device	PF
System library device	SYSTEM
Alternate system library device	ASR
System checkpoint file device	SCKP
Default family name	FAMILY
Removable device	REMOVE
Shared device	SHARE
Independent shared device	ISHARE
Load buffer controllers	LBC
Extended memory allocation	XM
UEM equipment initialization	UEMIN
Set access level limits	ACCESS
Set disk thresholds	THRESHOLD

NOS Mass Storage Concepts

Following are descriptions that define the NOS mass storage terminology and summarize the kinds of mass storage equipment assignments that can be specified in the EQPDECK. Table 3-3 summarizes the various functions that a particular mass storage device can serve. For example, if a device is an alternate system device (listed under the Alternate System column in table 3-3), then it cannot be a system device; it can contain temporary files, direct access files, and indirect access files; it can be a master device or a nonmaster device; it cannot be removable; it can be either an auxiliary device or a family device; and it can be a shared device or a link device.

Alternate System Device

Whereas a system device contains all the routines in the system library, an alternate system device contains copies of selected system library routines. The ASR entry in the EQPDECK (refer to ASR – Alternate System Library Device Assignment Entry later in this section) specifies which mass storage devices are to be alternate system devices; the *AD LIBDECK entry on the deadstart file specifies which system library routines are to reside on these alternate system devices. During system processing, the routines on the alternate system device are used instead of the ones on the system device.

This feature allows each routine in the system library to reside on the mass storage device that is most appropriate to the routine's use. For example, instead of using an 844 system device, a routine that is frequently used could use extended memory, which has a faster transfer rate, as an alternate system device.

NOTE

Use care when specifying alternate system devices. A mistake could result in no alternate system device defined. For example, suppose you specify the following in your EQPDECK:

```
ASR=7 .
ASR=11 .
```

and you specify the following in the corresponding LIBDECK:

```
*AD, 5, ABS/COMPASS
*AD, 11, PP/1SJ
```

Since the system could not match equipment numbers between the EQPDECK and the corresponding LIBDECK, the system library routine COMPASS will not be on equipment 7 or 5. However, 1SJ will be on equipment 11 as specified.

Table 3-3. Mass Storage Functions for Various Device Types

Other Possible Functions	Type¹	Type²	Type³	Type⁴	Type⁵	Type⁶	Type⁷	Type⁸	Type⁹
Alternate system device	No	–	Yes	Yes	Yes	Yes	Yes	Yes	Yes
System device	–	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Contain temporary files	Yes	Yes	–	Yes	Yes	Yes	Yes	Yes	Yes
Contain direct access files	Yes	Yes	Yes	–	Yes	Yes	Yes	Yes	Yes
Contain indirect access files	Yes	Yes	Yes	Yes	–	Yes	Yes	Yes	Yes
Master device	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Nonmaster device	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes
Removable device	No	No	No	Yes	Yes	Yes	No	Yes	No
Nonremovable device	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Auxiliary device	Yes	Yes	Yes	Yes	Yes	–	No	Yes	Yes
Family device	Yes	Yes	Yes	Yes	Yes	No	–	Yes	Yes
Shared device	Yes	Yes	Yes	Yes	Yes	Yes	Yes	–	–
Link device	Yes	Yes	Yes	Yes	Yes	Yes	Yes	–	–

1. System.
2. Alternate system.
3. Containing temporary files.
4. Containing direct access files.
5. Containing indirect access files.
6. Auxiliary.
7. Default family.
8. Shared.
9. Link.

Alternate Permanent File Family

More than one permanent file family can exist in a system: one default permanent file family and one or more alternate permanent file families. One permanent file family is defined as the default family by the EQPDECK FAMILY entry (refer to FAMILY – Default Family Name Assignment Entry later in this section). If another system's permanent file family is introduced, it is an alternate permanent file family; it can be added without interrupting the default permanent file family's operation.

This is a useful feature if a site has more than one system or has groups of installations. If one system fails, its permanent files can be accessed from another system.

As an example, a site with two systems might run with the mass storage configuration shown in table 3-4.

Table 3-4. Mass Storage Configuration for Two Systems at One Site

System	Ordinal	Device	Spindles	Access Used	Contents
X	7	844	2	A	Direct access files
Y	7	844	2	B	Direct access files

If system Y became inoperative, the B access could be connected to system X. This could be done without interrupting operations of system X.

The EQPDECK entries in system X would be:

Entries	Comments
EQ7=DI,ST=ON,UN=0/1,CH=0/3.	Defines access A.
EQ10=DI,ST=OFF,UN=0/1,CH=1/4.	Defines access B.
REMOVE=10.	Allows introduction of access B into system X during operation.

The EQPDECK entries in system Y would be:

Entries	Comments
EQ7=DI,ST=ON,UN=0/1,CH=1/4.	Defines access B.
EQ10=DI,ST=OFF,UN=0/1,CH=0/3.	Defines access A.
REMOVE=10.	Allows introduction of access A into system Y during operation.

To allow for introduction of an alternate permanent file family:

1. Define the equipment to be introduced or removed in the EQPDECKs for both systems (refer to the previous example).
2. Specify all of the equipment that may be introduced or removed during system processing as removable.
3. When you want to introduce the equipment into a system, use the ON operator command to indicate that the equipment that was set to the OFF position in the system in operation is now available. This introduces the alternate permanent file family.
4. Enable the validation files for the family by entering X.ISF(FM=familyname). Refer to section 20, System File Initialization.

Auxiliary Device

An auxiliary device is a mass storage device that is not part of a family. It is a supplementary permanent file storage device, which may be privately owned (PRIVATE) or may be shared by many users (PUBLIC). An auxiliary device resides on either a removable or nonremovable device. On the permanent file entry (PF) for an auxiliary device (for both a removable device and a fixed device without packs), a pack name is specified instead of a family name.

Refer to the NOS Version 2 Reference Set, Volume 3, for additional information about private and public auxiliary devices.

As an example, four 844 spindles to be used as a public auxiliary device could be defined as follows:

```
EQ6=DI , ST=ON , UN=1-4 , CH=2 .
PF=6 , X , name .
```

Private auxiliary devices can be created only after the system is up and running. An operator can make a public device a private device by entering the INITIALIZE command with the UN and TY=X parameters (refer to section 8, K-Display Utilities).

Family Device

A family device is a mass storage device that is part of a family. It can be either a removable device or a nonremovable device. The only difference between the two is that a nonremovable device containing permanent files can also contain a copy of the system library and/or temporary files. Refer to Alternate Permanent File Family earlier in this section.

On the PF entry, the family name is important if two systems' permanent files are to run on the same system. A user can only use one family of permanent files at one time; if the user does not specify one, the default FAMILY entry is used.

A family device can contain direct and indirect access files. These files are defined in the NOS Version 2 Reference Set, Volume 3. The files that are allowed are set by the device mask and secondary mask on the PF entry.

Link Device

Either extended core storage or extended semiconductor memory or STORNET is the medium through which several computer systems are linked to form a multimainframe operating environment (shared MMF). The link device contains the information necessary to manage the mass storage that can be shared by more than one mainframe. For a description of shared mass storage, refer to SHARE – Shared Device Entry and ISHARE – Independent Shared Device Entry later in this section.

Master Device

The master device contains all of the permanent file catalog entries, indirect access files, and file permits for a specific user. If permanent file access is required, the user's master device must be available on the system, unless all access is to be to an auxiliary device. The user index and family name uniquely describe a user's master device.

Each master device is organized into five logical sections.

1. Allocation information.

A master device, like all mass storage devices, maintains device labels and track reservation tables (TRTs).

The device label contains information describing the device, such as family name and user mask, as well as locations of permit and catalog information and indirect access files. Refer to the INITIALIZE – Initialization Entry later in this section.

The TRT is the key to allocating information on the master device and to describing the physical layout of data on the device. Refer to APRDECK later in this section.

2. Catalog information.

Catalog entries are used to determine the locations and attributes of permanent files. The catalogs for a master device are allocated to contain catalog entries for a specific group of user indexes. A particular catalog track may contain entries for many users, the number depending upon the number of catalog tracks defined for the device. The user index provides the mechanism for differentiating between user's files on a particular catalog track.

3. Permit information.

Users can explicitly or implicitly allow other users to access their permanent files. Refer to the PERMIT command in the NOS Version 2 Reference Set, Volume 3. Information describing the permission for all permanent files is in the permit file. Catalog entries contain a relative sector address within this permit file for permissions that have been granted for the file.

4. Indirect access files.

The master device contains all of the user's indirect access files. These files can be accessed by commands that generate working copies for manipulation by the user.

5. Direct access files.

Direct access files can reside either on the master device or on another device in the family, depending on the device masks specified on the PF entries. Direct access files are files that can be accessed at their location on mass storage. A working copy is not generated, so any updates or alterations made to the files are permanent.

Multispindle Device

To accommodate files that are larger than one device, you can specify multispindle device assignments. Up to eight spindles of 844 disk drives or up to three spindles of 885 disk drives can be included in the equipment definition of one logical device, when the device is first defined. All spindles must be available for access whenever the device is accessed.

Multispindle devices are treated as one logical device, having a track size equal to n times the single-spindle track size (n is the number of spindles in the device). The tracks of an n -spindle device are broken down into n equally-sized segments, each having a length equal to the single-spindle track size. Each segment is contained on a different physical unit.

844 Expander

A nonexpanded controller can have up to eight disk drives connected to it. Each of the connection paths is called a port and is identified by a port number ranging from 0 to 7. An expander (10304 extender) is a hardware device that can be connected between controllers and 844 disk drives to increase the number of disk drives that each controller can access.

The expander can be used only with 844-21 drives, although all equipment definitions and equipment driving software support the 64-drive addressing scheme for both 844-21 (DI/DK) and 844-41 (DJ/DL) type equipment.

Each expander consists of either two or four expansion elements. An expansion element connects to a single controller port and forms a connection path from that port to from one through eight disk drives. The connection paths between an expansion element and the eight possible disk drives are called ranks and are identified by a rank number ranging from 0 to 7. Two expanders with four expansion elements each can be connected to a single controller to allow that controller to access a maximum of 64 disk drives. Each expansion element, however, is logically independent and, as such, could be connected to any port of any controller.

A single controller maximum configuration can be visualized as an 8- by 8-square checkerboard with each square representing one of 64 disk drives (figure 3-1).

		Controller Port							
		0	1	2	3	4	5	6	7
Expander Rank	0	00	01	02	03	04	05	06	07
	1	10	11	12	13	14	15	16	17
	2	20	21	22	23	24	25	26	27
	3	30	31	32	33	34	35	36	37
	4	40	41	42	43	44	45	46	47
	5	50	51	52	53	54	55	56	57
	6	60	61	62	63	64	65	66	67
	7	70	71	72	73	74	75	76	77

Unit Numbers (octal)
 Port Digit
 Rank Digit

Figure 3-1. Expander Addressing Map

A column of squares in figure 3-1 represents all the drives that are accessed through a single controller port. A row of squares represents all the drives that have the same expander rank. Each disk drive that can be accessed by the controller is addressed by a 6-bit unit number. The rightmost 3 bits of this unit number select to which of the eight controller ports the drive is connected. The leftmost 3 bits of the unit number select to which of the eight ranks on an expansion element the drive is connected. This unit number is specified as a 2-digit octal number in the mass storage equipment EST entry (refer to EQ - Disk Equipment EST Entry later in this section). The right digit (port digit) of the unit number is the port number and the left digit (rank digit) is the rank of the unit in the particular expansion element.

If two disk drives are vertically adjacent on the expander addressing map (figure 3-1), their unit numbers are considered to be vertically ordered; that is, both drives are connected to the same expansion element, both have the same port number, and their rank numbers differ by one (refer to example 1).

If two disk drives are horizontally adjacent on the expander addressing map (figure 3-1), their unit numbers are considered to be horizontally ordered; that is, both drives have the same rank number, and their port numbers differ by one (refer to example 2). The special case of rank numbers of 0 for two horizontally adjacent drives is equivalent to the definition of consecutive unit numbers for other equipment.

All drives connected to a controller, either directly or through an expansion element, are supported as single-unit or multiunit logical devices. Unit numbers can range from 0 to 77_8 , rather than from 0 to 7, as for other equipment. Thus, a maximum of sixty-four 844 disk drives connected to a single controller can be addressed. However, a maximum of eight units can be specified per multiunit device. In addition, all units of a multiunit device must be connected to the same channel and, therefore, to the same controller.

Figure 3-2 illustrates a configuration in which two expansion elements and 20 disk drives are connected through one controller. An expansion element with eight drives is connected to port 0, an expansion element with six drives is connected to port 1, and six drives are connected to six ports (ports 2 through 7). Each disk drive is shown as a square with its appropriate unit number inside. This configuration is used in the following three examples to illustrate multiunit device assignments. The controller is assumed to be connected to channel 1. Refer to EQ – Disk Equipment EST Entry later in this section for specific information on assigning these devices.

		Controller Port							
		0	1	2	3	4	5	6	7
Expander Rank	0	00	01	02	03	04	05	06	07
	1	10	11						
	2	20	21						
	3	30	31						
	4	40	41						
	5	50	51						
	6	60							
	7	70							

Figure 3-2. 844 Expander Configuration With 20 Drives

NOTE

The following examples illustrate multiunit device assignment of devices. For a description of the mass storage EST entry in the EQPDECK, refer to EQ – Disk Equipment EST Entry later in this section.

Example 1:

Figure 3-3 illustrates a possible configuration for a three-unit vertically adjacent multiunit device. This device could be assigned in the EQPDECK, specifying EST ordinal 6, as:

EQ6=DI, ST=ON, UN=50/60/70, CH=1.




Figure 3-3. Vertically Adjacent Three-Unit Device

Example 2:

Figure 3-4 illustrates a possible configuration for a two-unit horizontally adjacent multiunit device. This device could be assigned in the EQPDECK, specifying EST ordinal 7, as:

EQ7=DI, ST=ON, CH=1, UN=40-41.

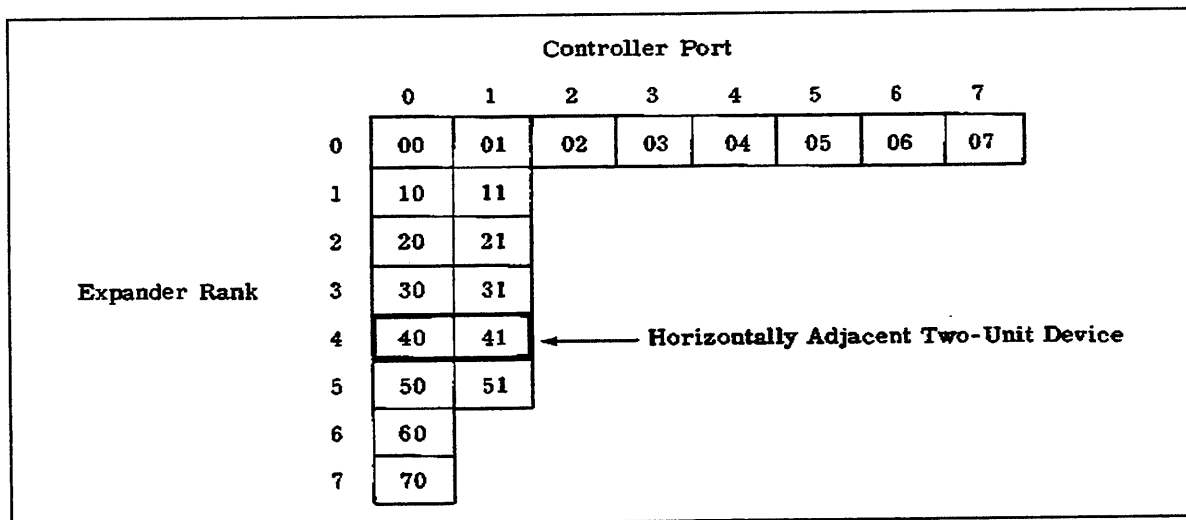


Figure 3-4. Horizontally Adjacent Two-Unit Device

Example 3:

Figure 3-5 illustrates a possible configuration of 20 disk drives into seven devices. These devices could be assigned in the EQPDECK as follows:

```
EQ6=DI,ST=ON,UN=60,CH=1.
EQ7=DI,ST=ON,UN=70,CH=1.
EQ10=DI,ST=ON,UN=50-51,CH=1.
EQ11=DI,ST=ON,UN=10/20/30/40,CH=1.
EQ12=DI,ST=ON,UN,11/21/31/41,CH=1.
EQ13=DI,ST=ON,UN=0-3,CH=1.
EQ14=DI,ST=ON,UN=4-7,CH=1.
```

Devices on ordinals 6 and 7 are defined as single units. These units may be initialized online into a two-unit device if they are also defined as removable.

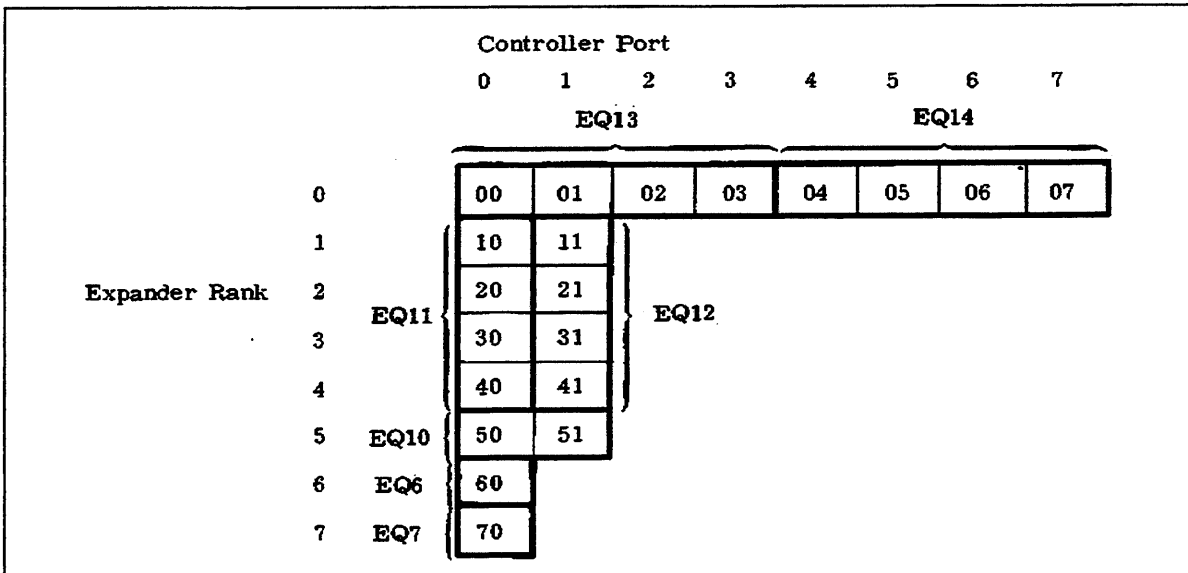


Figure 3-5. Hardware Configured into Seven Devices

Nonremovable Device

A nonremovable device cannot be physically removed during system operation. It can contain a copy of the system library, which means it is a system device; it can also be available for temporary files. It may or may not contain permanent files.

Removable Device

A removable device can be logically or physically added or removed during system operation without causing system malfunction.

A device is specified as removable with the REMOVE entry in the EQPDECK. During deadstart, a removable device is recovered just as is any other mass storage device, if the status is on. If the device is not available, then the status is displayed for the operator (E,M display).

Removable devices can contain permanent files but cannot contain the system library or temporary files, because a device containing active files (such as temporary or library files) cannot be removed from the system. A removable device can be either an auxiliary device or an alternate permanent file family device.

Shared Device

A shared device contains permanent files that can be accessed by more than one mainframe. To have these permanent files accessible to the mainframe, the device must be defined as shared in the mainframe. Refer to SHARE – Shared Device Entry and ISHARE – Independent Shared Device Entry later in this section.

A shared device can be removable. However, when unloading a shared device, it must be in global unload status before you remove it. Refer to the DSD commands UNLOAD and MOUNT in section 5 for more information concerning removable shared devices.

System Device

The system device is a nonremovable device on which the system library resides. It can also contain permanent and temporary files.

Temporary File Device

The temporary file device is a nonremovable device on which the following temporary system files reside:

- Library files
- Local files
- Queued files
- Rollout files
- System files
- Timed/event rollout files

Buffered Disks

Buffered disks are those disks for which the system provides an extended memory data buffer. The following buffered disks are supported on the mainframes indicated:

- The 819 Disk Storage Subsystem is available only on model 176. The buffer is kept in LCME.
- The 885-42 Disk Storage Subsystem is available on all models except CYBER 180-class machines and model 176. The buffer is kept in ESM.
- The 895 Disk Storage Subsystem is available on all CYBER 180-class machines except models 810, 815, 825, and 830. The buffer is kept in UEM.
- The 887, 9853, and 5830 Disk Storage Subsystems are available only on CYBER 860-, 960-, and 990-class mainframes with I4A IOU. The buffer is kept in UEM.
- Buffered disks are not available on models 810, 815, 825, and 830.

Using extended memory as a data buffer provides the following capabilities.

- The system treats the buffer as a disk cache so that multiple requests for a particular disk data block can potentially be satisfied by doing only one disk read.
- The system buffers data to maintain maximum transfer rates regardless of the user's buffer size.

To maintain maximum transfer rates, a CPU program that resides in CPUMTR controls the buffered disk I/O request processing. The system maintains the standard disk I/O interface to the user's programs.

The buffered disk error reporting process logs read/write errors in the binary maintenance log. An unrecovered read/write error is also logged in the system error log and an error message appears at the system control point. The binary maintenance log must be processed by the Hardware Performance Analyzer (HPA) to get detailed information concerning buffered device errors.

The amount of I/O buffer space in extended memory is important to system performance. Refer to XM – Declare Extended Memory Space for Buffers or User Access later in this section.

There may be situations when a particular user job runs slower with buffered disks. This happens if the job issues CIO reads when a buffer is full or issues CIO writes when a buffer is empty. (Both cases indicate inefficient programming techniques.) You can speed up the job by entering the DSD command ENABLE,LOGGING. The system will perform faster I/O and issue warning messages to the job dayfile.

Disk Array Subsystem (DAS)

The 5380 Disk Array Subsystem (DAS) is a disk storage subsystem that provides great storage capacity and operational speed. It consists of multiple disks in a storage cabinet. You can configure from one to three disks as a logical device; logical devices of two or more disks increase the storage capacity, operational speed, and/or data integrity.

The only disk units you can install in DAS are the 5832 Solid State Disk (SSD), the 5833 Sabre, and the 5838 Elite. You can configure SSD units as single units or as pairs of units driven in parallel. You can configure 5833 Sabre units in four ways:

- As a single unit
- As a single unit with another unit used for parity
- As a pair of units driven in parallel
- As a pair of units driven in parallel with another unit used for parity

You can configure 5838 Elite units in two ways:

- As a single unit
- As a single unit with another unit used for parity

Parallel units increase the storage capacity and operational speed. A parity unit ensures data integrity if a unit in the logical device fails.

When you configure a device with a parity unit, the system redundantly records data on both the parity unit and the nonparity unit(s). If a unit fails, the system continues to access data on the remaining unit(s). After the failing unit is repaired and/or replaced and you enter a RESTORE,est command, the system automatically restores data to the unit. While the data is being restored, users have access to the data on the device, although the effective transfer rate is somewhat degraded until the restoration is complete.

Dual-Channel Access for 5830 and 9853 Subsystems

Dual-channel access has several advantages:

- Increased reliability—access to the disks can continue even when one channel/controller is down.
- Increased performance—the system automatically distributes I/O requests to the two channels to maximize effective throughput.

To configure dual-channel access for the 5830 and 9853 disk subsystems, the channel addresses within the controller or control module must be set correctly. Refer to the appropriate controller reference manual for a description of how to set channel addresses for each subsystem.

For dual-channel access, both channels are connected to each controller, going from one controller to the next in a daisy-chain fashion. A request issued on a channel contains a channel address and is processed by one controller or the other based on the channel addresses set within the controller for that channel.

The system sets the channel address in the request based on the equipment controller number for the disk device as defined by the EQ keyword on the EQPDECK entry for that device. Because only one EQ value is specified for each device, the channel address contained in an I/O request for a dual-channel access device is the same regardless of the channel on which the request is issued.

To support this convention and to ensure that a controller does not receive requests on more than one channel at the same time, the channel addresses within the controller must be set in the following way:

EQPDECK

- Set the same channel address value for the first channel of the first controller and the second channel of the second controller.
- Set the remaining channel addresses (the second channel of the first controller and the first channel of the second controller) to some other value. You can use any value that does not conflict with the address used for the active channel ports for this pair of controllers or with the addresses of any additional controllers connected in daisy-chain fashion.

The following example illustrates a configuration for dual-channel access for a 5830 disk subsystem.

In this example, the logical equipment entry for a 5830 device is defined with the following EQPDECK entry:

```
EQ010=EC,ST=ON,EQ=0,UN=0,CH=1A/2A.
```

The 5830 subsystem is physically cabled as in figure 3-6.

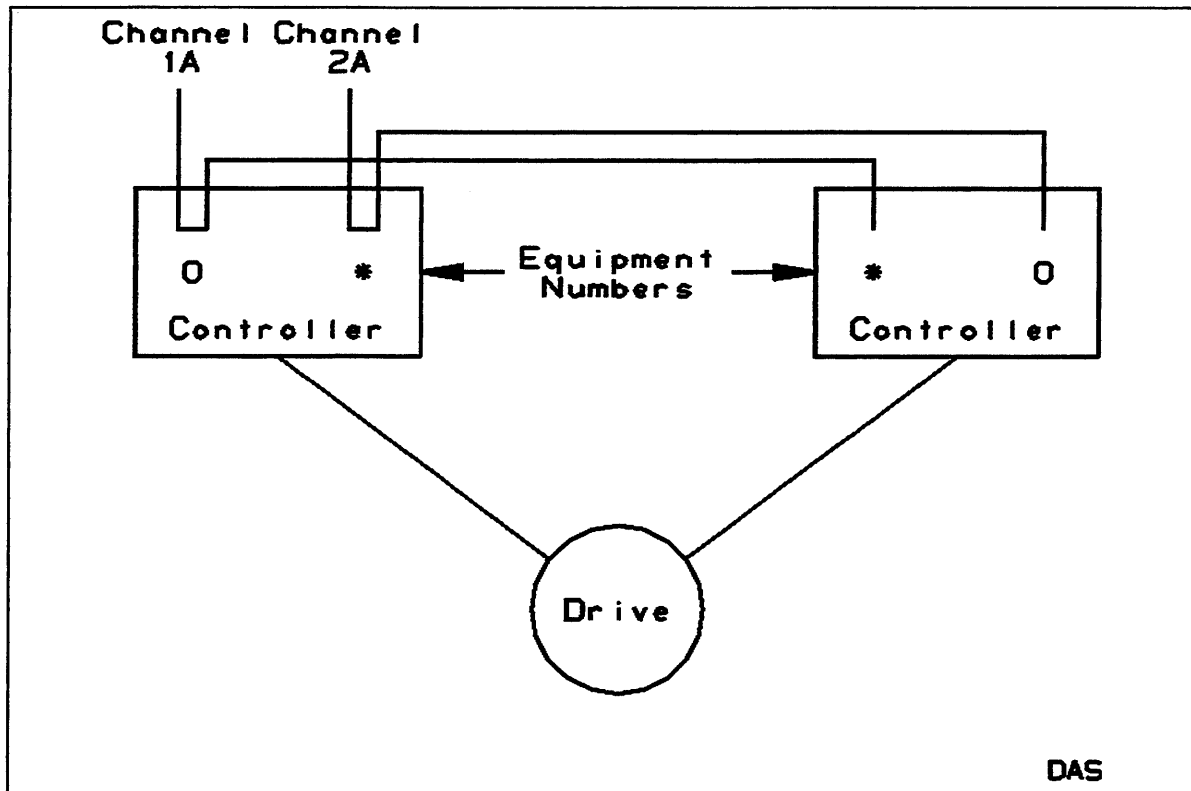


Figure 3-6. Dual-Channel Access to 5830 Drives

This configuration correctly configures only one channel per controller. The equipment number defined at each asterisk (*) can be any value other than zero (in this example) and does not appear in the EQPDECK.

Because two channels are defined, a separate driver is initiated for each channel. An I/O request for this disk can be processed by either driver, depending only on which driver sees the request first. Except for error recovery, NOS does not initiate a request on one access and complete that request on the alternate access. Error recovery can cause a disk request to be processed on the alternate channel if the request cannot be successfully completed on the original channel.

Extended Memory Overview

Extended memory (EM) refers to the five types of additional memory used by NOS to supplement central memory space and connect mainframes. This extended memory can be used to store programs and data, but not to execute programs. Shown below are the five types of extended memory and the mainframes that support each type. Also shown are the maximum EM size supported by each mainframe, how each mainframe connects to the extended memory, and what buffered devices are supported for each EM type.

Extended Memory	Mainframe	Maximum Size	Mainframe Connection	Buffered Device⁵
ECS	All	2MW	CPU port, ⁶ DDP	—
ESM	All	16MW	CPU port, ⁶ LSP	885-42
LCME	176	2MW	CPU port	819
STORNET	All	4MW	LSP	—
UEM	815,825,835,865,875	2MW	No connection, part of mainframe's main memory	895, 887,
	810,830	8MW		9853,
	840,845,850,855,860, 870,910	16MW		583x (DAS)
	990,995	8MW		

Physical Configuration

Extended core storage (ECS) is an external memory available for all mainframes. ECS connects to the mainframe using either a CPU port⁶ or a distributive data path (DDP) through a peripheral processor.

Extended semiconductor memory (ESM) is an external memory available for all mainframes. ESM connects to the mainframe using either a CPU port⁶ or a low-speed port (LSP) through a peripheral processor.

Large central memory extended (LCME) is an external memory available only for model 176 mainframes. LCME connects to the model 176 using a central processor unit (CPU) port.

STORNET is an external memory available for all CYBER 170 and CYBER 180 mainframes. STORNET connects to the mainframe using a low-speed port (LSP) through a peripheral processor.

5. The buffered devices shown are not necessarily supported on all of the mainframe types indicated.

6. A CPU port connection is not available on CYBER 180-class machines. CPU port on model 176 machines can only connect to LCME.

Unified extended memory (UEM) is a part of the mainframe's main memory and is separated from central memory by a partition defined at deadstart time by the site.

All extended memory except STORNET is accessed by special instructions. Instructions exist for multiple word transfers (FORTRAN: Level 1 variables; COMPASS: RE/WE instructions) and single word transfers (FORTRAN: Level 2 variables; COMPASS: RX/WX instructions). The multiple word transfers are available on all mainframe types; the single word transfers are available only with LCME and UEM on models 176, 865, 875, and CYBER 180-class machines.

Logical Configuration

Extended memory can be used for several purposes. The physical characteristics of the various extended memory types differ, which affects the ways in which extended memory can best be used. The possibilities are:

- To connect mainframes in a multimainframe environment (ESM, ECS, and STORNET).
- To store system routines and data files (all EM types).
- To supplement user central memory space (all EM types except STORNET).
- To provide a cache buffer for buffered I/O devices (LCME, ESM, and UEM).

ESM, ECS, or STORNET can be used as a link device to interconnect mainframes in a multimainframe environment. UEM cannot be used as the link device since it is part of a mainframe's main memory. For additional information concerning multimainframe operations, refer to section 13, Multimainframe Operations.

Extended memory can be used as a mass storage device. The following system files can be made extended memory resident as opposed to using central memory or disk space.

System Files	Description
Operating system	The alternate system residency EQPDECK entry ASR defines a device to be an alternate system device. LIBDECK directives specify which system library routines are to be placed on the alternate system device.
Temporary and rollout	The EQPDECK entry MSAL,T=est. or MSAL,R=est., where est is the EST ordinal of the EM device, causes the system to use extended memory as a mass storage device for temporary files or rollout files.
Secondary rollout	The EQPDECK entry MSAL,S=est., where est is the EST ordinal of the EM device, causes jobs that are less than a specified size to be rolled out to extended memory. The IPRDECK entry SRST defines the size of the rollout files allowed in extended memory.

System files should be placed in extended memory only if a performance improvement can be gained. Whether performance improves or not depends on your production environment, including memory and input/output requirements, available central and extended memory, and available disk units and channels.

The amount of memory available to a user job can be increased by defining user-accessible extended memory. NOS allows user data structures to reside in memory external to the job's field length. (Refer to COMPASS and FORTRAN reference manuals.)

A portion of extended memory is reserved for users by the XM EQPDECK entry. The user's memory is reserved in allocation units dependent on the amount of user-accessible extended memory defined at deadstart. The minimum allocation unit (user EM block size [M=million] UEBS) varies as follows. Service limits are defined for the user in the IPRDECK in UEBS units.

At least	But less than	Allocation unit (UEBS)
1000 words	1M words	1000 _g
1M words	2M words	2000 _g
2M words	4M words	4000 _g
4M words	8M words	10000 _g
8M words	16M words	20000 _g

This means that if you define 3 million words of user-accessible extended memory and a user executes RFL(EC=1), that user will be assigned 4000_g words of extended memory.

Extended memory provides a cache buffer for buffered I/O devices. Space is allocated by the system for buffered I/O with the XM EQPDECK entry. In most cases, system performance improves as the amount of extended memory defined for buffered I/O increases. However, ample space must be available for other uses.

CM and UEM Size Specification and Determination

When UEM is defined, its size is specified on the EQ EQPDECK entry for UEM. For models 865 and 875, the remaining memory is used for central memory. On CYBER 180-class machines, the remaining memory is divided between native mode (NOS/VE or EI) and central memory. If running in dual-state mode, the VE CMRDECK entry specifies the memory for NOS/VE. If running NOS only, central memory size equals the total memory size minus the EM size minus the EI size.

Simultaneous Use of Two Types of Extended Memory

NOS allows the simultaneous use of two types of extended memory in the following cases.

- Models 865 and 875 can access both UEM and ESM/ECS simultaneously. The recommended use of this combination of extended memory is to define ESM/ECS for system file residence and/or as a link device in a multimainframe environment and to define UEM for user-accessible extended memory. Since the main memory size for models 865 and 875 is limited to two million words, central memory space is given up for any UEM defined. Also, since ESM/ECS is an external device, fewer memory conflicts result when it is accessed rather than UEM.

In this case, an EQ entry is required in the EQPDECK for ESM/ECS as well as ASR and MSAL entries, as desired. User access is restricted to UEM via the XM entry when the EM parameter is specified.

- CYBER 180-class machines can also access both UEM and ESM/ECS/STORNET simultaneously. The purpose of the ESM/ECS/STORNET access is to allow connection of CYBER 180-class machines into a multimainframe environment. In this case, ESM/ECS is defined in the EQPDECK as the link device and UEM is defined for system files, I/O buffers, and user-accessible extended memory.

An EQ EQPDECK entry is required for both ESM/ECS/STORNET, but is not required for UEM. The ESM/ECS/STORNET EQ entry defines the extended memory as a link device via the LSP/DDP. The UEM EQ, if present, entry allows the use of extended memory as a mass storage device. An XM entry allocates extended memory for buffered I/O and user access.

Specifying Ranges of EST Ordinals

Several EQPDECK entries allow you to specify multiple EST ordinals in the same entry. For example, instead of entering

```
REMOVE=6,7,10,11,12.
```

you can enter:

```
REMOVE=6-12.
```

Ranges of ordinals can be mixed with individual ordinals, such as:

```
REMOVE=6,10-16,23.
```

Empty ordinals within a range are ignored. For example, specifying a range of 6 through 12 will ignore an unassigned ordinal 7. However, inappropriate devices within a range or specified individually (such as specifying a printer as a removable device) are diagnosed as errors.

EQ — Disk Equipment EST Entry

The purpose of the mass storage equipment (EQ) entries is to describe all mass storage peripheral equipment. This includes disk equipment and extended memory equipment. NOS requires that at least 6 million words of mass storage be available.

There can be up to 200 logical mass storage devices, and therefore, up to 200 mass storage EST entries (this number does not include ordinals 0 through 4, which are reserved for specific uses). An entry, however, can refer to more than one physical unit. For example, two 885 spindles can be defined as either two logical devices with two EQ entries or as one logical device with one EQ entry.

A unit is a dual-access unit if it is accessed by one mainframe through two different controller-channel access routes. To define a unit as a dual-access unit with its EQ entry, specify two channel parameters. The channels should be from two controllers. Only one channel of a dual-channel access controller is recommended for use on a single mainframe, since using both channels of the controller results in a performance degradation rather than an improvement. Therefore, if both channel accesses of a controller are physically connected to the same mainframe, you should define only one of them on an EQ entry.

NOTE

A device's EQ entry must precede any of the following entries for that device: ASR, MSAL, REMOVE, PF, SYSTEM, FAMILY, INITIALIZE, SHARE, ISHARE, ACCESS, and THRESHOLD. If you redefine a device's EQ entry, then you must also redefine those entries.

The format of the EQ entry for disk units is:

EQest=type,ST=status,UN=units,CH=ch₁pt₁/ch₂pt₂,AP=ap,IB=ib,EQ=eq.

Parameter	Description																																																						
est	EST ordinal of the disk units; from 5 to 777 _g .																																																						
type	Equipment type. The following disk units are supported by NOS.																																																						
	<table border="1"> <thead> <tr> <th>type</th> <th>Equipment</th> <th>Number of Units</th> </tr> </thead> <tbody> <tr> <td>DB⁷</td> <td>885-42, 7155-401</td> <td>1 to 3</td> </tr> <tr> <td>DC⁸</td> <td>895, 7165</td> <td>1 to 2</td> </tr> <tr> <td>DD⁹</td> <td>834, 7255-1</td> <td>1 to 8</td> </tr> <tr> <td>DF¹⁰</td> <td>887 (4K sector)</td> <td>1 to 3</td> </tr> <tr> <td>DG⁹</td> <td>836, 7255-1</td> <td>1 to 3</td> </tr> <tr> <td>DH¹⁰</td> <td>887 (16K sector)</td> <td>1 to 2</td> </tr> <tr> <td>DI</td> <td>844-21, 7054/7154 (half track)</td> <td>1 to 8</td> </tr> <tr> <td>DJ</td> <td>844-41/44, 7054/7154, 7155-1 (half track)</td> <td>1 to 8</td> </tr> <tr> <td>DK</td> <td>844-21, 7154 (full track)</td> <td>1 to 8</td> </tr> <tr> <td>DL</td> <td>844-41/44, 7154, 7155-1 (full track)</td> <td>1 to 8</td> </tr> <tr> <td>DM</td> <td>885-11/12, 7155-1 (half track)</td> <td>1 to 3</td> </tr> <tr> <td>DN¹⁰</td> <td>9853</td> <td>1</td> </tr> <tr> <td>DQ</td> <td>885-11/12, 7155-1 (full track)</td> <td>1 to 3</td> </tr> <tr> <td>DV¹¹</td> <td>819 (single-density)</td> <td>1</td> </tr> <tr> <td>DW¹¹</td> <td>819 (double-density)</td> <td>1</td> </tr> <tr> <td>EA¹⁰</td> <td>5380 Disk Array Subsystem; one 5832 SSD drive running in serial mode.</td> <td>1 to 8</td> </tr> <tr> <td>EB¹⁰</td> <td>5380 Disk Array Subsystem; two 5832 SSD drives running in parallel mode.</td> <td>1 to 6</td> </tr> </tbody> </table>	type	Equipment	Number of Units	DB ⁷	885-42, 7155-401	1 to 3	DC ⁸	895, 7165	1 to 2	DD ⁹	834, 7255-1	1 to 8	DF ¹⁰	887 (4K sector)	1 to 3	DG ⁹	836, 7255-1	1 to 3	DH ¹⁰	887 (16K sector)	1 to 2	DI	844-21, 7054/7154 (half track)	1 to 8	DJ	844-41/44, 7054/7154, 7155-1 (half track)	1 to 8	DK	844-21, 7154 (full track)	1 to 8	DL	844-41/44, 7154, 7155-1 (full track)	1 to 8	DM	885-11/12, 7155-1 (half track)	1 to 3	DN ¹⁰	9853	1	DQ	885-11/12, 7155-1 (full track)	1 to 3	DV ¹¹	819 (single-density)	1	DW ¹¹	819 (double-density)	1	EA ¹⁰	5380 Disk Array Subsystem; one 5832 SSD drive running in serial mode.	1 to 8	EB ¹⁰	5380 Disk Array Subsystem; two 5832 SSD drives running in parallel mode.	1 to 6
type	Equipment	Number of Units																																																					
DB ⁷	885-42, 7155-401	1 to 3																																																					
DC ⁸	895, 7165	1 to 2																																																					
DD ⁹	834, 7255-1	1 to 8																																																					
DF ¹⁰	887 (4K sector)	1 to 3																																																					
DG ⁹	836, 7255-1	1 to 3																																																					
DH ¹⁰	887 (16K sector)	1 to 2																																																					
DI	844-21, 7054/7154 (half track)	1 to 8																																																					
DJ	844-41/44, 7054/7154, 7155-1 (half track)	1 to 8																																																					
DK	844-21, 7154 (full track)	1 to 8																																																					
DL	844-41/44, 7154, 7155-1 (full track)	1 to 8																																																					
DM	885-11/12, 7155-1 (half track)	1 to 3																																																					
DN ¹⁰	9853	1																																																					
DQ	885-11/12, 7155-1 (full track)	1 to 3																																																					
DV ¹¹	819 (single-density)	1																																																					
DW ¹¹	819 (double-density)	1																																																					
EA ¹⁰	5380 Disk Array Subsystem; one 5832 SSD drive running in serial mode.	1 to 8																																																					
EB ¹⁰	5380 Disk Array Subsystem; two 5832 SSD drives running in parallel mode.	1 to 6																																																					

7. Not applicable for CYBER 180-class machines.

8. Applicable for CYBER 180-class machines except models 810, 815, 825, and 830. You must also specify the UEM EQ entry and XM entry to define the I/O buffer area in UEM.

9. Applicable for models 810, 815, 825, and 830 only. You must also specify the control module EQ entry.

10. Applicable only to CYBER 860-, 960-, and 990-class mainframes with I4A IOU.

11. Applicable to model 176 only.

Parameter	Description		
type (Continued)	Equipment type.		
	type	Equipment	Number of Units
	EC ¹²	5380 Disk Array Subsystem; one 5833 Sabre drive running in serial mode.	1 to 2
	ED ¹²	5380 Disk Array Subsystem; two 5833 Sabre drives: one for data and one for parity.	1 to 2
	EE ¹²	5380 Disk Array Subsystem; two 5833 Sabre drives running in parallel mode.	1
	EF ¹²	5380 Disk Array Subsystem; three 5833 Sabre drives: two for data and one for parity.	1
	EG ¹²	5380 Disk Array Subsystem; one 5838 Elite drive running in serial mode.	1
	EH ¹²	5380 Disk Array Subsystem; two 5838 Elite drives: one for data and one for parity.	1
	When running NOS/VE in dual state with NOS, disks DI, DJ, DK, DL, DM, and DQ must be defined in the NOS EQPDECK if NOS/VE is to use them.		

¹². Applicable only to CYBER 860-, 960-, and 990-class mainframes with I4A IOU.

Parameter	Description										
ST=status	Specifies whether or not the equipment is available for access; enter one of the following values:										
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DOWN</td> <td>All access to the equipment is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable, or if the equipment is to be used by NOS/VE in a dual-state environment (only applies to disks DI, DJ, DK, DL, DM, and DQ).</td> </tr> <tr> <td>IDLE</td> <td>New files are not assigned to a device with IDLE status unless no suitable alternative device exists, but you may continue to access files already on the device as if the device had a status of ON.</td> </tr> <tr> <td>OFF</td> <td>No user jobs can access the equipment; however, system utilities are permitted to access the device, so it can be dumped or loaded.</td> </tr> <tr> <td>ON</td> <td>The specified equipment is generally available.</td> </tr> </tbody> </table>	status	Description	DOWN	All access to the equipment is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable, or if the equipment is to be used by NOS/VE in a dual-state environment (only applies to disks DI, DJ, DK, DL, DM, and DQ).	IDLE	New files are not assigned to a device with IDLE status unless no suitable alternative device exists, but you may continue to access files already on the device as if the device had a status of ON.	OFF	No user jobs can access the equipment; however, system utilities are permitted to access the device, so it can be dumped or loaded.	ON	The specified equipment is generally available.
status	Description										
DOWN	All access to the equipment is inhibited for the operating system and user jobs. Specify DOWN if the equipment is malfunctioning and access is not desirable, or if the equipment is to be used by NOS/VE in a dual-state environment (only applies to disks DI, DJ, DK, DL, DM, and DQ).										
IDLE	New files are not assigned to a device with IDLE status unless no suitable alternative device exists, but you may continue to access files already on the device as if the device had a status of ON.										
OFF	No user jobs can access the equipment; however, system utilities are permitted to access the device, so it can be dumped or loaded.										
ON	The specified equipment is generally available.										

If the equipment is removable and is not available at deadstart, the system determines it is unavailable, even if its EQ status entry specifies ON. If INITIALIZE is entered, the equipment is not initialized until it is set to ON status. During system operation, the operator can initiate access to this device by entering the ON command.

If the equipment can be used with either one of two different systems (removable devices, not dual access), define the status of the EQ entry as ON in the system to which it is currently available for access; define the status of the EQ entry as OFF in the system to which it is not currently available for access.

Use the ST=DOWN parameter to prevent NOS from using disks that are to be used by NOS/VE. This applies only to disks DI, DJ, DK, DL, DM and DQ (these disks must be defined in the NOS EQPDECK if NOS/VE is to use them).

Parameter	Description
UN=units	<p>Defines unit number(s). Unit numbers are configuration dependent.</p> <p>For 834 or 836 units, unit numbers are generated as follows:</p> $\text{unit number} = (\text{control module physical equipment number}) * 10_8 + (834 \text{ or } 836 \text{ disk physical unit number})$ <p>For example, in UN=10 the 1 is from the EQ=1 parameter in the control module entry and the 0 is the physical unit number of the disk drive.</p> <p>For 819 units, unit numbers are either from 0 to 3, if on FLPP channels 2 and 3, or from 4 to 7, if on FLPP channels 6 and 7.</p> <p>For 844 units, unit numbers are from 0 to 77₈.</p> <p>For 885 units, unit numbers are from 40₈ to 77₈.</p> <p>For 887 units, unit numbers are from 0 to 7.</p> <p>For 895 units, unit numbers have the following format:</p> $\text{unit number} = \text{shxxxx}$ <p>where:</p> <p>s = storage control address (0 or 1).</p> <p>h = head of string controller address (0 or 1).</p> <p>xxxx = spindle number (0 to 17₈).</p> <p>For 583x units of type EA, EC, and EG, unit numbers are from 0 to 37₈. For 583x units of type EB, ED, EE, EF, and EH, unit numbers are from 0 to 7; the unit number specified is the unit number of the first unit of the device (that is, the unit in string zero). Refer to the 5830 Disk Array Subsystem Configuration Guide for more information on configuring 583x devices.</p> <p>For 9853 units, unit numbers are from 0 to 7.</p> <p>You can define the 834, 836, 844, 885, 887, 895, and 583x physical units with a separate EQ entry for each unit, or, if more contiguous storage is needed than is possible with one unit, you can define more than one physical unit as one logical device with one EQ entry. If the EQ entry defines more than one unit of a multispindle device, specify unit numbers as UN=un₁-un_n, if units are contiguous, or as UN=un₁/un₂/.../un_n, if units are noncontiguous. For 834, 836, and 385x units, all units must be on the same control module (controller).</p> <p>For example, three 885 units (0, 1, and 2) to be accessed as three units are defined as follows:</p> <pre>EQest=DM,ST=status,UN=0,CH=ch,AP=ap. EQest=DM,ST=status,UN=1,CH=ch,AP=ap. EQest=DM,ST=status,UN=2,CH=ch,AP=ap.</pre>

Parameter	Description								
UN=units	<p><i>(Continued)</i></p> <p>Three 885 units (0, 1, and 2) to be accessed as one logical unit are defined as follows:</p> <p>EQest=DM, ST=status, UN=0-2, CH=ch, AP=ap.</p> <p>or</p> <p>EQest=DM, ST=status, UN=0/1/2, CH=ch, AP=ap.</p> <p>An advantage to accessing the three units as one logical unit is that less space is used in CMR. A disadvantage is that if any of the units malfunctions or is destroyed, all units are affected.</p>								
CH=ch ₁	<p>For all devices except 819, 834, 836, 887, 583x, and 9853 the number of the channel or channels to which the controller is connected; from 0 to 13₈ and from 20₈ to 33₈.</p> <p>For 819 disks, the channel pair for input/output(I/O) multiplexer (primary access); enter one of these values:</p> <table border="1"> <thead> <tr> <th>ch₁</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>2</td> <td>Channels 2 and 3.</td> </tr> <tr> <td>4</td> <td>Channels 4 and 5.</td> </tr> <tr> <td>6</td> <td>Channels 6 and 7.</td> </tr> </tbody> </table> <p>For 834 and 836 disks, the channel number is specified on the control module EQ entry.</p> <p>For 887, 583x, and 9853 disks, the concurrent channel number from 0 to 11₈. The first and second channel numbers must be different.</p> <p>For an 895 disk, either a normal or concurrent channel; A concurrent channel is specified by Cch₁, or Cch₂. For example, CH=02/C03 defines the channel pair as nonconcurrent channel 02 and concurrent channel 03. Use of a concurrent channel frees up two nonconcurrent PPs for each channel that is moved to the concurrent I/O subsystem (CIO). If an 895 is used to deadstart, one channel access to the 895 must remain on a noncurrent channel.</p>	ch ₁	Description	2	Channels 2 and 3.	4	Channels 4 and 5.	6	Channels 6 and 7.
ch ₁	Description								
2	Channels 2 and 3.								
4	Channels 4 and 5.								
6	Channels 6 and 7.								
pt ₁	<p>For 887, 583x, and 9853 disks only, the channel port (A or B) for the first channel. Default is port A.</p>								

Parameter	Description
/ch₂	<p>Indicates dual-access unit; connect ch₁ and ch₂ to different controllers. The system balances channel activity on dual-access units.</p> <p>For 819 disks, the secondary access channel pair; one of the values described for ch₁.</p> <p>For 834 and 836 disks, dual access is indicated on the control module EQ entry.</p> <p>For 887, 583x, and 9853 disks, the concurrent channel number for the second channel; from 0 to 11₈. The first and second channel numbers must be different.</p>
pt₂	For 887, 583x, and 9853 disks only, the channel port (A or B) for the second access channel. Default is port A.
AP=ap	The 1- or 2-digit octal number that indicates which APRDECK to use. If AP=ap is omitted, the first APRDECK, APRD00, is assumed.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.
EQ=eq	Control module (controller) number for the 583x or 9853 disk. Control module numbers are from 0 to 7.

EQ — Control Module EST Entry

An 834 or 836 Disk Storage Subsystem consists of from one to eight 834 or from one to three 836 disk units and the corresponding control modules. To be operational, both the disk units and the control modules must be defined in the EQPDECK. The EQ entry for the 834 or 836 disk units is described under EQ – Disk Equipment EST Entry earlier in this section. The description of the EQ entry for a control module follows.

A control module for the 834 or 836 Disk Storage Subsystem is a controller that drives up to four 834 or three 836 disks and that interfaces with the operating system using one or two 7255-1 adapters (one per channel). The control module EQ entry is used to:

- Specify the channels to be used to access the 834 or 836 drives defined on the control module. Channel entries are not allowed on the 834 or 836 disk unit EQ entries.
- Specify what level of controlware is to be loaded into the control module during deadstart or when you use a LOADBC command.
- Allow maintenance access to a control module without affecting other control modules on the same channel(s).

The format of the control module entry is:

```
EQest=CM,EQ=eq,CH=ch1/ch2,CW=cw,IB=ib.
```

Parameter	Description										
est	EST ordinal of the control module; from 5 to 777 ₈ .										
EQ=eq	Physical equipment number of the control module; from 0 to 7. Each control module must be assigned a unique equipment number.										
CH=ch ₁	Number of the channel to which the control module is connected; from 0 to 13 ₈ and from 20 ₈ to 33 ₈ .										
/ch ₂	Indicates a dual-access control module; from 1 to 13 ₈ and from 20 ₈ to 33 ₈ .										
CW=cw	Indicates the control module type and whether controlware is to be installed in the control module during deadstart. The control module load sequence takes at least 15 seconds for each control module. It is advised that you load controlware only when necessary. cw can be one of these values:										
<table border="1"> <thead> <tr> <th>cw</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>CM</td> <td>Control module for 834 disk. Install controlware (default).</td> </tr> <tr> <td>NCM</td> <td>Control module for 834 disk. Do not install controlware.</td> </tr> <tr> <td>C2</td> <td>Control module for 836 disk. Install controlware.</td> </tr> <tr> <td>NC2</td> <td>Control module for 836 disk. Do not install controlware.</td> </tr> </tbody> </table>		cw	Description	CM	Control module for 834 disk. Install controlware (default).	NCM	Control module for 834 disk. Do not install controlware.	C2	Control module for 836 disk. Install controlware.	NC2	Control module for 836 disk. Do not install controlware.
cw	Description										
CM	Control module for 834 disk. Install controlware (default).										
NCM	Control module for 834 disk. Do not install controlware.										
C2	Control module for 836 disk. Install controlware.										
NC2	Control module for 836 disk. Do not install controlware.										
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.										

EQ — Extended Memory EST Entry

NOS supports the following types of extended memory.

- Large Central Memory Extended (LCME)
- Extended Core Storage (ECS)
- Extended Semiconductor Memory (ESM)
- Unified Extended Memory (UEM)
- STORNET

Refer to the Extended Memory Overview earlier in this section for an explanation of the various physical and logical configurations using extended memory. Also refer to the Examples of EQPDECK Entries for Extended Memory to see examples of how to use EQPDECK entries to define extended memory for both single mainframe and multimainframe configurations.

You can use extended memory as an alternate system residency for often accessed system routines. The faster access may result in a performance improvement. However, you should use caution when placing system routines in UEM to avoid possible performance degradation instead of an improvement. This can happen because a portion of central memory is reserved for UEM and therefore that portion is unavailable to users. Also, accessing the system routines that reside in UEM may require additional CPU overhead. Use the ACPD and PROBE utilities (refer to TRACER/PROBE Utilities, section 21) to determine the impact on system resources when using UEM.

NOTE

If extended memory is not included in the hardware configuration, do not make an extended memory EST entry.

The format of the entry is:

EQest=type, ST=status, MA=mode, ET=xmem/ddp/nc, SZ=size, CH=ch₁/ch₂, AP=ap, MC=mc, IB=ib.

Parameter Description

est	EST ordinal of the extended memory equipment; from 5 to 777 ₈ . If you are using extended memory as a link device in a multimainframe environment, est must be ordinal 5. Otherwise, extended memory can be defined as any EST ordinal in the given range.
type	Extended memory equipment type; enter one of these values:
	type Description
DE	DDP or LSP is not available.
DP	DDP or LSP is available (not applicable for UEM).

Parameter Description

ST=status Specifies whether extended memory is available for access; enter one of these values:

status	Description
---------------	--------------------

ON	Extended memory is available.
OFF	Extended memory is ignored during system operation.

MA=mode Maintenance mode (not applicable for UEM). If you specify MA=ON, online extended memory diagnostics are allowed to reference the half of extended memory that is placed in maintenance mode at the controller. The other half of extended memory is available to the system. The size of available physical extended memory is divided by 2 at deadstart.

When you initially place an extended memory device in maintenance mode, all mainframes using the extended memory must initialize it. When you place ECS in maintenance mode, you must also make the PRESET entry for multimainframe operation. Refer to INITIALIZE and PRESET later in this section. If you omit the MA=mode parameter, the default is MA=OFF.

If you use part of extended memory in maintenance mode and you have ESM, you must also use the MC=mc parameter to define which maintenance port to use.

ET=xmem/ ddp/nc Type of large memory and ddp. If you specify ddp, you must also specify xmem. If this parameter is not specified, E1/D1 is assumed.

xmem	Description
-------------	--------------------

E1	ECS I for all mainframes.
E2	ECS II for all mainframes.
LE	LCME only for model 176.
EM	UEM for models 865 and 875 and CYBER 180-class machines. Type must be DE. The system ensures that the sum of memory words specified by the MINCM CMRDECK entry and specified by this entry for UEM is present.
ES	ESM for all mainframes.
SN	STORNET for all CYBER 170 and CYBER 180 mainframes.

ddp	Description
------------	--------------------

D1	DC135 DDP. Default for xmem of E1 or E2.
D2	DC145 (parity enhanced) DDP, or ESM low-speed port (LSP). Default for xmem of ES and SN.

nc	Description
-----------	--------------------

NC	Indicates type DP does not have a CPU coupler. NC is not valid for xmem of DE. Default for xmem of SN.
----	--

Parameter Description

SZ=size Size of extended memory in words/1000₈. This value must not be larger than the size of the physical memory that is present.

- For ECS II and LCME, size can range from a minimum of 10₈ to a maximum of 10000₈.
- For STORNET, size can range from a minimum of 10₈ to a maximum of 20000₈.
- For ESM, size can range from a minimum of 10₈ to a maximum of 100000₈.
- For UEM, size can range from a minimum of 10₈ to a maximum of (100000₈-CM size).

NOTE

Since CIP uses a small amount of CM, you may need to make allowances for CIP when you enter a value for size.

- For ECS I, size is one of the values in the following table:

Size (Octal)	ECS I Available	Number of Banks
400	125K ¹³	1
1000	250K	2
2000	500K	4
4000	1000K	8
10000	2000K	16

The following table shows the value in EM words (both octal and decimal) and megabytes for a given value of size:

Size (Octal)	EM Words (Octal)	EM Words (Decimal)	EM Megabytes
10	10000	4096	
400	400000	131072	1 MB
1000	1000000	262144	2 MB
2000	2000000	524288	4 MB
4000	4000000	1048576	8 MB
10000	10000000	2097152	16 MB
20000	20000000	4194304	32 MB
40000	40000000	8388608	64 MB
100000	100000000	16777216	128 MB

13. K is 1000 60-bit words.

Parameter	Description
CH=ch ₁ /ch ₂	<p>Numbers of the channels to which the DDP or low speed port is connected; from 0 to 13₈ and from 20₈ to 33₈.</p> <p>If the equipment type is DE, do not specify a channel parameter. If a channel parameter is specified, the system recognizes the DE entry as a DP entry.</p> <p>If the equipment type is DP, specify either one or two channels. The second channel cannot be 0 (zero). If a DDP or low speed port is present, the loading of CPU programs residing in ECS or ESM still occurs via the CPU. A DDP or low speed port must be connected to a channel by itself.</p>
AP=ap	The 1- or 2-digit octal number that indicates which APRDECK to use. If you omit AP=ap, the first APRDECK, APRD00, is assumed.
MC=mc	Number of the channel to which the maintenance port is connected; from 1 to 13 ₈ and from 20 ₈ to 33 ₈ . Channel 0 cannot be used for the maintenance port. This channel is for ESM only and must be specified for proper reporting of ESM errors.
IB=ib	The 1- to 4-digit octal value; this value is entered in the installation byte for the specified EST ordinal. Refer to EST Formats in the NOS Version 2 Systems Programmer's Instant.

Examples of EQPDECK Entries for Extended Memory

The following examples show the EQPDECK entries that are necessary to define extended memory in various configurations.

Single Mainframe Examples

The EST ordinal of extended memory equipment ranges from 5 to 777₈. For single mainframe configurations, extended memory can be defined as any EST ordinal in the given range. For these examples, the machine identifier of the mainframe is assumed to be AA.

Example 1:

This example shows a CYBER 170 mainframe with ESM connected via the high-speed port. ESM can be used as a mass storage device, user-accessible extended memory, and for 885-42 disk I/O buffers.

EQ6=DP, ET=ES, ST=ON, SZ=10000, CH=32.	(Define 2 million words of ESM.)
MSAL, S=6.	(Define ESM as secondary rollout device.)
ASR=6.	(Define ESM as alternate system residency device.)
XM=AA, 3000, 1000.	(Allocate 3000000 octal words for 885-42 I/O buffers and 1000000 octal words for user-accessible extended memory.)

Example 2:

This example shows a CYBER 180-class machine with part of its mainframe memory defined as UEM, which can be used as a mass storage device, for user-accessible extended memory, and for 887, 895, or 9853 disk I/O buffers.

EQ7=DE, ST=ON, ET=EM, SZ=2000.	(Define 1/2 million words of memory as UEM.)
XM=AA, 1000, 200.	(Allocate 1000000 octal words for 895, 887, or 9853 I/O buffers and 200000 octal words for user-accessible extended memory.)

NOTE

If you are running a dual-state system, the VE CMRDECK entry affects the amount of memory available for NOS and thus for UEM.

Example 3:

This example shows a CYBER 180-class machine where UEM is used only for user-accessible extended memory, not as a mass storage device or for I/O buffers. No EQ entry for UEM is necessary in this case.

XM=AA, , 200 .

(Allocate 200000 octal words for user-accessible extended memory.)

Example 4:

This example shows a CYBER 170 model 865 or 875 mainframe with no ESM. Part of the mainframe memory is defined as UEM and is used as a mass storage device and for user-accessible extended memory. UEM cannot be used for 885-42 disk I/O buffers.

EQ11=DE, ST=ON, ET=EM, SZ=2000 .

(Define 1/2 million words of UEM.)

MSAL, S=11 .

(Define UEM as secondary rollout device.)

ASR=11 .

(Define UEM as alternate system residency device.)

XM=AA, , 200 .

(Allocate 200000 octal words for user-accessible extended memory.)

Example 5:

This example shows a CYBER 170 model 865 or 875 mainframe with no ESM (as in example 4) except that UEM is used only for user-accessible extended memory and not as a mass storage device. No EQ entry for UEM is necessary in this case.

XM=AA, , 200 .

(Allocate 200000 octal words for user-accessible extended memory.)

EQPDECK

Example 6:

This example shows a CYBER 170 model 865 or 875 mainframe with ESM connected via the high-speed port. ESM is used as a mass storage device, for 885-42 I/O buffers, and for user-accessible extended memory. No UEM is defined.

EQ10=DP,ST=ON,ET=ES,SZ=10000,CH=7.	(Define 2 million words of ESM.)
MSAL,S=10.	(Define ESM as secondary rollout device.)
ASR=10.	(Define ESM as alternate system residency device.)
XM=AA,3000,1000.	(Allocate 3000000 octal words for 885-42 I/O buffers and 1000000 octal words for user-accessible extended memory in ESM.)

Example 7:

This example shows a CYBER 170 model 865 or 875 mainframe with ESM connected via the high-speed port. ESM is used as a mass storage device and for 885-42 disk I/O buffers. Part of the mainframe's memory is defined as UEM and used for user-accessible extended memory.

You have the option on CYBER 170 model 865 and 875 mainframes to define user-extended memory either in ESM or UEM. If ESM is present, as indicated by an EQ entry for it, then user-accessible extended memory resides in ESM by default. However, you can force user-accessible extended memory to reside in UEM with the EM parameter on the XM entry. 885-42 disk I/O buffers must reside in ESM.

EQ10=DP,ST=ON,ET=ES,SZ=10000,CH=7.	(Define 2 million words of ESM.)
MSAL,S=10.	(Define ESM as secondary rollout device.)
ASR=10.	(Define ESM as alternate system residency device.)
XM=AA,3000,1000,EM.	(Allocate 3000000 octal words for 885-42 I/O buffers in ESM and 1000000 octal words for user-accessible extended memory in UEM.)

Multimainframe Examples

For all multimainframe configurations, the extended memory link device must be defined as the equipment with EST ordinal 5. The machine identifiers of the connected mainframes are assumed to be AA and AB. The following examples do not show the SHARE EQPDECK entries that are required to indicate multimainframe operation.

Example 1:

This example shows two CYBER 170 mainframes linked to ESM via high speed ports. In addition to being the link device, ESM can also be used as a mass storage device, for 885-42 disk I/O buffers, and for user-accessible extended memory.

EQ5=DP, ST=ON, ET=ES, SZ=10000, CH=5.	(Define 2 million words of ESM.)
XM=AA, 1000, 200.	(Allocate 1000000 octal words for 885-42 I/O buffers and 200000 octal words for user-accessible extended memory in ESM for machine AA.)
XM=AB, , 400.	(Allocate 400000 octal words for user-accessible extended memory in ESM for machine AB.)

The EQPDECK entries for the linked mainframe would be the same as shown above; the link device (ESM) is defined as equipment 5 and both XM entries must be present.

EQPDECK

Example 2:

This example shows two linked CYBER 180-class machines connected to ESM via low-speed ports. Part of the mainframe memory for machine AA is defined as UEM, which is used as a mass storage device, for 895, 887, or 9853 disk I/O buffers, and for user-accessible extended memory. On CYBER 180-class machines, user-accessible extended memory and 887 or 895 disk I/O buffers must reside in UEM; they cannot reside in ESM.

EQ5=DP, ST=ON, ET=ES/D2/NC, SZ=10000, CH=4. (Define ESM as the link device. NC indicates that there is no CPU access to ESM.)

EQ6=DE, ST=ON, ET=EM, SZ=2000. (Define 1/2 million words of UEM; without this EQ entry, UEM can be used only for user-accessible extended memory.)

MSAL, S=6. (Define UEM as secondary rollout device.)

ASR=6. (Define UEM as alternate system residency device.)

XM=AA, 1000, 200. (Allocate 1000000 octal words for 895, 887, or 9853 I/O buffers and 200000 octal words for user-accessible extended memory in UEM.)

The EQPDECK for the linked CYBER 180-class machine AB might look like the following. This mainframe has part of its memory defined as UEM used only for user-accessible extended memory.

EQ5=DP, ST=ON, ET=ES/D2/NC, SZ=10000, CH=4. (Define ESM as the link device.)

XM=AB, , 200. (Allocate 200000 octal words for user-accessible extended memory in UEM. Since no EQ entry is present for UEM, it cannot be used as a mass storage device or for I/O buffers.)

The XM entries appear only in the EQPDECK for the machine where the memory is being allocated. Do not place them in the EQPDECKs of any linked mainframes.

Example 3:

This example shows a CYBER 170 mainframe connected to ESM via a high-speed port linked to a CYBER 180-class machine connected to ESM via a low-speed port.

EQPDECK for CYBER 170

EQ5=DP, ST=ON, ET=ES, SZ=10000, CH=5.	(Define ESM as the link device.)
MSAL, S=5.	(Define ESM as secondary rollout device.)
ASR=5.	(Define ESM as alternate system residency device.)
XM=AA, 1000, 400.	(Allocate 1000000 octal words for 885-42 I/O buffers and 400000 octal words for user-accessible extended memory in ESM.)

EQPDECK for CYBER 180

EQ5=DP, ST=ON, ET=ES/D2/NC, SZ=10000, CH=4.	(Define ESM as the link device.)
EQ6=DE, ST=ON, ET=EM, SZ=2000.	(Define 1/2 million words of UEM; without this EQ entry, UEM can be used only for user-accessible extended memory.)
MSAL, S=6.	(Define UEM as secondary rollout device.)
ASR=6.	(Define UEM as alternate system residency device.)
XM=AA, 1000, 400.	(Linked 170 XM entry must be present to define user-accessible extended memory and 885-42 I/O buffers in ESM.)
XM=AB, 1000, 200.	(Allocate 1000000 octal words for 895, 887, or 9853 I/O buffers and 200000 octal words for user-accessible extended memory in UEM. This entry cannot be in the 170 EQPDECK.)

EQPDECK

Example 4:

This example shows a CYBER 170 model 865 or 875 connected to ESM via a low-speed port linked to a CYBER 180-class machine also connected to ESM via a low-speed port. This configuration allows continued system operation if the high-speed port to ESM from the CYBER 170 model 865 or 875 was down.

EQPDECK for CYBER 170 model 865 or 875

Note that ESM cannot be used for user-accessible extended memory or for 885-42 I/O buffers if a high-speed port is not present.

EQ5=DP, ST=ON, ET=ES/D2/NC, SZ=10000, CH=4.	(Define ESM as the link device; NC indicates that there is no CPU access.)
EQ11=DE, ST=ON, ET=EM, SZ=2000.	(Define 1/2 million words of UEM. This is allowed; in contrast to example 8 earlier, because ESM has no CPU access.)
MSAL, S=11	(Define UEM as secondary rollout device.)
ASR=11.	(Define UEM as alternate system residency device.)
XM=AA, , 200.	(Allocate 200000 octal words for user-accessible extended memory in UEM. This entry cannot be in the 180 EQPDECK.)

EQPDECK for CYBER 180

EQ5=DP, ST=ON, ET=ES/D2/NC, SZ=10000, CH=4.	(Define ESM as the link device.)
EQ6=DE, ST=ON, ET=EM, SZ=2000.	(Define 1/2 million words of UEM; without this EQ entry, UEM can be used only for user-accessible extended memory.)
MSAL, S=6.	(Define UEM as secondary rollout device.)
ASR=6.	(Define UEM as alternate system residency device.)
XM=AB, 1000, 200.	(Allocate 1000000 octal words for 895, 887, or 9853 I/O buffers and 200000 octal words for user-accessible extended memory in UEM. This entry cannot be in the 170 EQPDECK.)

MSAL — Mass Storage Allocation Control Entry

The MSAL entry assigns job files of the specified type to the mass storage devices defined by the specified EST ordinal.

The format of the MSAL entry is:

MSAL,t=est₁,est₂,...,est_n.

Parameter Description

t	File type; one of these values:
t	Description
B	LGO files.
D	Job dayfiles.
I	Input files. ¹⁴
L	Local files.
O	Output files. ¹⁴
P	Primary files.
R	Rollout files.
S	Secondary rollout files.
T	Temporary files.
est _i	EST ordinal of a nonremovable mass storage device, from 5 to 57 _g . Ranges of ordinals can be specified.

NOTE

In most cases, using ESM or UEM as a secondary rollout device improves system performance. On the other hand, depending on the mainframe type, the size of the jobs being rolled out, and the secondary rollout sector threshold (IPRDECK and DSD SRST command), CPU overhead may increase significantly. Using central memory for pseudo-control points instead of using the same memory as UEM for secondary rollout is normally preferable because of reduced CPU overhead. Use the ACPD and PROBE utilities (refer to TRACER/PROBE Utilities, section 21) to determine the best use of system resources.

¹⁴. Routing a file to queues changes a file's type, not its residency. Thus specifying an MSAL,I=est or MSAL,O=est entry does not necessarily force all input or output queued files to the specified devices.

EQPDECK

Secondary rollout files are rollout files whose size in sectors is smaller than a threshold specified by the IPRDECK or DSD entry SRST. These files are rolled out to devices specified by the MSAL,S= entry. All files selected for rollout that are equal to or greater in size than the threshold are rolled out to devices specified by the MSAL,R= entry. The default value of SRST is 0 (zero); thus, no secondary rollout files exist unless this value is changed. A possible use of this feature is:

Entries	Comments
EQPDECK	
:	
EQ5=DP, ST=ON, SZ=1000, CH=27.	Specify equipment 5 as extended memory with a DDP.
:	
MSAL, S=5.	Direct secondary rollout files to extended memory.
IPRDECK	
:	
SRST=20	Set threshold count so that no large rollout files are rolled out to extended memory.

DOWN — Down Channel Entry

The DOWN entry disables the use of channels at deadstart before the system attempts to use them to access devices connected to them. If a channel is the only remaining access to a device, it may not be downed unless that device is defined as DOWN. The DOWN entry does not apply to either the deadstart channel or the CIP channel.

The DOWN entry format is:

```
DOWN, CH=ch1, ch2, . . . , chn.
```

Parameter Description

ch_i Number of the channel; from 0 to 13₈ and 20₈ to 33₈. A concurrent channel is indicated by a C prefix (Cchi).

Example:

The following entry sets DOWN status for both nonconcurrent channel 23 and concurrent channel 4.

```
DOWN, CH=23, C4.
```

UP — Up Channel Entry

The UP entry allows resumption of normal use of channels that have been disabled by a previous DOWN entry.

The UP entry format is:

```
UP, CH=ch1, ch2, . . . , chn.
```

Parameter Description

ch_i Number of the channel; from 0 to 13₈ and 20₈ to 33₈. A concurrent channel is indicated by a C prefix (Cchi).

Example:

The following entry clears DOWN status of both nonconcurrent channel 23 and concurrent channel 4.

```
UP, CH=23, C4.
```

PF — Permanent Files Device Assignment Entry

Before initializing a mass storage device (with the INITIALIZE entry in the EQPDECK), add a PF entry¹⁵ for that device to the EQPDECK anywhere after the device's EQ entry. The PF entry information becomes part of the device's label when it is initialized during deadstart; this label is recovered during subsequent deadstarts. For subsequent deadstarts, it is not necessary that the PF entry be part of the EQPDECK on the deadstart file; if it is, it is ignored. Refer to the INITIALIZE entry for the default PF entries.

If the unit is a family device, the format is:

PF=est, type, dm, sm, name, device, nc.

If the unit is an auxiliary device, the format is:

PF=est, type, name, nc.

Parameter Description

Parameter	Description						
est	EST ordinal of the device; from 5 to 777 ₈ .						
type	Type of device; one of these values: <table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>F</td> <td>Family device. It can contain indirect access files if the dm parameter is from 1 to 377₈. It can contain direct access files if the sm parameter is from 1 to 377₈. It is a master device if the dm parameter is specified.</td> </tr> <tr> <td>X</td> <td>Auxiliary device, which can contain both direct and indirect access files. X must be specified on a unit's PF entry if any of the auxiliary device commands are to be used for the device.</td> </tr> </tbody> </table>	type	Description	F	Family device. It can contain indirect access files if the dm parameter is from 1 to 377 ₈ . It can contain direct access files if the sm parameter is from 1 to 377 ₈ . It is a master device if the dm parameter is specified.	X	Auxiliary device, which can contain both direct and indirect access files. X must be specified on a unit's PF entry if any of the auxiliary device commands are to be used for the device.
type	Description						
F	Family device. It can contain indirect access files if the dm parameter is from 1 to 377 ₈ . It can contain direct access files if the sm parameter is from 1 to 377 ₈ . It is a master device if the dm parameter is specified.						
X	Auxiliary device, which can contain both direct and indirect access files. X must be specified on a unit's PF entry if any of the auxiliary device commands are to be used for the device.						
dm	Specifies the unit's device mask; from 0 to 377 ₈ . Set according to information under Device Masks in section 17, PF Utilities. Omit this parameter if the device is an auxiliary device. The device mask for a permanent file device defines the groups of users whose catalogs reside on the device for a particular family.						
sm	Specifies the unit's secondary mask; from 0 to 377 ₈ . Set according to information under Device Masks in section 17, PF Utilities. Omit this parameter if the device is an auxiliary device. This parameter controls the residence of direct access files in the same way that dm controls the residence of indirect access files.						

¹⁵ Refer to table 3-5 for dependencies.

Table 3-5. PF Entry Parameter Settings (type, dm, name, device, and sm)

Type of PF Device	Files Permitted on Device	type	dm	name	device	sm
Auxiliary ¹	Indirect and/or direct	X	Omit	pack	Omit	Omit
Family	Direct only	F	0	family	1-77 ₈	1-377 ₈
	Indirect only (master device)	F	1-377 ₈	family	1-77 ₈	0
	Indirect and direct (master device)	F	1-377 ₈	family	1-77 ₈	1-377 ₈

1. If the device is to be a private auxiliary device, enter the INITIALIZE command after deadstart and specify the user name. Refer to section 8, K-Display Utilities, for the procedure.

Parameter	Description
-----------	-------------

name	<p>Designates either the name of the family to which the unit belongs or its pack name if it is an auxiliary device; from 1 to 7 alphanumeric characters. Do not use the family name 0: it is reserved.</p> <p>The family name describes the permanent file devices available to a user.¹⁶ A family may consist of from 1 to 63 logical devices; however, the master devices within the family must have device masks totaling 377_8 if all possible user indexes are to be accommodated.</p> <p>Usually a system runs with one family of permanent file devices available. But you can activate additional families on a system, in order to allow the users of these families to access their permanent files through an alternate system. This might be helpful if one system supplies backup service to another system. When more than one family is active on a system, users with matching user indexes access the same permanent files on a public auxiliary device. You can avoid this situation by predetermining a range of user indexes for each family running on a system. When a new family is introduced into a system, its user indexes should be checked against those of the family or families currently running and any matching indexes should be changed. Refer to the IPRDECK entry COMLIB later in this section.</p> <p>The pack name is the unique 7-character name associated with an auxiliary device. An auxiliary device is a self-contained permanent file device: all permanent files (whether direct or indirect access) represented by the catalogs on the device reside on that device. To access a file on an auxiliary device, users must specify the pack name as part of the permanent file request. The pack name is used instead of the usual algorithm for determining catalog location (user masks and family name). An auxiliary device can be private or public. Any user who knows the pack name and has the appropriate permissions and validations can access files on an auxiliary device. Only the owner user name can create files on a private auxiliary device (perform DEFINE, SAVE, or REPLACE requests).</p>
device	<p>Number of the device in the family; from 1 to 77_8. Omit this parameter if the device is an auxiliary device.</p> <p>A permanent file that does not reside on the master device has a device number in the catalog entry or on the master device. The device number specifies on which alternate device within the family the file resides.</p>

16. If not otherwise specified, the default family name becomes part of the tape label information. It is checked and verified if the user specifies the FA=A parameter on a command. Refer to the NOS Version 2 Reference Set, Volume 3 for a discussion of FA=A.

Parameter Description

nc Number of catalog tracks (optional) used only for master devices; from 1 to 200₈. This value must be a power of 2. If you do not specify nc, one of the following default values (based on the equipment type) is supplied.

Default nc	Equipment	Type
10	819	DV/DW
40	834	DD
40	836	DG
40	844-21	DI/DK
40	844-41/44	DJ/DL
10	885-11/12	DM/DQ
10	885-42	DB
10	887(4K)	DF
10	887(16K)	DH
10	895	DC
10	5832/5833/5838	EA-EH
10	9853	DN
4	Extended memory	DE
4	ECS or ESM with DDP	DP
1	Private device	

Examples:

PF=6, F, 125, 125, SYSTEM, 3, 200.
PF=17, X, PACK.

SYSTEM — System Library Device Assignment Entry

The SYSTEM entry specifies which mass storage devices are to contain copies of the NOS system library from the deadstart file. A system device can be any disk storage device as well as extended memory.

Throughput can be greatly improved by specifying more than one system device. For example, if two system devices are specified and they are on different channels, the time required to access system programs can be reduced. When the channel for one system device is busy, the other is accessed. Also, if hardware problems occur on one system device, the other system device can still be accessed. A good general rule is to have one copy of the NOS system library per pair of channels with a maximum of three copies. This allows alternate access to the system library while saving mass storage space.

The following restrictions apply.

- The EQ entry for a system device cannot have the status set to OFF.
- A REMOVE entry cannot exist in the EQPDECK for a device being specified as a system device.
- If more than one device is specified as a system device, all devices specified must be of the same type and have the same number of spindles. For example, if there are two system devices and the equipment EST ordinal for one of them specifies DI (single spindle), the equipment EST ordinal for the other one must also specify DI (single spindle).
- If no devices are specified as system devices, the system library resides on the first nonremovable mass storage device.
- An ASR entry cannot exist in the EQPDECK for a device being specified as a system device.

The SYSTEM entry format is:

```
SYSTEM=est1, est2, . . . , estn.
```

Parameter Description

est _i	EST ordinal of the device to contain a copy of the system library on the deadstart file; from 5 to 37 ₈ . The EQ entry must be set to ON status. One or more ordinals can be specified with one SYSTEM entry. Ranges of ordinals can be specified. The maximum number of system devices allowed depends upon the value of MXSY (refer to COMSMSC Parameters in the NOS Version 2 Installation Handbook).
------------------	---

ASR — Alternate System Library Device Assignment Entry

This entry specifies which mass storage devices are to be alternate system devices. An alternate system device is a mass storage device on which duplicate copies of system routines can be placed by the system, either for faster access than is possible from a system device or because they are frequently used programs. The following restrictions apply.

- The device must be a mass storage device, including extended memory.
- The device cannot be a removable device.
- The device cannot be a system device.

When loading a system routine, the system will access the routine from the alternate system device if practical. However, if the alternate system device is down or its access is more congested than the system device, the system will access the routine from the system device.

The procedure for selecting the records to be placed on the alternate device is in LIBDECK.

The ASR entry format is:

```
ASR=est1,est2,...,estn.
```

Parameter	Description
est _i	EST ordinal of mass storage device to be used as an alternative system device; from 5 to 37 ₈ . Ranges of ordinals can be specified.

SCKP — System Checkpoint File Device Assignment Entry

The SCKP entry specifies on which mass storage devices the checkpoint file is to reside. This entry is processed during a level 0, 1, or 2 deadstart. If no checkpoint devices are defined, the checkpoint file will reside on the first default mass storage device.

A level 0 deadstart clears the checkpoint status on all mass storage devices. To clear the checkpoint status for a specific mass storage device, use the RESET entry.

The SCKP entry format is:

```
SCKP=est1,est2,...,estn.
```

Parameter	Description
est _i	EST ordinal of the mass storage device to be defined as a system checkpoint file device; from 5 to 777 ₈ . Ranges of ordinals can be specified.

FAMILY — Default Family Name Assignment Entry

The FAMILY entry defines the default family. The family that is to be defined as the default family may reside on more than one device. The EST ordinal of any device within the family can be specified on the FAMILY entry, except in the following situation. If the member of the family whose device mask will have bit 2⁷ (200₈ in mask) set is being initialized, the FAMILY entry must specify the ordinal of this device. In all cases, the FAMILY entry must follow the EQ entry for the device specified.

The following restrictions apply.

- The status parameter for a default family device's EQ entry cannot be set to OFF.
- A REMOVE entry cannot exist in the EQPDECK for a device being specified as a default family device.

The FAMILY entry format is:

```
FAMILY=est.
```

Parameter Description

est	EST ordinal number of the mass storage device that the system automatically uses to determine your family when you do not specify a family name at login or job initiation; from 5 to 777 ₈ .
-----	--

REMOVE — Removable Device Assignment Entry

If a mass storage device is to be considered removable, you must specify it as such at deadstart with the REMOVE entry. This allows it to be introduced or removed during system operation. A device specified as removable cannot also have associated with it an ASR, SYSTEM, MSAL, FAMILY, DAYFILE, ACCOUNT, ERRLOG, or MAINLOG entry (refer to Dayfile Descriptions earlier in this section for the last four entries).

The format is:

```
REMOVE=est1,est2,...,estn.
```

Parameter Description

est _i	EST ordinal of mass storage device that is to be removable; from 5 to 777 ₈ . One or more ordinals may be specified with one REMOVE entry. Ranges of ordinals can be specified.
------------------	--

SHARE — Shared Device Entry

This entry identifies the rotating mass storage devices to be shared through a link device (ECS, ESM, or STORNET) by from two to four mainframes in a multimainframe complex. The tables necessary for the management of these devices (MST, TRT, MRT, and DAT) are maintained on the link device. The presence of the SHARE entry implies a multimainframe complex; shared status will be set for the link device. If the link device is the only device to be shared, it must be specified in the SHARE entry. If the SHARE entry is specified, the ISHARE entry must be omitted. Refer to PRESET – Preset the Link Device Entry later in this section.

Following is a list of the equipment types that can be shared.

Equipment	Type
ECS without DDP <i>or</i> ESM without LSP	DE
ECS with DDP <i>or</i> ESM with low-speed port (LSP) <i>or</i> STORNET	DP
844-21	DI/DK
844-41/44	DJ/DL
885-11/12	DM/DQ
5830 Disk Array Subsystem	EA, EB, EC, ED, EE, EF, EG, and EH

The format of SHARE is:

SHARE=est₁, est₂, . . . , est_n.

Parameter Description

Parameter	Description
est _i	EST ordinal of the mass storage device being shared; from 5 to 777 ₈ (EST ordinal 5 is reserved for the link device). Ranges of ordinals can be specified.

Refer to section 13, Multimainframe Operations, for suggestions on shared device configurations.

MAXSHARE — Maximum Shared Device Entry

This entry defines the maximum number of shared devices in a linked multiframe complex. This number is needed to determine the size of the tables to be allocated in extended memory when a PRESET is performed.

The MAXSHARE entry is only required if some shared devices are not configured in the EQPDECK for all the mainframes in the complex. If this entry is not specified, the default is the number of shared devices configured on this mainframe, as specified by the SHARE entry.

The format for the MAXSHARE entry is the following:

MAXSHARE=n, b.

Parameter Description

n	The total number of devices to be shared by any of the mainframes in the complex.
b	The total number of buffered devices to be shared by any of the mainframes in the complex. This parameter is optional.

ISHARE — Independent Shared Device Entry

This entry identifies the rotating mass storage devices that are to be independently shared by from 2 to 16 mainframes in a multimainframe complex. The tables necessary for the management of these devices (MST, TRT, MRT, and DIT) are maintained on the shared device. You cannot designate ECS or ESM in an ISHARE entry. When the ISHARE entry is specified, the SHARE entry must be omitted. Refer to PRESET – Preset the Independent Shared Device Entry later in this section.

Listed are the equipment types that can be independent shared devices.

Equipment Type

834	DD
836	DG
844-21	DI/DK
844-41/44	DJ/DL
885-11/12	DM/DQ

The ISHARE entry format is:

```
ISHARE=est1,est2, . . . ,estn
```

Parameter Description

est _i	EST ordinal of the mass storage device being shared; from 5 to 777 ₈ . Ranges of ordinals can be specified, as described under Specifying Ranges of EST Ordinals earlier in this section.
------------------	--

Refer to section 13, Multimainframe Operations, for suggestions on shared device configurations.

LBC — Load Buffer Controllers Entry

This entry identifies the type of controlware to be installed on the specified disk channels. Depending on the specified parameters, this entry can identify the channels as having half-track or full-track controlware but not install the controlware.

Unless you specify the LBC entry, the system examines the mnemonics of the device in the EQPDECK entry and causes the default version of controlware to be installed as follows:

Device Type	Controlware Version Number
DB ¹⁷	MA722
DC ¹⁸	MA464
DD ¹⁹	MA462
DG ¹⁸	MA462
DI	MA710
DJ	MA710
DK	MA401
DL	MA401
DM	MA721
DQ	MA721

Use the LBC entry to override these defaults. The LBC entry format is:

LBC, type, c₁, c₂, . . . , c_n.

Parameter Description

type	Description
type	Controlware to be installed; one of these values.
type	Description
CC	Install 7165 controller with full-track (MA464) controlware.
CN	Identify channel(s) as having full-track 7165 controller, but do not install the controlware.
FM	Install controller with full-track (MA721) controlware.
FT	Install controller with full-track (MA401) controlware.

17. Not applicable for CYBER 180-class machines.

18. Applicable for CYBER 180-class machines except models 810, 815, 825, and 830.

19. Applicable for models 810, 815, 825, and 830 only.

Parameter Description

type	Description
	<i>(Continued)</i>
type	Description
HT	Install controller with half-track (MA710) controlware.
ID	Install 7255-1 adapter with full-track (MA462) controlware.
M8	Install 7990 controller with MB466 controlware.
NF	Identify channel(s) as having full-track controller, but do not install the controlware.
NH	Identify channel(s) as having half-track controller, but do not install the controlware.
NI	Identify channel(s) as having full-track 7255-1 adapter, but do not install controlware.
NM	Identify channel(s) as having full-track 7155-1 controller, but do not install the controlware.
NN	Identify channel(s) as having NADs, but do not install the controlware.
NP	Identify channel(s) as having 7155-401 controller, but do not install the controlware.
NX	Identify channel(s) as having 5870 printer, but do not install the controlware.
N8	Identify channel(s) as having 7990 controller, but do not install the controlware.
PH	Install 7155-401 controller with MA722 controlware.
c_i	Disk channels; type determines if controlware is installed on these channels.

The controlware version number that can be loaded into the adapter or controller types is:

Adapter Controller	Controlware Version Number
-------------------------------	-----------------------------------

7054	MA710
7152	MA710, MA401
7154	MA710, MA401
7155-1	MA721
7155-401	MA722
7165	MA464
7255-1	MA462
7990	MB466

XM — Declare Extended Memory Space for Buffers or User Access

The XM entry reserves space for I/O buffers and causes initialization of tables for user access to extended memory.

NOTE

Except on models 865 and 875, the assignment of user access to extended memory forces jobs using user access to use CPU 0. This prevents CPUMTR from being locked out during large block transfers to or from extended memory on dual-CPU mainframes.

The XM entry is required if the system contains user-accessible extended memory or buffered devices. If UEM is to be used only for user-accessible extended memory, you do not need to specify an extended memory EST entry in the EQPDECK. To access UEM, it must be enabled (refer to the USER EXTENDED MEMORY IPRDECK entry later in this section).

The XM entry format is:

XM=id,iob,uec,EM.

Parameter	Description
id	Identifier of the mainframe that will access the reserved extended memory space; id is required. If an XM entry refers to shared extended memory (ECS or ESM as a link device), the XM entry must be present in the EQPDECKs of all linked mainframes. If an XM entry refers to nonshared extended memory (UEM), the XM entry must be present only in the EQPDECK for its own mainframe.
iob	Number of words/1000 ₈ reserved for buffers for DB, DC, DF, DH, DN, DV, or DW devices; from 40 ₈ to 37770 ₈ . This value is limited by the amount of memory available; either physically present or logically available (allocated in the CMRDECK and EQPDECK entries). The optimum number of words to reserve is the amount left after the alternate system library and user-accessible extended memory are taken into account. If no buffered devices are defined, enter 0.
uec	Number of words/1000 ₈ of extended memory to reserve for user access; maximum value is 77740 ₈ . This value is limited by the amount of memory available; either physically present or logically available (allocated in the CMRDECK and EQPDECK entries).
EM	For models 865 and 875 and CYBER 180-class machines, user-accessible extended memory can be allocated in UEM without allocating UEM as an equipment. If you want only user-accessible extended memory in UEM, clear or do not specify an extended memory EST entry in the EQPDECK and enter XM with the uec size and the EM keyword. For models 865 and 875, if you specify EM, user-accessible extended memory is allocated in UEM regardless of the presence of ECS or ESM; if you omit EM, user-accessible extended memory is allocated in the device defined in the extended memory EST entry.

UEMIN — UEM Equipment Initialization

The UEMIN entry enables full initialization for UEM equipment during a level 0 deadstart. It is equivalent to an INITIALIZE,AL,est entry. The UEMIN entry can be included as part of the EQPDECK on the deadstart file or it can be entered from the system console at deadstart time. If the UEMIN entry is being used, the EQ entry for the DE equipment is required. If you enter UEMIN when it's already set, it disables the automatic initialization of UEM. The UEMIN format is:

UEMIN.

ACCESS — Set Access Level Limits

The ACCESS entry allows you to specify equipment access level limits for mass storage, magnetic tape, two-port multiplexer, stimulator multiplexer, and unit record type equipment. This entry is invalid for other types of equipment. This entry determines the upper and lower limits for the range of access levels of the data allowed to be read from or written to the equipment. The default equipment access level limits are zero; no secure data can be read from or written on the equipment. Values for this entry should be supplied by a site security administrator.

If you want the equipment access level limits specified by this entry to become the device access level limits for mass storage equipment, that equipment must be initialized. Refer to the NOS Version 2 Security Administrator's Handbook for more information about equipment access and device access level limits.

This entry is ignored if the system is running in unsecured mode.

The format of the ACCESS entry is:

```
ACCESS, lower, upper, ordlist.
```

Parameter Description

lower	Access level name specified in deck COMSMLS (refer to the NOS Version 2 Installation Handbook) corresponding to the desired lower limit.
upper	Access level name specified in deck COMSMLS corresponding to the desired upper limit.
ordlist	One or more EST ordinals, separated by commas, or a range of ordinals as described in Specifying Ranges of EST Ordinals earlier in this section.

THRESHOLD — Set Disk Thresholds

The THRESHOLD entry enables you to set threshold values for the disk storage devices specified by EST ordinals. The system uses these threshold values as limits when monitoring disk verification failures, available disk space, and disk error processing. If a threshold value is exceeded, the system performs a corrective action, such as notifying the operator or restricting activity on the affected disk. The format of the THRESHOLD entry is:

```
THRESHOLD, type=value, ordlist.
```

Parameter Description

type	You can specify one of the following threshold types. Refer to the DSD THRESHOLD command in section 5 for additional information.												
	<table border="1"> <thead> <tr> <th style="text-align: left;">type</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>VF</td> <td>Verification failure threshold (default).</td> </tr> <tr> <td>RA</td> <td>Restricted activity threshold.</td> </tr> <tr> <td>LS</td> <td>Low space threshold.</td> </tr> <tr> <td>RE</td> <td>Recovered error threshold.</td> </tr> <tr> <td>UE</td> <td>Unrecovered error threshold.</td> </tr> </tbody> </table>	type	Description	VF	Verification failure threshold (default).	RA	Restricted activity threshold.	LS	Low space threshold.	RE	Recovered error threshold.	UE	Unrecovered error threshold.
type	Description												
VF	Verification failure threshold (default).												
RA	Restricted activity threshold.												
LS	Low space threshold.												
RE	Recovered error threshold.												
UE	Unrecovered error threshold.												
value	The threshold value can range from 0 to 3777 ₈ . If the value parameter is omitted, the following default values are used.												
	<table border="1"> <thead> <tr> <th style="text-align: left;">type</th> <th style="text-align: left;">value</th> </tr> </thead> <tbody> <tr> <td>VF</td> <td>0</td> </tr> <tr> <td>RA</td> <td>1/8 of the number of tracks on the device.</td> </tr> <tr> <td>LS</td> <td>1/16 of the number of tracks on the device.</td> </tr> <tr> <td>RE</td> <td>50₈</td> </tr> <tr> <td>UE</td> <td>0</td> </tr> </tbody> </table>	type	value	VF	0	RA	1/8 of the number of tracks on the device.	LS	1/16 of the number of tracks on the device.	RE	50 ₈	UE	0
type	value												
VF	0												
RA	1/8 of the number of tracks on the device.												
LS	1/16 of the number of tracks on the device.												
RE	50 ₈												
UE	0												
ordlist	Enter one or more EST ordinals, separated by commas, or a range of ordinals as described in Specifying Ranges of EST Ordinals earlier in this section.												

EQPDECK Entries Made Only During Deadstart

The following entries are valid only from the system console at deadstart time. They cannot be included as part of the EQPDECK on the deadstart file.

AUTOLOAD — Toggle Autoloading

The AUTOLOAD entry toggles the selection of buffer controller autoloading for all 7054/7154/7152/7155/7165 controllers and the 7255 adapter. The AUTOLOAD format is:

AUTOLOAD .

GRENADE — Clear Unit Reservations

The GRENADE entry causes unit reservations to be cleared on all 844 units physically connected to each 7054/7154/7152/7155 controller and any 834 or 836 units connected to a 7255 adapter. The GRENADE format is:

GRENADE .

INITIALIZE — Initialization Entry

To use a mass storage device that is defined with an EQ entry, it must have a label. A label is written on a device when you initialize it by using either the INITIALIZE command, during system operation, or the INITIALIZE entry in the EQPDECK, when it is displayed at the system console at deadstart time.

A mass storage device's label is contained on a logical track (usually track 0). It contains information about the allocation and characteristics of a device (and its units, if there is more than one unit on a device). This information is in the form of a label sector for the first unit, a TRT for the device, and a label sector for each unit.

Initialization does not automatically occur at each deadstart because mass storage device labels are recovered during all deadstarts. Therefore, initialize a device only in the following situations.

- To add a new mass storage device (no label exists on the device) use the INITIALIZE entry.
- If parts of the label on a permanent file device have been destroyed by maintenance operations (permanent files having been dumped to another device before diagnostics were run), use the INITIALIZE entry during deadstart to write a new label. Then reload the permanent files.
- If a device (usually a private auxiliary, public auxiliary, or alternate permanent file family device) is added to a system during operation, use the DSD INITIALIZE command to initialize it if it does not have a valid label on it when it is added to the system.
- When an extended memory device is initially placed in maintenance mode, all mainframes using extended memory must initialize it (the maintenance mode parameter is described under EQ – Extended Memory EST Entry earlier in this section). You must also enter the PRESET entry for multimainframe operation.

During a deadstart, the INITIALIZE entry has the following characteristics.

- During a level 0 deadstart, it can be entered at the system console only when the EQPDECK is displayed. It can be entered anywhere after the EQ entry for the device.
If it is placed in the deadstart file EQPDECK, the system issues the error message INCORRECT ENTRY when the EQPDECK is read from the tape.
- A total initialization (op=AL) assumes that no valuable information exists on the device and creates a new label. When the new label is created, all previously existing information on the device, except CTI, CDA, HIVS, and MSL, is lost.
- If the EQ status for the device is OFF when INITIALIZE is entered, initialization of the device occurs whenever the device is set to ON status by the operator with the DSD ON command during normal system operation.
- If the device is not a master device, INITIALIZE (op=AL) only writes a label; if it is a master device, then it also initializes the catalog track and writes EOIs at the beginning of the permit track, the indirect access track (data chain), and each catalog track.
- During a deadstart initialization (op=AL), all flaw reservations specified for a device are lost and must be reentered, except for 844 type devices with factory-formatted disk packs.
- During an initialize format pack (op=FP) for an 895 disk, NOS only formats those cylinders it plans to use in large record format. NOS does not format the disk cylinders occupied by CIP.
- During an initialize format pack (OP=FP) for a 583x device, the system first checks the current format of the device. If the device is formatted correctly, the system reformats the device only if FORCED FORMATTING is enabled. FORCED FORMATTING is enabled by default. To disable FORCED FORMATTING, enter the DISABLE, FORCED FORMATTING command either in the IPRDECK or from DSD.

NOTE

If a disk deadstart is in progress, the deadstart disk cannot be initialized. An attempt to initialize it will result in the system issuing the error message INCORRECT ENTRY.

The format of INITIALIZE is:

```
INITIALIZE,op,est1,est2,...,estn
```

Parameter Description

op	Level of initialization; one of these values:
op	Description
AF	Initialize account dayfile.
AL	Total initialization. For an 895 disk, the AL parameter is equivalent to the FP parameter if an initial install of CIP was done or if CIP was released.

Parameter	Description
-----------	-------------

op

(Continued)

op	Description
DF	Initialize system dayfile.
EF	Initialize error log.
FP	Initialize format pack (an automatic selection of AL also occurs). This option applies to 844, 895, and 583x disks only.
MF	Initialize binary maintenance log.
PF	Initialize permanent files.
QF	Initialize queued files.

est_iEST ordinal of mass storage device to be initialized; from 5 to 777₈.

If the ordinal refers to a family permanent file device, then family name, device number, and mask (if it is a master device) are specified on the PF entry.

If it is an auxiliary device, the pack name is specified on the PF entry.

Total initialization (op=AL or FP) is the only initialization that is independent of the content of the pack, if the initialization occurs during deadstart. If the initialization is done while the system is running, it is applied to the device after the check mass storage (CMS) routine has recovered it. If CMS cannot recover the device, the initialization is similar to a deadstart initialization (that is, all information on the device is lost).

The device number, family name, and device masks can only be changed during a total initialization. Since all devices may contain permanent files, you should include a PF entry for a device when performing a total initialization. If you do not, the device is assigned a default family name, device number, and device masks. It is possible that these parameters may conflict with other devices in the system. If a conflict occurs, resolve it by using PF entries.

If you initialize a nonremovable device without a PF entry, the device mask and secondary mask default to 377₈ for the smallest EST ordinal larger than 5. For all other equipment, including equipment 5, the default masks are set to 0 (zero). The default family name is SYSTid (where id is the machine identifier). The default device numbers begin at 1 and increase by 1 (starting with EST ordinal 5) for each device that you initialize without a PF entry.

The INITIALIZE entry operates in conjunction with the dayfile entries DAYFILE, ACCOUNT, ERRLOG, and MAINLOG (refer to Dayfile Descriptions earlier in this section) to determine where the dayfiles actually reside. The following examples illustrate the various cases. Assume that the system has three mass storage devices (EST ordinals 6, 7, and 10).

Example 1:

For this example, no dayfile entries are made and no previous dayfiles exist.

The following EQPDECK entry is made.

```
INITIALIZE,AL,6,7,10.
```

All dayfiles reside on ordinal 6.

Example 2:

In this example dayfile entries are made, but no previous dayfiles exist.

The following EQPDECK entries are made.

```
DAYFILE=6,200.
ACCOUNT=7,200.
ERRLOG=10.
MAINLOG=10,200.
INITIALIZE,AL,6,7,10.
```

In this case, the dayfiles reside on the indicated devices (system dayfile on ordinal 6, account dayfile on ordinal 7, error log and binary maintenance log on ordinal 10). The default buffer length is used for the error log buffer.

Example 3:

In this example, dayfile entries are made and previous dayfiles do exist.

Assume that the EQPDECK entries in example 2 are used.

Since a total initialization has been done on each device, no dayfiles are recovered. They reside on the indicated devices.

Example 4:

In this example, dayfile entries are made, previous dayfiles exist, but no dayfile initialization entries are made.

The following EQPDECK entries are made.

```
DAYFILE=6.
ACCOUNT=7.
ERRLOG=10.
INITIALIZE,PF,6.
```

The dayfiles may already reside on the specified devices, or they may reside on some combination of the possible devices. In either case, since no dayfile initialization entries are made, the old dayfiles are recovered. The residence of these dayfiles is governed by the residence of the old dayfiles. The PF initialization entry returns all permanent file space and relabels the device based on the recovered device parameters. The dayfiles and queued files on this device are not affected by this entry.

EQPDECK

Example 5:

In this example, dayfile entries are made, previous dayfiles exist, no dayfile initialization entries are made, and duplicate dayfiles are in existence.

Assume that the EQPDECK entries in example 4 are used.

For the dayfiles that do not have duplicates, the residence is defined by the current residence of the files, not the EQPDECK entries. But assume that an error log is recovered from ordinals 6 and 10. In this case, the most recent file becomes the active error log. Its previous residence overrides the EQPDECK entry. The other file becomes an inactive error log (an entry exists in the mass storage table of the device pointing to the inactive file, but the file is not in use by the system).

To produce an inactive error log, the site must run in the following manner.

1. Assume an 844 disk subsystem with two or more spindles is being used. Run with unit 1 equated to EQ6 and unit 0 unused.
2. Redeadstart, equate unit 0 to EQ6, and do not use unit 1.
3. Redeadstart, equate unit 0 to EQ6, and unit 1 to EQ7.

Since unit 0 has the most recent copy of the error log, this copy would become an active error log and the copy on unit 1 would become an inactive error log.

Example 6:

In this example, dayfile entries are made, the previous dayfiles from example 2 exist, and initialization entries are made.

The following EQPDECK entries are made.

```
DAYFILE=7.  
ACCOUNT=7.  
ERRLOG=10,300.  
MAINLOG=10.  
INITIALIZE,DF,6.  
INITIALIZE,QF,6.
```

In this case, the account dayfile is recovered and continued on ordinal 7. The binary maintenance log is recovered and continued on ordinal 10 with a CM buffer length of 100. The error log is recovered and continued on ordinal 10 with a CM buffer of 300₈ words. The system dayfile space on ordinal 6 (from example 2) is released and the new system dayfile starts on ordinal 7. The QF initialization entry releases all space reserved by queued files on ordinal 6.

The CM buffer length is not affected by dayfile recovery. It is always specified by the values defined in the EQPDECK entries. If no buffer length entries exist, the system default values are used.

PRESET — Preset the Link Device Entry

This entry defines allocation space and initializes the tables (MST, TRT, MRT, BAT, BRT, and BDT) on the link device that are required for management of shared multimainframe mass storage devices. The entry is valid only for level 0 deadstarts by the first mainframe in the multimainframe complex to deadstart.

Once PRESET is issued, the SHARE entry is disabled. Therefore, all SHARE entries must precede the PRESET entry. The PRESET entry has the format:

PRESET .

PRESET — Preset the Independent Shared Device Entry

This entry presets the independent shared devices in a multimainframe complex. The MST, TRT, MRT, and DIT are maintained on the mass storage device itself and are not affected by the PRESET entry. This entry is used in conjunction with the ISHARE entry. It is valid only on a level 0 deadstart by the first mainframe in the multimainframe complex to deadstart. All ISHARE entries must precede the PRESET entry.

The format is:

```
PRESET=est1,est2,...,estn.
```

Parameter Description

Parameter	Description
est _i	EST ordinal of the ISHARE device; from 5 to 777 ₈ . Ranges of ordinals can be specified, as described under Specifying Ranges of EST Ordinals earlier in this section.

Refer to section 13, Multimainframe Operations, for suggestions on shared device configurations.

WARNING

If a PRESET is entered for a device that is already in use by a second machine in a multimainframe complex, system failures may occur on both machines, and data on the device may be destroyed.

If a PRESET is entered for a device while a second machine in a multimainframe complex is down, and if that second machine later attempts to perform a level 3 deadstart and use that device, system failures may occur on both machines, and data on the device may be destroyed.

RESET — Reset Device Attributes

The RESET entry rescinds all device-related attributes resulting from entries such as REMOVE, SYSTEM, MSAL, and so on. It restores the values specified with the last EQest entry encountered. The format is:

```
RESET=est1,est2,...,estn.
```

Ranges of ordinals are not allowed; each ordinal must be entered individually.

APRDECK

The auxiliary mass storage parameter deck (APRDECK) is a text record on the deadstart file that is processed during system initialization. APRDECK entries identify areas of mass storage that are unusable (flawed areas) and prevent the system from accessing them. The system uses the information in the APRDECK entries to build the TRT for each device that resides in CMR and also in the mass storage device label.

You can place up to 64 APRDECKs on the deadstart file. Placing several APRDECKs on the same deadstart file allows you to use the same file to deadstart several configurations.

APRDECK Format

The first line in an APRDECK is the deck name. The format of the APRDECK name is:

```
APRDnn
```

Parameter	Description
nn	Number identifying the APRDECK; from 00 to 77 ₈ .

An APRDECK must have a name and may have flaw entries. The first APRDECK must contain the deck name APRD00 and nothing else. Subsequent APRDECKs must be numbered consecutively and can contain flaw entries.

The released version of the APRDECK contains no entries. You can enter flaws at three different times:

- During deadstart, after entering all EQPDECK modifications.
- During system operation, using the FLOW entry (refer to section 8, K-Display Utilities).
- During the configuration of a deadstart file.

If during deadstart you initialize a device and then enter NEXT, the system displays both the parameters on the device's EST entry and the APRDECK referenced by the EST entry. You can then change the flaws for the device. If the first APRDECK is referenced by the EST entry, the system displays the parameters on the device's EST entry and the APRDECK name, APRD00. You can then enter flaws for the device. These changes to the APRDECKs remain in effect until the next deadstart.

For example, in a EQPDECK, the EST entry for an 844-21 disk is:

```
EQ07=DI,ST=ON,EQ=0,CH=31/33,AP=5,UN=2.
```

After you initialize equipment 07 and enter NEXT, the following display appears.

EQ	TYPE	ST	EQ	UNITS	CHANNELS
07	DI-1	ON	0	02	31 33

```
APRD05
SLF=4173.
SLF=7062.
```

The APRDECK entries described in this section are those to be used for entering flaws during deadstart or during the configuration of the deadstart file. Use the entries as follows.

- Use the CAF entry to clear all flaw reservations on a device.
- Use the SPF entry to specify the physical address of a flaw in extended memory. (If a reservation for that physical address already exists, it remains in effect.)
- Use the SPF entry to specify the cylinder, track, and sector of a flaw in a disk. (If a reservation for that physical area already exists, it remains in effect.)
- Use the CPF entry to cancel a particular SPF entry.
- Use the SLF entry to specify the logical address of a flaw. (If a reservation for that logical address already exists, it remains in effect.)
- Use the CLF entry to cancel an SLF entry.

NOTE

All numeric values entered in the APRDECK must be entered in octal.

Either obtain flaw addresses from a customer engineer, or run the MST (mass storage test) on the device to determine the bad areas. MST specifies the physical address of flaws.

The system reads the flaw information recorded on the utility flaw map of an 881/883/885 disk pack during the initialization of 844/885 equipment and reserves the appropriate areas. For multiunit devices, the flaw reservation is the union of all utility flaw maps. This automatic flawing process occurs in addition to any APRDECK entries. However, you cannot clear areas recorded as flawed on the utility flaw map of an 881/883/885 disk pack with the CAF entry. Refer to appendix G for information on clearing these flaws.

You can list all APRDECKs on the deadstart file by accessing the system file SYSTEM with an ASSIGN or COMMON command, then using the T parameter on the CATALOG command. Refer to the NOS Version 2 Reference Set, Volume 3 for more information.

CAF — Clear All Flaw Reservations

The CAF entry clears all flaw reservations previously made with SLF or SPF entries. The format is:

CAF.

SLF — Set Logical Flaws on Any Mass Storage Device

Use this entry to specify the logical address of a flaw. If the track was previously reserved, that reservation remains in effect. The format is:

SLF=track.

Parameter Description

track	Logical track number. Refer to table 3-6 for valid ranges of track numbers.
-------	---

CLF — Clear Logical Flaws on Any Mass Storage Device

To cancel a reservation made with an SLF entry, enter the identical track number using the CLF entry. The format is:

CLF=track.

Parameter Description

track	Logical track number. Refer to table 3-6 for valid ranges of track numbers.
-------	---

SPF — Set Physical Extended Memory Track or Disk Area Flaws

Two formats exist for the SPF entry. One format prevents the system from using blocks (tracks) of extended memory. The other format prevents the system from using sectors on disks.

Use the following SPF format to prevent the system from using blocks (tracks) of extended memory.

SPF=Aaddress.

or

SPF=Aaddress₁-Aaddress₂.

Parameter Description

Aaddress	The 1- to 7-digit octal logical address in a track of extended memory; track containing the absolute address is reserved. The letter A must precede the address.
Aaddress ₁ Aaddress ₂	Lowest (Aaddress ₁) and highest (Aaddress ₂) addresses in a range of logical addresses in one or more tracks of extended memory. All tracks in the range are reserved. The letter A must precede the addresses; the hyphen is required.

Use the following SPF format to prevent the system from using sectors on disks.

SPF=Ccylinder, Ttrack, Ssector.

Refer to table 3-6 for the number of cylinders, tracks, and sectors for each device.

Parameter Description

Cylinder	Cylinder number; the letter C must precede the number.
Ttrack	Track number; the letter T must precede the number.
Ssector	Sector number; the letter S must precede the number.

CPF — Clear Physical Extended Memory Track or Disk Area Flaws

To cancel a reservation made with an SPF entry, enter the identical information with a CPF entry. The format is the same as for the SPF entries.

Table 3-6. Information for Setting or Clearing Flaw Areas

Device	Device Type	Logical Track Range ¹	Cylinders/ Device ¹	Tracks/ Cylinder ¹	Sectors/ Track ¹
819, single density	DV	4000-5465	633	12	24
819, double density	DW	4000-7153	1466	12	24
834	DD	4000-7135	1457	12	40
836	DG	4000-6565	1273	30	57
844-21, half track	DI	4000-7137	630	22	30
844-21, full track	DK	4000-7137	630	23	30
844-41/44, half track	DJ	4000-7147	1464	23	30
844-41/44, full track	DL	4000-7147	1464	23	30
885-11/12, half track	DM	4000-7221	1511	50	40
885-11/12, full track	DQ	4000-7221	1511	50	40
885-42	DB	4000-7221	1511	12	40
887, 4K sector	DF	4000-7343	1562	4	46
887, 16K sector	DH	4000-7343	1562	4	13
895	DC	4000-7351	1565	17	3
9853	DN	4000-7726	2601	23	25
Extended memory	DE/DP	4000-7620			
5832, 1X SSD	EA	4000-7745	1514	4	3
5832, 2X SSD	EB	4000-7745	1514	4	6
5833, 1X Sabre	EC	4000-7775	3135	7	6
5833, 1XP Sabre	ED	4000-7775	3135	7	6
5833, 2X Sabre	EE	4000-7761	3135	7	13
5833, 2XP Sabre	EF	4000-7761	3135	7	13
5838, 1X Elite	EG	4000-7750	5074	11	5
5838, 1XP Elite	EH	4000-7750	5074	11	5

1. Numbers are in octal.

IPRDECK

The IPRDECK contains the system installation parameters that determine the system's operation mode. From 1 to 64 IPRDECKs can exist on a deadstart file. The IPD entry in the CMRDECK specifies which IPRDECK to use. If you omit the IPD entry, the system uses the first IPRDECK on the deadstart file. IPRDECKs are named IPRDnn, where nn is from 00 to 77₈.

There are two IPRDECK console displays. The initial display, IPRINST, is an instruction display. It gives a brief description of all valid IPRDECK entries. The second display is the current IPRDECK. If either display overflows two screens, you can page the display.

You can modify the IPRDECK by entering the appropriate changes or additions from the console keyboard. Make these entries while either the IPRINST or IPRDECK is displayed. Each console entry supersedes the value currently specified in the IPRDECK.

NOTE

Changes made to the IPRDECK at deadstart time will be recovered across a level 3 deadstart; but not across a level 0, level 1, or level 2 deadstart. To make these changes permanent, you must create a new deadstart file with these changes incorporated into the IPRDECK.

You can list all IPRDECKs on your system by accessing the system file SYSTEM with a COMMON command, then using the T parameter on the CATALOG command. Refer to the NOS Version 2 Reference Set, Volume 3 for more information concerning these commands.

Most of the IPRDECK entries are also valid DSD commands that can be used to make changes during system operation. Changes to the IPRDECK using DSD commands are retained after a level 3 deadstart, but not after a level 0, level 1, or level 2 deadstart.

Job Control Information

The QUEUE, SERVICE, and DELAY entries in the IPRDECK relate to job control. General information concerning job control follows.

Job Scheduling

Job scheduling is the control of jobs in the input (IN), executing (EX), and output (OT) queues for each service class. Scheduling in the input and output queues is based on the priority of a queue entry relative to all queue entries in the system. The priority of a queue entry depends both upon how long the entry has been waiting in the queue and upon the parameters specified on the QUEUE entry in the IPRDECK. The following formula shows how the system computes the priority; all values are octal.

$$p = \frac{(ct-et)}{wf} + lp$$

Variable	Description
p	Priority; $LP \leq p \leq UP$. LP (a parameter on the QUEUE entry) is the lowest priority and UP is the highest priority.
wf	Weighting factor; WF parameter on the QUEUE entry.
ct	Current time in seconds.
et	Time in seconds at which the job entered the queue.
lp	Lowest priority; LP parameter on the QUEUE entry.

When an input or output queue entry is created, its priority is the lowest priority (LP) for its service class. The queue priority of the queue entry increases as time passes. The rate at which the priority increases depends upon the weighting factor (WF). The larger the weighting factor, the slower the priority increases. (The queue priority of an entry with a WF of 10_8 increases eight times slower than an entry with a WF of 1.) The queue priority increases either until the queue entry is selected for processing or until the queue priority reaches the highest priority (UP). If the queue priority of an entry reaches UP, it remains at UP until the entry is selected for processing. If the queue priority is zero, the job or file is never selected by the job scheduler and stays in the queue until the operator either enters a DROP command or resets the priority to a nonzero number.

Job scheduling for executing jobs determines how much execution time a job gets. The amount of execution time depends both on the job's scheduling priority (which is parameters specified on the QUEUE and SERVICE entries in the IPRDECK) and on the scheduling priorities of other jobs that may be waiting for execution.

When a job in the input queue is selected for execution, it gets an initial priority (IP; a parameter on the QUEUE entry) for the execution queue. Scheduling priority increases for a job in the execution queue in the same way as in the input and output queues. After the job is scheduled to a control point, its scheduling priority does not age; it is set to fixed values that are changed only when specific events occur. Initially, the job's scheduling priority is set to the upper priority bound for its service class (the value of the UP parameter on the QUEUE entry). The job's scheduling priority remains at this fixed value until the job exceeds either its control point time slice or its central memory time slice.

If the job exceeds its control point time slice (the CT parameter on the SERVICE entry), its scheduling priority is lowered to the control point slice priority (the CP parameter on the SERVICE entry). If the job exceeds its central memory time slice (the CM parameter on the SERVICE entry) for the first time, its scheduling priority is lowered to the initial lower priority (the IL parameter on the QUEUE entry); if the job exceeds its central memory time slice for a second or subsequent time, its scheduling priority is lowered to the lower priority bound (the LP parameter on the QUEUE entry).

Exceeding the CP or CM slice does not, by itself, force a job to roll out. A job remains at a control point until the job scheduler determines that it needs this job's control point and/or memory for a job with a higher priority. If a job rolls out, either to mass storage or to a pseudo-control point, its scheduling priority starts at the current value and begins to age upward (in the same fashion as input and output queue files). When the job scheduler again selects this job to run at a control point, its scheduling priority is again set to UP.

For interactive jobs, there is an additional execution queue priority, TP (a parameter on the SERVICE entry). It is assigned to the execution queue entry of a job restarting after terminal I/O. The value of TP aids response time to program prompts. Also, for interactive jobs, the initial priority (IP) has an added significance. In addition to being the priority at which jobs are scheduled from the input queue to the execution queue, IP is the priority assigned when a terminal command is entered. Using separate TP and IP parameters allows the system to give faster responses to users interacting with a job. To achieve this, a value for the TP parameter slightly greater than the value for the IP parameter is recommended (refer to tables 3-7 and 3-8 and figure 3-6).

The relative values of the QUEUE and SERVICE parameters, both among service classes and within a service class, affect system performance. For an example of ranges of service class priorities, refer to figure 3-6. For an example set of specific entries for the QUEUE and SERVICE parameters, refer to tables 3-7 and 3-8.

Control Points

A job leaves a central memory control point when any of the following conditions are present:

- A job completes, aborts, or is suspended.
- A system request causes a job to be rolled out to wait for a resource to become available or for a specified time period to elapse.

Such system requests include a job request for a tape or disk pack, a ROLLOUT command, and execution of the ROLLOUT macro.

- Terminal input/output is required.

A job leaves central memory when the following conditions occur:

- The system requests terminal input and typeahead input is not available.
- You request terminal output and the recall parameter is specified on the request.
- You issue a RECALL macro after a request for terminal output that omitted the recall parameter.
- The control point is made available for a higher priority job.

This ensures reasonable service to all users in the system. The operating system controls the amount of time each type of job can be at a control point. This ensures that one job does not monopolize system resources.

When a job is rolled out, the priority increases as time passes, giving that job a better chance to be selected for execution again. When the job is selected and rolled in, its priority is changed to the value of the UP parameter on the QUEUE entry, and the job scheduling priority cycle begins again. This description on job rollout applies to local batch, remote batch, and interactive jobs that are not doing interactive I/O. For interactive jobs that do terminal I/O within a time slice, scheduling priority is slightly different. When I/O is complete and input, for example, is available, the system assigns the rolled-out job the terminal I/O scheduling priority (TP parameter on the SERVICE entry). The TP parameter can be used to give the job a priority equal to the priority of jobs still within their initial time slice, an advantage over jobs in a second time slice, and a larger advantage over jobs in a third or higher time slice.

If a job at a control point exceeds the central memory time slice and it is not a subsystem, the scheduling priority is set to the initial lower or lower bound priority (the IL or LP parameter on the QUEUE entry) for its service class. Thus, any job in the queue with a higher priority forces the executing job with the lower priority to be rolled out. The rolled-out job ages normally until its priority is higher than the priorities of either the jobs in the input queue or a job that is executing; then it is again scheduled to a control point.

Once a job is scheduled, it is desirable to use the resources allocated before another job forces it out. If a job maintained its scheduling priority when it was assigned to a control point, another job could age past that job and force it to be rolled out before it had an opportunity to use its time slice. For this reason, when a job is assigned to a control point and its priority is within the queue aging range, it is given a priority equal to the highest priority (the UP parameter in the QUEUE entry) for its service class.

Pseudo-control Points

If pseudo-control points are defined, a job that is forced out of a control point may remain in central memory at a pseudo-control point, rather than being rolled out to a mass storage device.

- If the job is preempted by a higher priority job, the scheduling priority stays the same.
- If the job's control-point time slice has expired, the scheduling priority drops to the control-point slice priority.

Jobs with the following characteristics are excluded from using pseudo-control points:

- Connection to a system control point.
- Accumulator overflow flags set.
- Assigned to the first control point or the last control point.
- Error flag set.

Number of Control Points and Pseudo-control Points Available

Selecting the number of control points available on the system depends on the amount of memory space available, the job mix, and the mode in which the system is being run. Up to 34_g control points can be defined. Each control point needs 200_g words of CMR space. For example, if an installation is running only TAF, then four or five control points may suffice. On the other hand, if the system is running a large number of interactive terminals with heavy permanent file activity, 20 or more control points may be needed. You may need to study memory and control point use in order to correctly determine the setting of this option. The DSD W,R display and/or the ACPD and PROBE utilities may be useful in making this determination.

- If memory use is high and control point use is low, select fewer control points.
- If control point use is high and memory use is low, select more control points.

Selecting the number of pseudo-control points depends on the same variables that are used to determine the number of control points. Up to 34_g pseudo-control points can be defined. The pseudo-control points provide more effective use of central memory in environments where an insufficient number of control points create a system bottleneck. Like a control point, each pseudo-control point needs 200_g words of CMR space.

- If memory use is high and control point use is lower, do not select any pseudo-control points.
- If control point use is high and memory use is low, select pseudo-control points.

CPU Scheduling

Job scheduling can be managed by controlling access to memory (by rollin and rollout) and control points (by moving jobs between control points and pseudo-control points). It is also possible to directly control the allocation of CPU time to jobs by specification of service class CPU scheduling parameters. This method may become necessary when many jobs can fit into memory simultaneously (jobs may not roll out when their central memory slice expires) or when there are enough control points for all contending jobs (jobs may not move to pseudo-control points when their control point slice expires).

The system maintains a CPU wait queue of all jobs that are ready to run, sorted in descending order by CPU priority. Whenever the job currently executing in a CPU must give up the CPU for some reason, the first job on the queue that can use the CPU is selected for execution. Reasons for giving up the CPU include:

- Going into recall status to wait for I/O to complete.
- Preemption by a higher priority job.
- Expiration of the jobs' CPU slice.

The CPU slice is the method of distributing CPU service among jobs at the same CPU priority. A CPU slice is a time value expressed in milliseconds. When a job has completed its CPU slice, it is inserted into the CPU wait queue behind jobs with the same CPU priority that have time remaining in their CPU slices. When all jobs with the same CPU priority have completed their CPU slices, jobs with completed CPU slices are again selected for execution.

The CPU slice does not have to be a continuous span of time. For example, if a job is giving up the CPU frequently to wait for I/O, the job's CPU slice is typically interleaved with that of other jobs. Jobs performing I/O will be given preference in obtaining the CPU under these circumstances:

- If a job gives up the CPU to wait for I/O, that job is recalled when the I/O is complete.
- A job that has been recalled and has not completed its CPU slice is inserted into the CPU wait queue before those jobs at the same CPU priority that have not been recalled.
- A job that has been recalled and has completed its CPU slice is inserted into the CPU wait queue behind those jobs at the same CPU priority that have not completed their CPU slice. However, such a job is inserted into the wait queue before jobs at the same CPU priority that have completed their CPU slice but have not been recalled.

The US (unextended CPU slice) SERVICE parameter allows specification of different CPU slice values by service class. If there is an intensive CPU-bound job mix, jobs in service classes with longer CPU slices receive proportionately more CPU service than jobs of the same priority in service classes with shorter CPU slices.

The SE (CPU slice extension) SERVICE parameter defines an additional maximum amount of time that may be added, if needed, to the CPU slice of jobs performing I/O and using significant amounts of CPU time. This parameter improves the I/O transfer rate that the job achieves and may improve overall system throughput. However, this improvement is at the expense of a more exactly proportional allocation of CPU service to CPU intensive jobs.

The RS (recall CPU slice) SERVICE parameter defines the minimum guaranteed amount of time that a recalled job keeps the CPU before it can be preempted by another recalled job at the same CPU priority. This parameter is not involved in determining the amount of CPU time allocated to a job and normally should be set to the same value for all service classes. However, under unusual circumstances in which the I/O behavior of jobs in a service class is consistent and well understood, it may be possible to optimize their performance with a nonstandard RS value.

Memory Control

You can control the maximum memory allowed for job types and for service classes with the parameters you specify on the SERVICE and DELAY entries in the IPRDECK.

These parameters specify these lengths:

- Maximum field length divided by 100₈ for a job in a service class.
- Maximum field length divided by 100₈ for all jobs of the specified service class.
- Maximum extended memory length in words divided by 1000₈ for a job in a service class.
- Maximum extended memory length in words divided by 1000₈ for all jobs of the specified service class.
- Amount of central or extended memory to leave between job field length, if possible.

Initially, the scheduler attempts to find the highest priority job that meets the memory constraints. However, if the scheduler is unable to schedule a job and has explicitly rejected one or more jobs because the total field length or the total extended memory field length for a service class would be exceeded, it attempts to schedule a job a second time if flexible partitions are enabled (refer to FLEXIBLE PARTITIONS in this section). During this second attempt, any job that requires other jobs to be rolled out is not scheduled; otherwise, the constraints (service class total field length and total extended memory field length) are ignored, and the job is scheduled at the lower bound priority, LP (a parameter on the QUEUE entry). This means that the constraints are applied as long as there are enough jobs of each service class. However, if central memory is unused and no other jobs are available, the scheduler attempts to schedule the jobs without the constraints.

All of these parameters can be changed by using the SERVICE, QUEUE, and DELAY entries.

Example of Job Control Parameters

An example set of entries for the job control parameters is shown in tables 3-7 and 3-8. These entries fall within the example of ranges of service class priorities shown in figure 3-7. Neither the specific entries nor the ranges are recommended; they are strictly examples to aid you in selecting QUEUE and SERVICE entry parameters.

The following discussion indicates the significance of the values chosen and how they relate to each other.

The entry (lowest) priority (LP parameter in the QUEUE entry) of the system service class input queue is higher than all entry priorities, except the network supervision and subsystem entry priorities, because it is assumed that an operator-initiated job should receive prompt attention. A system job rolls out any batch job. The entry priority (LP) of the network supervision service class input queue is set high to ensure adequate response time from network programs and facilities such as CS, NS, and NVF.

The queue priorities for local batch, remote batch, and detached jobs are similar. The time slice for detached jobs is shorter than for the local batch and remote batch jobs. The assumption is that detached jobs need less CPU time than either local batch or remote batch jobs.

The queue priorities are explained under Job Scheduling earlier in this section.

The time slices for the various service classes reflect the following objectives:

- To keep system jobs with their high entry priority from monopolizing system resources.
- To keep at a minimum rollout activity caused by diagnostics running as maintenance service class jobs.
- To allow most interactive jobs to compile, load, and begin execution in one time slice.
- To give batch jobs a large time slice, because little is gained from rolling out batch jobs. There is no problem with the time slices for batch jobs compared to interactive jobs, because, with the priorities shown, an interactive job generally causes a batch job to roll out.
- To ensure prompt service to all interactive users, without employing an excessive number of rollouts, by setting the time slices for interactive jobs low. The time slice parameters are critical to good interactive performance. In some cases, depending on the system load, job size, and so forth, it may be desirable to change these parameters during operation.

Interactive performance is sensitive to the relative values of the QUEUE and SERVICE scheduling parameters, both within a service class and among service classes. Tables 3-7 and 3-8 show almost no overlaps of values except that local batch, remote batch, and detached jobs could age slightly past interactive jobs if they remained rolled out for a very long time (about 26 minutes). For the interactive service class, the range between entry (lowest) priorities and the highest priority is wide, so that few jobs are at the highest priority simultaneously. If many jobs reach the highest priority, their priorities are the same, and the order in which the jobs entered the queue is lost. The job scheduler selects jobs with equal queue priority in a random manner. The terminal I/O priority (TP) is set slightly higher than the initial priority (IP) to reduce response time for the user interacting with a job as compared to the user initiating a new job step. This parameter setting improves the perceived responsiveness of the system for the interactive users.

The CPU priorities reflect the following objectives.

- The maintenance service class jobs are run at the lowest priority. This handles the background CPU and memory diagnostics.
- All other jobs, except network supervision service class jobs, run at the same priority. It is generally not desirable to run one class of jobs at a higher priority than another because the system would roll in jobs that occupy memory without executing, until they exceed the central memory time slice.
- The network supervision service class is set high to ensure adequate performance from network programs and facilities such as CS, NS, and NVF.

The CPU slice parameters in the example are the same for all service classes. In this case, all jobs that contend for the CPU with the same CPU priority receive equal service. An individual job performing I/O may receive more CPU service if it has the non-zero CPU slice extension than CPU-bound jobs with the same CPU priority.

If sufficient memory is available to allow many jobs to remain in central memory at control points and pseudo-control points, it may be necessary to set differing CPU slice values (using the SERVICE command US parameter) to achieve the desired allocation of CPU time to jobs in different service classes. Job rollout is less frequent and may not deny service to jobs in service classes with relatively small central memory slices. Setting the CPU slice (SERVICE command US parameter) larger for one service class than for another results in jobs in the first service class receiving proportionately more CPU time than jobs in the second class when all the jobs are at control points and contending for the CPU.

Table 3-7. Example Set of Job Control Parameters, Part I

Service Class⁵	LP¹	UP¹	WF¹	LP²	UP²	WF²	IP²	IL²	TP³
SY (system)	7770	7776	1	2000	7000	1	7000	4000	–
BC (local batch)	10	4000	1	1000	4004	1	2000	2000	–
RB (remote batch)	10	4000	1	1000	4004	1	2000	2000	–
CT (communication task)	7770	7776	1	3000	7000	1	7000	4000	–
TS (interactive)	7000	7770	1	3700	7000	1	4004	3770	4024
NS (network supervision)	7770	7776	1	7770	7776	1	7772	7772	–
DI (detached)	10	4000	1	1000	4000	1	2000	2000	–
SS (subsystem)	7770	7776	1	7770	7776	1	7772	7772	–
MA (maintenance)	1	10	1	1	10	1	10	4	–
I _n (Installation) ⁴	10	4000	1	1000	4004	1	2000	2000	–

1. Input queue QUEUE parameter.
2. Execution queue QUEUE parameter.
3. Interactive job initial scheduling TP SERVICE parameter.
4. Installation classes I0, I1, I2, and I3 all have release values the same as class BC.
5. All values are octal; DELAY parameters are JS=1, CR=30, AR=1750, and MP=400.

Table 3-8. Example Set of Job Control Parameters, Part II

Service Class⁶	LP¹	UP¹	WF¹	CT²	CM²	CP³	NJ⁴	TD⁵
SY (system)	7000	7776	1	10	20	6770	7777	-
BC (local batch)	1	4000	1	20	200	3770	7777	-
RB (remote batch)	1	4000	1	20	200	3770	7777	-
CT (communication task)	7000	7776	1	10	200	6770	7777	-
TS (interactive)	1	7000	1	10	10	6770	7777	113 ⁸
NS (network supervision)	1	7000	1	10	200	7770	7777	-
DI (detached)	1	7000	1	20	20	3770	7777	341 ⁹
SS (subsystem)	7400	7776	1	10	20	7770	7777	-
MA (maintenance)	7000	7776	1	10	20	1	7777	-
I _n (installation ⁷)	1	7000	1	10	200	3770	7777	-

1. Output queue QUEUE parameter.
2. Time slice SERVICE parameters.
3. Control point slice priority CP SERVICE parameter.
4. Number of jobs NJ SERVICE parameter.
5. Time-out delay TD SERVICE parameter.
6. All values are octal; DELAY parameters are JS=1, CR=30, AR=1750, and MP=400.
7. Installation classes I0, I1, I2, and I3 all have release values the same as class BC.
8. A TS suspended job times out after 10 minutes when the time-out delay is 113.
9. A DI suspended job times out after 30 minutes when the time-out delay is 341.

Table 3-9. Example Set of Job Control Parameters, Part III

Service Class	PR¹	RS²	SE³	US⁴
SY (system)	30	4	20	20
BC (local batch)	30	4	20	20
RB (remote batch)	30	4	20	20
CT (communication task)	30	4	20	20
TS (interactive)	30	4	20	20
NS (network supervisor)	74	4	20	20
DI (detached)	30	4	20	20
SS (subsystem)	70	4	20	20
MA (maintenance)	2	4	20	20
In (installation) ⁵	30	4	20	20

1. CPU priority PR SERVICE parameter.
2. Recall CPU slice RS SERVICE parameter.
3. CPU slice extension SE SERVICE parameter.
4. Unextended CPU slice US SERVICE parameter.
5. Installation classes I0, I1, I2 and I3 all have release values the same as class BC.




Figure 3-7. Example of Ranges of Service Class Priorities

IPRDECK Entries Made Only During Deadstart

You can enter the following IPRDECK entries only during deadstart. They cannot be entered as DSD commands. Changes to the IPRDECK are not retained across deadstart unless a new deadstart file is created to reflect those changes. The entries are described in alphabetical order.

COMLIB

Format:

COMLIB, username, family, D.

Default:

None.

Significance:

This entry specifies a user name and family for a library of permanent files that can be accessed by users on all other families. This feature eliminates the need to maintain duplicate copies of the files on the families that require access to the files. This entry can also be used to reduce validation file accesses for user names such as LIBRARY in a single family environment. The number of COMLIB entries that can be specified is controlled by the CMRDECK entry CLT.

Parameter Description

username	The 1- to 7-character user name to be associated with the files to be accessed. This parameter should be unique to all families in the system. It cannot be associated with more than one family.
family	Family name to be associated with username. If 0 (zero) is specified, the system default family name is used.
D	If D is specified, the associated username and family are deleted from the common library table.

CPM

Format:

CPM, $s_1=n_1$, $s_2=n_2$.

Default:

System selection.

Significance:

This entry alters the central processor multiplier of type s_i , which is used in SRU calculations. The s_i parameters are either 0 or 1 to indicate the multipliers S0 or S1, respectively. Entering 0= n obtains a multiplier to be used for S0 and entering 1= n obtains a multiplier to be used for S1. (Refer to the NOS Version 2 Administration Handbook for a discussion of multiplier use.) The values of n_i range from 1 to 47₈ and are used as indexes to values defined in COMSSRU in order to determine the multiplier value. The default values are listed in table 3-10.

Table 3-10. CPM Default Values for n_i

n_i (Octal)	COMSSRU Name	Mainframe Model	COMSSRU Default Multiplier Value
1	CP62	6200	1.0
2	CP64	6400	1.0
3	CP65	6500	1.0
4	CP66	6600	1.0
5	CP67	6700	1.0
6	CP71	71	1.0
7	CP72	72	1.0
10	CP73	73	1.0
11	CP74	74	1.0
12	C171	171	1.0
13	C172	172	1.0
14	C173	173	1.0
15	C174	174	1.0
16	C175	175	1.0
17	C176	176	1.0
20	C720	720	1.0
21	C730	730	1.0

(Continued on next page)

Table 3-10. CPM Default Values for n_i

n_i (Octal)	COMSSRU Name	Mainframe Model	COMSSRU Default Multiplier Value
<i>(Continued from previous page)</i>			
22	C740	740	1.0
23	C750	750	1.0
24	C760	760	1.0
25	C810	810	1.0
26	C815	815	1.0
27	C825	825	1.0
30	C830	830	1.0
31	C835	835	1.0
32	C840	840	1.0
33	C845	845	1.0
34	C850	850	1.0
35	C855	855	1.0
36	C860	860/870	1.0
37	C865	865	1.0
40	C875	875	1.0
41	C961	960-11	1.0
42	C963	960-31/32	1.0
43	C990	990/994/995	1.0
44	ICM1	Model on which you are installing.	1.0
45	ICM2	Model on which you are installing.	2.0
46	ICM3	Model on which you are installing.	3.0
47	ICM4	Model on which you are installing.	4.0
50	ICM5	Model on which you are installing.	5.0

CSM**Format:**

CSM=csm.

Default:

64

Significance:

This entry sets the operating system character set mode. To change the character set mode for the products, a change must be made in IPARAMS, and the products must be reassembled.

csm	Description
63	63-character set.
64	64-character set.

The system assumes a 64-character set if there is no CSM entry in the current IPRDECK.

NOTE

Unpredictable and possibly serious problems occur if the operating system is operating in one character set and the products are operating in another. Therefore, ensure that all installed products and the operating system are in the same mode.

DISK VALIDATION**Format:**

ENABLE,DISK VALIDATION.
DISABLE,DISK VALIDATION.

Default:

Enabled.

Significance:

NOS performs hardware verification on each mass storage device during the deadstart process before users are allowed to access the device. The hardware verification sequence includes writing data to the disk, reading the data from the disk, and finally comparing the read data with the write data to ensure integrity. For 887 and 9853 disks, the hardware verification includes running in-line diagnostics. Since the diagnostics verify the quality of the disk media more completely, it is advantageous to run them. These diagnostics take approximately 3 minutes to execute, during which time the 887 or 9853 disk is not available for use. In order to keep deadstart time to a minimum (such as, during testing), disk validation can be disabled using the DISABLE,DISK VALIDATION command.

IPRDECK

DSD

Format:

DSD,level,command₁#command₂#...#command_n

Default:

None.

Significance:

This entry specifies the initial command(s) to be executed by the DSD program when the deadstart is complete. The commands can be the minimum number of characters recognizable by DSD but must include terminating characters. Only the last DSD entry on the IPRDECK for a given deadstart level is processed; other DSD entries for the same level are ignored. The DSD entry cannot exceed one line and a maximum of 60 characters can be specified.

Parameter Description

level Level of deadstart (0, 1, or 2).

command_i DSD command to be executed for the level of deadstart specified.

Several commands can be specified by separating them with the # (6-bit display code 60) or % (6-bit display code 63) character. These characters may misposition parts of the console display of IPRDECK if they appear as the upper 6 bits in a byte.

If you specify a series of commands including DSD display selection commands, the display selection commands must be the last commands. Failure to do so results in the system ignoring any commands following the DSD display selection commands.

Examples:

```
DSD,0,MAI%X.QREC(P0=N)
DSD,0,X.QREC(P0=N)#SET,AEIQ.#QB.
```

EI**Format:**

EI=nn₁, nn₂, . . . , nn_x.

Default:

None.

Significance:

The EI entry should be used to ensure loading of the correct environment interface for the operating system. This entry specifies up to six 2-character environment interface names, one of which must be loaded while deadstarting a CYBER 180-class machine. The deadstart will not be allowed to continue if the environment interface loaded by CTI is not one of those specified by the EI entry. You will have to load the correct environment interface or modify the EI entry in the IPRDECK. If an EI entry is not present, you will not be warned at deadstart.

EXTENDED STACK PURGING**Formats:**

ENABLE, EXTENDED STACK PURGING.
DISABLE, EXTENDED STACK PURGING.

Default:

Disabled.

Significance:

These entries specify the default action for instruction-stack purging for nonsystem-origin jobs on CYBER 180-class machines. Refer to the MODE macro in the NOS Version 2 Reference Set, Volume 4 for a description of instruction-stack purging.

HARDWARE FAULT INJECTION**Formats:**

ENABLE, HARDWARE FAULT INJECTION.
DISABLE, HARDWARE FAULT INJECTION.

Default:

Disabled.

Significance:

The entries enable and disable the simulation of hardware faults by normal jobs executing special instructions. This simulation is intended for test purposes only, and should not be enabled on a production system.

IPRDECK

KEYPM

Format:

KEYPM=cc.

Default:

26

Significance:

This entry specifies the keypunch mode to be assumed during system operation.

cc	Description
26	O26 keypunch mode.
29	O29 keypunch mode.

This entry is used for all batch jobs submitted if the keypunch mode is not specified on the job command. This does not apply to RBF.

MEMORY CLEARING

Formats:

ENABLE, MEMORY CLEARING.
DISABLE, MEMORY CLEARING.

Default:

Disabled.

Significance:

When memory clearing is enabled, central and extended memory are cleared when released from a job (that is, when a job is rolled out, terminates, or reduces its field length). When memory clearing is disabled, memory is cleared only when a job requests additional memory.

MICRO**Format:**

MICRO=name₁, name₂, . . . , name₇.

Default:

None.

Significance:

The MICRO entry should be used to ensure loading of the correct microcode for the operating system. This entry specifies up to seven 7-character microcode names, one of which must be loaded while deadstarting a CYBER 180-class machine. The deadstart will not be allowed to continue if the microcode loaded by CTI is not one of those specified by the MICRO entry. You will have to load the correct microcode or modify the MICRO entry in the IPRDECK. If a MICRO entry is not present, you will not be warned at deadstart.

NAMIAF**Format:**

NAMIAF=maxt .

Default:

200₈

Significance:

This entry specifies the number of network terminals that can be connected to IAF at one time. The maximum number of terminals IAF can support is 1039. However, this value is dependent on the number of network terminals, multiplexer ports, and stimulator ports specified in the deadstart decks.

Parameter Description	
maxt	Total number of network terminals; maxt can range from 1 to 1440 ₈ .

PROBE**Formats:**

ENABLE, PROBE .
DISABLE, PROBE .

Default:

Disabled.

Significance:

These entries enable and disable the data gathering facility of CPUMTR.

IPRDECK

SCP

Formats:

ENABLE, SCP.
DISABLE, SCP.

Default:

Disabled.

Significance:

These entries specify whether to use the system control point facility. You must enable SCP if CDCS, IAF, MAP, MCS, MSE, NAM, NVE, PLA, RBF, RHF, SMF, SSF, or TAF will be used. If none of these will be used, leave SCP disabled so that more CMR space is available.

SCRSIM

Formats:

ENABLE, SCRSIM.
DISABLE, SCRSIM.

Default:

Disabled.

Significance:

These entries enable or disable the simulation of the status/control register using the interlock register on CYBER 70 Computer Systems (refer to appendix D for information on the SCRSIM simulator).

SECCATS**Format:**

SECCATS=cat₁,cat₂, . . . ,cat_n.

Default:

All categories are enabled.

Significance:

This entry specifies the security access categories that will be allowed in the system for processing when the system is in a secured mode. The values should be supplied by a site security administrator. Refer to the NOS Version 2 Security Administrator's Handbook.

Parameter Description	
cat _i	Category names specified in deck COMSMLS corresponding to the desired access categories. Initially, all categories are enabled. The first SECCATS entry clears all categories, then sets the specified categories. Subsequent SECCATS entries set additional categories:
SECCATS=ALL.	Enables all 32 categories.
SECCATS=NUL.	Disables all 32 categories.

SPC**Format:**

SPC,d=lines.

Defaults:

64 lines at 6 lines per inch.

85 lines at 8 lines per inch.

Significance:

This entry specifies the charge in number of lines for a page of printed output at the specified print density on non-PFC printers.

Parameter Description	
d	Density in lines per inch; d can be 6 or 8 lines per inch, but must be specified in octal (6 or 10 ₈).
lines	Lines per page; lines can range from 16 to 255, but must be specified in octal (20 ₈ to 377 ₈).

IPRDECK

SPD

Format:

SPD=d.

Default:

6

Significance:

This entry specifies the assumed density for printed output.

d	Description
----------	--------------------

6	6 lines per inch.
---	-------------------

10 ₈	8 lines per inch.
-----------------	-------------------

SPL

Format:

SPL=length.

Default:

60

Significance:

This entry specifies the assumed page length in number of lines for printed output; length can range from 16 to 255, but must be specified in octal (20₈ to 377₈).

SPW

Format:

SPW=width.

Default:

136

Significance:

This entry specifies the assumed page width in number of characters for printed output; width can range from 40 to 136, but must be specified in octal (50₈ to 210₈).

SUBCP**Formats:**

ENABLE, SUBCP.
 DISABLE, SUBCP.

Default:

Disabled.

Significance:

These entries specify whether CPUMTR is to be initialized to handle subcontrol point (TAF) processing.

If SUBCP is disabled, CPUMTR is not initialized to handle subcontrol point processing. If you are not running TAF and if no user applications use subcontrol point processing, disable SUBCP so that CPUMTR uses less central memory.

TCVM**Format:**

TCVM=mode.

Default:

AS

Significance:

This entry sets the tape to be assumed during system operation.

mode	Description
AS	ASCII 9-track conversion.
US	ANSI (previously known as USASI) 9-track conversion (same as AS).
EB	EBCDIC 9-track conversion.

IPRDECK

TDEN

Format:

TDEN=density.

Default:

HY for 7-track tapes.

PE for 9-track tapes.

Significance:

This entry sets the system tape density. When the density is set, any tape unit accessed is automatically set to this density unless specified otherwise by a magnetic tape request. Two TDEN entries may be present, one for 7 track and one for 9 track.

density	Description
LO	200 cpi (7 track).
HI	556 cpi (7 track).
HY	800 cpi (7 track).
HD	800 cpi (9 track).
PE	1600 cpi (9 track).
GE	6250 cpi (9 track).

TDTY

Format:

TDTY=devicetype.

Default:

NT

Significance:

This entry sets the default tape device type.

devicetype	Description
AT	Automated Cartridge Subsystem (ACS).
CT	Cartridge Tape Subsystem (CTS).
MT	7-track tape device.
NT	9-track tape device.

TMS**Formats:**

ENABLE, TMS .
 DISABLE, TMS .

Default:

Disabled.

Significance:

These entries enable and disable the Tape Management System (TMS). These entries are meaningless if the TMS binaries are not installed in the deadstart file.

TMSTO**Format:**

TMSTO, N=to, S=to.

Default:

N=TC, S=FC

Significance:

This entry specifies the default values for the TMS tape options parameter for the LABEL command and the LABEL macro. The N parameter specifies the default for nonsystem origin jobs and the S parameter specifies the default for system origin jobs.

to	Description
TC	Sets the default to TO=TC (TMS processing with catalog error checking).
TE	Sets the default to TO=TE (TMS processing without catalog error checking).
FC	Sets the default to TO=FC (non-TMS processing with catalog error checking).
FE	Sets the default to TO=FE (non-TMS processing without catalog error checking).

This entry is meaningless if TMS is disabled.

IPRDECK

TRACE

Formats:

ENABLE, TRACE.
DISABLE, TRACE.

Default:

Disabled.

Significance:

The entries enable or disable specification of monitor functions and other system data to be traced, the TRACE and TRAP capabilities. To use these capabilities, ENABLE,SYSTEM DEBUG. must also be entered.

IPRDECK Entries

You can enter the following IPRDECK entries during deadstart and online during system operation using the L display or DSD commands. The reason for entering them online is to change the system's operation between deadstarts. Entries made during system operation are not retained across deadstarts. Entries made during deadstart are not retained across deadstarts unless a new deadstart file is created to reflect the changes.

A description of the enabling and disabling of subsystems follows. The remaining IPRDECK entries are described in alphabetical order.

Subsystems

You can initiate a subsystem by either of the following methods:

- Enter an explicit call to the subsystem procedure using a DSD command. The command format is:

```
subffff
```

Variable	Description
sub	3-character mnemonic for the name of the subsystem and subffff is the name of the subsystem procedure (refer to Subsystem Control Commands in section 5).

- Enter the DSD command AUTO or MAINTENANCE. This will initiate calls to all subsystems that are currently enabled. The calls will be to subsystem procedures with the same name as the name of the subsystem. You can choose to have either the AUTO command or the MAINTENANCE command issued automatically at deadstart time by specifying the command on the DSD entry in the IPRDECK.

NOTE

You cannot use either of the preceding methods to initiate the ATF, MCS, and RBF subsystems. The NAM subsystem automatically initiates them if they are selected in the NAM startup file.

A subsystem procedure may be either a record in the deadstart file or a permanent file under system user name SYSTEMX (user index 377777g). If a subsystem procedure with the same name is present in both places, the permanent file copy will be used.

NOTE

Before you initiate a subsystem that uses the system control point facility (such as NAM or CDCS), you must enable the system control point facility with the ENABLE,SCP IPRDECK entry. Refer to SCP earlier in this section.

All subsystems, except BIO and MAG, are disabled by default. You can enable or disable a subsystem by using IPRDECK entries or by using the SUBSYST L-display utility. (For information about the SUBSYST L display, refer to the NOS Version 2 Operations Handbook.)

The IPRDECK entries that enable and disable subsystems may also assign a required control point for a subsystem. Thus, even if you do not want to enable a subsystem, you may want to explicitly disable a subsystem with an IPRDECK entry to assign a required control point for the subsystem.

NOTE

If you choose to specify a control point for any subsystem, it is recommended that you specify a unique control point for each subsystem. Otherwise, a subsystem may not be able to initiate successfully because its designated control point is already occupied by another subsystem. When this happens, the subsystem being initiated waits indefinitely for the other subsystem to give up the control point.

Use these IPRDECK entries to enable or disable a subsystem:

ENABLE, subsystem, cp.
 DISABLE, subsystem, cp.

Parameter Description

subsystem Three characters that select the desired subsystem; one of these values:

subsystem	Description
ATF	Automated Tape Facility. The ATF subsystem must be active before users can access the tape units for the Automated Cartridge Subsystem (ACS).
BIO	Local batch I/O for central site line printers, card readers, and card punches.
CDC	CYBER Database Control System.
CYB	CYBIS-NAM Interface.
IAF	Interactive Facility. Do not specify the cp parameter on the ENABLE entry for IAF.
MAG	Magnetic Tape Subsystem. Even if your users do not use magnetic tapes, enable MAG if removable auxiliary packs are used. Disabling MAG frees a control point for other use.
MAP	MAP III or MAP IV.
MCS	Message Control System.
MSE	Mass Storage Extended Subsystem.
NAM	Network Access Method.
NVE	NOS/VE Dual-State Interface.
RBF	Remote Batch Facility.

Parameter Descriptionsubsystem *(Continued)***subsystem Description**

RDF	Remote Diagnostic Facility. Do not specify the CP parameter on the ENABLE entry for RDF.
RHF	Remote Host Facility.
SMF	Screen Management Facility.
SSF	NOS-SCOPE 2 Station Facility.
STM	Interactive Stimulator.
TAF	Transaction Facility.

NOTE

Since ATF, MCS and RBF are started by the NAM subsystem (if they are selected in the NAM startup file), the ENABLE and DISABLE commands establish only the control points at which the subsystems will run.

If an ENABLE is made for NVE and the subsystem procedure name is also NVE, use the DOWN,CH EQPDECK entry to automatically down any channels that NOS/VE uses.

If STIMULA is used as the interactive stimulator, you must enable the STM subsystem at the last control point.

cp Control point where the subsystem will reside. If you omit cp, the current control point is used. If you enter 0 (zero) as the control point, the subsystem may reside at any available control point. If you enter -n as the control point, the subsystem's control point will be relative to the system control point (-1 would be the last control point before the system control point).

Omit cp for IAF and RDF.

ACS TAPE PF STAGING**Formats:**

ENABLE,ACS TAPE PF STAGING.
DISABLE,ACS TAPE PF STAGING.

Default:

Disabled.

Significance:

These entries specify whether permanent files that reside on Automated Cartridge Subsystem (ACS) tape alternate storage are automatically staged to disk. If disabled, jobs that attempt to access permanent files residing only on ACS tapes are aborted. Contrast this entry with the CARTRIDGE PF STAGING entry, which controls staging from MSE or optical disk, and the TAPE PF STAGING entry, which controls staging from 9-track tapes and CTS tapes.

AUTORESTART

Formats:

ENABLE, AUTORESTART.
DISABLE, AUTORESTART.

Default:

Enabled.

Significance:

These entries enable or disable the AUTORESTART installation parameter. If AUTORESTART is enabled and if step mode is set during an environmental shutdown, the system automatically clears step mode when the power environment returns to normal. All jobs waiting in the input and rollout queues are initiated by performing the processing associated with the DSD AUTO command. This recovery process requires no operator interaction.

If AUTORESTART is disabled and step mode is set during an environmental shutdown, the operator must manually enter the UNSTEP and AUTO commands when the message POWER/ENVIRONMENT NORMAL is displayed at the system control point.

CARTRIDGE PF STAGING

Formats:

ENABLE, CARTRIDGE PF STAGING.
DISABLE, CARTRIDGE PF STAGING.

Default:

Disabled.

Significance:

These entries specify whether permanent files that reside on either MSE or optical disk alternate storage are staged to disk. If disabled, jobs that attempt to access permanent files residing only on MSE or optical disk alternate storage are aborted. Contrast this entry with the ACS TAPE PF STAGING entry, which controls staging from ACS tapes, and the TAPE PF STAGING entry, which controls staging from 9-track tapes and CTS tapes.

CLASS

Format:

CLASS,ot,sc₁,sc₂,...,sc_n.

Default:

Origin Type	Service Class
Batch	BC
Remote batch	RB
Interactive	TS

Significance:

Specifies the valid service classes for each origin type. The system stores this information in the service class table (refer to the NOS Version 2 Systems Programmer's Instant).

Parameter Description

ot	Origin type; must be BC (batch), RB (remote batch), or IA (interactive). This parameter is required and order dependent.
sc _i	Service classes. Each class selected causes validation for that service class for the origin type specified by ot. Entering a service class that already has validation clears validation for that service class. You can select from 1 to 36 service classes for each origin type. This command does not accept class SS (subsystem).

Service class is one of these values:

sc	Description
BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class n ($0 \leq n \leq 8$).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SY	System.
TS	Interactive.
ALL	Causes validation of all service classes except subsystem (SS) and deadstart (DS).
NUL	Clears validation for all service classes.

IPRDECK

CPTT

Format:

CPTT=n.

Default:

100₈ sectors.

Significance:

This entry specifies the CPUPFM transfer threshold (CPTT). The value of CPTT determines whether an indirect access file transfer will be processed entirely within the PP program PFM or whether the CPU program CPUPFM will be called to perform the transfer. Since CPUPFM can transfer large files faster than PFM, CPUPFM transfers files that exceed the threshold value.

Depending on the configuration, a site can adjust this value up or down. A site with buffered I/O devices, for example, may find it advantageous to set a lower value. If a site wishes to disable CPUPFM transfers altogether, a value of zero (CPTT=0) may be specified.

Parameter	Description
n	CPUPFM transfer threshold value that will be defined in common deck COMSPFM; $0 \leq n \leq 7777_8$ PRUs.

DEBUG

Format:

DEBUG.

Default:

Disabled.

Significance:

This entry selects or clears debug mode, depending upon the current status.

If enabled, debug mode is selected. The message DEBUG appears in the header of the left screen display. Debug mode provides system origin privileges to validated users and allows modifications to be made to the running system.

If disabled, debug mode is cleared. It is recommended that debug mode not be allowed in a normal production environment.

On a secured system, this entry is ignored. While the console is in security unlock status, debug mode can be set using DSD command DEBUG.

DELAY**Format:**

DELAY, ARar, CRcr, JQjq, JSjs, MPmp.

Default:

None.

Significance:

This entry specifies the system delay parameters. Refer to table 3-7 for an example set of parameter entries. Refer to figure 3-6 for an example of ranges of service class priorities.

Parameter Description

Parameter	Description
ARar	PP recall interval in milliseconds. This parameter specifies the default time interval after which a peripheral processor program in the PP recall queue will be recalled. ar ranges from 1 to 7777 ₈ .
CRcr	CPU recall period in milliseconds. This parameter specifies the amount of time a job remains in recall (X status) when an RCL request is placed in RA+1. cr ranges from 1 to 7777 ₈ .
JQjq	Exponent used to determine the input file scheduling interval in seconds. jq ranges from 0 to 14 ₈ . The interval in seconds between scheduling of input files is calculated as follows: $\text{interval} = 2^{**jq}$
JSjs	Job scheduler interval in seconds. This parameter specifies the interval in which the job scheduler is called. The scheduler may also be called at other times. js ranges from 1 to 7777 ₈ .
MPmp	Memory padding value expressed as 100 ₈ word blocks. mp ranges from 0 to 777 ₈ . This parameter specifies how much additional (unassigned) memory should be allocated between the end of the newly assigned job field length and the beginning of the next job. Increasing this value reduces the probability that a job will be storage moved in response to a request for more memory.

DDP ROLLOUT PATH

Formats:

ENABLE,DDP ROLLOUT PATH.
DISABLE,DDP ROLLOUT PATH.

Default:

Disabled.

Significance

These entries determine which path a rollout to extended memory will take (CPU or PP). The required supporting hardware (EM or DDP) must be present.

DFPT

Format:

DFPT=dt_n.

Default:

DI1 (see the DFPT installation parameter in common deck COMSPFM).

Significance:

The DFPT entry resets the system default removable pack type. When accessing a removable auxiliary device with a permanent file command, the system checks that the equipment type and pack name of the device match the equipment type (R parameter) and pack name (PN parameter) on the command. If R is not specified, the system uses the equipment type specified by DFPT.

Parameter Description

dt	Disk device type of the removable pack.
n	Number of spindles in the removable pack; $1 \leq n \leq 8$.

ENGR**Formats:**

ENABLE, ENGR.
DISABLE, ENGR.

Default:

Disabled.

Significance:

These entries enable or disable engineering mode. If enabled, the ENGR message appears in the header of the left screen display. Engineering mode allows the peripheral processing unit (PPU)/hardware diagnostics and the 881/883 pack reformatting utility FORMAT to run while the system is in operation.

On a secured system, these entries are ignored. When the console is in security unlock status, engineering mode can be enabled using the DSD command ENABLE,ENGR.

FLEXIBLE PARTITIONS**Formats:**

ENABLE, FLEXIBLE PARTITIONS.
DISABLE, FLEXIBLE PARTITIONS.

Default:

Enabled.

Significance:

This entry enables or disables flexible memory partitioning. The job scheduler attempts to use memory space to the greatest extent possible when flexible partitions are enabled. Some service classes may be allotted more total memory space than memory partitioning constraints normally allow. If flexible partitions are disabled, the total memory used by jobs of a given service class are never allowed to exceed the maximum specified on the SERVICE command (AM and EM parameters).

FORCED FORMATTING

Formats:

ENABLE, FORCED FORMATTING.
DISABLE, FORCED FORMATTING.

Default:

Enabled.

Significance:

These commands are applicable only to 583x devices. When FORCED FORMATTING is enabled and the INITIALIZE,FP option for a device is also specified, the driver always formats the device. The driver does not check whether the device is already formatted correctly.

When FORCED FORMATTING is disabled and the INITIALIZE,FP option is specified for a device, the driver firsts checks whether the device is correctly formatted. If the device is already formatted correctly, the driver does not format the device. However, the driver does issue the clustering commands to the controller when necessary.

If FORCED FORMATTING is enabled when the system detects that a parity drive that was previously offline is now online, the system automatically formats the drive before restoring data.

If FORCED FORMATTING is disabled when the system detects that a parity drive that was previously offline is now online, the system first checks whether the drive is formatted correctly. The driver formats the drive only when necessary.

LOCK**Format:**

LOCK.

Default:

Unlocked.

Significance:

This entry specifies the system is locked. This software function prevents entry of restricted commands; all other DSD commands can be entered. The console is normally locked when the system is being used in a production environment.

LOGGING**Formats:**

ENABLE, LOGGING.
DISABLE, LOGGING.

Default:

Disabled.

Significance:

These entries specify whether dayfile messages intended for system analysis are logged in the dayfile. Typical messages deal with informing the user that the program is making inefficient CIO calls, such as reading to a full buffer or writing from an empty buffer. The dayfile messages are documented in an appendix of the NOS Version 2 Operations Handbook.

MASTER MSE

Formats:

ENABLE, MASTER MSE.
DISABLE, MASTER MSE.

Default:

Disabled.

Significance:

These entries specify whether the MSE executive program (SSEXEC), when initialized, is to run in master (enabled) or slave (disabled) mode.

MS VALIDATION

Formats:

ENABLE, MS VALIDATION.
DISABLE, MS VALIDATION.

Default:

Disabled.

Significance:

This entry enables or disables mass storage validation. If enabled, CMR is increased by 48 words, and the system verifies that, for each mass storage device, the sum of the counts of unreserved tracks and preserved files equals values specified in the device's mass storage table.

If the device is a master device (contains user catalogs), the system also verifies these conditions:

- The device's track reservation table (TRT) specifies that the first tracks of the indirect access file chain and the permit area are reserved and preserved.
- The label track is linked to the first catalog track.
- The number of catalog tracks is a power of 2.
- The catalog chain is reserved, of correct length, and contiguous if flagged as such in the device's MST.

To enable/disable mass storage validation with a DSD command entry, enable the validation in the IPRDECK during a level 0, 1, or 2 deadstart.

OQSH**Format:**

OQSH=level.

Default:

None.

Significance:

This entry specifies the value of the output queue special handling (OQSH) level. Output files with an access level equal to or greater than the output queue special handling level will remain in the queue until released by the DSD command RELEASE. When no OQSH level or access level name LVLO is selected, all files will be processed. The value for this entry is supplied by the site security administrator.

Parameter Description

level	Access level name specified in deck COMSMLS (refer to the NOS Version 2 Installation Handbook) that corresponds to the desired output queue special handling level.
-------	---

PCLASS**Format:**

CLASS, sc₀, sc₁, . . . , sc₇.

Default:

None.

Significance:

Specifies the service class associated with each priority level (P0 through P7) for selection on the Job command (refer to the NOS Version 2 Reference Set, Volume 3).

Parameter Description

sc _i	The 2-character service class symbol. Because parameters are positional, a comma must appear for any null parameter. The default for a null parameter is that no service class will be associated with the priority level represented by the null parameter's position in the string. This command does not accept the SS (subsystem) service class. Refer to the SERVICE IPRDECK entry later in this section for a list of service classes.
-----------------	--

PF VALIDATION

Formats:

ENABLE, PF VALIDATION.
DISABLE, PF VALIDATION.

Default:

Disabled.

Significance:

These entries enable or disable preserved file (PF) validation. If enabled, the system aborts an attach of a direct access permanent file if its end-of-information (EOI) was altered during recovery of the file. If NA (no abort) is specified on the attach request, the system attaches the file.

If mass storage validation is also enabled, TRT verification of preserved files takes place during a level 3 deadstart as follows:

- For all files, the system ensures that all tracks are reserved and that no circular linkage exists.
- For all queued, permanent direct-access, and fast-attach files, the system also ensures that the first track is preserved.

If mass storage validation is enabled on a level 1 or 2 deadstart, TRT verification takes place automatically, regardless of the status of PF VALIDATION.

PRIVILEGED ANALYST MODE

Formats:

ENABLE, PRIVILEGED ANALYST MODE.
DISABLE, PRIVILEGED ANALYST MODE.

Default:

Disabled.

Significance:

These entries enable or disable privileged analyst mode operations. If enabled, a user validated with AW=CPAM (refer to the NOS Version 2 Administration Handbook for information on MODVAL validation) is permitted to read system status information (such as the system dayfile, account file and error log) using a nonsystem-origin job. On a secured system, this entry is ignored. Refer to the DSD ENABLE/DISABLE command in section 5 for more details.

PRIVILEGED RDF**Formats:**

ENABLE, PRIVILEGED RDF.
 DISABLE, PRIVILEGED RDF.

Default:

Disabled.

Significance:

These entries enable or disable privileged mode of RDF. If enabled, a user's commands are checked to ensure that a maintenance function is being performed.

QUEUE**Format:**

QUEUE, sc, qt, ILil, IPip, LP1p, UPup, WFwf.

Default:

None.

Significance:

This entry specifies the queue priorities associated with the input, executing, and output queues for each job service class. Refer to table 3-7 for an example set of parameter entries and to figure 3-6 for an example of ranges of service class priorities.

Parameter Description

sc Service class.

sc Description

BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class n ($0 \leq n \leq 8$).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SS	Subsystem.
SY	System.
TS	Interactive.

Parameter Description

qt	Description
qt	Job queue type.
qt	Description
EX	Executing jobs.
IN	Input queued files.
OT	Output queued files.
ILil	Priority a job receives when it initially exhausts its CM time slice (refer to the SERVICE command CM parameter). Thereafter, whenever the CM time slice is exhausted the job's priority will be set to the value of lp. il ranges from 0 to 7777 ₈ , but must be in the range of values for lp and up. This parameter is valid only for executing jobs (EX).
IPip	Initial priority only for an executing batch job or for an interactive job when a terminal command is entered. Online interactive jobs with terminal I/O available are scheduled at tp priority (refer to the SERVICE command TP parameter). ip ranges from 0 to 7777 ₈ , but must be in the range of values for lp and up.
LPlp	Lowest priority. For an input or output queue file, the priority assigned to a file or job entering the queue. For an executing job, the priority assigned to job which has exceeded a non-initial CM slice. The IL value is used for the first CM slice. lp ranges from 0 to 7777 ₈ , but must be less than the value of up.
UPup	Highest priority. For input and output queues, this is the highest priority a file can reach in that queue; aging stops when this priority is reached. For the execution queue, this is the priority assigned to a job when initially assigned to a control point. up ranges from 0 to 7777 ₈ , but must be greater than the value of lp.
WFwf	Weighting factor for queue priority calculation; wf must be a power of 2, from 1 to 4000 ₈ . The smaller the weighting factor, the faster the queue entry reaches its highest priority.

REMOVABLE PACKS

Formats:

ENABLE, REMOVABLE PACKS.
DISABLE, REMOVABLE PACKS.

Default:

Enabled.

Significance:

These entries enable or disable automatic label checking for mass storage devices that are defined as removable.

If enabled, automatic label checking occurs. This status must be available to perform label verification before removable devices can be accessed.

If disabled, any removable devices introduced into the system will not be recognized.

RESIDENT RDF

Formats:

ENABLE, RESIDENT RDF.
DISABLE, RESIDENT RDF.

Default:

Disabled.

Significance:

These entries enable or disable resident mode of RDF. While in resident mode, RDF remains active, regardless of terminal inactivity, until RDF is disabled. When resident mode is disabled, RDF becomes inactive if no one is logged on at the remote diagnostic terminal for 15 minutes.

SECONDARY USER COMMANDS

Formats:

ENABLE, SECONDARY USER COMMANDS.
DISABLE, SECONDARY USER COMMANDS.

Default:

Disabled.

Significance:

The enable option allows jobs to issue more than one USER command. These entries are ignored on a secured system since secondary USER commands are not allowed on a secured system.

SECURES**Format:**

SECURES, ot, LA=1a, UA=ua.

Default:

The security access limits of the system origin type (SY) are set to access level name LVLO. All security access limits for other origin types are set to system limits.

Significance:

This entry specifies system limits for the system origin type (ot equals SY) or limits for other origin types (ot equals BC, IA, or RB). System limits define the highest and lowest levels allowed in the system. No job may execute and no file may be accessed or created at a level outside this range.

Limits for origin types other than SY must be within the system limits. If origin type SY is specified, limits for all origin types are set to the selected values.

Both 1a and ua must be entered; they can be set to the same value, restricting system access or a particular origin type access to a single level. These values are supplied by a site security administrator.

This entry is not meaningful in an unsecured system.

Parameter Description

Parameter	Description										
ot	Origin type.										
	<table border="1"> <thead> <tr> <th>ot</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>SY</td> <td>System.</td> </tr> <tr> <td>BC</td> <td>Batch.</td> </tr> <tr> <td>IA</td> <td>Interactive.</td> </tr> <tr> <td>RB</td> <td>Remote batch.</td> </tr> </tbody> </table>	ot	Description	SY	System.	BC	Batch.	IA	Interactive.	RB	Remote batch.
ot	Description										
SY	System.										
BC	Batch.										
IA	Interactive.										
RB	Remote batch.										
1a	Access level name specified in deck COMSMLS (refer to the NOS Version 2 Installation Handbook) corresponding to the desired lower access level limit.										
ua	Access level name specified in deck COMSMLS corresponding to the desired upper access level limit.										

SERVICE**Format:**

SERVICE, sc, AMam, CMcm, CPcp, CScs, CTct, DSds, DTsc, ECec, EMem, FCfc, FLfl, FSfs, NJnj, PRpr, RSrs, SEse, TDtd, TPtp, USus.

Default:

None.

Significance:

This entry specifies the service limits associated with each service class. Refer to table 3-7 for an example set of parameter entries and to figure 3-6 for an example of ranges of service class priorities.

Parameter	Description																						
sc	Service class; one of these values: <table border="1" data-bbox="451 758 1346 1199"> <thead> <tr> <th>sc</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>BC</td> <td>Local batch.</td> </tr> <tr> <td>CT</td> <td>Communication task.</td> </tr> <tr> <td>DI</td> <td>Detached interactive.</td> </tr> <tr> <td>In</td> <td>Installation class n ($0 \leq n \leq 8$).</td> </tr> <tr> <td>MA</td> <td>Maintenance.</td> </tr> <tr> <td>NS</td> <td>Network supervisor.</td> </tr> <tr> <td>RB</td> <td>Remote batch.</td> </tr> <tr> <td>SS</td> <td>Subsystem.</td> </tr> <tr> <td>SY</td> <td>System.</td> </tr> <tr> <td>TS</td> <td>Interactive.</td> </tr> </tbody> </table>	sc	Description	BC	Local batch.	CT	Communication task.	DI	Detached interactive.	In	Installation class n ($0 \leq n \leq 8$).	MA	Maintenance.	NS	Network supervisor.	RB	Remote batch.	SS	Subsystem.	SY	System.	TS	Interactive.
sc	Description																						
BC	Local batch.																						
CT	Communication task.																						
DI	Detached interactive.																						
In	Installation class n ($0 \leq n \leq 8$).																						
MA	Maintenance.																						
NS	Network supervisor.																						
RB	Remote batch.																						
SS	Subsystem.																						
SY	System.																						
TS	Interactive.																						
AMam	Maximum field length divided by 100_8 for all jobs of the specified service class. This parameter partitions central memory by limiting the field length available to each service class. For example, if scheduling a job to a control point would cause the AM value for its service class to be exceeded for its service class, it may not be scheduled until the required field length is available. This means that a lower priority job from a different service class may be scheduled first. However, a job that would cause the AM value for its service class to be exceeded, can be scheduled to a control point if not enough jobs in other service classes exist to fill central memory and FLEXIBLE PARTITIONS is enabled. The system attempts to use central memory to its greatest capacity. am ranges from 0 to 77777_8 .																						

Parameter	Description																		
CMcm	Central memory time slice in seconds. This parameter specifies the maximum amount of time a job of the specified service class can remain in central memory either at a control point or at a pseudo-control point before its priority is set to the lower boundary. When the job initially exhausts its CM time slice, this lower boundary is the value of il. Subsequently, when the job exhausts its CM time slice, this lower boundary is the value of lp (refer to the QUEUE command IL and LP parameters). cm ranges from 0 to 7777 ₈ .																		
CPcp	Control point slice priority. This parameter specifies the value of the scheduling priority which will be set after a job has been at a control point for ct seconds. The priority specified must be greater than or equal to the lower bound and less than or equal to the upper bound execution queue scheduling priorities specified by the CB parameter (lp ≤ cp ≤ up). Refer to the QUEUE command LP and UP parameters.																		
CScs	Cumulative size in PRUs allowed for all indirect access permanent files. cs indicates a limit value for the cumulative size.																		
	<table border="1"> <thead> <tr> <th>cs</th> <th>Limit Value</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Unlimited</td> </tr> <tr> <td>1</td> <td>1000₈</td> </tr> <tr> <td>2</td> <td>5000₈</td> </tr> <tr> <td>3</td> <td>50000₈</td> </tr> <tr> <td>4</td> <td>100000₈</td> </tr> <tr> <td>5</td> <td>200000₈</td> </tr> <tr> <td>6</td> <td>400000₈</td> </tr> <tr> <td>7</td> <td>Unlimited</td> </tr> </tbody> </table>	cs	Limit Value	0	Unlimited	1	1000 ₈	2	5000 ₈	3	50000 ₈	4	100000 ₈	5	200000 ₈	6	400000 ₈	7	Unlimited
cs	Limit Value																		
0	Unlimited																		
1	1000 ₈																		
2	5000 ₈																		
3	50000 ₈																		
4	100000 ₈																		
5	200000 ₈																		
6	400000 ₈																		
7	Unlimited																		
CTct	Control point time slice in seconds. This parameter specifies the maximum amount of time that a job of the specified priority can remain at a control point before its priority is changed to the control point slice priority (cp). ct ranges from 1 to 7777 ₈ .																		

Parameter	Description																		
DSds	Size in PRUs allowed for individual direct access permanent files. ds indicates the limit value for the size of the files.																		
	<table border="1"> <thead> <tr> <th>ds</th> <th>Limit Value</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Unlimited</td> </tr> <tr> <td>1</td> <td>1000₈</td> </tr> <tr> <td>2</td> <td>5000₈</td> </tr> <tr> <td>3</td> <td>50000₈</td> </tr> <tr> <td>4</td> <td>100000₈</td> </tr> <tr> <td>5</td> <td>200000₈</td> </tr> <tr> <td>6</td> <td>400000₈</td> </tr> <tr> <td>7</td> <td>Unlimited</td> </tr> </tbody> </table>	ds	Limit Value	0	Unlimited	1	1000 ₈	2	5000 ₈	3	50000 ₈	4	100000 ₈	5	200000 ₈	6	400000 ₈	7	Unlimited
ds	Limit Value																		
0	Unlimited																		
1	1000 ₈																		
2	5000 ₈																		
3	50000 ₈																		
4	100000 ₈																		
5	200000 ₈																		
6	400000 ₈																		
7	Unlimited																		
DTsc	Service class to which a detached job is assigned if the job is detached. The default value for sc is DI (detached interactive).																		
ECec	Maximum user-accessible extended memory field length in words divided by UEBS ²⁰ for any job of the specified service class. This parameter performs the same function for extended memory field length that the FL parameter does for central memory field length. ec ranges from 0 to 3777 ₈ .																		
EMem	Maximum extended memory length in words divided by UEBS ¹⁹ for all jobs of the specified service class. This parameter performs the same function for extended memory field length that the AM parameter does for central memory field length. em ranges from 0 to 3777 ₈ .																		

²⁰. Refer to Extended Memory Overview earlier in this section to determine the value of UEBS.

Parameter	Description																		
FCfc	Number of permanent files allowed. fc indicates a limit value that is the maximum number of permanent files allowed.																		
	<table border="1"> <thead> <tr> <th>fc</th> <th>Limit Value</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Unlimited</td> </tr> <tr> <td>1</td> <td>10_g</td> </tr> <tr> <td>2</td> <td>40_g</td> </tr> <tr> <td>3</td> <td>100_g</td> </tr> <tr> <td>4</td> <td>200_g</td> </tr> <tr> <td>5</td> <td>1000_g</td> </tr> <tr> <td>6</td> <td>4000_g</td> </tr> <tr> <td>7</td> <td>Unlimited</td> </tr> </tbody> </table>	fc	Limit Value	0	Unlimited	1	10 _g	2	40 _g	3	100 _g	4	200 _g	5	1000 _g	6	4000 _g	7	Unlimited
fc	Limit Value																		
0	Unlimited																		
1	10 _g																		
2	40 _g																		
3	100 _g																		
4	200 _g																		
5	1000 _g																		
6	4000 _g																		
7	Unlimited																		
FLfl	Maximum field length divided by 100 _g for any job of the specified service class. If more memory is requested, the job is aborted. You typically use this parameter to limit the memory requirement for jobs of a specific service class during certain hours of the day. For example, you may use the FL parameter to specify a maximum field length for all batch service class jobs between the hours of 2 p.m. and 4 p.m. fl ranges from 0 to 7777 _g .																		
FSfs	Size in PRUs allowed for individual indirect access permanent files. fs indicates a limit value for the size of the files.																		
	<table border="1"> <thead> <tr> <th>fs</th> <th>Limit Value</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Unlimited</td> </tr> <tr> <td>1</td> <td>10_g</td> </tr> <tr> <td>2</td> <td>30_g</td> </tr> <tr> <td>3</td> <td>100_g</td> </tr> <tr> <td>4</td> <td>300_g</td> </tr> <tr> <td>5</td> <td>1000_g</td> </tr> <tr> <td>6</td> <td>2000_g</td> </tr> <tr> <td>7</td> <td>Unlimited</td> </tr> </tbody> </table>	fs	Limit Value	0	Unlimited	1	10 _g	2	30 _g	3	100 _g	4	300 _g	5	1000 _g	6	2000 _g	7	Unlimited
fs	Limit Value																		
0	Unlimited																		
1	10 _g																		
2	30 _g																		
3	100 _g																		
4	300 _g																		
5	1000 _g																		
6	2000 _g																		
7	Unlimited																		

Parameter	Description
NJnj	Maximum number of jobs. For each service, this parameter specifies the number of jobs that can be executing in the system. nj ranges from 0 to 7777 ₈ .
PRprS	CPU priority. When more than one job is waiting for the CPU, the jobs with the highest CPU priority access the CPU first. Jobs with lower CPU priority access the CPU only when higher-priority jobs no longer wish to use the CPU. pr ranges from 2 to 77 ₈ .
RSrs	Recall CPU slice in milliseconds. The minimum guaranteed amount of time that a recalled job keeps the CPU before it can be preempted by another recalled job at the same CPU priority. rs ranges from 1 to 7777 ₈ , but the value specified must be less than or equal to the value specified for the unextended CPU slice, us.
SEse	CPU slice extension in milliseconds. The maximum additional amount of CPU time that a job performing I/O gets after its unextended CPU slice expires before the CPU slice expires (and the CPU is switched to an equal-priority job). This is the maximum extension. The slice is actually extended from the amount already used by the value of the recall CPU slice, rs, each time that the job is recalled. se ranges from 0 to 7777 ₈ .
TDtd	Suspension time-out delay. A suspended job will not be timed out for td*10 ₈ seconds. The maximum delay is approximately 9 hours. td ranges from 0 to 7777 ₈ .
TPtp	Terminal job priority assigned to a terminal job that is rolled out waiting for terminal input. tp ranges from 0 to 7777 ₈ , but must be in the range of values for the execution queue parameters lp and up (refer to the QUEUE command LP and UP parameters).
USus	Unextended CPU slice in milliseconds. The amount of CPU time that a CPU-bound job gets before the CPU is switched to an equal-priority job. us ranges from 1 to 7777 ₈ , but the value specified for us must be greater than or equal to the value specified for the recall CPU slice, rs.

SPINDOWN

Formats:

ENABLE, SPINDOWN.
DISABLE, SPINDOWN.

Default:

Enabled.

Significance:

These entries enable or disable the spindown of 834, 836, 887, and 9853 disk units. When spindown is enabled, entering a CHECK POINT SYSTEM command will cause all 834, 836, 887, and 9853 disk units that are on and not globally unloaded to automatically spin down. When spindown is disabled, entering a CHECK POINT SYSTEM command will not spin down the 834, 836, 887, and 9853 disk units. System checkpoints that are initiated by mainframe errors do not affect the state of these disk units.

SRST

Format:

SRST=n.

Default:

0

Significance:

This entry specifies the secondary rollout sector threshold. Any rollout file smaller than n sectors ($0 \leq n \leq 7777_8$) is considered a secondary rollout file for the purpose of equipment selection.

NOTE

The size of the rollout file for any job must be at least seven sectors larger than the combined size in sectors of the job's central memory and extended memory field lengths.

SYSTEM DEBUG

Formats:

ENABLE, SYSTEM DEBUG.
DISABLE, SYSTEM DEBUG.

Default:

Disabled.

Significance:

These entries enable or disable the system debug mode of operation. When the system is in system debug mode, it is less tolerant of system errors; that is, it is more likely to hang upon experiencing errors. When the system is not in system debug mode, it rates system errors as critical or noncritical. For critical errors, the system partially or totally interrupts system operation to tend to the errors. For noncritical errors, the system logs them in the binary maintenance log (BML) and inasmuch as possible allows system operation to proceed. You can initiate the system debug mode with the DSD ENABLE command or the corresponding IPRDECK entry.

There is another system state called debug mode, which is conceptually different from that of system debug mode. Debug mode is the state of the system where a user with system origin privileges can make modifications to the running system. You can initiate this mode of operation with the DSD command DEBUG. The left screen header of the system console indicates whether the system is in debug mode or not.

TAPE PF STAGING

Formats:

ENABLE,TAPE PF STAGING.
DISABLE,TAPE PF STAGING.

Default:

Disabled.

Significance:

These entries specify whether permanent files that reside on 9-track and CTS tape alternate storage are staged to disk. If disabled, jobs that attempt to access permanent files residing only on 9-track tapes or CTS tapes are aborted. Contrast this entry with the ACS TAPE PF STAGING entry, which controls staging from ACS tapes, and the CARTRIDGE PF STAGING entry, which controls staging from MSE or optical disk.

UNLOCK

Format:

UNLOCK.

Default:

Unlocked.

Significance:

This entry specifies the system console is unlocked. All DSD commands can be entered when the console is unlocked. The console is usually locked when the system is being used in a production environment. Refer to LOCK earlier in this section.

USER EXTENDED MEMORY

Formats:

ENABLE,USER EXTENDED MEMORY.
DISABLE,USER EXTENDED MEMORY.

Default:

Disabled.

Significance:

These entries enable or disable scheduling of jobs that access user extended memory.

LIBDECK

LIBDECK is a SYSEDIT directive record on the deadstart file. SYSEDIT reads LIBDECK during the system load. LIBDECK specifies program residence, field length, record type, and parameter format.

Up to 64 LIBDECK records can be placed on the deadstart file. LIBDECKs are named LIBDnn, where nn is 00 to 77₈. A specific record can be selected with a LIB=n entry in CMRDECK. The multiple deadstart deck capability enables the use of a single deadstart file on virtually any system configuration.

You can list all LIBDECKs on the deadstart file by accessing the system file SYSTEM with a COMMON command, then using the T parameter on the CATALOG command. Refer to the NOS Version 2 Reference Set, Volume 3 for more information concerning these commands.

The following list provides brief descriptions of SYSEDIT directives acceptable in LIBDECK. Complete descriptions of all SYSEDIT directives are in section 19, SYSEDIT. A list of valid record types follows the directives.

Directive Format	Significance
*AD,nn,ty ₁ /rec ₁ ,ty ₂ /rec ₂ ,...,ty _n /rec _n	Specifies the alternate device to be used in addition to the system device(s) for storing ABS, OVL, PP, and REL type records. nn is either the EST ordinal or the equipment type, ty _i is the record type, and rec _i is the record name.
*CM,ty ₁ /rec ₁ ,ty ₂ /rec ₂ ,...,ty _n /rec _n	Defines record rec _i of type ty _i as being central memory resident; valid only for record types ABS, OVL, or PP. Like any other job in NOS, SYSEDIT has a field length restriction of 376000 ₈ central memory words. You cannot use more than approximately 326000 ₈ central memory words (376000 ₈ -50000 ₈ words for SYSEDIT's program FL) when defining records to be central memory resident.
*FL,ty ₁ /rec ₁ -fl ₁ ,ty ₂ /rec ₂ -fl ₂ ,..., ty _n /rec _n -fl _n	Record rec _i of type ty _i is loaded with a field length specified by fl _i (fl _i is field length divided by 100 ₈).
*MS,ty ₁ /rec ₁ ,ty ₂ /rec ₂ ,...,ty _n /rec _n	Defines record rec _i of type ty _i as being mass storage resident. This is the default residence for routines with no storage area specified in LIBDECK.

Directive Format	Significance
*PROC,rec ₁ ,rec ₂ ,...,rec _n	Defines record rec _i of type PROC as a procedure file.
*SC,ty ₁ /rec ₁ ,ty ₂ /rec ₂ ,..., ty _n /rec _n	Defines record rec _i of type ty _i as product set format commands. The command parameters are processed in product set format (refer to the NOS Version 2 Reference Set, Volume 3).
*/ comments	Defines comment lines.

The following record types may be specified in SYSEEDIT directives. Some directives do not allow all types.

Type (ty _i)	Description
ABS	Multiple entry point overlay.
CAP	Fast dynamic load capsule.
OPL	Modify old program library deck.
OPLC	Modify old program library common deck.
OPLD	Modify old program library directory.
OVL	Central processor overlay.
PP	12-bit peripheral processor program.
PPL	16-bit peripheral processor program (CYBER 180 models only).
PPU	First-level peripheral processor (FLPP) program (model 176 only).
PROC	Procedure.
REL	Relocatable central processor program.
TEXT	Unrecognizable as one of the other types.
ULIB	User library program.

DIS Operations

- DIS Job Dayfile Display (A) 4-3
- DIS Job Status Display (B) 4-4
- DIS Memory Displays (C, D, F, G) 4-6
- DIS Exchange Package Display (X) 4-8
- DIS Directory Display (Z) 4-9
- Console Operation 4-10
- Display Selection Commands 4-13
- DIS Keyboard Entries 4-14
- Memory Entry Commands 4-18
- PP Call Commands 4-20

DIS displays information about a single job. Under DIS, the X display shows the exchange package area for the job. Central memory addresses relative to the job's reference address are used for the data and program displays.

Initiate DIS at a control point to monitor the progress of a job by using any of the following methods:

- You can call DIS by entering X.DIS,fl (fl=field length desired) or X.DIS (field length of 60000₈ assumed by default). This brings DIS to an empty control point to initiate utility programs.
- A job can execute a command in the form DIS (the job must be system origin or have system origin privileges).
- You can call DIS to an existing job by entering the command:

```
DIS, jsn.
```

Parameter	Description
jsn	Job sequence name to which you want DIS assigned.

On either a secured or unsecured system, before you can call DIS to an existing job, the console must be unlocked (refer to the DSD UNLOCK command in section 5). For a secured system, the system must be set to security unlock status by the security administrator.

You can toggle between DSD and DIS. DIS permanently returns control to DSD when you enter DROP.; the job is not dropped unless no commands remain.

When DIS is called to a control point, automatic command processing stops and the A and B displays for DIS appear on the left and right display screens, respectively. Keyboard entry is necessary to begin processing subsequent commands. Unless automatic command processing is reenabled,¹ the job is stopped after each command is processed. That is, only one command can be processed at a time. Under DIS, the B display shows only the condition of the job to which it is assigned, including upcoming commands. When the job is not using the central processor, a copy of its exchange package is displayed on the X display. Displays available under DIS are selected in the same manner as DSD displays. Refer to Console Operation later in this section for information concerning display selection commands and to DIS Keyboard Entries for information about other DIS commands.

1. You can initiate automatic command processing by entering either a period (.) or the RCS command (refer to Console Operation and to DIS Keyboard Entries later in this section).

The displays available under DIS are:

Display Identifier	Description
A	Job dayfile. Messages and files (local FNT entries) attached to the job.
B	Job status. Individual job status, equipment assigned, current messages, and command buffer.
C,D	Central memory. Contents of central memory words (selectable 8-word groups) in five columns of four octal digits with display code equivalents (same as the DSD C and D displays).
F	Central memory. Contents of central memory words (selectable 8-word groups) in four columns of five octal digits with display code equivalents.
G	Central memory. Contents of central memory words (selectable 8-word groups) in four columns of five octal digits with COMPASS instruction mnemonic equivalents.
H	File status. All files assigned to the job as well as equipment assigned to files.
M	Extended memory. Contents of 60-bit words of extended memory (selectable 8-word groups) in five columns of four octal digits with display code equivalents (same as the DSD M display).
N	Blank screen. You may wish to select the N display or one screen to reduce the screen filter on the other screen.
T,U	Text display. Displays text from central memory in coded lines (up to 60 characters per line). The display terminates after 256 words have been displayed.
V	Central memory buffer. Displays directly from central memory. The display terminates after 512 words have been displayed.
X	Exchange package. Breakpoint address and the exchange package.
Y	Monitor functions. Displays mnemonics and the values of all monitor functions (same as the DSD Y display).
Z	Directory. DIS displays available.

NOTE

Although all displays listed may appear on the left screen, only the B, C, D, and N displays may appear on the right screen. If you attempt to bring any other display to the right screen, the message INCORRECT COMMAND is issued to the job dayfile and is displayed in the message buffer of the B display.

DIS Job Dayfile Display (A)

Figure 4-1 shows the DIS job dayfile display (A). The figure shows the job dayfile messages for the control point to which DIS is currently assigned and as many files attached to that control point as will fit in the display. All files attached to the job can be observed by using the file status display (H).

DIS A. JOB DAYFILE.								
Day file messages	09.46.34.RFL(060000)							
	09.46.34.DIS.							
	09.47.31.USER,ABC1234,.							
	09.47.31.ABSC, S.							
	09.48.43.PACKNAM,721C,R=DL.							
	09.49.39.NOS.							
	09.50.06. GET,NOS/PN=PACKC,R=DJ,UN=KRONMOD.							
	09.50.06. DEVICE UNAVAILABLE.							
	09.50.29. EXIT.							
	09.50.41. GET,NOS/PN=PACKV2,R=DJ,UN=KRONMOD.							
	09.50.41. DEVICE UNAVAILABLE.							
	09.55.21.COMMON,SYSTEM.							
	FMT	NAME	TYPE	EST	TRACK	FS	STATUS	LEVEL
	0	INPUT=	IN=	7	4004	ND	5	LVL0
5	SYSTEM	LI=	6	4010		1	LVL0	
10	ZZZZC0	LI=	7	4213	ND	5	LVL0	
11	ZZZZC2	LO	6	5652	ND	307	LVL0	

Files (local FMT entries) attached to the job

Figure 4-1. DIS Job Dayfile Display (A)

The level field, shown in the figure, is displayed only on a secured system.

DIS Job Status Display (B)

The DIS B display shows the status of a specific job executing at a control point. It also shows: any equipment assigned exclusively to the job by EST ordinal, message 1 and message 2 from the control point area, and the current command buffer, allowing you to anticipate future job requirements.

Figure 4-2 shows the DIS job status display (B).

```

DIS B.  JOB STATUS.

JSN = AAAC           EJTO =      5           P = 2046
                   SRUA =      1           RA = 1741
                   SRUL = 777777         FL = 26
                   CS = BATCH           RAE = 160
                   CONN =      0           FLE = 0
                   LEVEL = LVL0         CPA = 600
CPU = X
EST = 1              }   Equipment assigned
MS1 = DEVICE UNAVAILABLE.
MS2 =                }   Current messages
NOS, PACKV2.
RETURN, NOS.
REVERT, NOLIST.
EXIT.
REVERT, ABORT. PACKC OR PACKV2 NOT FOUND.
REVERT, NOLIST.
EXIT.
REVERT. CCL
EXIT. CCL
REVERT, ABORT. CCL
}   Command buffer

```

Figure 4-2. DIS Job Status Display (B)

The job status portion of the display shows three columns of information. Each item has the form item=value. The items are described next in the order that they appear on the display.

Item	Description
JSN	Job sequence name of the DIS job.
UI	User index.
FM	Current family name.
PN	Current pack name.
CPU	CPU status.
EJTO	Executing job table ordinal.
SRUA	System resource units accumulator.
SRUL	Account block limit for system resource units.
CS	Connection status.
CONN	Connection number (interactive jobs only).
LEVEL	Job access level (secured systems only).
P	P register address from exchange package.
RA	Central memory reference address.
FL	Central memory field length.
RAE	Extended memory reference address.
FLE	Extended memory field length.
CPA	Control point area address.

DIS Memory Displays (C, D, F, G)

Figure 4-3 shows the DIS data storage display (F). The contents of each central memory word is displayed in columns of five octal digits along with the display code equivalent. Only the memory locations currently assigned to the job can be displayed. The message ****SECURED AREA**** is displayed for all other locations.

The DIS C and D displays have the same format as the DSD C display, except that the DIS memory displays may show a managed table bias word preceding the groups of central memory words.

DIS F. CENTRAL MEMORY.									
00000060	00000	00000	00000	00000					
00000061	00000	00000	00000	00000					
00000062	00000	00000	00000	00000					
00000063	00000	00000	00000	00000					
00000064	03171	51517	16000	00001	COMMON	A			
00000065	40000	00000	00010	00507	5	A	EG		
00000066	40000	00000	00020	00131	5	B	AY		
00000067	40000	00000	00000	00000	5				
00000070	03171	51517	16562	33123	COMMON, SYS				
00000071	24051	55700	00000	00000	TEM.				
00000072	00000	00000	00000	00000					
00000073	00000	00000	00000	00000					
00000074	00000	00000	00000	00000					
00000075	00000	00000	00000	00000					
00000076	00000	00000	00000	00000					
00000077	00000	00000	00000	00000					

Figure 4-3. DIS Data Storage Display (F)

Figure 4-4 shows the DIS program storage display (G). The G display shows the contents of central memory and the COMPASS mnemonic translation.

DIS G. CENTRAL MEMORY.						
0000000	0000	0000	0000	0000	PS	PS
0000001	05160	42000	00000	00000	NE BB	PS
0000002	23312	32405	15000	00001	AX XB DXX+X	BXX+--X PS
0000003	00000	00000	00000	00000	PS	PS
0000004	00000	00000	00000	00000	PS	PS
0000005	00000	00000	00000	00000	PS	PS
0000006	00000	00000	00000	00000	PS	PS
0000007	00000	00000	00000	00000	PS	PS
0000010	00000	00000	00000	00000	PS	PS
0000011	00000	00000	00000	00000	PS	PS
0000012	00000	00000	00000	00000	PS	PS
0000013	00000	00000	00000	00000	PS	PS
0000014	00000	00000	00000	00000	PS	PS
0000015	00000	00000	00000	00000	PS	PS
0000016	00000	00000	00000	00000	PS	PS
0000017	00000	00000	00000	00000	PS	PS
Address	Memory contents				COMPASS mnemonic translation	

Figure 4-4. DIS Program Storage Display (G)

DIS Exchange Package Display (X)

The DIS X display shows the breakpoint address (BKP=addr) if a breakpoint was set at address addr. The breakpoint address is followed by the job's exchange package.

Figure 4-5 shows the DIS exchange package display (X).

```

DIS X.  EXCHANGE PACKAGE.
  BKP =   100.

  P   =    452      A0 =    600      B0 =      0
  RA  =   174100    A1 =      1      B1 =      1
  FL  =     600     A2 =      0      B2 =  777776
  EM  =  74070000  A3 =      0      B3 =      1
  RAE =      0     A4 =      0      B4 =      3
  FLE =      0     A5 =      2      B5 =      1
  MA  =     600     A6 =      1      B6 =      0
  EEA =      0     A7 =    375      B7 =      0

X0 = 7777 7777 7777 7700 0000
X1 = 0000 0000 0000 0000 0000
X2 = 0000 0000 0000 0000 0375          C
X3 = 0000 0000 0000 0000 0000
X4 = 0000 0000 0000 0000 0003          C
X5 = 2331 2324 0515 0000 0001  SYSTEM  A
X6 = 0516 0420 0000 0000 0000  ENDP
X7 = 0000 0000 0000 0000 0000

+ AND - WILL SET *BKP* TO (P) +/- 1.

```

Figure 4-5. DIS Exchange Package Display (X)

DIS Directory Display (Z)

Figure 4-6 shows the DIS directory display (Z). The Z display lists all displays available under DIS control. If the letter entered to select the left screen display is not a valid display identifier, the Z display is selected automatically.

```
DIS Z.  DIRECTORY.  LEFT SCREEN ONLY DISPLAY
      A  JOB DAYFILE.
      F  CENTRAL MEMORY.  FOUR GROUPS OF FIVE.
      G  CENTRAL MEMORY.  FOUR GROUPS OF FIVE.
      H  FILE STATUS.
      M  EXTENDED MEMORY.
      T  TEXT DISPLAY.
      U  TEXT DISPLAY.
      V  CM BUFFER.
      X  EXCHANGE PACKAGE.
      Y  MONITOR FUNCTIONS.
      Z  DIRECTORY.

      LEFT AND RIGHT SCREEN DISPLAYS

      B  JOB STATUS.
      C  CENTRAL MEMORY.  FIVE GROUPS OF FOUR.
      D  CENTRAL MEMORY.  FIVE GROUPS OF FOUR.
      N  BLANK SCREEN.
```

Figure 4-6. DIS Directory Display (Z)

Console Operation

Unlike DSD, DIS is not interpretive. You must complete every entry and signal DIS to act upon the message by pressing CR or NEXT. The following rules apply to all DIS commands.

- For input, spaces in an octal field are ignored but can be inserted for readability.
- For output, all octal fields are right-justified with leading zero fill; excess octal digits are ignored.

In addition to the command entries, the following keys have special meaning to DIS when entered as the first character. The corresponding special keys for the CC545, CC598B, and CC634B console types are listed along with a description of the use of each key.

CC545	CC598B	CC634B	Description
*	*	F15	If DSD has relinquished the main display console to DIS, this key acts as a quick hold, and DIS drops the display channel so that DSD can use it.
+	Up arrow, +, or grey +	+	Pages the left screen forward for the A, C, D, F, G, H, M, T, and U displays. Increments the breakpoint address for the X display.
-	Down arrow, -, or grey -	-	Pages the left screen backward for the C, D, F, G, M, T, and U displays. Decrements the breakpoint address for the X display. Resets the H display to the beginning. On the A display, the - converts to a BEGIN,.
(PgUp or ((Pages the right screen forward for the C and D displays.
)	PgDn or))	Pages the right screen backward for the C and D displays.
/	/	/	Advances the left screen memory display address by the value in the lower 18 bits of the first word displayed (applicable only to memory displays C, D, F, G, and M).
.	.	.	Sets auto mode (initiates automatic command processing). This key performs the same function as the RCS command described under DIS Keyboard Entries later in this section.

CC545	CC598B	CC634B	Description
8	8	8	Advances the pointer indicating the first address of managed tables for the left screen (applicable only to memory displays C, D, F, and G).
9	9	9	Decrements the pointer indicating the first address of managed tables for the left screen (applicable only to memory displays C, D, F, and G).
Right blank (display key)	Tab	→	Advances the left screen display sequence established by the SET,screen. command (refer to Display Selection Commands later in this section).
CR (carriage return)	Enter/Return	NEXT	Sets the repeat entry flag; the message REPEAT ENTRY is displayed on the error message line of the left screen. The subsequent command entry is processed but is not erased after completion. That command is processed each time this key is pressed. To clear the repeat entry mode, press the left blank (erase) key on the CC545 or the ← key on the CC634B.

The following keys are interpreted as control characters by DIS.

CC545	CC598B	CC634B	Description
Left blank (erase)	Esc	← or ERASE	Clears current DIS keyboard entry and any resultant error message; auto mode (automatic command processing) is also cleared.
BKSP (backspace)	Back Space	←	Deletes last character displayed and clears error message (if one exists).
CR (carriage return)	Enter/Return	NEXT	Initiates processing of an entered command.

The following keyboard messages may appear above your entry.

Message	Description
AUTO MODE.	The command buffer is read automatically. Automatic command processing can be selected by the RCS command or by entering a period.
COMMAND BUFFER FULL.	The ELS command has been entered and there is insufficient room in the command buffer to add the characters requested.
COMMAND TOO LONG.	The command you entered is too long.

Message	Description
COMMANDS ON FILE.	The ELS command has been entered and commands are currently located in a file instead of in a command buffer. It is not possible to add a command to the file.
DIRECT CPU INPUT.	The N command has been entered and all data entered from the keyboard is being passed directly to central memory.
DISK BUSY.	DIS is waiting for an overlay to be loaded from a mass storage device.
EXTENDED MEMORY NOT AVAILABLE.	You entered the ENFLE command, but user extended memory is not defined.
INCORRECT ENTRY.	The command cannot be processed.
INCORRECT FL REQUEST.	You entered an ENFL command to set the field length. The requested field length must be greater than 10000 ₈ and less than 377700 ₈ including negative field length.
INCORRECT PRIORITY.	The priority you entered using the ENPR command is not correct. CPU priority must be greater than 1 and less than 71 ₈ .
INSUFFICIENT FIELD LENGTH.	Your field length is not long enough to support either extended memory or O26.
JOB ACTIVE.	The previous request has not completed. The command must be reentered when the job is not active.
OUT OF RANGE.	The memory entry address is greater than the field length.
PP BUSY.	DIS is waiting for a PP to be assigned in order to process a keyboard entry.
REPEAT ENTRY.	The command in the command buffer is repeated each time you press CR, Enter/Return, or NEXT. This can be cleared by pressing the left blank, Esc, or ← key.
REQUEST EXCEEDS MAXIMUM FLE.	You requested more extended memory than your system can support.
STORAGE NOT AVAILABLE.	The amount of central memory requested by the ENFL command or required by the O26 file editor is not available.
WAITING FOR EXTENDED MEMORY.	DIS is waiting for extended memory after you entered the ENFLE command.
WAITING FOR STORAGE.	DIS is waiting for central memory after you entered the ENFL or O26 command.

Display Selection Commands

Specify the displays you want shown on the console's left and right screens by using the following commands. Follow the entry with a carriage return.

Command	Description								
xy.	Brings the x and y displays to the left and right screens, respectively. Note that although all DIS displays may appear on the left console screen, only the B, C, D, and N displays may appear on the right screen. If you attempt to bring any other display to the right screen, the entry is interpreted as a command and the message INCORRECT ENTRY appears. In addition, if the letter entered to select the left screen display (x) is not a valid display identifier, the Z display is selected automatically.								
xz,loc.	Brings specified memory display to the left screen, if not currently selected, and provides display modifications as follows: <table border="1" data-bbox="412 709 1341 1205"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>x</td> <td>Display identifier (C, D, F, G, or M).</td> </tr> <tr> <td>z</td> <td>Type of display modifications: 0-3, 4, 5, or 6 <ul style="list-style-type: none"> Use 0-3 to display specified group (8 words) starting at location loc. Use 4 to display all 8-word groups in contiguous locations starting at location loc. Use 5 to advance the display by loc locations. Use 6 to decrement the display by loc locations. </td> </tr> <tr> <td>loc</td> <td>Location parameter (maximum of eight digits for central memory address or seven digits for extended memory address).</td> </tr> </tbody> </table>	Parameter	Description	x	Display identifier (C, D, F, G, or M).	z	Type of display modifications: 0-3, 4, 5, or 6 <ul style="list-style-type: none"> Use 0-3 to display specified group (8 words) starting at location loc. Use 4 to display all 8-word groups in contiguous locations starting at location loc. Use 5 to advance the display by loc locations. Use 6 to decrement the display by loc locations. 	loc	Location parameter (maximum of eight digits for central memory address or seven digits for extended memory address).
Parameter	Description								
x	Display identifier (C, D, F, G, or M).								
z	Type of display modifications: 0-3, 4, 5, or 6 <ul style="list-style-type: none"> Use 0-3 to display specified group (8 words) starting at location loc. Use 4 to display all 8-word groups in contiguous locations starting at location loc. Use 5 to advance the display by loc locations. Use 6 to decrement the display by loc locations. 								
loc	Location parameter (maximum of eight digits for central memory address or seven digits for extended memory address).								
x,addr.	If x specifies one of the memory displays (C, D, F, G, or M), addr is the address used to obtain the bias address for the managed table display. (The bias address is the lower 18 bits of the word at addr.)								
SET,xxxx.	Sets the left screen display sequence; xxxx consists of one to four display identifiers. Pressing the right blank key on the CC545, the Tab key on the CC598B, or the → key on the CC634B after this command is entered causes the first display to appear on the left console screen. Pressing the key again selects the second display. The next display in the specified sequence appears on the left screen each time the key is pressed, for example, SET,ACFD.								

DIS Keyboard Entries

You can enter the following commands while in DIS.

If a job is currently active (CPU active, waiting, on recall, or PP active), many commands are not accepted and the message JOB ACTIVE is displayed.

Command	Description
BEGIN,pname,pfile.	Sets auto mode and calls the procedure pname that is on file pfile.
BKP,addr.	Breakpoint to address addr in the program. Central processor execution begins at the current value of P and stops when P=addr, and DIS is the only PP active at the control point.
BKPA,addr.	Breakpoint to address addr in the program with assigned PPs. Central processor execution begins at the current value of P and stops when P=addr. PPs attached to the control point can still be active. DIS clears addr to stop the program at that point. The breakpoint may be cleared by setting the breakpoint address to a new value.
CEF.	Clears the skip-to-exit flag. This allows command processing to continue with the next command instead of skipping to the EXIT command after an error.
DCP.	Drops the central processor and displays the exchange package area on the X display.
DDF.	Calls the display disk file (DDF) utility to the control point.
DIS.	Reloads the main DIS overlay.
DROP.	Drops DIS, but normal processing of the job continues (it does not drop the job until all commands are processed).
ELS.commandstring.	Allows entry of the commandstring command after the last command in the command buffer, if there is space. This command is valid only when auto mode is not set.
ENAi,addr.	Sets register Ai=addr in the exchange package area.
ENBi,addr.	Sets register Bi=addr in the exchange package area.
ENEM,n.	Sets CPU program exit mode to n ($0 \leq n \leq 7$).
ENFL,f.	Sets central memory field length FL=f in the exchange package area ($0 \leq f \leq 777777_8$). f must be at least 10000_8 if user extended memory is assigned.

Command	Description
ENFLE, fle.	Sets extended memory field length FLE to fle000 in the exchange package area ($1 \leq fle \leq 7777_8$). If user extended memory is assigned (fle=0), central memory FL, set by the ENFL command, must be greater than or equal to 10000_8 .
ENP, addr.	Sets P=addr (next instruction address).
ENPR, pr.	Sets the CPU priority to pr ($2 \leq pr \leq 67_8$).
ENPR, *.	Sets the CPU priority to the value defined for the job's service class.
ENS.commandstring.	Allows entry of the commandstring command as the next unprocessed command in the command buffer. The command can then be processed using RNS, RSS, or DROP. Use of ENS with CCL procedure files produces unexpected results. This command is valid only when auto mode is not set.
ENTER./command1/ command2.	Allows entry of the command1 and command2 commands from the keyboard and sets auto mode.
ENTL, timlmt.	Sets the job time limit to timlmt (77777_8 specifies no limit).
ENXi, cont.	Sets register Xi=cont in the exchange package area.
ENXi, Lcont.	Sets register Xi=cont, left-justified, in the exchange package area.
ENXi, Dcharacters.	Sets register Xi to characters in display code.
ENXi, b, value.	Sets byte b of register Xi to value.
ERR.	Sets forced error flag (FSE), terminates program execution, and clears auto mode if set.
GO.	Restarts a program that has paused.
HOLD.	DIS relinquishes the display console, but the job is held at the present status. The console must be reassigned to continue use of DIS.
M.characters.	Enters characters as a CPU program command. Data is stored at RA+CCDR.
N.characters.	Sets direct CPU input mode. Characters entered from keyboard are passed one character at a time, right-justified, directly into central memory at RA+CCDR. Pressing the left blank key on the CC545, the Esc key twice on the CC598B, or the ← key twice on the CC634B clears direct CPU input mode.

Command	Description
OFFSWs.	Turns off sense switch s for the job ($1 \leq s \leq 6$).
ONSWs.	Sets sense switch s for the job ($1 \leq s \leq 6$).
O26.	Calls the O26 file editor to a control point. Refer to the NOS Version 2 Systems Programmer's Instant for a description of the editing commands.
RCP.	Requests central processor. Depending on job priority, execution begins at the next program address for a job suspended by a DCP request.
RCS.	Sets auto mode, which initiates automatic command processing. All succeeding commands are read from the command buffer and processed automatically until an SCS command, an erase function, or an error is encountered. A period (.) may also be used to initiate automatic command processing.
RNS.	Reads and processes the next command in the DIS command buffer.
ROLLOUT.	Allows the job to roll out. This command should be issued when the message ROLLOUT REQUESTED appears (or the * or F15 key may be used).
ROLLOUT,spr.	Places job in rollout status for spr job scheduler delay intervals. The job is automatically rolled back in after this period of time. If a number greater than 777_8 is specified for spr, 777_8 is used.
RSS.	Reads the next command from the command buffer and stops prior to CPU execution. This is used to initiate breakpointing of a program.
RSS,commandstring.	Reads the commandstring command and stops prior to execution. The action is similar to ENS followed by RSS except that the command buffer is not cleared.
SCS.	Clears auto mode, which stops automatic command processing.
SUI,userindex.	Allows access to a user index above AUIMX (377700_8). Any permanent file activity that is to be done on such user indices must be done through system origin jobs. This command is not accepted by a secured system.
T,addr.	Changes the T display to start at address addr.

Command	Description
U,addr.	Changes the U display to start at address addr.
UCC=c	Sets the uppercase character c. This command does not terminate with a period.
V,addr.	Changes the V display to start at address addr.
X.commandstring.	Processes commandstring as the next command. Only the first 50 characters following X are used. This may be used to enter a leading slant or a command that is the same as DIS display.
* commandstring.	If an asterisk (*) followed by a blank and commandstring is encountered during automatic command processing (auto mode), commandstring is interpreted as a direct DIS display selection command. For example, * C4,100. will set the left screen display to the central memory C display at address 100. Using this feature, it is possible to set up procedure files that use DIS to breakpoint a program to a desired stopping point.
commandstring.	Processes commandstring as a command if it is not a recognizable DIS command.
,	Calls the O26 file editor to a control point. Refer to the NOS Version 2 Systems Programmer's Instant for a description of the editing commands. (This command is the same as the O26 command.)

Memory Entry Commands

The following commands are used in conjunction with the C, D, F, G, and M memory displays to change the contents of central memory and extended memory. Only locations relative to the reference address (RA) of the job to which DIS is assigned can be changed. When changing the contents of memory relative to a job, the negative field length area of the job can be accessed by specifying a negative address. For example, to change the content of RA-3, enter the address 777775_8 .

On a secured system no memory entry commands are allowed unless the security unlock status is set. The memory display shows the message

```
****SECURED AREA****
```

instead of the contents of the memory locations to prevent you from examining these locations; you may not alter the contents of these locations.

Character values or numeric data can replace the current word contents. Either one 12-bit byte, one 15-bit parcel, one 30-bit parcel, or 60 bits can be changed. A single byte can be changed by inserting the byte number after the location to be changed; bytes are numbered 0 through 4 from left to right. The address and contents are assembled right-justified with leading zero fill. Leading zeros may be omitted in the entry. Only words within the field length of the job may be changed.

CAUTION

Improper use of these commands may result in damage to the system or to user jobs.

Formats and descriptions of the memory entry commands follow.

Command	Description
addr,cont. or addr+cont.	Changes the contents of memory location addr to cont. The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next memory location. ²
addr,b,cont. or addr+b,cont.	Changes the contents of byte b at memory location addr (eight digits) to cont. Each location consists of five 12-bit bytes, numbered 0 through 4 from left to right. The contents are octal characters. The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next memory location. ²

2. If the message REPEAT ENTRY is displayed above the entry line, the cont field is not cleared and may be entered in successive memory locations as many times as desired by pressing CR or NEXT. The repeat entry mode is enabled by pressing CR or NEXT before initial entry of the command. This is also applicable to the b and n fields of the second, fifth, and seventh commands.

Command	Description
addr,Dcont. or addr+Dcont.	Changes the contents of memory location addr (eight digits) to display code characters cont (left-justified and zero-filled). The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next memory location. ³
addr,Lcont. or addr+Lcont.	Changes the contents of memory location addr (eight digits) to left-justified cont. The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next memory location.
addr,In,cont. or addr+In,cont.	Changes the contents of instruction n (0 through 3 from left to right) at memory location addr (eight digits) to cont; cont may be a 15- or 30-bit instruction. However, one or more bits must be set in the upper 15 bits of a 30-bit instruction or the entry will be treated as a 15-bit instruction. The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next memory location. ³
Eaddr,cont. or Eaddr+cont.	Changes the contents of extended memory location addr to cont. The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next extended memory location.
Eaddr,b,cont. or Eaddr+b,cont.	Changes the contents of byte b at extended memory location addr to cont. Each location consists of five 12-bit bytes, numbered 0 through 4 from left to right. The contents are four octal characters. The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next extended memory location.
Eaddr,Dcont. or Eaddr+Dcont.	Changes the contents of extended memory location addr to display code characters cont (left-justified and zero-filled). The second form of the command performs essentially the same function but leaves the address at addr+1, allowing immediate entry for the next extended memory location. ³

3. If the message REPEAT ENTRY is displayed above the entry line, the cont field is not cleared and may be entered in successive memory locations as many times as desired by pressing CR or NEXT. The repeat entry mode is enabled by pressing CR or NEXT before initial entry of the command. This is also applicable to the b and n fields of the second, fifth, and seventh commands.

PP Call Commands

Any PP program having a name that begins with a letter may be initiated by DIS. However, before entering any of these commands, it is necessary to have a working knowledge of the PP program to be called. This ensures correct use of the specified program.

CAUTION

Improper use of these commands may result in damage to the system or to user jobs.

In table 4-1, prg denotes the name of the PP program and n is the control point number.

Table 4-1. PP Call Formats

Command	Description	Format of PP Call Initiated
prg.	Calls PP program prg to the control point.	18/3Lprg,6/n,36/0
prg,p1.	Calls PP program prg to the control point; p1 is an octal parameter required by prg.	18/3Lprg,6/n,18/0,18/p1
prg,p1,p2.	Calls PP program prg to the control point; p1 and p2 are octal parameters required by prg.	18/3Lprg,6/n,18/p1,18/p2

Display Selection Commands	5-2
Dayfile Commands	5-3
Queued File Utility Commands	5-4
Job Processing Control Commands	5-4
General Job Control Commands	5-4
Interactive Job Control Commands	5-8
Job Communication Commands	5-9
Job Flow Commands	5-10
Peripheral Equipment Control Commands	5-24
Subsystem Control Commands	5-50
System Control Commands	5-56
Secured System Control Commands	5-70
Memory Entry Commands	5-72
Channel Control Commands	5-74
Extended Memory Flag Register Commands	5-75
Breakpoint Package Commands	5-75
PP Breakpoint Commands	5-75
V Display	5-77
PP Breakpoint Precautions	5-79
CPUMTR Breakpoint Commands	5-80
CPUMTR Breakpoint Display (C Display)	5-81
CPUMTR Breakpoint Precautions	5-83
Debugging Commands	5-84

After the system has been deadstarted successfully, you can enter the DSD commands necessary to provide optimum performance and reliability for users. The following general categories of DSD commands are available for this purpose.

Category	Description
Display selection	Selects DSD displays.
Dayfile	Dumps the system, account, or error log dayfile to a specified device.
Queued file utility	Provides control over selected queued files.
Job processing control	Provides added control over job scheduling and processing.
Peripheral equipment control	Controls the peripheral equipment available to the system.
Subsystem control	Schedules a subsystem to a control point or terminates a current subsystem.
System control	Maintains system integrity in a normal production environment or debugs a system that is in an abnormal state.
Memory entry	Changes the contents of central memory and extended memory.
Channel control	Controls activity on a specified data channel in circumstances where abnormal hardware and/or system operation is detected.
Extended memory flag register	Clears and sets bits in the extended memory flag register.
Breakpoint package	Provides control over PP breakpoint processing and CPUMTR breakpoint processing.
Debugging commands	Traps certain conditions and traces selected pool PP to CPUMTR and MTR to CPUMTR functions.

Although all DSD commands are generally available, many of them are seldom used in a normal production environment. Many DSD commands are used only by the system analyst for maintenance or debugging. These commands include all memory entry and channel control commands as well as several commands in the other categories listed on an unsecured system. Memory entry and several other commands are restricted on a secured system (refer to the UNLOCK command later in this section).

When unusual problems arise, do not attempt corrective action unless you have considerable experience relating to the current problem. Misguided attempts to correct a system problem can often destroy information required to successfully analyze the problem.

To assist customer engineers and software analysts in tracking problems, the system enters the first characters (up to 25) of the following commands into the error log after each execution.

DOWN,param.
FORM,param.
IDLE,param.
INITIALIZE,param.
LOG,param.
OFF,param.
ON,param.
REDEFINE,param.
99.
Memory entry commands
Channel control commands
Any other commands that you specify at installation time

These commands are prefixed by the characters DS in the error log but otherwise appear exactly as they are entered. This feature can be enabled and disabled using the 99 command (refer to System Control Commands later in this section).

The manner in which the DSD commands are entered and the use of special keyboard characters are described in the NOS Version 2 Operations Handbook. Command formats are fixed field; that is, the fields in the command format must be specified as shown. Embedded blanks are allowed in octal fields. Leading spaces in command entries are not allowed.

Display Selection Commands

The system display program DSD generates system-oriented displays. You can select any of the DSD displays with the following DSD command:

xy.

where x and y represent the letter designation (screen identifier) of the displays.

Display x appears on the left screen and display y appears on the right screen. If x and y are identical, both screens display the same information, except for the B and P displays when using the CC598B or CC634B console. Refer to the NOS Version 2 Operations Handbook for details concerning the DSD display formats and keyboard operating instructions.

You can specify a sequence of DSD displays that you want displayed on the left screen. To preselect the left screen display sequence, enter the following DSD command:

SET,screens.

where screens represents a string of four letters (screen identifiers) designating any four DSD displays. Usually you specify four different displays although DSD accepts any four valid screen identifiers. Refer to the NOS Version 2 Operations Handbook for DSD displays and keyboard operating instructions.

Dayfile Commands

The system saves messages in five types of dayfiles.

- Account dayfile
- Binary maintenance log
- Error log
- Job dayfile
- System dayfile

The account dayfile keeps a record of all resources charged to a job. This dayfile can be used for customer billing and other accounting purposes. The binary maintenance log records the information used in Control Data maintenance in binary format. The error log records system error messages, such as disk errors. Job dayfiles keep entries for individual jobs. The system dayfile keeps a history of all commands for all jobs processed.

The following commands dump the account dayfile, system dayfile, or error log to a system-defined mass storage device. The resultant mass storage file is put in the output queue for printing. The system automatically prints the job dayfile at the end of the job's output.

Command	Description
X.AFD.	Requests that the account dayfile be dumped to a system-defined mass storage device. The resultant mass storage file is put in the output queue for printing.
X.DFD.	Requests that the system dayfile be dumped to a system-defined mass storage device. The resultant mass storage file is put in the output queue for printing.
X.ELD.	Requests that the error log be dumped to a system-defined mass storage device. The resultant mass storage file is put in the output queue for printing.

Refer to section 18, Queue/Dayfile Utilities, for more information on dayfile dumps.

The binary maintenance log is designed to be processed through an interpreter program, and therefore is normally dumped to tape or disk.

Refer to the NOS Version 2 Operations Handbook for descriptions of dayfile displays.

Queued File Utility Commands

The following commands are used to control queued input and output files.

Command	Description
X.QREC	Deactivates or activates selected queued files and purges selected inactive queued files.
X.QMOVE	Moves queued files from one mass storage device to another.

Refer to section 18, Queue/Dayfile Utilities, for more information on queued file utility commands.

Job Processing Control Commands

Under normal circumstances, control over job processing is performed automatically by the system. Although the following commands may not be used frequently, they provide an added measure of control over job processing.

General Job Control Commands

Several of the commands described here change internal system parameters which control job scheduling and processing. Give careful consideration to their use since job flow and overall system performance can be affected. Refer to the individual command descriptions for further information.

Command	Description
CKP,jsn.	Checkpoints the job with job sequence name jsn. The checkpoint information includes a copy of the job's field length, the system information used for job control, and the name and contents of all local files currently assigned to the control point. It is the responsibility of the user's job to establish a magnetic tape or mass storage permanent file to receive the checkpoint information. Otherwise, checkpoint information is automatically written to a local file named CCCCCC and is not available if a restart becomes necessary (restart is user-controlled and is accomplished using the RESTART command). Refer to the NOS Version 2 Reference Set, Volume 3 for supplementary information concerning the checkpoint/restart feature available to users.

NOTE

If the current job command has secure system memory (SSM) status set, the job cannot be checkpointed. SSM status is set for certain jobs to prevent dumping of the job's field length. Subsystems cannot be checkpointed.

Command	Description																
DROP,jsn,qt,ujn.	<p>Drops the job with job sequence name jsn from the queue qt where it currently resides. If a jsn is not specified, a 1- to 7-character user job name ujn can be specified after the queue type.</p> <p>The console must be unlocked to use this form of the DROP command. If no jsn or ujn is specified, all jobs in the specified queue type are dropped. If the queue type is not specified, the default is the executing queue.</p> <p>The DROP command cannot be used to terminate a subsystem. Also the DROP command will not terminate an interactive user's session (this can be done only with a KILL command).</p> <p>The queue type is one of the following:</p> <table border="1"> <thead> <tr> <th>qt</th> <th>Queue Type</th> </tr> </thead> <tbody> <tr> <td>ALL</td> <td>All jobs and queued files.</td> </tr> <tr> <td>EX</td> <td>All jobs in the executing queue (including the rolled out jobs).</td> </tr> <tr> <td>IN</td> <td>All jobs in the input queue.</td> </tr> <tr> <td>PL</td> <td>All jobs in the plot queue.</td> </tr> <tr> <td>PR</td> <td>All jobs in the print queue.</td> </tr> <tr> <td>PU</td> <td>All jobs in the punch queue.</td> </tr> <tr> <td>WT</td> <td>All jobs in the wait queue.</td> </tr> </tbody> </table>	qt	Queue Type	ALL	All jobs and queued files.	EX	All jobs in the executing queue (including the rolled out jobs).	IN	All jobs in the input queue.	PL	All jobs in the plot queue.	PR	All jobs in the print queue.	PU	All jobs in the punch queue.	WT	All jobs in the wait queue.
qt	Queue Type																
ALL	All jobs and queued files.																
EX	All jobs in the executing queue (including the rolled out jobs).																
IN	All jobs in the input queue.																
PL	All jobs in the plot queue.																
PR	All jobs in the print queue.																
PU	All jobs in the punch queue.																
WT	All jobs in the wait queue.																
ENPR,jsn,pr.	Sets the CPU priority for the job with job sequence name jsn to pr ($2 \leq pr \leq 67_8$).																
ENPR,jsn,*.	Sets the CPU priority for the job with job sequence name jsn to the value defined for the job's service class.																
ENQP,jsn,pr.	Enters queue priority of pr for a queue type file with job sequence name jsn. The value of pr can range from the minimum to the maximum for the job's service class. If the priority is zero, an input file is not scheduled back to a control point automatically. The value specified overrides the current queue priority for the file. The current queue priority can either be increased or decreased using this command.																

Command	Description
KILL,jsn.	<p>Drops the job with job sequence name jsn from the executing job table (EJT) without exit processing. If you want the job to proceed with exit processing, use the DROP command. The KILL command cannot be used to drop a subsystem.</p> <p>KILL is the only command that immediately terminates an interactive user's session. The DROP and OVERRIDE commands only terminate the current job step and do not log the user out of the system. There will be no recoverable job after you issue the KILL command.</p> <p>NOTE</p> <hr/> <p>Before initiating the command, ensure that the correct job sequence name has been specified.</p> <hr/> <p>In some cases, a KILL command will be intercepted by a job's reprieve processing. If the job does not terminate after finishing its reprieve processing, issue another KILL command to terminate the job.</p>
OVERRIDE,jsn.	<p>Certain types of job processing are unaffected by the DROP, KILL, or STOP commands. These include operations such as setting permanent file device interlocks, interlocking files/tracks, clearing VSN entries, and waiting for certain types of tape/PP activity to end. The OVERRIDE command terminates this type of processing and drops the job with job sequence name jsn from the executing job table regardless of queue priority. The OVERRIDE command will not terminate an interactive user's session (this can be done only with a KILL command). Unlock the console (refer to the UNLOCK command later in this section) to enter this command. Use of this command is recorded in the error log.</p> <p>NOTE</p> <hr/> <p>Exercise extreme caution in using the OVERRIDE command. Undesirable situations (such as interlocks being left set, VSN entries remaining uncleared, or certain tape/PP activities left outstanding) could occur that would cause potentially damaging system activity. Never use this command during normal operations.</p> <hr/>

Command	Description
RERUN,jsn.	Terminates the job with job sequence name jsn, then reruns the job from the beginning. The job must be in rerun status as set by the RERUN command or macro.
ROLLIN,jsn,L.	Allows the job defined by job sequence name jsn to be scheduled to an available control point. L is optional; if omitted, the job can be selected by the scheduler for rollout.
ROLLOUT,jsn,sd.	<p>Removes the currently executing job with job sequence name jsn and makes it a rolled out job. A subsystem cannot be rolled out. sd is the number of scheduler intervals before the job can be scheduled again. The acceptable range for sd is between 0 and 777₈. If sd is not present or is 0, the job is not scheduled back to a control point automatically. That is, your action is required to return the job to a control point. This can be done by using the ROLLIN command.</p> <p>The amount of time required for one job scheduler interval is initially set in the IPRDECK but may be changed by using the DELAY command (JS parameter) described later in this section. Normally, it is a 1-second interval.</p>

Interactive Job Control Commands

The following commands apply only to interactive online jobs. The interactive facility subsystem must be active at control point 1. Refer to the CDCNET Network Operations Manual for information about sending messages to terminal users that are connected through CDCNET.

Command	Description
DIAL,jsn,messagetext. DIAL,*,messagetext.	<p>Sends message messagetext (48 characters maximum) to the terminal currently assigned to the job with job sequence name jsn (when jsn=*, the message is sent to all active terminals). Examine the T display to determine the appropriate job sequence name. The message is sent to the terminal immediately except when output is being sent to the terminal. In that case, the message follows the output data.</p>
WARN,messagetext.	<p>Sends message messagetext (48 characters maximum) to all terminals currently logged into IAF. The message is received at a terminal upon completion of the current command or at the end of a job step. Each subsequent terminal that logs into IAF also receives this message. This continues until either a new message is entered or the message is cleared by using the WARN command. In addition, the current message also appears at the IAF Subsystem control point on the B,O display.</p> <p>When sent to an interactive terminal, the message messagetext is always preceded by the statement</p> <pre>hh.mm.ss. WARNING</pre> <p>specifying the time (hours.minutes.seconds) when you entered the WARN command.</p> <p>For example, if you enter</p> <pre>WARN,SYSTEM SHUTDOWN AT 1500.</pre> <p>the following information would be transmitted to all interactive terminals.</p> <pre>hh.mm.ss. WARNING SYSTEM SHUTDOWN AT 1500.</pre> <p>This command is typically used to notify interactive users of an interruption in service or system shutdown.</p>
WARN.	<p>Clears the message entered by the WARN,messagetext command. Unless this command is entered, the existing message (if any) continues to be transmitted to each new terminal that logs into IAF.</p>

Job Communication Commands

The following job communication commands are used to communicate with a job currently in the executing job table.

Command	Description
CFO,jsn,messagetext.	Sends a message messagetext (36 characters maximum) from the console to the job with job sequence name jsn. Bit 14 of the job's RA must be set before the CFO command is accepted. The message is placed in locations RA+70 ₈ through RA+74 ₈ of the program's field length.
COMMENT,jsn. messagetext	Enters comment messagetext (48 characters maximum) in the dayfile for the job with job sequence name jsn.
GO,jsn.	Clears the pause bit of the job with job sequence name jsn. A job may set the pause bit if an error is encountered or if an operator response is required. If jsn is not specified, the command applies to the system control point.
OFFSW,jsn,s ₁ ,s ₂ ,...,s ₆ .	Turns off one or more sense switch(es) s _i (1 ≤ s _i ≤ 6) of the job with job sequence name jsn. Refer to Subsystem Control Commands later in this section for definitions of the sense switches that can be set for the BIO, IAF, MSE, and TAF Subsystems.
ONSW,jsn,s ₁ ,s ₂ ,...,s ₆ .	Turns on sense switch(es) s _i (1 ≤ s _i ≤ 6) of the job with job sequence name jsn. Refer to Subsystem Control Commands later in this section for definitions of the sense switches that can be set for the BIO, IAF, MSE, and TAF Subsystems.
PAUSE,jsn.	Sets the pause bit of the job with job sequence name jsn. If jsn is not specified, the command applies to the system control point.

Job Flow Commands

The following job flow commands affect scheduling and execution of jobs in the system. Improper use of these commands can drastically hamper job flow as well as system performance. In certain cases, jobs may be lost.

The priority associated with each priority parameter in the following commands is established by an entry in the IPRDECK for each service class. The value of each priority parameter for each service class is listed on the system control display(s) (refer to the NOS Version 2 Operations Handbook for more information on the S display).

Command	Description						
CLASS,ot,sc₁,sc₂,...,sc_n	Defines the authorized service class for each origin type. Refer to the SCTD L display in section 9, L-Display Utilities, for further information.						
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ot</td> <td>Origin type (must be batch, remote batch, or interactive). This parameter is order dependent. Types with brief descriptions follow: BC Batch. RB Remote batch. IA Interactive.</td> </tr> <tr> <td>sc_i</td> <td>Authorizes service classes for the specified origin type. Entering a service class that already has validation clears validation for that service class. Service classes with brief descriptions (in parentheses) follow: BC (Local batch). CT (Communication task). DI (Detached interactive). In [Installation class n (0 ≤ n ≤ 8)]. MA (Maintenance). NS (Network supervisor). RB (Remote batch). SY (System). TS (Interactive). ALL (Sets all classes except SS and DS). NUL (Clears all service classes).</td> </tr> </tbody> </table>	Parameter	Description	ot	Origin type (must be batch, remote batch, or interactive). This parameter is order dependent. Types with brief descriptions follow: BC Batch. RB Remote batch. IA Interactive.	sc _i	Authorizes service classes for the specified origin type. Entering a service class that already has validation clears validation for that service class. Service classes with brief descriptions (in parentheses) follow: BC (Local batch). CT (Communication task). DI (Detached interactive). In [Installation class n (0 ≤ n ≤ 8)]. MA (Maintenance). NS (Network supervisor). RB (Remote batch). SY (System). TS (Interactive). ALL (Sets all classes except SS and DS). NUL (Clears all service classes).
Parameter	Description						
ot	Origin type (must be batch, remote batch, or interactive). This parameter is order dependent. Types with brief descriptions follow: BC Batch. RB Remote batch. IA Interactive.						
sc _i	Authorizes service classes for the specified origin type. Entering a service class that already has validation clears validation for that service class. Service classes with brief descriptions (in parentheses) follow: BC (Local batch). CT (Communication task). DI (Detached interactive). In [Installation class n (0 ≤ n ≤ 8)]. MA (Maintenance). NS (Network supervisor). RB (Remote batch). SY (System). TS (Interactive). ALL (Sets all classes except SS and DS). NUL (Clears all service classes).						

Command	Description
CPTT=n	<p>Changes the CPUPFM transfer threshold (CPTT) to the value n ($0 \leq n \leq 7777_8$ PRUs). The value of CPTT determines whether an indirect access file transfer will be processed entirely within the PP program PFM or whether the CPU program CPUPFM will be called to perform the transfer. Since CPUPFM can transfer large files faster than PFM, CPUPFM will transfer files that exceed the threshold value.</p> <p>The default value is 100_8 sectors, as defined in common deck COMSPFM. Depending on the configuration, you may want to adjust this value up or down. A site with buffered I/O devices, for example, may find it advantageous to set a lower value. To disable CPUPFM transfers altogether, enter CPTT=0. The current value of CPTT is displayed on the W,M display.</p>
DELAY,p ₁ ,p ₂ ,...,p _n	<p>Alters current system delay parameters. Examine the W,M display to determine the current delay parameter values.</p>

p _i	Description
ARar	PP recall interval in milliseconds. This parameter specifies the default time interval after which a peripheral processor program in the PP recall queue will be recalled. ar ranges from 1 to 7777_8 .
CRcr	CPU recall period in milliseconds. This parameter specifies the amount of time a job remains in recall (X status) when an RCL request is placed in RA+1. cr ranges from 1 to 7777_8 .
JQjq	Exponent used to determine the input file scheduling interval in seconds. jq ranges from 0 to 14_8 . The interval in seconds between scheduling of input files is calculated as follows:
	interval = 2^{**jq}
JSjs	Job scheduler interval in seconds. This parameter specifies the interval in which the job scheduler is called. The scheduler may also be called at other times. js ranges from 1 to 7777_8 .

Command	Description
DELAY,p ₁ ,p ₂ ,...,p _n	(Continued)

p _i	Description
MPmp	Memory padding value expressed as 100 ₈ word blocks. mp ranges from 0 to 777 ₈ . This parameter specifies how much additional (unassigned) memory should be allocated between the end of the newly assigned job field length and the beginning of the next job. Increasing this value reduces the probability that a job will be storage moved in response to a request for more memory.

The value for each system delay parameter may be established by using a DELAY command entry in the IPRDECK. If no DELAY entry is present, default values are provided. Figure 5-1 provides space to record the original values (specified in the W,M display) in the event that any are altered temporarily. For additional information concerning the DELAY command, refer to section 3, Deadstart Decks.

Delay	Values
AR	
CR	
JQ	
JS	
MP	

M04486

Figure 5-1. Record of Original Values in W,M Display (Delay Values)

Command	Description
PCLASS, <i>sc</i> ₀ , <i>sc</i> ₁ ,..., <i>sc</i> ₇ .	Selects the service class associated with each priority level (<i>p</i> ₀ through <i>P</i> ₇) on the Job command. This parameter is order dependent; a comma must appear for any parameters not specified. The default for any parameters not specified is 0 (zero). Subsystem (SS) service class is not allowed with this command. Refer to the SCTD L display in section 9, L-Display Utilities, for more information.

*sc*_{*i*} is a 2-character service class parameter. Entries are listed and briefly described below.

<i>sc</i>_{<i>i</i>}	Description
BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class <i>n</i> ($0 \leq n \leq 8$).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SY	System.
TS	Interactive.

Command	Description
QUEUE,sc,qt,qp ₁ , qp ₂ ,...,qp _n	Alters the queue priorities qp _i associated with the input, executing, and output queues qt for each service class sc. Examine the L display SDSPLAY to determine the priority values currently associated with each service class.

sc	Description
BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class n (0 ≤ n ≤ 8).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SS	Subsystem.
SY	System.
TS	Interactive.

There are three job queue types.

qt	Description
EX	Executing.
IN	Input.
OT	Output.

Command	Description												
QUEUE,sc,qt,qp ₁ , qp ₂ ,...,qp _n	(Continued) There are five priorities.												
	<table border="1"> <thead> <tr> <th>qp_i</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ILil</td> <td>Priority a job receives when it initially exhausts its CM time slice (refer to the SERVICE command CM parameter). Thereafter, whenever the CM time slice is exhausted, the job's priority will be set to the value of lp. il ranges from 0 to 7777₈, but must be in the range of values for lp and up. This parameter is valid only for executing jobs (EX).</td> </tr> <tr> <td>IPip</td> <td>Initial scheduling priority for an executing batch job or for an interactive job when a terminal command is entered. Online interactive jobs with terminal I/O available are scheduled at tp priority (refer to the SERVICE command TP parameter). ip ranges from 0 to 7777₈, but must be in the range of values for lp and up.</td> </tr> <tr> <td>LPip</td> <td>Lowest priority. For an input or output queue file, the priority assigned to a file or job entering the designated queue. For an executing job, the priority assigned to a job which has exceeded a non-initial CM slice. The IL value is used for the first CM slice. lp ranges from 0 to 7777₈, but must be less than the value of up.</td> </tr> <tr> <td>UPup</td> <td>Highest priority. For input and output queues, this is the highest priority a file can reach in that queue; aging stops when this priority is reached. For the execution queue, this is the priority assigned to a job when initially assigned to a control point. up ranges from 0 to 7777₈, but must be greater than the value of lp.</td> </tr> <tr> <td>WFwf</td> <td>Weighting factor wf for queue priority calculation. wf must be a power of 2, from 1 to 4000₈. The smaller the weighting factor, the faster the queue entry reaches its highest priority.</td> </tr> </tbody> </table>	qp _i	Description	ILil	Priority a job receives when it initially exhausts its CM time slice (refer to the SERVICE command CM parameter). Thereafter, whenever the CM time slice is exhausted, the job's priority will be set to the value of lp. il ranges from 0 to 7777 ₈ , but must be in the range of values for lp and up. This parameter is valid only for executing jobs (EX).	IPip	Initial scheduling priority for an executing batch job or for an interactive job when a terminal command is entered. Online interactive jobs with terminal I/O available are scheduled at tp priority (refer to the SERVICE command TP parameter). ip ranges from 0 to 7777 ₈ , but must be in the range of values for lp and up.	LPip	Lowest priority. For an input or output queue file, the priority assigned to a file or job entering the designated queue. For an executing job, the priority assigned to a job which has exceeded a non-initial CM slice. The IL value is used for the first CM slice. lp ranges from 0 to 7777 ₈ , but must be less than the value of up.	UPup	Highest priority. For input and output queues, this is the highest priority a file can reach in that queue; aging stops when this priority is reached. For the execution queue, this is the priority assigned to a job when initially assigned to a control point. up ranges from 0 to 7777 ₈ , but must be greater than the value of lp.	WFwf	Weighting factor wf for queue priority calculation. wf must be a power of 2, from 1 to 4000 ₈ . The smaller the weighting factor, the faster the queue entry reaches its highest priority.
qp _i	Description												
ILil	Priority a job receives when it initially exhausts its CM time slice (refer to the SERVICE command CM parameter). Thereafter, whenever the CM time slice is exhausted, the job's priority will be set to the value of lp. il ranges from 0 to 7777 ₈ , but must be in the range of values for lp and up. This parameter is valid only for executing jobs (EX).												
IPip	Initial scheduling priority for an executing batch job or for an interactive job when a terminal command is entered. Online interactive jobs with terminal I/O available are scheduled at tp priority (refer to the SERVICE command TP parameter). ip ranges from 0 to 7777 ₈ , but must be in the range of values for lp and up.												
LPip	Lowest priority. For an input or output queue file, the priority assigned to a file or job entering the designated queue. For an executing job, the priority assigned to a job which has exceeded a non-initial CM slice. The IL value is used for the first CM slice. lp ranges from 0 to 7777 ₈ , but must be less than the value of up.												
UPup	Highest priority. For input and output queues, this is the highest priority a file can reach in that queue; aging stops when this priority is reached. For the execution queue, this is the priority assigned to a job when initially assigned to a control point. up ranges from 0 to 7777 ₈ , but must be greater than the value of lp.												
WFwf	Weighting factor wf for queue priority calculation. wf must be a power of 2, from 1 to 4000 ₈ . The smaller the weighting factor, the faster the queue entry reaches its highest priority.												

The priority associated with each queue is established by using QUEUE command entries in the IPRDECK for each service class. These entries normally reflect the ideal queue priorities for the job mix of the particular installation. The values specified in the IPRDECK are considered critical to optimum system performance and are not normally altered. Figure 5-2 provides space to record the original values specified in the L display SDSPLAY. For additional information on the QUEUE command, refer to section 3, Deadstart Decks.

Job Service Class	Priorities for Each Queue Type										
	Input Queue			Executing Queue					Output Queue		
	LP	UP	WF	IL	IP	LP	UP	WF	LP	UP	WF
SY											
BC											
RB											
TS											
DI											
NS											
SS											
MA											
CT											
10											
11											
12											
13											
14											
15											
16											
17											
18											

06557M

Figure 5-2. Record of Original Values in L Display SDSPLAY

Command	Description																						
SERVICE , <i>sc,p₁,p₂,...,p_n</i> .	Alters the service limits <i>p_i</i> associated with each service class <i>sc</i> .																						
	<table border="1"> <thead> <tr> <th>sc</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>BC</td> <td>Local batch.</td> </tr> <tr> <td>CT</td> <td>Communication task.</td> </tr> <tr> <td>DI</td> <td>Detached interactive.</td> </tr> <tr> <td>In</td> <td>Installation class <i>n</i> ($0 \leq n \leq 8$).</td> </tr> <tr> <td>MA</td> <td>Maintenance.</td> </tr> <tr> <td>NS</td> <td>Network supervisor.</td> </tr> <tr> <td>RB</td> <td>Remote batch.</td> </tr> <tr> <td>SS</td> <td>Subsystem.</td> </tr> <tr> <td>SY</td> <td>System.</td> </tr> <tr> <td>TS</td> <td>Interactive.</td> </tr> </tbody> </table>	sc	Description	BC	Local batch.	CT	Communication task.	DI	Detached interactive.	In	Installation class <i>n</i> ($0 \leq n \leq 8$).	MA	Maintenance.	NS	Network supervisor.	RB	Remote batch.	SS	Subsystem.	SY	System.	TS	Interactive.
sc	Description																						
BC	Local batch.																						
CT	Communication task.																						
DI	Detached interactive.																						
In	Installation class <i>n</i> ($0 \leq n \leq 8$).																						
MA	Maintenance.																						
NS	Network supervisor.																						
RB	Remote batch.																						
SS	Subsystem.																						
SY	System.																						
TS	Interactive.																						
	<table border="1"> <thead> <tr> <th>p_i</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AMam</td> <td>Maximum field length/100₈ for all jobs of the specified service class. This parameter is used to partition central memory by limiting the total field length available to each service class. For example, if scheduling a job to a control point would cause the AM value to be exceeded for its service class, it may not be scheduled until the required field length is available. This means that a lower priority job from a different service class may be scheduled first. However, a job that would cause the AM value for its service class to be exceeded can be scheduled to a control point if there are not enough jobs in other service classes to fill the central memory and if FLEXIBLE PARTITIONS is enabled. The system always attempts to use central memory to its greatest capacity. am ranges from 0 to 77777₈.</td> </tr> <tr> <td>CMcm</td> <td>Central memory time slice in seconds. This parameter specifies the maximum amount of time a job of the specified service class can remain in central memory, either at a control point or at a pseudo-control point, before its priority is set to the lower boundary. When the job initially exhausts its CM time slice, this lower boundary is the value of il. Subsequently, when the job exhausts its CM time slice, this lower boundary is the value of lp (refer to the QUEUE command IL and LP parameters). cm ranges from 0 to 7777₈.</td> </tr> </tbody> </table>	p_i	Description	AMam	Maximum field length/100 ₈ for all jobs of the specified service class. This parameter is used to partition central memory by limiting the total field length available to each service class. For example, if scheduling a job to a control point would cause the AM value to be exceeded for its service class, it may not be scheduled until the required field length is available. This means that a lower priority job from a different service class may be scheduled first. However, a job that would cause the AM value for its service class to be exceeded can be scheduled to a control point if there are not enough jobs in other service classes to fill the central memory and if FLEXIBLE PARTITIONS is enabled. The system always attempts to use central memory to its greatest capacity. am ranges from 0 to 77777 ₈ .	CMcm	Central memory time slice in seconds. This parameter specifies the maximum amount of time a job of the specified service class can remain in central memory, either at a control point or at a pseudo-control point, before its priority is set to the lower boundary. When the job initially exhausts its CM time slice, this lower boundary is the value of il. Subsequently, when the job exhausts its CM time slice, this lower boundary is the value of lp (refer to the QUEUE command IL and LP parameters). cm ranges from 0 to 7777 ₈ .																
p_i	Description																						
AMam	Maximum field length/100 ₈ for all jobs of the specified service class. This parameter is used to partition central memory by limiting the total field length available to each service class. For example, if scheduling a job to a control point would cause the AM value to be exceeded for its service class, it may not be scheduled until the required field length is available. This means that a lower priority job from a different service class may be scheduled first. However, a job that would cause the AM value for its service class to be exceeded can be scheduled to a control point if there are not enough jobs in other service classes to fill the central memory and if FLEXIBLE PARTITIONS is enabled. The system always attempts to use central memory to its greatest capacity. am ranges from 0 to 77777 ₈ .																						
CMcm	Central memory time slice in seconds. This parameter specifies the maximum amount of time a job of the specified service class can remain in central memory, either at a control point or at a pseudo-control point, before its priority is set to the lower boundary. When the job initially exhausts its CM time slice, this lower boundary is the value of il. Subsequently, when the job exhausts its CM time slice, this lower boundary is the value of lp (refer to the QUEUE command IL and LP parameters). cm ranges from 0 to 7777 ₈ .																						

Command	Description								
SERVICE,sc,p ₁ ,p ₂ ,...,p _n .	<i>(Continued)</i>								
	<table border="1"> <thead> <tr> <th>p_i</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>CPcp</td> <td>Control point slice priority. This parameter specifies the value of the scheduling priority which will be set after a job has been at a control point for ct seconds. The priority specified must be greater than or equal to the lower bound and less than or equal to the upper bound execution queue scheduling priorities specified by the CB parameter ($lp \leq cp \leq up$). Refer to the QUEUE command LP and UP parameters.</td> </tr> <tr> <td>CScs</td> <td>Cumulative size in PRUs allowed for all indirect access permanent files; cs indicates a limit value for the cumulative size: <ul style="list-style-type: none"> 0 Unlimited 1 1000_g 2 5000_g 3 50000_g 4 100000_g 5 200000_g 6 400000_g 7 Unlimited </td> </tr> <tr> <td>CTct</td> <td>Control point time slice in seconds. This parameter specifies the maximum amount of time that a job of the specified priority can remain at a control point before its priority is changed to the control point slice priority (cp). ct ranges from 1 to 7777_g.</td> </tr> </tbody> </table>	p_i	Description	CPcp	Control point slice priority. This parameter specifies the value of the scheduling priority which will be set after a job has been at a control point for ct seconds. The priority specified must be greater than or equal to the lower bound and less than or equal to the upper bound execution queue scheduling priorities specified by the CB parameter ($lp \leq cp \leq up$). Refer to the QUEUE command LP and UP parameters.	CScs	Cumulative size in PRUs allowed for all indirect access permanent files; cs indicates a limit value for the cumulative size: <ul style="list-style-type: none"> 0 Unlimited 1 1000_g 2 5000_g 3 50000_g 4 100000_g 5 200000_g 6 400000_g 7 Unlimited 	CTct	Control point time slice in seconds. This parameter specifies the maximum amount of time that a job of the specified priority can remain at a control point before its priority is changed to the control point slice priority (cp). ct ranges from 1 to 7777 _g .
p_i	Description								
CPcp	Control point slice priority. This parameter specifies the value of the scheduling priority which will be set after a job has been at a control point for ct seconds. The priority specified must be greater than or equal to the lower bound and less than or equal to the upper bound execution queue scheduling priorities specified by the CB parameter ($lp \leq cp \leq up$). Refer to the QUEUE command LP and UP parameters.								
CScs	Cumulative size in PRUs allowed for all indirect access permanent files; cs indicates a limit value for the cumulative size: <ul style="list-style-type: none"> 0 Unlimited 1 1000_g 2 5000_g 3 50000_g 4 100000_g 5 200000_g 6 400000_g 7 Unlimited 								
CTct	Control point time slice in seconds. This parameter specifies the maximum amount of time that a job of the specified priority can remain at a control point before its priority is changed to the control point slice priority (cp). ct ranges from 1 to 7777 _g .								

Command	Description
SERVICE,sc,p ₁ ,p ₂ ,...,p _n	<i>(Continued)</i>
	p_i
	Description
DSds	Size in PRUs allowed for individual direct access permanent files; ds (0 through 7) is used to specify the corresponding values shown next: 0 Unlimited 1 1000 ₈ 2 5000 ₈ 3 50000 ⁸ 4 100000 ₈ 5 200000 ₈ 6 400000 ₈ 7 Unlimited
DTsc	Service class to which a detached job is assigned. The default value for sc is DI (detached interactive).
ECec	Maximum user-accessible extended memory field length in words divided by UEBS ¹ for any job of the specified service class. This parameter performs the same function for extended memory field length that the FL parameter does for central memory field length. ec ranges from 0 to 3777 ₈ .
EMem	Maximum extended memory divided by UEBS ¹ for all jobs of the specified service class. This parameter performs the same function for extended memory field length that the AM parameter does for central memory field length. em ranges from 0 to 3777 ₈ .

1. Refer to Extended Memory Overview in section 3 to determine the value of UEBS.

Command	Description
SERVICE,sc,p₁,p₂,...,p_n <i>(Continued)</i>	
p_i	Description
FCfc	<p>Number of permanent files allowed; fc (0 through 7) is used to specify the corresponding values (limit value in octal) shown next:</p> <p style="margin-left: 40px;">0 Unlimited</p> <p style="margin-left: 40px;">1 10₈</p> <p style="margin-left: 40px;">2 40₈</p> <p style="margin-left: 40px;">3 100₈</p> <p style="margin-left: 40px;">4 200₈</p> <p style="margin-left: 40px;">5 1000₈</p> <p style="margin-left: 40px;">6 4000₈</p> <p style="margin-left: 40px;">7 Unlimited</p>
FLfl	<p>Maximum field length divided by 100₈ for any job of the specified service class. If more memory is requested, the job is aborted. You typically use this parameter to limit the memory requirement for jobs of a specific service class during certain hours of the day. For example, you may use the FL parameter to specify a maximum field length for all batch service class jobs between the hours of 2 p.m. and 4 p.m. fl ranges from 0 to 7777₈.</p>
FSfs	<p>Size in PRUs allowed for individual indirect access permanent files; fs (0 through 7) is used to specify the corresponding values (limit value in octal) shown next:</p> <p style="margin-left: 40px;">0 Unlimited</p> <p style="margin-left: 40px;">1 10₈</p> <p style="margin-left: 40px;">2 30₈</p> <p style="margin-left: 40px;">3 100₈</p> <p style="margin-left: 40px;">4 300₈</p> <p style="margin-left: 40px;">5 1000₈</p> <p style="margin-left: 40px;">6 2000₈</p> <p style="margin-left: 40px;">7 Unlimited</p>

Command	Description
SERVICE,sc,p ₁ ,p ₂ ,...,p _n .	(Continued)
p_i	Description
NJnj	Maximum number of jobs. For each service class, this parameter specifies the number of jobs that can be executing in the system. nj ranges from 0 to 7777 ₈ .
PRpr	CPU priority. When more than one job is waiting for the CPU, jobs with the highest CPU priority access the CPU first. Jobs with lower CPU priority access the CPU only when higher-priority jobs no longer wish to use the CPU. pr ranges from 2 to 77 ₈ .
RSrs	Recall CPU slice in milliseconds. The minimum guaranteed amount of time that a recalled job keeps the CPU before it can be preempted by another recalled job at the same CPU priority. rs ranges from 1 to 7777 ₈ , but the specified value must be less than or equal to the value of the unextended CPU slice, us.
SEse	CPU slice extension in milliseconds. The maximum additional amount of CPU time that a job performing I/O gets after its unextended CPU slice expires before the CPU slice expires (and the CPU is switched to an equal-priority job). This is the maximum extension. The slice is actually extended by the value of the recall CPU slice, rs, each time the job is recalled. se ranges from 0 to 7777 ₈ .
TDtd	Suspension time-out delay. A suspended job will not be timed out for td*10 ₈ seconds. td ranges from 0 to 7777 ₈ . The maximum delay is approximately 9 hours.
TPtp	Terminal job priority assigned to a terminal job that is rolled out waiting for terminal input. tp ranges from 0 to 7777 ₈ , but must be in the range of values for the execution queue parameters lp and up (refer to the QUEUE command LP and UP parameters).
USus	Unextended CPU slice in milliseconds. The amount of CPU time that a CPU-bound job gets before the CPU is switched to an equal-priority job. us ranges from 1 to 7777 ₈ , but the specified value must be greater than or equal to the value specified for the recall CPU slice, rs.

Command	Description
SERVICE,sc,p ₁ ,p ₂ ,...,p _n .	<p data-bbox="542 226 691 252"><i>(Continued)</i></p> <p data-bbox="542 279 1360 621">The service limits associated with each service class are established by using SERVICE command entries in the IPRDECK. These entries normally reflect the ideal service limits for the job mix of the particular installation. The values specified in the IPRDECK are important to optimum system performance and are not normally altered. However, when changes are necessary they are usually temporary and the original values will be reset. Figure 5-3 provides space to record the original values specified in the L display SDSPLAY. For additional information concerning the SERVICE command, refer to section 3, Deadstart Decks.</p>
SRST=t.	<p data-bbox="542 642 1360 730">Changes the secondary rollout sector threshold to the value t($0 \leq t \leq 7777_8$). Any rollout file smaller than t sectors is considered a secondary rollout file for the purpose of equipment selection.</p>
NOTE	
<p data-bbox="542 827 1300 915">The size of the rollout file for any job must be at least seven sectors larger than the combined size in sectors of the job's central memory and extended memory field lengths.</p>	

Job Service Class	Service Limits																			
	CP	CT	CM	NJ	TD	FL	AM	TP	AJ	DT	EC	EM	DS	FC	CS	FS	PR	SE	RS	US
SY																				
BC																				
RB																				
TS																				
DI																				
NS																				
SS																				
MA																				
CT																				
I0																				
I1																				
I2																				
I3																				
I4																				
I5																				
I6																				
I7																				
I8																				

04487-1M

Figure 5-3. Record of Original Values in L Display SDSPLAY

Peripheral Equipment Control Commands

The commands described in this category provide control of the channel-connected peripheral equipment available to the system.² To control unit record equipment that is connected through communications ports, use the RBF K display described in section 8 and the PSU K display described in the NOS Version 2 Operations Handbook. You should become familiar with the following DSD displays, which are closely associated with the use of these and other commands described throughout this section.

- Equipment status display (E,A).
- Disk configuration display (E,C).
- Disk errors display (E,E).
- Family status display (E,F).
- Disk thresholds display (E,H).
- Disk status display (E,M).
- Resource requests display (E,P).
- Tape status display (E,T).
- BIO status display (I).

A complete description of each of these displays is given in the NOS Version 2 Operations Handbook.

Command	Description
ASSIGN,jsn,est.	Assigns equipment defined by EST ordinal est (normally a tape unit) to the job with job sequence name jsn. This command is entered in response to a flashing REQUEST message. Use of this command for assignment of a tape unit should not normally be required because tape assignment is performed automatically when a volume serial number (VSN) is specified in the job request. However, if a VSN is not specified in the job request for a labeled or unlabeled tape, the REQUEST message appears at the job's control point (on the B,O display), and the ASSIGN command must be entered to assign a tape unit to the job.
	<p>NOTE</p> <p>The mnemonic DS (assigns the console EST entry) or NE (assigns the null equipment EST entry) can be used instead of the EST ordinal for these equipments.</p>
BKSP,est,rr.	Backspaces rr octal logical records on the print file for the BIO equipment defined by EST ordinal est. When rr is not specified, the default is one record.

² Operation of peripheral equipment is described in the NOS Version 2 Operations Handbook.

Command	Description
BKSPF,est,ff.	Backspaces ff octal files on the print file for the BIO equipment defined by EST ordinal est. When ff is not specified, the default is one file.
BKSPRU,est,ss.	Backspaces ss physical record units (PRUs) on the print file for the BIO equipment defined by EST ordinal est. The PRU count, ss, must be specified. There is no default setting. Printing resumes at the beginning of a line.
CONTINUE,est.	Resumes printing on the BIO equipment defined by EST ordinal est.
DOWN,CH=ch, EQ=est. or DOWN,CCH=cch, EQ=est. or DOWN,MCH=ch, EQ=est. or DOWN,EQ=est.	<p>The first two forms of the command disallow the use of channel ch or concurrent channel cch for the mass storage or device specified by EST ordinal est. If the EQ=est DOWN,CCH=cch,parameter is omitted, channel ch or cch is disallowed for all devices in the EST. The EQ=est parameter cannot be orspecified if the CH=ch or CCH=cch parameter specifies a magnetic tape channel. Use of all ports of a concurrent channel is disallowed.</p> <p>If channel ch or cch is the only channel available to a mass storage device, its use is disallowed providing the device is down. If channel ch or cch is</p> <ul style="list-style-type: none"> • not defined as a magnetic tape or mass storage channel • the last active channel on an equipment that is not globally unloaded (status N on the E,M display) • a nonimpact printer channel and any equipment on the channel is logically ON or assigned to a control point <p>the command is ignored and the following message appears on the left screen.</p> <p style="text-align: center;">INCORRECT ENTRY</p> <p>Correct the entry and retry the command.</p> <p>The third form of the command disallows the use of MUX channel ch for the equipment specified by EST ordinal est (this form applies only to model 176). If the EQ=est parameter is omitted, MUX channel ch is disallowed for all equipments in the EST.</p> <p>The fourth form of the command disallows the use of equipment with EST ordinal est for all channels. It disallows all access to the equipment by the operating system and user jobs. This device state is used to allow diagnostic routines exclusive use of the device for problem isolation and for preventing access to a device that is being repaired.</p>

Command	Description
DOWN,CH=ch, EQ=est.	<p data-bbox="537 226 683 254"><i>(Continued)</i></p> <p data-bbox="537 281 1365 499">This command also clears the SUSPCT state for a tape unit that is currently unavailable (but not OFF or DOWN) because an unrecovered hardware or tape I/O error has occurred on the tape or unit. This command can be entered followed by the ON,EQ=est command to restore the unit to service after a customer engineer has checked it out, or the unit has been tested following Control Data's recommended procedures:</p> <p data-bbox="537 527 1365 554">You receive an error message if you attempt to DOWN equipment:</p> <ul data-bbox="537 581 1365 800" style="list-style-type: none"> • That already has a DOWN status. • That is assigned to a job. • That is the last system device not DOWN.³ • That contains the system dayfile, account file error log, or maintenance log.³ <p data-bbox="537 827 1110 854">In these cases, the following message appears.</p> <pre data-bbox="586 877 802 905">INCORRECT ENTRY</pre> <p data-bbox="537 926 1057 953">Correct the entry and retry the command.</p>
DOWN,MID=mid.	<p data-bbox="537 982 618 1010">NOTE</p> <hr data-bbox="532 1014 1360 1020"/> <p data-bbox="537 1045 1279 1108">This command should be used with caution since it directly affects the operation of system peripheral equipment.</p> <hr data-bbox="532 1119 1360 1125"/> <p data-bbox="537 1157 1279 1283">Confirms that a machine running in low-speed port multimainframe mode with machine identifier mid is down. Enter this command in response to a request at the system control point of the form:</p> <pre data-bbox="586 1304 1045 1331">IF mid DOWN ENTER *DOWN,MID=mid*</pre> <p data-bbox="537 1352 1289 1415">Make this entry if the machine is down; otherwise, clear the message by entering:</p> <pre data-bbox="586 1436 683 1463">GO, SYS.</pre>

³. Applies only to mass storage equipment.

Command	Description
END,est,rc.	<p>Terminates the current operation on the BIO equipment defined by EST ordinal est. If est defines a line printer or card punch, BIO then assigns the next available file to that equipment. If est defines a card reader that is actively reading cards when END,est is entered, the job terminates at the last card read. The next card is treated as the beginning of a new job. If another card deck follows the end-of-information card (multipunch 6/7/8/9), it is processed normally.</p> <p>If est defines a card reader that is stopped because of a compare error when END,est. is entered, perform the following steps:</p> <ol style="list-style-type: none"> 1. Remove the remainder of the card deck, except the end-of-information card, from the card reader input hopper. 2. Ready the card reader to read the end-of-information card. <p>The job terminates; and if another card deck follows the end-of-information card, it is processed normally.</p> <p>The rc parameter cancels a portion of the repeat count specified for that equipment (refer to the REPEAT command later in this section). For example, if the current operation on equipment est had been set to be repeated five times (operation performed six times), entering a value of 4 for rc would permit the operation to be performed just two times. If the repeat count is zero, this command performs the END operation once.</p>
FORM,est,fc.	<p>Assigns a 2-character forms code fc to the BIO line printer or card punch defined by EST ordinal est. Only those files in the output queue assigned the forms code fc are directed to equipment est. A user can assign a forms code to an output file using the ROUTE command. (For a description of the ROUTE command, refer to the NOS Version 2 Reference Set, Volume 3.) The value of the forms code can range from AA to 99. If the forms code is not present, the current forms code field is cleared (value is null).</p>
ID,est,id.	<p>Assigns a numeric identifier id to the card punch, card reader, or printer defined by EST ordinal est. The value of the identifier can range from 00 to 67₈. Only those files in the proper queue with an identifier equal to id are directed to the card punch or printer defined by est. Refer to the LOAD command in this section to assign an identifier to a queue type file. All subsequent jobs loaded from card reader est are assigned the identifier id.</p>

Command	Description
IDLE,EQ=est.	<p>Use this command when you want to continue using the disk storage device specified by EST ordinal est even though it is failing occasionally. When the IDLE command is used in conjunction with the THRESHOLD command, it allows some degree of flexibility in deciding when to repair such a device. New files will not be assigned to a device with IDLE status unless no suitable alternative device exists, but files already on the device may continue to be accessed as if the device had a status of ON. If the device is a spun down 834, 836, 887, or 9853 drive, it will be automatically spun up. It may take up to 3 minutes to spin up the device.</p> <p>Suppose a failing disk storage device causes the system to set the device's status to OFF. Suppose also that the device is failing infrequently, for example, once or twice per hour, and contains several permanent files that must remain available during peak production periods. In this situation, you may use the THRESHOLD command to increase the verification failure threshold and then use the IDLE command to allow continued use of the device. This device could then be repaired after the peak production period has passed, causing less disturbance to system users.</p>

NOTE

The decision to continue using a failing device should be made under the advisement of a customer engineer. If there is a good chance of data corruption due to the nature of the failure, it may be best to leave the device OFF until it can be repaired.

If you enter an IDLE command for a device with a status of OFF or DOWN, the system performs a device verification to ensure the hardware is usable and the disk pack contains valid data. If this verification fails and the verification failure threshold is exceeded, the device state remains unchanged. However, if the device verification finds no errors the device's state changes to IDLE, the verification failure count (displayed on the E,H display) is set to zero, and any jobs that were rolled out waiting for the device are allowed to continue processing.

Command	Description
IDLE, EQ=est.	<p data-bbox="516 243 667 268"><i>(Continued)</i></p> <p data-bbox="516 296 1341 384">If you attempt to change a device's state to IDLE, the message INCORRECT ENTRY appears if one of the following conditions is true:</p> <ul data-bbox="521 411 1325 541" style="list-style-type: none"> <li data-bbox="521 411 919 436">• The device is already IDLE. <li data-bbox="521 464 1105 489">• The device is DOWN and assigned to a job. <li data-bbox="521 516 1325 541">• The device is DOWN and no channels are UP for the device. <p data-bbox="516 569 1304 657">If you enter IDLE for a device with a status of ON, the verification process is not performed; the device state is simply changed to IDLE.</p> <p data-bbox="516 684 1252 743">If you enter IDLE for a device other than a disk, the IDLE command is ignored.</p>
IDLEFAMILY, est.	<p data-bbox="516 764 1252 823">This command performs one of the two following functions, depending on the status of the family.</p> <ul data-bbox="521 850 1341 1150" style="list-style-type: none"> <li data-bbox="521 850 1341 1035">• If the family is active, the IDLEFAMILY command deactivates it by aborting all new jobs and USER commands for the family on the equipment specified by EST ordinal est. Jobs in progress are allowed to complete. ISF and permanent file utilities, such as PFDUMP, can be used on the inactive family. <li data-bbox="521 1062 1230 1150">• If the family is inactive, the IDLEFAMILY command activates it. Jobs can then access the family on the equipment specified by EST ordinal est. <p data-bbox="516 1178 1341 1234">The FAMC column of the E,F display shows the number of jobs in progress on the equipment.</p>

Command	Description
INITIALIZE,op,est ₁ , est ₂ ,...,est ₅ .	Reverses current setting of initialize option op for mass storage devices defined by EST ordinals est (maximum of five devices). Examine the E,A display to determine correct EST ordinals.
op	Level of Initialization
AF	Initialize inactive account dayfile.
AL	Total initialization.
DF	Initialize inactive system dayfile.
EF	Initialize inactive error log.
FP	Format pending. ⁴
FT	Total initialization as full-track device. ⁵
HT	Total initialization as half-track device. ⁵
MF	Initialize binary maintenance log.
PF	Initialize permanent files.
QF	Initialize inactive queue files.

This command provides the capability to initialize and flaw tracks on any mass storage device during normal system operation. Each time this command is entered it is logged in the error log.

If local unload (L) status is set for the device, the INITIALIZE command is ignored and the following message appears on the left console screen.

INCORRECT ENTRY

However, the INITIALIZE command is allowed to execute while local unload status is set if an error code (CE, IL, LE) is set.

4. Sets format pending bit in the MST for 881/883 pack reformatting. Refer to the description of the FORMAT program in appendix G.

5. Device must be removable.

Command	Description
INITIALIZE,op,est ₁ , est ₂ ,...,est ₅ .	(Continued)

Entry of this command does not in itself initialize the specified device. It merely sets initialize status for the device so that it may be initialized. However, if fast attach files (special system files) are active on the specified device and initialization level of AL or PF is specified, initialize status cannot be set until these files are returned. In this case, the message

FAST ATTACH FILES ON DEVICE.

appears at the system control point on the system status display (B,O). Refer to the description of this message in the NOS Version 2 Operations Handbook for additional information.

The procedure involved in initializing a mass storage device is described in section 8, K-Display Utilities, under the heading Initialize K Display. The following describes system activity when initialization occurs.

If the device is shared in a multimainframe environment, initialization does not proceed until all other mainframes sharing the device have processed an UNLOAD command for the device and user counts on all machines are zero. If all machines have not unloaded the device, this control point message is displayed:

EQest BUSY ON ID=id.

Variable	Description
est	EST ordinal.
id	Machine identifier of the first machine found without unload status set.

Command	Description
INITIALIZE,op,est ₁ , est ₂ ,...,est ₅ .	<p data-bbox="537 226 686 252"><i>(Continued)</i></p> <p data-bbox="537 281 1369 369">If initialize status is set on this device for another mainframe, the INITIALIZE command is ignored and this message appears at the system control point on the system status display (E,O):</p> <p data-bbox="591 399 1081 420">INITIALIZE PENDING ON THIS DEVICE.</p> <p data-bbox="537 449 1369 951">When the AL initialization option is specified, the label on the device to be initialized is either updated or a new label is created. If the label on the device is bad or cannot be recognized, the new label is created and all current data on the device is lost. If the label is found to be good, it is updated and all permanent file information is cleared. In this case, system library or temporary files (local, rollout, and so forth) residing on the device are not disturbed. If the device being initialized is a master device, the system scans all other mass storage devices in the family that contain direct access files and releases the space for files with catalogs on this device. If the device being initialized contains direct access files, the system scans all other master devices and sets the catalog entries on those devices to indicate that the files were purged. All or part of the permanent file system can be initialized and then reloaded if necessary (refer to the description of permanent file utilities in section 17, Permanent File Utilities).</p>
LOAD,est,id.	<p data-bbox="537 968 1321 1123">Requests that a job be loaded from equipment defined by EST ordinal est (normally tape unit). The job is assigned a numeric identifier id ranging from 00 to 67₈. If id is not specified, it defaults to 00. Examine the E,A display to determine the EST ordinal of the equipment to be used.</p>

Command	Description
MOUNT,est,P.	<p data-bbox="521 239 1351 394">Clears the local unload (L) and global unload (N) status for a mass storage device and reactivates the device. If the device is a spun down 834 or 836 drive, it will be automatically spun up. The device is defined by EST ordinal est (examine the E,A display to determine the EST ordinal).</p> <p data-bbox="521 417 1351 541">When you specify P in the MOUNT command for an independent shared device in a multimainframe environment, the system presets the device with EST ordinal est. The preset (P) option can be specified only on the first mainframe to access the device.</p> <p data-bbox="521 564 1351 657">If the device defined by EST ordinal is not a mass storage device, the MOUNT command is ignored and this message appears on the left console screen:</p> <p data-bbox="574 682 862 709" style="text-align: center;">INCORRECT EQUIPMENT.</p> <p data-bbox="521 732 1351 888">If the device is shared in a multimainframe environment and another mainframe has an unsatisfied initialize request pending for that device, the MOUNT command is ignored and this message appears at the system control point on the system status display (B,O):</p> <p data-bbox="574 913 1062 940" style="text-align: center;">INITIALIZE PENDING ON THIS DEVICE.</p>
MOUNT,est,vsn.	<p data-bbox="521 957 1351 1173">Allows an operator to mount a cartridge on the Automated Cartridge Subsystem (ACS). The tape drive is defined by EST ordinal est (examine the E,A display to determine the EST ordinal). The cartridge is defined by its external VSN. If a cartridge mounted using the MOUNT command has not been assigned to a job for an installation-defined period of time, the subsystem unloads and dismounts the cartridge.</p>

Command	Description
MSAL,t=est ₁ , est ₂ ,...,est _n .	Assigns job files of type t to mass storage devices defined by EST ordinals est _i . The specified mass storage device must be nonremovable. A limit of 47 devices may be specified on a single entry. Examine the disk status display (E,M) to determine if the device is nonremovable.

t	Description
B	LGO.
D	Job dayfile.
I	Input.
L	Local.
O	Output.
P	Primary.
R	Rollout.
S	Secondary rollout.
T	Temporary.

If no devices are specified for a file type, the system selects a temporary device.

Secondary rollout files do not exist until the command SRST=t (where t is the size in sectors of the file, from 0 to 7777₈) defines a threshold; the default value of a threshold is 0. All rollout files smaller than the threshold are secondary rollout files. These files are assigned to the devices specified with the MSAL,S command.

All files greater than or equal to the threshold are assigned to the rollout file devices (MSAL,R command).

The following example illustrates a use for secondary rollout files. Equipment 5 is extended memory via DDP. Secondary rollout files to extended memory is assigned by

```
MSAL, S=5.
```

Command	Description
MSAL,t=est ₁ , est ₂ ,...,est _n .	<p data-bbox="521 241 667 268"><i>(Continued)</i></p> <p data-bbox="521 289 1331 415">The threshold count is set by the SRST command so that rolled out files under 20 sectors long are assigned to extended memory. For additional information concerning the SRST command, refer to section 3, Deadstart Decks.</p> <pre data-bbox="574 443 987 787"> EQPDECK : EQ005=DP,ST=ON,SZ=4000,CH=30. : MSAL,S=5. : IPRDECK : SRST=20. : </pre>
NEXTVSN,est,vsn.	<p data-bbox="521 808 1331 934">Specifies that the tape with VSN vsn is to follow the tape currently mounted on the magnetic tape unit defined by EST ordinal est. This command is allowed only in response to the following prompt on the E,P display:</p> <pre data-bbox="574 955 967 982"> NEXTVSN,EST,VSN TO CONTINUE </pre> <p data-bbox="521 1003 1331 1192">This message is posted when an end-of-reel condition occurs on a tape file and the user does not specify the VSN of the next reel. The EST ordinal of the affected tape unit is displayed under the UNIT header on the E,P display. After entry of the NEXTVSN command, a request to mount the specified VSN is posted on the E,P display.</p>

Command	Description
OFF,EQ=est.	<p data-bbox="532 226 1360 506">Logically turns off the device defined by EST ordinal est. This command allows you to logically remove a device from the operating environment so that no user access of the device is permitted. If a user job attempts to access the device, it is rolled out to wait for the device status to become ON or IDLE. However, system utilities are permitted to access the device, so it can be dumped or loaded. If the device is a spun down 834, 836, 887, or 9853 drive, it will be automatically spun up. It may take up to 3 minutes to spin up the device.</p> <p data-bbox="532 533 1360 684">If you enter the OFF command for a device with a status of DOWN, the system performs a verification of the device label. This provides assurance that the correct pack is mounted on the drive and the controller is accessing the right drive after a repair action.</p> <p data-bbox="532 714 1360 800">If you attempt to change a device's state to OFF, the message INCORRECT ENTRY appears if one of the following conditions is true:</p> <ul data-bbox="532 829 1360 961" style="list-style-type: none"> • The device is already OFF. • The device is DOWN and assigned to a job. • The device is DOWN and no channels are UP for the device. <p data-bbox="532 989 1360 1331">Examine the E,A display to determine the EST ordinal and current status (ON, OFF, IDLE, or DOWN) of the device. If est specifies a disk storage device and the system library or temporary files (local, rollout, and so forth) reside on that device, it should not be turned off. Examine the disk status display (E,M) to determine which disk storage devices have system residency or allow system allocation of temporary files. In addition, if an MSAL entry is currently specified for a disk storage device, it is cleared when that device is turned off. The MSAL designation is not reset automatically when the device is turned back on and must be reset manually (if necessary) using the DSD command MSAL.</p> <p data-bbox="532 1358 1360 1575">This command also clears the SUSPCT state for a tape unit that is currently unavailable (but not OFF or DOWN) because an unrecovered hardware or tape I/O error has occurred on the tape or unit. This command can be entered followed by the ON,EQ=est command to restore the unit to service after a customer engineer has checked it out, or the unit has been tested following Control Data's recommended procedures.</p>

Command	Description
ON, EQ=est.	<p>Logically turns on the device defined by EST ordinal est. This command allows you to activate a device currently having OFF status in the EST. If the device is a spun down 834, 836, 887, or 9853 drive, it will be automatically spun up. Examine the E,A display to determine the EST ordinal and current status (ON, OFF, IDLE, or DOWN) of the device.</p> <p>Use this command to bring a disk storage device back into use following its repair. You can also use the ON command after repairing a controller that provides the only access to a disk storage device.</p> <p>When you enter ON for a device with a status of OFF or DOWN, the device's state does not change to ON immediately. First the system attempts to verify that the drive and channels accessing the device are usable and the data on the disk pack is in the proper format. If this verification process fails and the verification failure threshold is exceeded, the system sets the device's state to OFF. If the verification finds no errors, the device's state changes to ON, the verification failure count (displayed on the E,H display) is set to zero, and any jobs that were rolled out waiting for the device are allowed to continue processing.</p> <p>If you enter ON for device with a status of IDLE, the verification process is not performed; the device's state simply changes to ON.</p> <p>If you attempt to change a device's state to ON, the message INCORRECT ENTRY appears if one of the following conditions is true:</p> <ul style="list-style-type: none"> • The device is already ON. • The device is DOWN and assigned to a job. • The device is DOWN and no channels are UP for the device.
PRSIZE, est, ps.	<p>Sets the paper status ps to short (S) or long (L) paper for the printer with EST ordinal est. Short paper has a form length of 8.5 inches and long paper has a form length of 11 inches.</p>
REPEAT, jsn, rc. ⁶	<p>Alters the repeat count for the queue file specified by job sequence name jsn. If the file is still in the queue, the repeat count for the file is set to rc. If the file is already being processed by BIO, the remaining repeat count for the file is increased by rc. The maximum value that can be entered for rc is 77₆. The default value for rc is 1.</p>

6. When the current BIO operation is repeated, maximum line and card limits are reinitialized prior to printing or punching of the file being processed. User control limits apply individually to each output file copy produced.

Command	Description
REPRINT,est,pr.	Terminates the current operation on the BIO printer equipment defined by EST ordinal est and reenters the job in the print queue with a queue priority specified by pr00 (service class minimum \leq pr00 \leq service class maximum; refer to the S display in the NOS Version 2 Operations Handbook; pr value is multiplied by 100 _g internally). If pr is not specified, the service class default priority is assigned.
REPUNCH,est,pr.	Terminates the current operation on the BIO card punch equipment defined by EST ordinal est and reenters the job in the punch queue with a queue priority specified by pr00 (service class minimum \leq pr00 \leq service class maximum; refer to the S display in the NOS Version 2 Operations Handbook; pr value is multiplied by 100 _g internally). If pr is not specified, the service class default priority is assigned.
RESTORE,est.	Initiates the restoration of parity protection for the 583x (DAS) device defined by EST ordinal est. This command is valid only for a device with a device type of ED, EF, or EH.
RETRY,est.	<p data-bbox="537 810 1365 869">Reissues a tape operation that previously aborted with one of the following load point errors appearing in the E,P display:</p> <p data-bbox="594 894 829 982" style="text-align: center;"> CLEANER FAULT READ ID BURST WRITE ID BURST </p> <p data-bbox="537 1010 1365 1167">est represents the EST ordinal of the tape unit where the load point error occurred. Examine the E,P display to determine the EST ordinal. The tape unit must be ready before you enter the RETRY command. Examine the E,T display to ensure that the tape is loaded and the tape unit is ready.</p> <p data-bbox="537 1188 1365 1470">The RETRY command is to be used after the system detects a load point error on a tape. Upon detecting such an error, the system issues an error message to the E,P display and unloads the tape. The operator then has the opportunity to fix the failing tape or tape drive. Upon completing this step, the operator reloads the tape on the same tape unit and enters the RETRY command. The system then reissues the tape operation. If the failure cannot be fixed at this time, the operator should enter the TERMINATE command to abort the recovery attempt.</p>

Command	Description
SCRATCH,est.	<p>Declares the tape mounted on an unassigned magnetic tape unit, defined by EST ordinal est, to be a scratch tape. This command enables a tape to be available to satisfy scratch VSN requests and still be assigned by its original VSN. Thus, the VSN defined on the tape (in VOL1 label) is not redefined as scratch although the VSN will appear as SCRATCH on the tape status display (E,T).</p> <p>Scratch status is retained for only one job assignment. This allows a tape to be used for scratch purposes on a temporary basis. For example, a job requests a tape mounted on the tape unit defined in this command by specifying the current VSN for that tape in the request. The tape is then assigned to the job as a scratch tape (the original VSN is retained and not made scratch). When that job releases the tape, SCRATCH status is cleared, and unless this command is entered again, that tape would not be assigned as a scratch tape in future requests. To determine if SCRATCH status is in effect for a tape, monitor the tape status display (E,T).</p>
SECUREQ,est, LA=lowerlevel, UA=upperlevel.	<p>Changes the equipment access level limits for the unit record equipment with EST ordinal est. On a secured system, only files that have access levels within these limits can be printed on the specified equipment. The original limits are set during deadstart by the ACCESS command entries in the EQPDECK. The parameters LA=lowerlevel and UA=upperlevel specify the lower and upper access level limits for the equipment.</p>
NOTE	
<p>On a secured system this command is accepted only if a security administrator has placed the system in security unlock status.</p>	
SKIP,est,rr.	<p>Skips forward rr octal logical records on the print file for the BIO equipment defined by EST ordinal est. When rr is not specified the default is one record.</p>
SKIPF,est,ff.	<p>Skips forward ff octal files on the print file for the BIO equipment defined by EST ordinal est. When ff is not specified the default is one file.</p>

Command	Description
SKIPRU,est,ss.	Skips forward ss PRUs on the print file for the BIO equipment defined by EST ordinal est. All parameters must be specified; there are no default settings. The PRU count, ss, is limited to 10 ₈ PRUs (the current buffer size) plus the number of PRUs remaining in the buffer. If the buffer was empty, ss would be limited to 20 ₈ PRUs. If ss is larger than the number of PRUs remaining in the buffer, the buffer is filled again and the additional PRUs are skipped on the new print file. For example, if five PRUs remained in the print file and ss was specified as 10, the remaining five PRUs would be skipped, the buffer filled again, and five additional PRUs skipped. Printing resumes at the beginning of a line.
SPINDOWN,est.	Spins down the 834, 836, 887, or 9853 disk storage device defined by EST ordinal est. To enter this command, the console must be unlocked (refer to the UNLOCK command later in this section).
SPINUP,est.	Spins up the 834, 836, 887, or 9853 disk storage device defined by EST ordinal est.
NOTE	
Spinning down an 834, 836, 887, or 9853 disk storage device that does not have global unload status (N), can cause mass storage device status errors or permanent file errors when the device is spun up.	
STOP,est.	Stops printing on the BIO equipment defined by EST ordinal est.

Command	Description
SUPPRESS,est.	<p>Suppresses automatic printer carriage control on the BIO line printer defined by EST ordinal est. This command stops the page eject function on the line printer to provide a continuous listing for the current job.</p>
TEMP=est ₁ , est ₂ ,...,est _n .	<p>Reverses current set or clear condition of temporary file status for mass storage devices defined by EST ordinals est_i. When temporary file status is set, the system can use the specified device for allocation of temporary files. This command is not valid if the device specified is defined as removable. Examine the disk status display (E,M) to determine:</p> <ul style="list-style-type: none"> • The EST ordinal of the device. • If the device is defined as removable. • If temporary file status is currently selected (set) for the device.
TERMINATE,est.	<p>Aborts a tape load point error recovery attempt (refer to the RETRY command in this section). est represents the EST ordinal of the tape unit where the load point error occurred. Examine the E,P display to determine the EST ordinal.</p> <p>The TERMINATE command also clears the error message in the E,P display and issues the error message to the system dayfile.</p>

Command	Description												
THRESHOLD, type=value, est ₁ ,est ₂ ,...,est _n .	You can use the THRESHOLD command to set disk storage device thresholds for the devices defined by EST ordinals est _i . You must specify at least one EST ordinal. Use this command in conjunction with the IDLE command. You can specify one of the following threshold types. If no threshold type is specified, VF is assumed.												
	<table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>VF</td> <td> <p>Verification failure threshold. Each time an unrecovered error occurs on a device, the system performs a device verification. If the error occurs again during the verification process and the error is in the device or in the only available channel or in controlled access to the device, the verification failure count is incremented and compared with the verification failure threshold. If the threshold is exceeded, the system sets the device's state to OFF.</p> <p>If the threshold is not exceeded, the device's state is set to IDLE to restrict allocation of new files on the device. Because a disk storage error can cause a job to abort, it is important to understand that the higher this threshold is set the more jobs that may abort before the system sets the device's state to OFF. By setting the verification failure threshold to a value greater than zero, you can continue to use a failing device. Do this only when the failures are relatively infrequent, the failures do not cause corruption of permanent files, and the device is essential during the next several hours of production. After the peak hours of production, the device can be repaired during a planned corrective maintenance period.</p> </td> </tr> <tr> <td>RA</td> <td>Restricted activity threshold. When the available space on a device falls below this threshold, new files are not assigned to the device unless no other suitable device is available.</td> </tr> <tr> <td>LS</td> <td>Low space threshold. When the available space on a device falls below this threshold, the operator is notified on the A,OPERATOR display.</td> </tr> <tr> <td>RE</td> <td>Recovered error threshold. Each time a recovered error occurs on a disk storage device a count is incremented and checked against the recovered error threshold for that device. If the threshold is exceeded, the operator is notified on the A,OPERATOR display.</td> </tr> <tr> <td>UE</td> <td>Unrecovered error threshold. Each time an unrecovered error occurs on a disk storage device a count is incremented and checked against the unrecovered error threshold for that device. If the threshold is exceeded, the operator is notified on the A,OPERATOR display.</td> </tr> </tbody> </table>	type	Description	VF	<p>Verification failure threshold. Each time an unrecovered error occurs on a device, the system performs a device verification. If the error occurs again during the verification process and the error is in the device or in the only available channel or in controlled access to the device, the verification failure count is incremented and compared with the verification failure threshold. If the threshold is exceeded, the system sets the device's state to OFF.</p> <p>If the threshold is not exceeded, the device's state is set to IDLE to restrict allocation of new files on the device. Because a disk storage error can cause a job to abort, it is important to understand that the higher this threshold is set the more jobs that may abort before the system sets the device's state to OFF. By setting the verification failure threshold to a value greater than zero, you can continue to use a failing device. Do this only when the failures are relatively infrequent, the failures do not cause corruption of permanent files, and the device is essential during the next several hours of production. After the peak hours of production, the device can be repaired during a planned corrective maintenance period.</p>	RA	Restricted activity threshold. When the available space on a device falls below this threshold, new files are not assigned to the device unless no other suitable device is available.	LS	Low space threshold. When the available space on a device falls below this threshold, the operator is notified on the A,OPERATOR display.	RE	Recovered error threshold. Each time a recovered error occurs on a disk storage device a count is incremented and checked against the recovered error threshold for that device. If the threshold is exceeded, the operator is notified on the A,OPERATOR display.	UE	Unrecovered error threshold. Each time an unrecovered error occurs on a disk storage device a count is incremented and checked against the unrecovered error threshold for that device. If the threshold is exceeded, the operator is notified on the A,OPERATOR display.
type	Description												
VF	<p>Verification failure threshold. Each time an unrecovered error occurs on a device, the system performs a device verification. If the error occurs again during the verification process and the error is in the device or in the only available channel or in controlled access to the device, the verification failure count is incremented and compared with the verification failure threshold. If the threshold is exceeded, the system sets the device's state to OFF.</p> <p>If the threshold is not exceeded, the device's state is set to IDLE to restrict allocation of new files on the device. Because a disk storage error can cause a job to abort, it is important to understand that the higher this threshold is set the more jobs that may abort before the system sets the device's state to OFF. By setting the verification failure threshold to a value greater than zero, you can continue to use a failing device. Do this only when the failures are relatively infrequent, the failures do not cause corruption of permanent files, and the device is essential during the next several hours of production. After the peak hours of production, the device can be repaired during a planned corrective maintenance period.</p>												
RA	Restricted activity threshold. When the available space on a device falls below this threshold, new files are not assigned to the device unless no other suitable device is available.												
LS	Low space threshold. When the available space on a device falls below this threshold, the operator is notified on the A,OPERATOR display.												
RE	Recovered error threshold. Each time a recovered error occurs on a disk storage device a count is incremented and checked against the recovered error threshold for that device. If the threshold is exceeded, the operator is notified on the A,OPERATOR display.												
UE	Unrecovered error threshold. Each time an unrecovered error occurs on a disk storage device a count is incremented and checked against the unrecovered error threshold for that device. If the threshold is exceeded, the operator is notified on the A,OPERATOR display.												

Command	Description																		
THRESHOLD, type=value, est ₁ ,est ₂ ,...,est _n .	<p>(Continued)</p> <p>The threshold value can range from 0 to 3777₈. Specify a threshold value of * to obtain the default value for the specified type, for example, THRESHOLD,RA=*,21. The following default values are used.</p> <table border="1"> <thead> <tr> <th>type</th> <th>value</th> </tr> </thead> <tbody> <tr> <td>VF</td> <td>0</td> </tr> <tr> <td>RA</td> <td>1/8 of the number of tracks on the device</td> </tr> <tr> <td>LS</td> <td>1/16 of the number of tracks on the device</td> </tr> <tr> <td>RE</td> <td>50₈</td> </tr> <tr> <td>UE</td> <td>0</td> </tr> </tbody> </table>	type	value	VF	0	RA	1/8 of the number of tracks on the device	LS	1/16 of the number of tracks on the device	RE	50 ₈	UE	0						
type	value																		
VF	0																		
RA	1/8 of the number of tracks on the device																		
LS	1/16 of the number of tracks on the device																		
RE	50 ₈																		
UE	0																		
TRAIN,est,t.	<p>Assigns or changes print train identifier t of the line printer defined by EST ordinal est. This command can set the identification if it was not specified in the EQ entry of the EQPDECK, or change an identification previously included in the EQPDECK. An LR designation in the EQ entry indicates a 580-12 line printer, LS is a 580-16 line printer, LT is a 580-20 line printer, and LX is a 5870 printer. Print trains supported for the 580 printers are 595-1/596-1, 595-5/596-5, and 595-6/596-6. The print train supported for the 5870 printer is t = 7. The t field specifies the print train.</p> <table border="1"> <thead> <tr> <th>t</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>595-1/596-1 (CDC graphic 63/64-character set); default.</td> </tr> <tr> <td>1</td> <td>595-1/596-1 (CDC graphic 63/64-character set).</td> </tr> <tr> <td>2</td> <td>Reserved for future use.⁷</td> </tr> <tr> <td>3</td> <td>Reserved for future use.⁷</td> </tr> <tr> <td>4</td> <td>595-6/596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).</td> </tr> <tr> <td>5</td> <td>595-6/596-6 (ASCII graphic 63/64-character set).</td> </tr> <tr> <td>6</td> <td>595-6/596-6 (ASCII graphic 95-character set).</td> </tr> <tr> <td>7</td> <td>595-6/596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).</td> </tr> </tbody> </table>	t	Description	0	595-1/596-1 (CDC graphic 63/64-character set); default.	1	595-1/596-1 (CDC graphic 63/64-character set).	2	Reserved for future use. ⁷	3	Reserved for future use. ⁷	4	595-6/596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).	5	595-6/596-6 (ASCII graphic 63/64-character set).	6	595-6/596-6 (ASCII graphic 95-character set).	7	595-6/596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).
t	Description																		
0	595-1/596-1 (CDC graphic 63/64-character set); default.																		
1	595-1/596-1 (CDC graphic 63/64-character set).																		
2	Reserved for future use. ⁷																		
3	Reserved for future use. ⁷																		
4	595-6/596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).																		
5	595-6/596-6 (ASCII graphic 63/64-character set).																		
6	595-6/596-6 (ASCII graphic 95-character set).																		
7	595-6/596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).																		

7. These values are allowed but will default to 595-1/596-1.

Command	Description																					
TRAIN,est,t.	<p><i>(Continued)</i></p> <p>If an invalid external characteristic (EC)⁸ is specified, the queued file processor cannot output the file. The following shows which files will print and which files will not print for a given print train selection.</p> <table border="1"> <thead> <tr> <th>Print Train Selected</th> <th>Will Print File With the Specified EC</th> <th>Will Not Print File With the Specified EC</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>None, B4, B6</td> <td>A4, A6, A9</td> </tr> <tr> <td>1</td> <td>None, B4, B6</td> <td>A4, A6, A9</td> </tr> <tr> <td>4</td> <td>None, A4, A6, A9</td> <td>B4, B6</td> </tr> <tr> <td>5</td> <td>None, A4, A6</td> <td>B4, B6, A9</td> </tr> <tr> <td>6</td> <td>A9</td> <td>None, B4, B6, A4, A6</td> </tr> <tr> <td>7</td> <td>None, B4, B6, A4, A6, A9</td> <td></td> </tr> </tbody> </table>	Print Train Selected	Will Print File With the Specified EC	Will Not Print File With the Specified EC	0	None, B4, B6	A4, A6, A9	1	None, B4, B6	A4, A6, A9	4	None, A4, A6, A9	B4, B6	5	None, A4, A6	B4, B6, A9	6	A9	None, B4, B6, A4, A6	7	None, B4, B6, A4, A6, A9	
Print Train Selected	Will Print File With the Specified EC	Will Not Print File With the Specified EC																				
0	None, B4, B6	A4, A6, A9																				
1	None, B4, B6	A4, A6, A9																				
4	None, A4, A6, A9	B4, B6																				
5	None, A4, A6	B4, B6, A9																				
6	A9	None, B4, B6, A4, A6																				
7	None, B4, B6, A4, A6, A9																					

8. Refer to the NOS Version 2 Reference Set, Volume 3 for a discussion of the ROUTE command EC parameter.

Command	Description
UNLOAD,est.	Physically unloads a tape or logically removes a removable disk storage device from the operating system. The device to be unloaded is defined by EST ordinal est (examine the E,A display to determine the EST ordinal). Also, in a multimainframe environment, the UNLOAD command must be issued if another mainframe wants to initialize a shared mass storage device, whether the device is removable or nonremovable (refer to the INITIALIZE command earlier in this section). If a magnetic tape is specified, the tape is physically unloaded. If a removable disk storage device is specified, you can dismount the disk pack on that device after all activity on the device has been completed.

NOTE

If a nonremovable shared mass storage device is to be specified, the console must be unlocked (refer to the UNLOCK command later in this section).

Magnetic tape units:

If a tape is currently assigned to a job, it cannot be unloaded. If you attempt to unload it, the UNLOAD command is ignored and this message appears on the left console screen:

```
UNIT NOT AVAILABLE
```

Examine the tape status display (E,T) before entering the UNLOAD command to determine if the tape to be unloaded is currently assigned to a job. If the tape is not currently assigned, entering this command unloads the specified tape.

Mass storage devices:

The UNLOAD command is valid for any shared mass storage device in a multimainframe environment for the purpose of initialization. Otherwise, the command is valid only for removable devices. (Only removable devices may be physically removed by unloading.) After entering the UNLOAD command, monitor the disk status display (E,M). Execution of this command immediately causes local unload (L) status to appear in the STATUS field for that device. While L status is displayed, no new users are permitted to access files on the device. A user currently accessing files on the device can continue while at least one direct access file from the device is attached to the job. When the user count is zero and there are no checkpoint requests pending, one of the following two actions occurs.

Command	Description
UNLOAD,est.	<p data-bbox="526 222 675 254"><i>(Continued)</i></p> <ul data-bbox="526 275 1305 401" style="list-style-type: none"> <li data-bbox="526 275 1305 401">• If the device is removable and the L status is set in all machines accessing the device, global unload (N) status is displayed. This indicates that the device may now be physically dismantled. <p data-bbox="578 432 667 464">NOTE</p> <hr data-bbox="578 464 1351 468"/> <p data-bbox="578 495 1321 558">If a multispindle family is mounted on a single spindle device, only the first device shows the global unload status.</p> <hr data-bbox="578 558 1351 562"/> <ul data-bbox="526 615 1317 741" style="list-style-type: none"> <li data-bbox="526 615 1317 741">• If an initialize is pending on the device and all other machines accessing the device have L status set, the initialization proceeds. However, initialization cannot take place if the device has been unloaded. <p data-bbox="578 772 667 804">NOTE</p> <hr data-bbox="578 804 1351 808"/> <p data-bbox="578 835 1351 898">A device should be physically dismantled only if global unload status (N) is displayed on all machines accessing the device.</p> <hr data-bbox="578 898 1351 903"/> <p data-bbox="526 957 1235 1020">If a removable pack is dismantled before the N status is displayed, the following may occur.</p> <ul data-bbox="526 1041 1351 1335" style="list-style-type: none"> <li data-bbox="526 1041 1003 1073">• Mass storage device status errors. <li data-bbox="526 1094 1317 1157">• Permanent file errors when the pack is remounted at some later date. <li data-bbox="526 1178 1351 1335">• If another pack has been mounted, accesses made by a previously attached user may destroy information on the new pack or the user may retrieve information from the new device which he is not necessarily privileged to access. Mass storage device status errors are also possible in this situation. <p data-bbox="578 1367 667 1398">NOTE</p> <hr data-bbox="578 1398 1351 1402"/> <p data-bbox="578 1430 1333 1558">If the Mass Storage Extended Subsystem (MSE) is active, it must be idled before unloading a removable family pack which has MSE files. After dismantling the family pack, MSE can be initialized again.</p> <hr data-bbox="578 1558 1351 1562"/>

Command	Description
UNLOAD,est.	<p data-bbox="521 243 667 268"><i>(Continued)</i></p> <p data-bbox="521 296 748 321">Optical disk units:</p> <p data-bbox="521 348 1333 436">If an optical disk is currently assigned to a job, it cannot be unloaded. If you attempt to unload it, the UNLOAD command is ignored and this message appears on the left console screen:</p> <pre data-bbox="573 468 829 489">UNIT NOT AVAILABLE</pre> <p data-bbox="521 516 1300 604">To determine if the optical disk to be unloaded is currently assigned to a job, examine the optical disk status display (E,S) before entering the UNLOAD command.</p> <p data-bbox="521 632 1235 688">If the optical disk is not currently assigned, entering this command unloads the specified optical disk.</p> <p data-bbox="220 709 399 894">UP,CH=ch. or UP,MCH=ch.⁹ or UP,CCH=cch, EQ=est.</p> <p data-bbox="521 709 1341 798">The first two forms of the command allow resumption of normal use of channel ch by magnetic tape units or mass storage devices, reversing the effects of the DOWN command.</p> <p data-bbox="521 825 1341 947">The third form of the command allows resumption of normal use of both ports of concurrent channel cch for the mass storage device specified by EST ordinal est. If the EQ=est parameter is omitted, channel cch resumes normal use for all devices in the EST.</p> <p data-bbox="521 974 1325 1031">If you enter UP for a channel (ch or cch) that is already UP, the following message appears on the console:</p>
VALIDATE,est.	<pre data-bbox="573 1058 789 1079">INCORRECT ENTRY</pre> <p data-bbox="521 1106 1260 1131">If the channel is already up, no further action is necessary.</p> <p data-bbox="521 1159 1308 1276">Causes validation of mass storage tables associated with the equipment defined by EST ordinal est. The equipment must be available mass storage and the MS VALIDATION option must have been selected in the IPRDECK at deadstart.</p>
VSN,est.	<p data-bbox="521 1297 1333 1386">Reverses the effect of a previous VSN,est,vsn command, and forces the magnetic tape subsystem to read the label on the tape. This command is not allowed on ACS tape units.</p>

⁹. Model 176 only.

Command	Description				
VSN,est,.	Declares the tape mounted on an unassigned magnetic tape unit, defined by EST ordinal est, to be a scratch tape. This command is similar in function to the SCRATCH command in that it enables a tape to be available to satisfy scratch VSN requests. However, if the tape is labeled and a write function is performed, the VSN specified in the VOL1 label will be rewritten as a scratch VSN, destroying the original VSN and making the tape available for future scratch VSN requests. The VSN also appears as ***est (est is the ordinal of the est) on the tape status display (E,T). Refer to the VSN,est,vsn command, next, for a discussion of the INCORRECT ENTRY message, which is also applicable to VSN,est,. This command is not allowed on ACS tape units.				
VSN,est,vsn. ¹⁰	<p>Assigns VSN vsn to an unassigned magnetic tape unit defined by EST ordinal est. This command allows you to specify a 1- to 6-character VSN for a mounted, unlabeled tape so it may be assigned and referenced automatically. For example, when a job specifies a VSN in the request for an unlabeled tape, an entry for that job appears in the resource requests display (E,P). This display indicates the job sequence name of the job; the type of tape unit, 7-track (MT), 9-track (HD, PE, or GE), or CTS cartridge (CT), on which the tape is to be mounted; the required VSN; the user name of the job; and the required write ring status (IN or OUT). If the correct tape is not currently mounted, mount the tape on an available unit (ensuring that track type and write ring status are correct), ready the unit, and enter this command. The system equates the VSN you enter with that specified by the job and assigns the tape automatically upon demand. This command is not allowed on ACS tape units.</p> <p>If the tape mounted on the tape unit defined by EST ordinal est is a labeled tape, has already had a VSN assigned by console command, or has not yet been checked for a label by the magnetic tape subsystem, this command is ignored. The message</p> <p style="text-align: center;">INCORRECT ENTRY</p> <p>appears on the left console screen. To change a VSN previously assigned by this command, clear the first VSN by entering:</p> <p style="text-align: center;">VSN,est.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>est</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>est</td> <td>EST ordinal of the tape unit.</td> </tr> </tbody> </table>	est	Description	est	EST ordinal of the tape unit.
est	Description				
est	EST ordinal of the tape unit.				

¹⁰. Special characters cannot be entered using this command. If a special character is encountered in vsn, the VSN entered is truncated at the character preceding the special character.

Command	Description						
VSN,est,vsn.	<p><i>(Continued)</i></p> <p>This command can then be entered:</p> <p style="padding-left: 40px;">VSN, est, vsn.</p> <table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>est</td> <td>EST ordinal of the tape unit.</td> </tr> <tr> <td>vsn</td> <td>New VSN.</td> </tr> </tbody> </table>	Parameter	Description	est	EST ordinal of the tape unit.	vsn	New VSN.
Parameter	Description						
est	EST ordinal of the tape unit.						
vsn	New VSN.						

(The preceding message and procedure for changing a VSN are also applicable to the VSN,est,. command.)

If a job specifies a VSN in the request for a labeled tape, assignment occurs automatically, without your intervention, unless the correct tape is not mounted. In this case, an entry is formed in the resource requests display (E,P), which describes the tape to be mounted. When the tape is mounted and the tape unit made ready, assignment occurs automatically without additional intervention by you. For multivolume files, automatic tape assignment occurs only if the tape units on which the tapes are mounted are similar and on the same channel(s). That is, if the first volume of the file is on a 669 tape unit on channels 13 and 33, all subsequent volumes must be on a 669 unit on channels 13 and 33. When assigning tapes, models 679-2, 679-3, and 679-4 drives (800/1600 cpi) are similar. Also models 679-5, 679-6, and 679-7 drives (1600/6250 cpi) are similar. All CTS tape units are similar and all ACS tape units are similar.

If two or more unassigned tapes having identical VSNs are mounted on units of the same tape type, the message

REQUEST, dt, vsn

appears on the B,O and B,A displays. The dt field is either a device type (CT or MT) or a density requirement (HD, PE, or GE); vsn is the VSN required.

You must assign one of the tapes using the ASSIGN command. If the duplicate VSNs are SCRATCH, the resource executive routine assigns one automatically.

NOTE

It is not possible to specify a VSN of SCRATCH with this command since only six characters may be used to define a VSN. To define a scratch tape (used to satisfy scratch VSN requests), refer to the description of the SCRATCH command earlier in this section.

Subsystem Control Commands

The commands that follow provide control over which subsystems are to be used. When a system deadstart is performed, parameters specified in the IPRDECK determine which subsystems initially are available. Scheduling other subsystems to a control point or terminating a current subsystem depends on your action.

When a subsystem is scheduled to a specific control point, any job currently assigned to that control point is rolled out if it is not another subsystem or special system job. If the job currently assigned to the control point is another subsystem or cannot be rolled out, the new subsystem waits for the control point to become free. In this case, either terminate the job (if the subsystem requires that control point) or specify another control point for the subsystem using the SUBSYST L-display utility (refer to section 9, L-Display Utilities, for more information).

Command	Description						
BIO.	<p>Calls the Batch Input/Output (BIO) Subsystem. BIO must be active in order to use any of the local batch peripheral devices. BIO equipment activity can be monitored on the I display (refer to the NOS Version 2 Operations Handbook).</p> <p>Response to line printer errors is controlled by setting sense switch 1 or 2 as follows:</p> <table border="1"> <thead> <tr> <th>Entry</th> <th>Response</th> </tr> </thead> <tbody> <tr> <td>ONSW,BIO,1.</td> <td>Lines causing print error messages are not flagged or retried.</td> </tr> <tr> <td>ONSW,BIO,2.</td> <td>Lines causing print error messages are flagged and retried. The printer halts and is placed on hold until you correct the error and enter CONTINUE,est. The line reprints and printing continues.</td> </tr> </tbody> </table>	Entry	Response	ONSW,BIO,1.	Lines causing print error messages are not flagged or retried.	ONSW,BIO,2.	Lines causing print error messages are flagged and retried. The printer halts and is placed on hold until you correct the error and enter CONTINUE,est. The line reprints and printing continues.
Entry	Response						
ONSW,BIO,1.	Lines causing print error messages are not flagged or retried.						
ONSW,BIO,2.	Lines causing print error messages are flagged and retried. The printer halts and is placed on hold until you correct the error and enter CONTINUE,est. The line reprints and printing continues.						
CDCffff. ¹¹	Calls the CYBER Database Control System (CDCS). CDCffff calls a procedure file named CDCffff under the system user index to initiate the CDCS Subsystem.						
CYBffff. ¹¹	Calls the CYBIS-NAM Interface (CYB) Subsystem to the control point specified in the IPRDECK or by the SUBSYST L-display utility. CYBffff calls a procedure file named CYBffff under the system user index to initiate the CYB Subsystem.						
IAFffff. ¹¹	Calls the Interactive Facility (IAF) Subsystem to control point 1. IAFffff calls a procedure file named IAFffff under the system user index to initiate the IAF Subsystem. The IAF Subsystem must be active before jobs can be entered from an interactive terminal. Use the T display to monitor activity at the IAF Subsystem control point.						

¹¹. The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

Command	Description														
IAFffff.	<p><i>(Continued)</i></p> <p>Several options are available to control the operation of IAF. Select these options by setting sense switches after the subsystem is initiated. Although normal operation does not require these sense switches to be set, sense switch 5 (ONSW,IAF,5) is set by default when the IAF Subsystem is initiated. Once set, the sense switches may be turned off again by entering OFFSW,IAF,s (s is the sense switch number.)</p> <table border="1"> <thead> <tr> <th>Entry</th> <th>Response</th> </tr> </thead> <tbody> <tr> <td>ONSW,IAF,1.</td> <td>When the IAF Subsystem is terminated (by an IDLE,IAF command), enter users into a detached state and inhibit restarting operations.</td> </tr> <tr> <td>ONSW,IAF,2.</td> <td>The IAF Subsystem will not reload after an abort. This option releases the output file to the output queue.</td> </tr> <tr> <td>ONSW,IAF,3.</td> <td>Aborts the IAF Subsystem on all abnormal conditions. This ensures that the subsystem does not continue to operate if an internal malfunction occurs. This may be used with sense switch 5 to enable an analyst to determine the problem. Some conditions cause the IAF Subsystem to abort even though sense switch 3 is not set.</td> </tr> <tr> <td>ONSW,IAF,4.</td> <td>Enables a dump on normal termination. This option should be activated whenever a dump is desired after the IDLE,IAF command.</td> </tr> <tr> <td>ONSW,IAF,5.</td> <td>Dumps information to the output file after the IAF Subsystem is dropped or aborted. This option provides a listing that may assist an analyst in determining the problem that existed when the subsystem dropped or aborted. This sense switch is set by default when the IAF Subsystem is activated.</td> </tr> <tr> <td>ONSW,IAF,6.</td> <td>Releases the output file containing dump information written after the IAF Subsystem is dropped or aborted (refer to the ONSW,IAF,5 option) immediately after the dump is taken. When sense switch 6 is not set, the output file is not released until the IAF Subsystem is terminated. It is then printed with the job dayfile. This switch applies only when sense switch 5 is set.</td> </tr> </tbody> </table>	Entry	Response	ONSW,IAF,1.	When the IAF Subsystem is terminated (by an IDLE,IAF command), enter users into a detached state and inhibit restarting operations.	ONSW,IAF,2.	The IAF Subsystem will not reload after an abort. This option releases the output file to the output queue.	ONSW,IAF,3.	Aborts the IAF Subsystem on all abnormal conditions. This ensures that the subsystem does not continue to operate if an internal malfunction occurs. This may be used with sense switch 5 to enable an analyst to determine the problem. Some conditions cause the IAF Subsystem to abort even though sense switch 3 is not set.	ONSW,IAF,4.	Enables a dump on normal termination. This option should be activated whenever a dump is desired after the IDLE,IAF command.	ONSW,IAF,5.	Dumps information to the output file after the IAF Subsystem is dropped or aborted. This option provides a listing that may assist an analyst in determining the problem that existed when the subsystem dropped or aborted. This sense switch is set by default when the IAF Subsystem is activated.	ONSW,IAF,6.	Releases the output file containing dump information written after the IAF Subsystem is dropped or aborted (refer to the ONSW,IAF,5 option) immediately after the dump is taken. When sense switch 6 is not set, the output file is not released until the IAF Subsystem is terminated. It is then printed with the job dayfile. This switch applies only when sense switch 5 is set.
Entry	Response														
ONSW,IAF,1.	When the IAF Subsystem is terminated (by an IDLE,IAF command), enter users into a detached state and inhibit restarting operations.														
ONSW,IAF,2.	The IAF Subsystem will not reload after an abort. This option releases the output file to the output queue.														
ONSW,IAF,3.	Aborts the IAF Subsystem on all abnormal conditions. This ensures that the subsystem does not continue to operate if an internal malfunction occurs. This may be used with sense switch 5 to enable an analyst to determine the problem. Some conditions cause the IAF Subsystem to abort even though sense switch 3 is not set.														
ONSW,IAF,4.	Enables a dump on normal termination. This option should be activated whenever a dump is desired after the IDLE,IAF command.														
ONSW,IAF,5.	Dumps information to the output file after the IAF Subsystem is dropped or aborted. This option provides a listing that may assist an analyst in determining the problem that existed when the subsystem dropped or aborted. This sense switch is set by default when the IAF Subsystem is activated.														
ONSW,IAF,6.	Releases the output file containing dump information written after the IAF Subsystem is dropped or aborted (refer to the ONSW,IAF,5 option) immediately after the dump is taken. When sense switch 6 is not set, the output file is not released until the IAF Subsystem is terminated. It is then printed with the job dayfile. This switch applies only when sense switch 5 is set.														

Command	Description
IDLE,sub.	Sets idledown status for subsystem sub.
	<p>NOTE</p> <hr/> <p>Do not use this command to idle the NAM subsystem. NAM is idled by using the procedure shown in NAM Shutdown (refer to section 15, Network Operations).</p> <hr/> <p>The subsystem terminates when idledown conditions are met. MAG terminates when no tapes are assigned. MSE terminates when no requests are outstanding and no MSE utilities are connected. BIO terminates when no active equipments remain. TAF terminates when no active transactions are outstanding. For all other subsystems, there are no idledown conditions; they terminate immediately.</p> <p>The system does not initiate new activity, such as assigning tapes and beginning print jobs, when idledown status is set. It is recommended that you use this command for terminating all subsystems.</p>
MAGfff. ¹²	<p>Calls the Magnetic Tape Subsystem (MAG) to the control point specified in the IPRDECK or by the SUBSYST L-display utility. MAG must be active in order for tape processing and resource allocation to take place and to make the resource requests available. Tape use can be monitored on the tape status (E,T) and resource requests (E,P) displays. MAGfff executes a procedure file named MAGfff under the system user index to initiate the MAG Subsystem.</p> <p>NOTE</p> <hr/> <p>You must not change unit numbers or power down tape units after MAG is initialized. Also, MAG should not be stopped (via the STOP command) while tapes are assigned. The IDLE command should be used instead.</p> <hr/>
MAPfff. ¹²	<p>Calls the Matrix Algorithm Processor (MAP) Subsystem to the control point specified in the IPRDECK or by the SUBSYST L-display utility. MAPfff calls a procedure file named MAPfff under the system user index to initiate the MAP Subsystem.</p>
MCSfff. ¹²	<p>Calls the Message Control System (MCS) to the control point specified in the IPRDECK or by the SUBSYST L-display utility. MCSfff calls a procedure file named MCSfff under the system user index to initiate the MCS Subsystem. NAM must be active at a control point to use MCS.</p>

¹². The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

Command	Description						
MSEfff.13	<p>Calls the Mass Storage Extended Subsystem (MSE) processing program to the control point specified in the IPRDECK or by the SUBSYST L-display utility. MSEfff calls a procedure file named MSEfff under the system user index to initiate the MSE.</p> <p>You can turn the trace mode on or off by using the following sense switch entries after the subsystem is initiated.</p> <table border="1"> <thead> <tr> <th>Entry</th> <th>Response</th> </tr> </thead> <tbody> <tr> <td>ONSW,MSE,1.</td> <td>Turn the trace mode on. The trace mode issues all MSE messages to the job dayfile.</td> </tr> <tr> <td>OFFSW,MSE,1.</td> <td>Turn the trace mode off. Default is trace mode off.</td> </tr> </tbody> </table>	Entry	Response	ONSW,MSE,1.	Turn the trace mode on. The trace mode issues all MSE messages to the job dayfile.	OFFSW,MSE,1.	Turn the trace mode off. Default is trace mode off.
Entry	Response						
ONSW,MSE,1.	Turn the trace mode on. The trace mode issues all MSE messages to the job dayfile.						
OFFSW,MSE,1.	Turn the trace mode off. Default is trace mode off.						
NAMfff.13	<p>Calls the network access method (NAM) to the control point specified in the IPRDECK or by the SUBSYST L-display utility. NAMfff calls a procedure file named NAMfff under the system user index to initiate several NAM supervisor modules. These modules appear momentarily at different control points as needed during NAM activity.</p> <p>The release versions are NAM and NAMNOGO. Use NAM if you want to initialize the network with no operator intervention. Use NAMNOGO if you want to alter the memory file. Refer to section 15, Network Operations, for more information.</p> <p>Additional commands that you can enter during NAM initiation are also described in section 15, Network Operations.</p>						
RDFfff.13	<p>Calls the Remote Diagnostic Facility (RDF) Subsystem to control point 1. RDFfff calls a procedure file named RDFfff or IAFfff (depending upon whether the site has the IAF enabled) under the system user index to initiate the RDF Subsystem.</p> <p>If your site uses IAF, ensure that NAM is active at a control point (refer to the B,O or B,A display) before you enter the RDFfff command.</p>						

13. The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

Command	Description
RHFfff. ¹⁴	Calls the Remote Host Facility (RHF) Subsystem to the control point specified in the IPRDECK or by the SUBSYST L-display utility. RHFfff calls a procedure file named RHFfff under the system user index to initiate the RHF Subsystem.
SMFfff. ¹⁴	Calls the Screen Management Facility (SMF) Subsystem to the control point specified in the IPRDECK or by the SUBSYST L-display utility. SMFfff calls a procedure file name SMFfff under the system user index to initiate the SMF Subsystem. IAF must be active to use SMF.
SSFfff. ¹⁴	Calls the SCOPE 2 Station Facility (SSF) Subsystem to the control point specified in the IPRDECK or by the SUBSYST L-display utility. SSFfff calls a procedure file named SSFfff under the system user index to initiate the SSF Subsystem.
STMfff. ¹⁴	Calls the stimulator to the last control point. STMfff calls a procedure file named STMfff under the system user index to initiate the stimulator.
STOP,sub.	Drops (terminates) the subsystem sub. Any acceptable three-letter subsystem name can be specified. Unlock the console to enter this command. This command can cause termination errors in the subsystem being dropped. The IDLE,sub command does the same thing without the termination errors.

¹⁴. The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

Command	Description										
TAFfff. ¹⁵	<p>Calls the Transaction Facility (TAF) Subsystem to the control point specified in the IPRDECK or by the SUBSYST L-display utility. TAF must be active for transaction processing to take place.</p> <p>TAFfff calls a procedure file named TAFfff under the system user index to initiate the TAF Subsystem. NAM must be active at a control point for interactive access to TAF. Batch jobs can access TAF without NAM.</p> <p>Response to an operator drop or fatal error is controlled by setting sense switches as follows:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Entry</th> <th style="text-align: left;">Response</th> </tr> </thead> <tbody> <tr> <td>ONSW,TAF,3.</td> <td>Attempt recovery after TAF is dropped or aborted. To successfully drop the transaction subsystem once this entry has been made, the command OFFSW,TAF,3 must be entered prior to the IDLE,TAF command.</td> </tr> <tr> <td>ONSW,TAF,4.</td> <td>Automatically restart TAF after a recovery.</td> </tr> <tr> <td>ONSW,TAF,5.</td> <td>Dump the entire field length and release the output file after the transaction subsystem is dropped or aborted.</td> </tr> <tr> <td>ONSW,TAF,6.</td> <td>Print the job dayfile upon termination.</td> </tr> </tbody> </table> <p>Activity at the transaction control point can be monitored on the O displays (refer to the NOS Version 2 Operations Handbook).</p>	Entry	Response	ONSW,TAF,3.	Attempt recovery after TAF is dropped or aborted. To successfully drop the transaction subsystem once this entry has been made, the command OFFSW,TAF,3 must be entered prior to the IDLE,TAF command.	ONSW,TAF,4.	Automatically restart TAF after a recovery.	ONSW,TAF,5.	Dump the entire field length and release the output file after the transaction subsystem is dropped or aborted.	ONSW,TAF,6.	Print the job dayfile upon termination.
Entry	Response										
ONSW,TAF,3.	Attempt recovery after TAF is dropped or aborted. To successfully drop the transaction subsystem once this entry has been made, the command OFFSW,TAF,3 must be entered prior to the IDLE,TAF command.										
ONSW,TAF,4.	Automatically restart TAF after a recovery.										
ONSW,TAF,5.	Dump the entire field length and release the output file after the transaction subsystem is dropped or aborted.										
ONSW,TAF,6.	Print the job dayfile upon termination.										

Refer to the AUTO and MAINTENANCE commands in this section and the SUBSYST L display in the NOS 2 Operations Handbook for additional information concerning subsystem control.

¹⁵. The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

System Control Commands

The following DSD commands control the operating system as well as the subsystems that run under the system. Several of these commands are typically used only for debugging purposes when the system is in an abnormal state. You may use others frequently to maintain system integrity in a normal production environment. Appropriate cautions are included with individual command descriptions in cases where you do not normally use the command or it is not recommended for use in a production environment.

Command	Description
AUTO.	<p>Calls enabled subsystems to control points, initiates automatic job processing, and spins up any 834, 836, 887, or 9853 disk drive that is ON and not unloaded. The IPRDECK used at deadstart time determines which subsystems will be activated by default. However, any of those subsystems not currently assigned to a control point can be disabled or others enabled through use of the SUBSYST L display (refer to the NOS Version 2 Operations Handbook). Individual subsystems can also be called to a control point or removed independent of the AUTO command by using the Subsystem Control Commands described earlier in this section. For additional information concerning the AUTO command, refer to Initiating Job Processing in section 2, Deadstart.</p>
CHECK POINT SYSTEM.	<p>Provides for termination of job processing and writes the contents of central memory tables to mass storage. The SCKP EQPDECK entry specifies which mass storage devices will contain a copy of the checkpoint file. This command is typically entered in preparation for recovery deadstart. The following sequence of operations takes place:</p> <ol style="list-style-type: none"> <li data-bbox="532 1108 1360 1234">1. Sense switch 1 is set for the IAF Subsystem, which causes all users to be placed in detached job status when an operator drop occurs. When all users are in detached job status, the IAF Subsystem is dropped and the checkpoint continues. <li data-bbox="532 1255 1360 1314">2. All job scheduling is inhibited. (This has the same effect as if the IDLE command was entered.) <li data-bbox="532 1335 1360 1436">3. Jobs with a zero subsystem identifier are rolled out. All of these jobs are recovered on a level 1 or level 2 recovery deadstart. <li data-bbox="532 1457 1360 1516">4. The system moves the system dayfile buffers maintained in CMR to disk. <li data-bbox="532 1537 1360 1598">5. All subsystems except the magnetic tape executive (MAG) are aborted.

Command	Description
CHECK POINT SYSTEM.	(Continued)
	<p>6. MAG is rolled out when no other jobs are active. The rolling out of MAG allows recovery of all tape files associated with jobs rolled out if the tapes are not repositioned prior to the level 1 or level 2 recovery deadstart. Because of this, if a tape unit is used for deadstart, it should be unassigned prior to the checkpoint command.</p> <p>7. The system is left in an idle state. If SPINDOWN is enabled, all 834, 836, 887, or 9853 disk drives that are ON and not unloaded are spun down. Normal processing may be continued with an AUTO command. If this is done, no attempt should be made to later perform a level 1 or level 2 recovery deadstart unless another checkpoint command is performed.</p>

During the processing of the checkpoint, the message

PROCESSING SYSTEM CHECKPOINT

is issued at the system control point. It is also possible that under certain circumstances the checkpoint routine will not be able to properly abort a job (such as one that has NOEXIT selected and has a nonzero subsystem identifier). If such a job continues processing after it has been aborted by the system, the operator must then abort the job for the checkpoint to continue. A more desirable approach to this situation is to ensure that jobs such as permanent file dumps are finished prior to the checkpoint. For additional information concerning the CHECK POINT SYSTEM command, refer to Preparing for Recovery Deadstart in section 2, Deadstart.

DATE.yy/mm/dd. Changes the current system date. Unlock the console before entering this command (refer to the UNLOCK command later in this section).

Parameter	Description
yy	Year; 70 through 99 (corresponds to the years 1970 through 1999) or 00 through 33 (corresponds to the years 2000 through 2033).
mm	Month; 01 through 12.
dd	Day; 01 through 31.

Command	Description
DEBUG.	<p>Reverses the current status of the debug mode. When the debug mode is set, the message DEBUG appears in the header of the left screen display. The debug mode provides system origin privileges to validated users and allows modifications to be made to the running system. Unlock the console before entering this command (refer to the UNLOCK command later in this section). Use of the debug mode is not commonly allowed in a normal production environment.</p> <hr/> <p>NOTE</p> <hr/> <p>On a secured system, the console must be in security unlock status to accept this command (refer to the UNLOCK,username,password command later in this section).</p> <hr/>
DIS,jsn.	<p>Calls the job display package (DIS) to the job with job sequence name jsn. The A and B displays for DIS automatically appear on the left and right console screen, respectively. Unlock the console before entering this command (refer to the UNLOCK command later in this section). Refer to section 4, DIS Operations, for complete information concerning the DIS package.</p> <hr/> <p>NOTE</p> <hr/> <p>On a secured system, the console must be in security unlock status to accept this command (refer to the UNLOCK,username,password command later in this section).</p> <hr/>

Command	Description										
DISABLE,op. or ENABLE,op.	Disables or enables option op. If you enter the ENABLE command and op is currently enabled, the system ignores the command. The system also ignores the DISABLE command if you enter it and op is already disabled. op is one of the following options:										
	<table border="1"> <thead> <tr> <th>op</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ACS TAPE PF STAGING</td> <td>Enables or disables the staging of Automated Cartridge Subsystem (ACS) tape alternate storage resident files to disk. If disabled, jobs attempting to access permanent files that reside only on ACS tapes are aborted. This entry controls staging of files from ACS tapes, whereas the TAPE PF STAGING entry controls staging of files from 9-track tapes and CTS tapes.</td> </tr> <tr> <td>AUTORESTART</td> <td>Enables or disables the AUTORESTART installation parameter. If AUTORESTART is enabled and if step mode is set during an environmental shutdown, the system automatically clears step mode when the power environment returns to normal. All jobs waiting in the input queue are initiated by performing the processing associated with the AUTO command described earlier in this section. This recovery process requires no operator action.</td> </tr> <tr> <td>CARTRIDGE PF STAGING</td> <td>Enables or disables the staging of MSE or optical disk permanent files to disk. Disabling CARTRIDGE PF STAGING aborts jobs that attempt to access files residing only on either MSE or optical disk. Enabling CARTRIDGE PF STAGING allows both MSE and optical disk files to be staged to disk when accessed. Contrast this entry with the ACS TAPE PF STAGING entry, which enables or disables staging from ACS tapes, and the TAPE PF STAGING entry, which enables or disables staging from 9-track tapes and CTS tapes.</td> </tr> <tr> <td>DDP ROLLOUT PATH</td> <td>Enables or disables the use of the DDP for job rollout and rollin. If DDP or EM equipment is not present, this entry has no effect.</td> </tr> </tbody> </table>	op	Description	ACS TAPE PF STAGING	Enables or disables the staging of Automated Cartridge Subsystem (ACS) tape alternate storage resident files to disk. If disabled, jobs attempting to access permanent files that reside only on ACS tapes are aborted. This entry controls staging of files from ACS tapes, whereas the TAPE PF STAGING entry controls staging of files from 9-track tapes and CTS tapes.	AUTORESTART	Enables or disables the AUTORESTART installation parameter. If AUTORESTART is enabled and if step mode is set during an environmental shutdown, the system automatically clears step mode when the power environment returns to normal. All jobs waiting in the input queue are initiated by performing the processing associated with the AUTO command described earlier in this section. This recovery process requires no operator action.	CARTRIDGE PF STAGING	Enables or disables the staging of MSE or optical disk permanent files to disk. Disabling CARTRIDGE PF STAGING aborts jobs that attempt to access files residing only on either MSE or optical disk. Enabling CARTRIDGE PF STAGING allows both MSE and optical disk files to be staged to disk when accessed. Contrast this entry with the ACS TAPE PF STAGING entry, which enables or disables staging from ACS tapes, and the TAPE PF STAGING entry, which enables or disables staging from 9-track tapes and CTS tapes.	DDP ROLLOUT PATH	Enables or disables the use of the DDP for job rollout and rollin. If DDP or EM equipment is not present, this entry has no effect.
op	Description										
ACS TAPE PF STAGING	Enables or disables the staging of Automated Cartridge Subsystem (ACS) tape alternate storage resident files to disk. If disabled, jobs attempting to access permanent files that reside only on ACS tapes are aborted. This entry controls staging of files from ACS tapes, whereas the TAPE PF STAGING entry controls staging of files from 9-track tapes and CTS tapes.										
AUTORESTART	Enables or disables the AUTORESTART installation parameter. If AUTORESTART is enabled and if step mode is set during an environmental shutdown, the system automatically clears step mode when the power environment returns to normal. All jobs waiting in the input queue are initiated by performing the processing associated with the AUTO command described earlier in this section. This recovery process requires no operator action.										
CARTRIDGE PF STAGING	Enables or disables the staging of MSE or optical disk permanent files to disk. Disabling CARTRIDGE PF STAGING aborts jobs that attempt to access files residing only on either MSE or optical disk. Enabling CARTRIDGE PF STAGING allows both MSE and optical disk files to be staged to disk when accessed. Contrast this entry with the ACS TAPE PF STAGING entry, which enables or disables staging from ACS tapes, and the TAPE PF STAGING entry, which enables or disables staging from 9-track tapes and CTS tapes.										
DDP ROLLOUT PATH	Enables or disables the use of the DDP for job rollout and rollin. If DDP or EM equipment is not present, this entry has no effect.										

(Continued on next page)

Command	Description
<i>(Continued from previous page)</i>	
DISABLE,op. or ENABLE,op.	<i>(Continued)</i>
	op
	Description
	ENGR¹⁶
	Enables or disables engineering mode. When engineering mode is set, the message ENGR appears in the header of the left screen display. Engineering mode allows PP/hardware diagnostics and FORMAT/FDP to be executed if the user has system origin privileges. Unlock the console before entering this command (refer to the UNLOCK command later in this section).
	FLEXIBLE PARTITIONS
	Enables or disables flexible memory partitioning. The job scheduler attempts to use memory space to the greatest extend possible when flexible partitions are enabled. Some service classes may be allotted more total memory space than memory partitioning constraints normally allow. If flexible partitions are disabled, the total memory used by jobs of a given service class is never allowed to exceed the maximum specified on the SERVICE command (AM and EM parameters).
	FORCED FORMATTING
	Enables or disables forced formatting of 583x devices. When FORCED FORMATTING is enabled, the driver reformats the devices when either you specify the INITIALIZE,FP option for the device, or the system detects that a previously offline parity drive is now online. The driver does not determine whether the device is already formatted correctly.
	When FORCED FORMATTING is disabled, the driver formats the device only when necessary.
	LOGGING
	Enables or disables logging of dayfile messages concerned with program efficiency.

(Continued on next page)

16. On a secured system, the console must be in security unlock status to accept this command (refer to the UNLOCK.username,password command later in this section).

Command	Description
<i>(Continued from previous page)</i>	
DISABLE,op. or ENABLE,op.	<i>(Continued)</i>
	op
	Description
	MASTER MSE
	Enables or disables master mainframe mode for MSE processing. When MSE is brought to a control point, the SSEXEC program runs if master mainframe mode is enabled. The SSSLV program runs if master mainframe mode is disabled. This entry has no effect unless MSE processing is activated.
	MS VALIDATION
	Enables or disables automatic verification of mass storage tables. This command cannot be used unless the MS VALIDATION option has been selected in the IPRDECK used at deadstart. The validation that occurs for each level of recovery deadstart is described under Preparing for Recovery Deadstart in section 2, Deadstart.
	PF VALIDATION
	Enables or disables verification of BOI/EOI on preserved files. Selection of this option causes BOI/EOI information to be verified when a direct access file is attached or purged, and when an existing local file is defined as a direct access permanent file.
	PRIVILEGED ANALYST MODE
	Enables or disables PRIVILEGED ANALYST MODE operations. When enabled, a user validated with AW-CPAM (refer to the NOS Version 2 Administration Handbook for information on MODVAL validation) is permitted to read system status information (such as the system dayfile, account file, and error log) using a nonsystem-origin job. Unlock the console (refer to the UNLOCK command later in this section) to enable or disable this option. This command is not accepted on a secured system.

(Continued on next page)

Command	Description
<i>(Continued from previous page)</i>	
DISABLE,op. or ENABLE,op.	<i>(Continued)</i>
	op
	Description
	PRIVILEGED RDF
	Enables or disables privileged mode of the Remote Diagnostic Facility (RDF). When enabled, a user's commands are checked to ensure that a maintenance function is being performed.
	REMOVABLE PACKS
	Enables or disables automatic label checking for mass storage devices defined as removable. Examine the disk status display (E,M) (refer to the NOS Version 2 Operations Handbook) to determine which mass storage devices (if any) are defined as removable. If this option is disabled, all removable devices subsequently introduced into the system cannot be accessed. This option must be enabled to perform label verification before those devices can be accessed.
	RESIDENT RDF
	Enables or disables resident mode of RDF. While in resident mode, RDF remains active, regardless of terminal activity, until you disable RDF. When resident mode is disabled (the default condition), RDF becomes inactive if no one logs into the remote diagnostic terminal for 15 minutes. To allow terminal activity to resume, you must reactivate RDF with a call to RDFffff.
	SECONDARY USER COMMANDS
	Enables or disables the acceptance of more than one USER command in a job. If disabled, any USER command encountered after the first one causes the job to be aborted. Unlock the console (refer to the UNLOCK command later in this section) to enable or disable this option. This command is not accepted on a secured system.
	SPINDOWN
	Enables or disables the spinning down of 834, 836, 887, and 9853 disk storage devices during a system checkpoint. System checkpoints that are initiated by mainframe errors do not affect the state of these devices.

(Continued on next page)

Command	Description
<i>(Continued from previous page)</i>	
DISABLE,op. or ENABLE,op.	<i>(Continued)</i>
	op
	SYSTEM DEBUG
	Description
	<p>Enables or disables the system debug mode of operation. Unlock the console (refer to the UNLOCK command later in this section) to enable or disable this option.</p> <p>When the system is in system debug mode, it is less tolerant of system errors; that is, it is more likely to hang upon experiencing errors. When the system is not in system debug mode, it rates system errors as critical or noncritical.</p> <p>For critical errors, the system partially or totally interrupts system operation to tend to the errors. For noncritical errors, the system logs them in the binary maintenance log (BML) and inasmuch as possible allows system operation to proceed. You can initiate the system debug mode with the DSD ENABLE command or the corresponding IPRDECK entry.</p> <p>There is another system state called debug mode which is conceptually different from that of system debug mode. Debug mode is the state of the system where a user with system origin privileges can make modifications to the running system. You can initiate this mode of operation with the DSD command DEBUG. The left screen header of the system console indicates whether the system is in debug mode or not.</p>
<i>(Continued on next page)</i>	

Command	Description
<i>(Continued from previous page)</i>	
DISABLE,op. or ENABLE,op.	<i>(Continued)</i>
	op
	Description
	TAPE PF STAGING Enables or disables the staging of tape alternate storage resident files to disk. The media for tape alternate storage is 9-track tapes and CTS tapes. Disabling TAPE PF STAGING aborts jobs that attempt to access files residing only on tape alternate storage. This entry controls staging of files from 9-track tapes and CTS tapes. Contrast this entry with the ACS TAPE PF STAGING entry, which controls staging of files from ACS tapes, and the CARTRIDGE PF STAGING entry, which controls staging of files from MSE and optical disk.
	USER EXTENDED MEMORY Enables or disables use of the user accessible area of extended memory. If disabled, no job can access the user area of extended memory and, other than subsystems, all jobs currently accessing the area are rolled out.
IDLE.	Prevents any new jobs from being scheduled to a control point but does not terminate the jobs currently assigned. If a job is rolled out while this command is in effect, it will not be scheduled back to a control point until the AUTO, MAINTENANCE, or SCHEDULE command is entered.
K.messagetext.	Allows entry of data messagetext in the user- or system-defined CPU buffer for control when the K display is active. Refer to section 8, K-Display Utilities, for information about the K display.
L.messagetext.	Allows entry of data messagetext in the user- or system-defined CPU buffer for control when the L display is active. Refer to section 9, L-Display Utilities, for information about the L display.
LOCK.	Locks the console keyboard. This command prevents entry of restricted commands (refer to the UNLOCK command later in this section for a list of restricted commands). All other DSD commands can be entered when the console is locked. The console is normally locked when the system is being used in a production environment.

(Continued on next page)

Command	Description
<i>(Continued from previous page)</i>	
MAINTENANCE.	This command performs the same functions as the AUTO command but additionally starts several maintenance routines. Refer to Initiating Job Processing in section 2, Deadstart, for complete information concerning this command.
SCHEDULE.	Initiates automatic job processing (reversing the effect of a previous IDLE command), but does not initiate subsystems as the AUTO command does.
STEP.	<p>Sets the monitor in step mode. Setting the monitor in step mode stops all central memory I/O operations and prevents the system from processing PP requests when the next monitor function is encountered. Pressing the space bar releases the present step and stops again for each subsequent monitor function.¹⁷ If a DSD command is entered while the system is in step mode and all PPs are assigned, it is possible the request cannot be satisfied. In that case, the screen(s) blinks until you correct the condition by clearing the entry.</p> <p>When step mode is set, the message STEP appears in the header of the left screen display. Unlock the console before entering this command (refer to the UNLOCK command later in this section).</p> <p>This command is generally used for debugging purposes and should not be used in a normal production environment. In addition, the system may set step mode automatically upon detection of a main power failure or abnormal environmental condition.</p>

(Continued on next page)

¹⁷. Monitor functions are described in the NOS Version 2 Systems Programmer's Instant.

Command	Description
---------	-------------

(Continued from previous page)

STEP,jsn,ff,b,v. or STEP,,ff,b,v.	Sets step mode for the job with job sequence name jsn (first form of the command) or all jobs (second form of the command) at the next occurrence of monitor function ff when output register byte b has value v.
---	---

If jsn is specified (first form of the command), the monitor function ff, the output register byte number b, and the output register value v can be optionally specified for more precise control of when the step mode is executed. The values for b can be from 0 to 4; the values for v can be from 0 to 7777₈.

If jsn is not specified (second form of the command), the monitor function ff must be specified. The output register byte number b and value v can be optionally specified for more precise control of when the step mode is executed. The values for b can be from 0 to 4; the values for v can be from 0 to 7777₈.

For either form of the command, pressing the spacebar releases the current step and steps the job or system at the next occurrence of the monitor function ff, if specified, or the next monitor function from the job if the first form is used and no monitor function is specified.

Using the second form of the command may stop all central memory I/O operations and prevent the system from processing PP requests. If a DSD command is entered while the system is in step mode and all PPs are assigned, it is possible the request cannot be satisfied. In this case, the screens blink until you clear the entry. The message STEP followed by the monitor function number ff (if specified) appears in the header of the left screen displays while this command is in effect. The console must be unlocked before entry of this command is permitted (refer to the UNLOCK command later in this section).

These commands are generally used only for debugging purposes. Do not enter these commands if the system has automatically set step mode because of a power failure or abnormal environmental condition.

TIME.hh.mm.ss.	Changes the current system time. Unlock the console before entering this command (refer to the UNLOCK command later in this section).
----------------	---

Parameter	Description
hh	Hour; 00 through 23.
mm	Minute; 00 through 59.
ss	Second; 00 through 59

(Continued on next page)

Command	Description
---------	-------------

(Continued from previous page)

UNLOCK.	<p>Unlocks the console keyboard. When this command is active, the message UNLOCK appears in the header of the left screen display. Although all DSD commands can be entered when the console is unlocked, the following commands are restricted to entry only when the console is unlocked.</p> <p>DATE.yy/mm/dd.¹⁸</p> <p>DEBUG.</p> <p>DISABLE,PRIVILEGED ANALYST MODE.¹⁹</p> <p>DISABLE,SECONDARY USER COMMANDS.¹⁹</p> <p>DISABLE,SYSTEM DEBUG.</p> <p>DROP,,qt,ujn.</p> <p>ENABLE,ENGR.¹⁸</p> <p>ENABLE,PRIVILEGED ANALYST MODE.¹⁹</p> <p>ENABLE,SECONDARY USER COMMANDS.¹⁹</p> <p>ENABLE,SYSTEM DEBUG.</p> <p>OVERRIDE,jsn.</p> <p>SPINDOWN,est.</p> <p>STEP.</p> <p>STEP,jsn,ff,b,v.</p> <p>STOP,sub.</p> <p>TIME.hh.mm.ss.¹⁸</p> <p>UNLOAD,est. (est specifies a nonremovable shared mass storage device)</p> <p>All memory entry commands¹⁸</p> <p>All channel control commands</p> <p>All extended memory flag register commands</p> <p>All breakpoint package commands¹⁸</p>
---------	---

(Continued on next page)

¹⁸. Refer to Secured System Control Commands later in this section for information on unlocking the console keyboard on a secured system.

¹⁹. Not allowed on secured systems.

Command	Description				
<i>(Continued from previous page)</i>					
UNLOCK.	<i>(Continued)</i> Always lock the console when the system is being used in a production environment. However, you can unlock the console to enter the STEP command if a system failure requiring a recovery deadstart occurs (refer to Preparing for Recovery Deadstart in section 2, Deadstart). Refer to Secured System Control Commands later in this section for information on unlocking the console keyboard on a secured system.				
UNSTEP.	Clears step mode. This command clears the effect of any format of STEP command. If the system has set step mode because of a main power failure or abnormal environmental condition, do not enter this command until the conditions that caused the automatic setting have returned to normal.				
X.MDD(p).	Initiates the monitor display driver (MDD). MDD is a PP program and is independent of the operating system. Refer to the CYBER Initialization Package User's Handbook for more information.				
<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>p</td> <td>Port number parameter. If you enter 1, the first port is connected. If you enter 2, the second port is connected. If you do not specify this parameter, the second port is connected by default.</td> </tr> </tbody> </table>		Parameter	Description	p	Port number parameter. If you enter 1, the first port is connected. If you enter 2, the second port is connected. If you do not specify this parameter, the second port is connected by default.
Parameter	Description				
p	Port number parameter. If you enter 1, the first port is connected. If you enter 2, the second port is connected. If you do not specify this parameter, the second port is connected by default.				
X.name. or X.name(parameters). or X.name,fl.	Calls a system program or utility specified by name to an available control point. If parameters are to be passed to the program, the second form of the command is used. In both the first and second form of the command, the field length specified in the library for the command is used. If no field length is specified in the library, a value of 60000 ₈ is assumed. If a field length different from the default is required, the third form of the command is used. The field length is specified by fl. Only the first 58 characters following X. are used.				

(Continued on next page)

Command	Description
----------------	--------------------

(Continued from previous page)

99.	Disables or enables syntax overlay processing and logging of DSD commands in the system dayfile and error log. That is, depending upon current status, syntax overlay processing and logging are reversed. When this mode is enabled, 99 appears on the left screen header.
-----	---

When syntax overlay processing and logging of DSD commands is disabled, DSD does not allow you to enter any command requiring the loading of an overlay from mass storage. Disable overlay loading only when the system is in an abnormal state to prevent PPs from being requested when they cannot perform the necessary tasks (for example, when a system disk channel is hung). A 99 command that enables logging will be logged itself, but a 99 command which disables logging will not be logged.

This command is normally used only for debugging purposes. When the system is in an abnormal state, the commands that require entry of the 99 command depend on which syntax overlays for DSD are in central memory.

Secured System Control Commands

The following commands are allowed only on a secured system.

Command	Description
OQSH=level.	<p>Specifies the output queue special handling (OQSH) level on a secured system. The (OQSH) level is set initially during deadstart by the OQSH IPRDECK entry. The current OQSH level appears on the DSD S display. The OQSH command can be entered at any time from the system console to change the current level. Output files with an access level greater than or equal to the OQSH level specified in this command are not printed but remain in the queue until released by the operator (refer to the RELEASE command, next). If level is set to the lowest access level or no level is specified, no files are held in the queue.</p>
RELEASE,jsn.	<p>Allows the operator to release a file from the output queue whose access level is equal to or above the OQSH level on a secured system. Output queue files and their access levels can be examined using the DSD Q display. The output file with job sequence name jsn is released from the output queue and is processed by the Batch Input/Output Subsystem. The RELEASE command can be entered at any time from the system console. Other restrictions based on device access levels and file access levels set by your site continue to apply.</p>
SECURES,ot, LA=lowerlevel, UA=upperlevel.	<p>Sets the security access level limits for the system (ot=SY) or the security access level limits for a particular origin type (ot=IA, BC, or RB). The system limits determine the range of access levels allowed in the system; jobs may not execute at an access level outside this range, and files may not be created or accessed at an access level outside this range. Origin type limits determine the range as it applies to jobs of one particular origin type. By default, all origin type limits are the same as the system limits.</p> <p>The options available for using this command to change the system limits (that is, whether they can be raised, lowered, or both) are controlled by the OPSECM CMRDECK entry (refer to section 3, Deadstart Decks). The origin type limits for ot=IA, BC, and RB can be changed to any values within the system limits. When the system limits are changed (ot=SY), all origin type limits are reset to the new system limits.</p> <p>The parameters LA and UA specify the lower and upper access level limits. Both parameters must be specified. The value corresponding to the lower access level must be less than or equal to the value corresponding to the upper access level.</p>

Command	Description
UNLOCK,username, password.	<p>Unlocks the console keyboard for special commands when the system is running in secured mode. When this command is entered, the message SECURITY UNLOCK appears in the header on the left screen display. If the system is running in secured mode, the following commands are restricted to entry only when the console is in security unlock status.</p> <p>DATE.yy/mm/dd.</p> <p>X.DDF.</p> <p>DEBUG.</p> <p>DIS,jsn.</p> <p>ENABLE,ENGR.</p> <p>QDSPLAY,jsn.</p> <p>SECURES,ot,LA=lowerlevel,UA=upperlevel.</p> <p>SECUREQ,est,LA=lowerlevel,UA=upperlevel.</p> <p>TIME.hh.mm.ss.</p> <p>All memory entry commands.</p> <p>All PP breakpoint commands.</p> <p>All CPUMTR breakpoint commands.</p> <p>In addition, the memory displays are partially disabled on a secured system unless the console is in security unlock status. If you attempt to display anything outside of the system tables, a memory display shows the following message instead of the contents of the memory locations.</p> <p>****SECURED AREA****</p>

Memory Entry Commands

The following commands are used to change the contents of central memory and extended memory. Either absolute locations or those relative to a specific job's reference address (RA) can be changed.

CAUTION

Use these commands with extreme caution to avoid damage to the system or to user jobs.

Before entering these commands you must unlock the console keyboard (refer to the UNLOCK command earlier in this section). On a secured system, the console must be in security unlock status. To change central memory, you must bring a central memory display (C, D, F, or G) to the left screen. To change extended memory, you must bring the M display to the left screen. It is this display that controls whether the memory entry commands change absolute memory locations or relative memory locations.

For example, if you enter

C, .

to bring up the C display, any memory entry commands entered make changes to absolute memory locations. If you enter

C, jsn.

where jsn is a valid job sequence name, any memory entry commands entered make changes to memory locations relative to the job's RA.

Character values or numeric data can replace the current word contents. Either one 12-bit byte or a 60-bit word can be changed. A single byte can be changed by inserting the byte number after the location; bytes are numbered 0 through 4 from left to right. The address and contents are assembled right-justified with zeros filling unused leading positions. Leading zeros may be omitted in the entry.

When you are changing the contents of memory relative to a specific job (a valid job sequence name was specified when the memory display was brought to the left screen), the negative field length area of the job can be accessed. This area is accessed by specifying a negative address in the memory entry command. For example, to change the contents of word RA-3, you enter the address 77777775₈.

On an unsecured system when you are changing a memory location relative to a specific job's RA, the system checks for field length violations. If you specify an address larger than your job's RA plus field length, you cannot change its contents.

On a secured system no memory entry commands are allowed unless the security unlock status is set. The memory display shows the message

****SECURED AREA****

instead of the contents of the memory locations to prevent you from examining these locations; you may not alter the contents of these locations. When you are displaying absolute memory you can see only the system tables.

Formats and descriptions for these commands are:

Command	Description
addr,cont. or addr+cont. ²⁰	Changes the contents of central memory location addr (maximum of eight digits) to cont (maximum of 20 digits). The second form of the command performs essentially the same function but is used when it is necessary to change successive memory locations.
addr,b,cont. or addr+b,cont. ²⁰	Changes the contents of byte b at central memory location addr (maximum of eight digits) to cont. Each location consists of five 12-bit bytes, numbered 0 through 4 from left to right. The contents are four octal characters. The second form of the command performs essentially the same function but is used when it is necessary to change successive memory locations.
addr,Dcont. or addr+Dcont. ²⁰	Changes the contents of central memory location addr (maximum of eight digits) to display code characters cont (left-justified and zero-filled). The second form of the command performs essentially the same function but is used when it is necessary to change successive memory locations.
Eaddr,cont. or Eaddr+cont. ²⁰	Changes the contents of extended memory location addr (maximum of seven digits) to cont (maximum of 20 digits). The second form of the command performs essentially the same function but is used when it is necessary to change successive extended memory locations.
Eaddr,b,cont. or Eaddr+b,cont. ²⁰	Changes the contents of byte b at extended memory location addr (maximum of seven digits) to cont. Each location consists of five 12-bit bytes, numbered 0 through 4 from left to right. The contents are four octal characters. The second form of the command performs essentially the same function but is used when it is necessary to change successive extended memory locations.
Eaddr,Dcont. or Eaddr+Dcont. ²⁰	Changes the contents of extended memory location addr (maximum of seven digits) to display code characters cont (left-justified and zero-filled). The second form of the command performs essentially the same function but is used when it is necessary to change successive extended memory locations.

²⁰. If the + sign is specified, the address is incremented by one location (addr+1) after the initial entry is processed; the - sign causes the address to be decremented by one location (addr-1). This allows immediate entry for the next (or previous) memory location. If the message REPEAT ENTRY is displayed above the entry, the cont field is not cleared and can be entered in successive memory locations as many times as desired.

Channel Control Commands

The following commands enable control activity on a specified data channel in circumstances where abnormal hardware and/or system operation is detected. Extreme caution must be exercised if any of these commands are entered during normal system operation. In addition, the console keyboard must be unlocked before entry of any of these commands is permitted (refer to the UNLOCK command earlier in this section). DSD does not reserve the channel specified in any of the channel control commands. The channels are numbered 0 to 13₈ in a 10 PP system and 0 to 13₈, 20₈ to 33₈ in a 20 PP system.

CAUTION

Extreme caution must be exercised when using the following channel control commands.

Command	Description
ACN,cc.	Activates channel cc. This command alerts and prepares peripheral equipment on channel cc for the exchange of data.
DCH,cc.	Drops channel cc. This is a software function to release the current reservation of channel cc.
DCN,cc.	Deactivates channel cc. As a result, peripheral equipment on channel cc stops and any current I/O operation is terminated.
FCN,cc,func.	Outputs a function code func to channel cc. This releases all equipment selections on that channel. If func is not specified, a zero function code (no activity) is output.
IAN,cc.	Inputs data to the pseudo A register from channel cc.
LDC,nnnn.	Loads the pseudo A register with nnnn (normally a peripheral equipment function code). The current value of nnnn is shown in the channel status display (W,C).
MCH,cc.	Master clears and removes all 3000 Computer Systems peripheral equipment selections on channel cc (6681 function code 1700 ₈ is issued).
OAN,cc.	Outputs the contents of the pseudo A register to channel cc.

Extended Memory Flag Register Commands

The extended memory flag register commands are used to set or clear specific bits in the extended memory flag register. These commands are typically used when debugging a problem in a multiframe environment. The console must be unlocked (refer to the UNLOCK command earlier in this section).

CAUTION

Extreme caution must be exercised when using these commands.

Command	Description
CFR,bb.	Clears the extended memory flag register bit bb ($0 \leq bb \leq 17$).
SFR,bb.	Sets the extended memory flag register bit bb ($0 \leq bb \leq 17$).

Breakpoint Package Commands

The breakpoint package consists of a set of DSD commands that controls PP breakpoint processing and CPUMTR breakpoint processing.

PP Breakpoint Commands

The following DSD commands control PP breakpoint processing. These commands and the DSD V display enable you to communicate with the PP breakpoint package. The V display is described following the command descriptions. The PP breakpoint commands require that the system console be unlocked (refer to the UNLOCK command earlier in this section); on a secured system, the console must be in security unlock status.

CAUTION

Use these commands with extreme caution to avoid damage to the system or to user jobs. Incorrect usage may cause a system to hang up and require a deadstart.

NOTE

Breakpoint can be set for only one PP program at a time. Before setting breakpoint for another program, it must be cleared on the first program.

Command	Description
BKP.	Clears an existing PP breakpoint.
BKP,prog,cp.	<p>Sets breakpoint for PP program prog at control point cp. The control point parameter cp is optional. The PP program prog can be any PP program or overlay except the following:</p> <ul style="list-style-type: none"> • PPR overlays (1DD, 0SE, etc.). • Mass storage drivers. • Mass storage error processors. • Any PP program that overwrites PP memory greater than 7500_g with instructions or data. The breakpoint program uses PP memory locations 7500_g to 7777_g. If the PP program is allowed to modify those locations on load, execution, or overlay load, unpredictable results may occur. Refer to the PP breakpoint precautions in the following section. • Concurrent (16-bit) PP programs. <p>The PP program need not be a main program, and it can reside in central memory, disk, or extended memory (ECS/ESM/UEM).</p>
PPnn.A,val.	Sets the A register to the value val for PP number nn. Up to 18 bits can be specified for the value parameter val.
PPnn.C,addr.	Causes the V display for PP number nn to display 100 _g bytes of PP memory, starting at base address addr, as the C block.
PPnn.C,addr,val.	Changes the contents of PP memory word addr for PP number nn to value val. The specified address must point to one of the PP memory words displayed in the C block of the V display.
PPnn.D,addr.	Causes the V display for PP number nn to display 100 _g bytes of PP memory, starting at base address addr, as the D block.
PPnn.D,addr,val.	Changes the contents of PP memory word addr for PP number nn to value val. The specified address must point to one of the PP memory words displayed in the D block of the V display.
PPnn.EXR.	Resumes execution in PP number nn with a breakpoint set in the PP-resident subroutine EXR (after the next overlay load completes, but before execution of the overlay begins).
PPnn.FTN.	Resumes execution in PP number nn with a breakpoint set in the PP-resident subroutine FTN+1 (to trap the next monitor function before it is issued to MTR/CPUMTR).

Command	Description
PPnn.G.	Resumes normal PP program execution for PP number nn with no breakpoint set.
PPnn.G,addr.	Resumes PP program execution for PP number nn with a breakpoint set at address addr. The address addr must be within executable code and lie on a PP instruction boundary, but the validity of the address is not checked.
PPnn.P,addr.	Resets the P register to address addr for PP number nn. When the PP program resumes execution, it will start at P=addr.
PPnn.S.	Steps the PP program one instruction for PP number nn. Indexed jump instructions, channel jump instructions, and instructions that jump to the current program address or current program address minus one cannot be stepped. If you try to step these instructions, the following message appears on the system console: CAN-T STEP.
PPnn.SX.	Steps the PP program one instruction for PP number nn. This command performs the same function as the PPnn.S. command except that if an RJM instruction is encountered, a breakpoint is set to the instruction after the RJM and the entire subroutine is executed.
V,nn.	Brings up the DSD V display for PP number nn.
V,HNG.	Brings up the DSD V display for a hung PP.
V,PS.	Brings up the DSD V display for the pseudo PP.

V Display

The DSD V display complements the PP breakpoint commands to provide an analyst's tool for PP breakpoint processing. The information displayed is for the PP number specified in the V,nn command. Figure 5-4 shows the format of the V display. The display shows the PP communications area for the specified PP (10₈ CM words), the contents of the P and A registers, and two 100₈ word blocks of PP memory (C and D). The base addresses of the C and D blocks are set by PPnn.C,addr and PPnn.D,addr commands.

If the specified PP is not at breakpoint when the V display is brought up, the display shows only the communications area and the following message:

NOT AT BREAK.

Breakpoint Package Commands

```

V,nn.  PP BREAKPOINT.
7420   0310 1702 0000 0000 1234
       0000 0001 0000 0000 0000
       xxxx xxxx xxxx xxxx xxxx
       xxxx xxxx xxxx xxxx xxxx      P = 0460
       xxxx xxxx xxxx xxxx xxxx
       xxxx xxxx xxxx xxxx xxxx      A = 220346
       xxxx xxxx xxxx xxxx xxxx
       xxxx xxxx xxxx xxxx xxxx

C  (1100) 0200 2754 2000 3465 5400 3472 1006 3401
   (1110) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (1120) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (1130) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (1140) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (1150) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (1160) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (1170) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx

D  (0000) 0000 0002 0056 4567 0000 0003 4157 0234
   (0010) 0000 0001 0000 0000 0000 0000 0000 0000
   (0020) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (0030) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (0040) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (0050) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (0060) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx
   (0070) xxxx xxxx xxxx xxxx xxxx xxxx xxxx xxxx

```

Figure 5-4. PP Breakpoint V Display

PP Breakpoint Precautions

The following precautions and information are important to the successful use of the PP breakpoint package.

- The breakpoint must be set in executable code and must lie on a PP instruction boundary.
- When a breakpoint is set, two words of PP memory are replaced with the following:

```
RJM BKP
```

When the breakpoint is hit, the two replaced words are immediately restored. However, 16-bit instructions are not restored properly since they revert to 12-bit instructions.

- The PP breakpoint package uses PP locations 7500₈ through 7777₈. If these locations are altered by the running PP, the system will probably fail.
- If the PP program being breakpointed encounters a mass storage error, the mass storage driver's error processor will overlay the PP breakpoint code, and will probably crash the operating system.
- If the PP program being breakpointed performs code modification on the instruction at the address at which the breakpoint has been set, unpredictable results may occur.
- The PP breakpoint package cannot be used on DIS or O26. However, it can be used on other PPs at the same control point as DIS or O26.
- The PP breakpoint information (BKP,prog) remains set until BKP is entered or a SYSEEDIT or any level of deadstart is performed.

CPUMTR Breakpoint Commands

The following DSD commands control CPUMTR breakpoint processing. These commands and the DSD C display enable you to communicate with the CPUMTR breakpoint package. The C display is described following the command descriptions. The CPUMTR breakpoint commands require that the system console be unlocked (refer to the UNLOCK command earlier in this section); on a secured system, the console must be in security unlock status.

CAUTION

Use these commands with extreme caution to avoid damage to the system or to user jobs. Incorrect usage may cause a system to hang up and require a deadstart.

NOTE

Only one CPUMTR breakpoint can be set at a time. Any existing breakpoint is cleared when you enter a CPB,BKP,... command.

Command	Description
CPB,BKP,addr.	Sets breakpoint at absolute address addr. Restarts execution, if the CPU is waiting for a GO.
CPB,BKP,//addr.	Sets breakpoint at CPUMTR relative address addr. Restarts execution, if the CPU is waiting for a GO.
CPB,BKP,/bbbb/addr.	Sets breakpoint at address addr in CPUMTR block bbbb. Restarts execution, if the CPU is waiting for a GO.
CPB,DSP.	Brings up the CPUMTR breakpoint variant of the C display on the left screen if the C display is not already active. Causes the memory block in the C display to begin at the breakpoint address.
CPB,DSP,addr.	Brings up the CPUMTR breakpoint variant of the C display on the left screen if the C display is not already active. Causes the memory block in the C display to begin at absolute address addr.
CPB,DSP,//addr.	Brings up the CPUMTR breakpoint variant of the C display on the left screen if the C display is not already active. Causes the memory block in the C display to begin at CPUMTR relative address addr.
CPB,DSP,/bbbb/addr.	Brings up the CPUMTR breakpoint variant of the C display on the left screen if the C display is not already active. Causes the memory block in the C display to begin at address addr in CPUMTR block bbbb.
CPB,DSP+	Brings up the CPUMTR breakpoint variant of the C display on the left screen if the C display is not already active. Advances each of the memory block addresses by 10 ₈ .
CPB,DSP-	Brings up the CPUMTR breakpoint variant of the C display on the left screen if the C display is not already active. Decrements each of the memory block addresses by 10 ₈ .

Command	Description
CPB,GO.	Clears any breakpoint that is set and allows CPUMTR to restore its registers and continue execution.
CPB,P+	Sets breakpoint at the address following the current breakpoint address and restarts execution, if the CPU is waiting for a GO.
CPB,P-	Sets breakpoint at the address preceding the current breakpoint address and restarts execution, if the CPU is waiting for a GO.

CPUMTR Breakpoint Display (C Display)

The CPUMTR breakpoint display is a variant of the DSD C display that is formatted for use in CPUMTR breakpoint processing. This display is initiated by entering one of the CPB,DSP commands. Figure 5-5 shows the format of the CPUMTR breakpoint display. The display consists of the 20₈-word exchange package from the current or last breakpoint (displayed in 5 groups of 4 octal digits per line). All exchange package information is displayed except RA, FL, RAE, FLE, and MA. The P address shown is the actual breakpoint address.

The exchange package is followed by two 10₈-word blocks of central memory (displayed in 4 groups of 5 octal digits per line). The address of the memory block to be displayed can be specified using one of the CPB,DSP commands. Even if the address specified in the CPB,DSP command is a relative address, the addresses shown in the display are absolute; thus, giving the actual location in memory of the displayed data.

C. CENTRAL MEMORY.										
00062662	54726		1		0					
00062663	0		106		1					
00062664	0	55730		75230						
00062665	0	43		55713						
00062666	0	55731		4						
00062667	0	1022		0						
00062670	0	54533		55676						
00062671	0	43		7200						
}										Exchange Package
00062672	7777	7777	7777	7777	0000					
00062673	0000	0000	0000	0000	0003				C	
00062674	3423	1200	0000	0000	0000	1SJ				
00062675	0000	0000	0000	0000	0001				A	
00062676	0000	0000	0000	0000	0003				C	
00062677	0000	0000	0002	0000	0000			B		
00062700	3423	1200	0002	0000	0000	1SJ	B			
00062701	4000	0002	0001	0040	0344	5	B A 5C9			
}										
00054726	51200	00052	43324	20360		(P)8Z7C			
00054727	11232	51300	00111	20303		ISUK	IJCC			
00054730	03330	54777	61000	46000		COE*	D			
00054731	43724	20760	51700	55010		8	7G (E/H			
00054732	02000	64037	61000	46000		B	F54 D			
00054733	10322	43405	51200	00057		HZT1E(P	.			
00054734	76770	20451	20735	15664		BD(P	(.			
00054735	20251	12667	11242	43422		PUIV	ITT1R			
}										Memory Block
00054736	13227	03020	54776	46000		KR	XP=			
00054737	03030	54777	61000	46000		CCE*	D			
00054740	61600	54746	07410	55735			E*-G6E.2			
00054741	51200	00144	51140	55010		(P	A9(LE/H			
00054742	21244	73110	03020	54776		QT*	YHCBE*			
00054743	03010	55735	43301	53110		CAE.	28XMYH			
00054744	20352	15114	36713	15773		P2Q(L3	Y.			
00054745	04000	71362	61000	46000		D	GK D			

Figure 5-5. CPUMTR Breakpoint Display

CPUMTR Breakpoint Precautions

The following precautions and information are important to the successful use of the CPUMTR breakpoint package.

- The breakpoint must be set in executable code.
- When a breakpoint is set, DSD stores RJ BKP at the specified address. When the breakpoint address is hit, CPUMTR saves all of the registers in the display area, clears the breakpoint, and then waits to receive a GO. The message CPUMTR BREAKPOINT HIT is displayed at the system control point.
- From the time that the breakpoint is hit until the CPU continues, the system clock stops. Therefore, you should reset the system clock (with the DSD TIME command) after you have finished breakpoint processing.
- A CPUMTR breakpoint can be set in monitor mode or program mode. However, alternating between the modes is not advised. Going from monitor to program mode is probably safe, but going from program to monitor mode is almost certain to fail causing CPUMTR to abort with the following message issued to the system console:

```
CPUMTR ERROR EXIT.
```

- Since the C display is resident in RPL, changing displays is allowed (before the breakpoint is hit) because you can always return to the C display. However, once the breakpoint is hit until the CPU continues, CPUMTR is hung. This means that all activity in the system stops until the CPU restarts. If DSD attempts to request a monitor function, the following message appears in the B display:

```
NO MONITOR RESPONSE.
```

This is normal since CPUMTR is hung in the breakpoint code and the message may be ignored. In this case, the DSD request that required the monitor function is not processed.

- If you attempt to reference a CPUMTR block that is not loaded, the following message appears on the system console:

```
INCORRECT ENTRY.
```

- The C display retains the last addresses displayed. For example, if you change to another display and then come back to the C display it will still show the same absolute addresses.
- Once a breakpoint has been hit and CPUMTR is waiting for a CPB,GO (or another CPB,BKP) command, the contents of any register or memory location can be changed by using the normal memory entry commands. To change a register, simply change the corresponding memory location on the C display (as initiated by the CPB,DSP command). The exchange package shown on the C display will be reloaded to display the actual registers when execution resumes.

NOTE

Since CPUMTR has been halted, any memory entries made at this time must be preceded by a 99 command to disable DSD syntax overlay processing and the logging of the command to the error log.

Debugging Commands

The following commands trap certain conditions and at the same time trace selected PP and MTR to CPUMTR functions.

TRACE

TRAP

These debugging commands can be used to solve specific system problems. Care must be exercised in using the commands in a production environment as performance may be impacted significantly.

The following entry must be made in the IPRDECK at deadstart.

ENABLE, TRACE.

In addition, system debug mode must be enabled at the system console (the console must be unlocked to do this) in order to execute the TRAP and TRACE commands.

Command	Description				
TRACE,type, parameters.	Permits specification of monitor functions and other system data. Entering the command without any parameters (TRACE,PPU. or TRACE,MTR.) clears the trace function. The following types may be traced.				
	<table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>PPU</td> <td> Traces PP to CPUMTR requests. Each trace buffer entry consists of the following words: <ul style="list-style-type: none"> • Time stamp. • PP input register. • PP output register. • PP message buffer (6 words). • Absolute address 1 (default is SMRL). • Absolute address 2 (default is CMCL). The format of the time stamp word is: <pre>***dhms111</pre> In this format, dhms is bits 23-0 (day, hour, minute, and second) of PDTL and 111 is bits 17-0 of RTCL. <p>The trace buffer can hold 20 entries. The oldest entry in the buffer is overwritten by a new entry once the buffer is filled. The DSDI directive TBDUMP,PPU displays the buffer, as well as the functions traced, the absolute addresses saved, and the buffer IN pointer.</p> </td> </tr> </tbody> </table>	type	Description	PPU	Traces PP to CPUMTR requests. Each trace buffer entry consists of the following words: <ul style="list-style-type: none"> • Time stamp. • PP input register. • PP output register. • PP message buffer (6 words). • Absolute address 1 (default is SMRL). • Absolute address 2 (default is CMCL). The format of the time stamp word is: <pre>***dhms111</pre> In this format, dhms is bits 23-0 (day, hour, minute, and second) of PDTL and 111 is bits 17-0 of RTCL. <p>The trace buffer can hold 20 entries. The oldest entry in the buffer is overwritten by a new entry once the buffer is filled. The DSDI directive TBDUMP,PPU displays the buffer, as well as the functions traced, the absolute addresses saved, and the buffer IN pointer.</p>
type	Description				
PPU	Traces PP to CPUMTR requests. Each trace buffer entry consists of the following words: <ul style="list-style-type: none"> • Time stamp. • PP input register. • PP output register. • PP message buffer (6 words). • Absolute address 1 (default is SMRL). • Absolute address 2 (default is CMCL). The format of the time stamp word is: <pre>***dhms111</pre> In this format, dhms is bits 23-0 (day, hour, minute, and second) of PDTL and 111 is bits 17-0 of RTCL. <p>The trace buffer can hold 20 entries. The oldest entry in the buffer is overwritten by a new entry once the buffer is filled. The DSDI directive TBDUMP,PPU displays the buffer, as well as the functions traced, the absolute addresses saved, and the buffer IN pointer.</p>				

Command	Description								
TRACE,type, parameters.	(Continued)								
	<table border="1"> <thead> <tr> <th>type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>MTR</td> <td> <p>Traces MTR to CPUMTR functions. Each trace buffer entry consists of the following words:</p> <ul style="list-style-type: none"> • Time stamp. • Function. • Absolute address 1 (default is SMRL). • Absolute address 2 (default is CMCL). • *MCT* <p>*MCT* expands to the following for each entry in the memory control table (MCT):</p> <ul style="list-style-type: none"> • MCT word 1. • MCT word 2. • CPA/PCPA address 1 (default is STSW). • CPA/PCPA address 2 (default is FLSW). • Word TFSW from the CPA/PCPA. • Word JSNE from the EJT. <p>Depending on the number of pseudo-control points defined, there is room for approximately 10 entries in the trace buffer. The oldest entry in the buffer is overwritten by a new entry once the buffer is filled.</p> <p>The DSDI directive TBDUMP,MTR displays the buffer, as well as the functions traced, the absolute and CPA/PCPA addresses saved, and the buffer IN pointer.</p> </td> </tr> <tr> <td>SET,ABS</td> <td>Sets the absolute addresses of the CM words to be saved in the trace buffer entry.</td> </tr> <tr> <td>SET,CPA</td> <td>Sets the CPA/PCPA addresses of the CM words to be saved in the trace buffer entry.</td> </tr> </tbody> </table>	type	Description	MTR	<p>Traces MTR to CPUMTR functions. Each trace buffer entry consists of the following words:</p> <ul style="list-style-type: none"> • Time stamp. • Function. • Absolute address 1 (default is SMRL). • Absolute address 2 (default is CMCL). • *MCT* <p>*MCT* expands to the following for each entry in the memory control table (MCT):</p> <ul style="list-style-type: none"> • MCT word 1. • MCT word 2. • CPA/PCPA address 1 (default is STSW). • CPA/PCPA address 2 (default is FLSW). • Word TFSW from the CPA/PCPA. • Word JSNE from the EJT. <p>Depending on the number of pseudo-control points defined, there is room for approximately 10 entries in the trace buffer. The oldest entry in the buffer is overwritten by a new entry once the buffer is filled.</p> <p>The DSDI directive TBDUMP,MTR displays the buffer, as well as the functions traced, the absolute and CPA/PCPA addresses saved, and the buffer IN pointer.</p>	SET,ABS	Sets the absolute addresses of the CM words to be saved in the trace buffer entry.	SET,CPA	Sets the CPA/PCPA addresses of the CM words to be saved in the trace buffer entry.
type	Description								
MTR	<p>Traces MTR to CPUMTR functions. Each trace buffer entry consists of the following words:</p> <ul style="list-style-type: none"> • Time stamp. • Function. • Absolute address 1 (default is SMRL). • Absolute address 2 (default is CMCL). • *MCT* <p>*MCT* expands to the following for each entry in the memory control table (MCT):</p> <ul style="list-style-type: none"> • MCT word 1. • MCT word 2. • CPA/PCPA address 1 (default is STSW). • CPA/PCPA address 2 (default is FLSW). • Word TFSW from the CPA/PCPA. • Word JSNE from the EJT. <p>Depending on the number of pseudo-control points defined, there is room for approximately 10 entries in the trace buffer. The oldest entry in the buffer is overwritten by a new entry once the buffer is filled.</p> <p>The DSDI directive TBDUMP,MTR displays the buffer, as well as the functions traced, the absolute and CPA/PCPA addresses saved, and the buffer IN pointer.</p>								
SET,ABS	Sets the absolute addresses of the CM words to be saved in the trace buffer entry.								
SET,CPA	Sets the CPA/PCPA addresses of the CM words to be saved in the trace buffer entry.								

Command	Description
---------	-------------

TRACE,type,
parameters.

(Continued)

The following parameters are used with TRACE,MTR and TRACE,PPU. Trace data is collected by CPUMTR immediately before the selected function is processed.

parameters	Description
------------	-------------

- | | |
|------------------------|---|
| * | Traces all functions of the selected type. |
| F_1, F_2, \dots, F_3 | Traces up to five functions. Functions are specified by their symbolic names. |

The following parameters are used with TRACE,SET,ABS and TRACE,SET,CPA.

parameters	Description
------------	-------------

- | | |
|------------------|--|
| $addr_1, addr_2$ | Sets the absolute (TRACE,SET,ABS) or relative (TRACE,SET,CPA) addresses of the words to be saved in each trace buffer entry. The addresses must be octal digits. |
|------------------|--|

The following examples illustrate various forms of the TRACE command.

Example	Description
TRACE, PPU, ROCM, JACM.	Causes a trace entry to be created each time a JACM or ROCM monitor function is issued.
TRACE, MTR, ACSF.	Causes a trace entry to be created each time an ACSF function is issued.
TRACE, SET, ABS, 21, 152.	Causes words 21 ₈ and 152 ₈ to be saved in each trace buffer entry.
TRACE, SET, CPA, 21, 57.	Causes words 21 ₈ and 57 ₈ of each CPA and PCPA to be saved for each MTR trace buffer entry.

Command	Description
TRAP,type,parameters.	Halts the system immediately (in monitor mode) upon detection of a specified condition. Only one type of trap can be active at once. Entering the command without any parameters (TRAP.) clears the trap.

The following types do not accept parameters:

type	Description
NFL	Checks for erroneous data in negative field length of executing jobs.
MCT	Checks for invalid data in the memory control table.
EJT	Checks for invalid data in the executing job table.
DJB	Checks for jobs that are no longer executing, but are failing to advance.

The following types accept parameters:

type	Description
CPA	Checks a specified control point area field or pseudo-control point area field for a given value.
MEM	Checks a field specified by an absolute address (central memory only) for a given value.

The following parameters, which are associated with CPA and MEM, indicate the field and value being checked:

parameters	Description
W	Absolute address (type=MEM) or offset (type=CPA).
V	Value (octal). If V=N, any nonzero field value halts the system.
H	Left-most bit of field. Decimal value ranges from 0 to 59.8
L	Right-most bit of field. Decimal value ranges from 0 to 59.

The following examples illustrate various forms of the TRAP command:

Example	Description
TRAP,CPA,20,N,47,36.	Checks for any nonzero value in bits 47 to 36 of word 20 in each control point area.
TRAP,MEM,37461,36,7,3	Checks for a value of 36 in bits 7 to 3 of 37461.

Express Deadstart Dump Interpreter (DSDI)

6

Calling the Express Deadstart Dump Interpreter	6-3
Input Directives	6-6
Directive Format	6-7
List Control Directives	6-8
EJ — Force Page Eject	6-8
EJOFF — Turn Off Auto Page Eject	6-8
EJON — Turn On Auto Page Eject	6-8
PD — Reset Print Line Density	6-8
* — Comment in Subtitle Line	6-9
File Manipulation and Control Directives	6-9
DISPOSE — Dispose Alternate List File to Print Queue	6-9
OUTPUT — Assign Output to Alternate List File	6-9
READ — Read Alternate Directives File	6-10
REWIND — Rewind File	6-10
Central Memory/Extended Memory Dump Directives	6-11
Dump Control Directives	6-11
ALLMEM — Extend Central Memory Dumps	6-11
CM — Set Memory Type to Central Memory	6-11
EC — Set Memory Type to Extended Memory	6-12
RA — Reset Reference Address	6-12
RAC — Reset Reference Address to RA of Control Point	6-13
UEC — Set Memory Type to User-Access Extended Memory	6-13
Memory Dump Directives	6-14
C — Dump Memory in Instruction Parcel Format	6-14
D — Dump Memory in Byte Format	6-14
E — Dump Memory in Word Format	6-15
I — Dump 64-Bit Memory in Instruction Parcel Format	6-15
W — Dump 64-Bit Memory in Word Format	6-16
PP Dump Directives	6-17
AP — Dump Analysis of PP and PP Memory in Octal Line Format	6-17
MPP — Move PP	6-18
P — Dump PP Memory in Octal Block Format	6-18
PF — Dump FLPP Memory in Octal Block Format	6-19
PMS — Read PP Select Switches	6-19
PO — Dump 16-Bit PP Memory in Octal Block Format	6-20
PX — Dump 16-Bit PP Memory in Hexadecimal Block Format	6-20
Q — Dump PP Memory in Octal Line Format	6-21
QF — Dump FLPP Memory in Octal Line Format	6-21
QOA — Dump 16-Bit PP Memory in Octal/ASCII Line Format	6-22
QOD — Dump 16-Bit PP Memory in Octal/Display Line Format	6-22
QXA — Dump 16-Bit PP Memory in Hexadecimal/ASCII Line Format	6-23
QXD — Dump 16-Bit PP Memory in Hexadecimal/Display Line Format	6-23
CMR Dump Directives	6-24
ACCOUNT — Dump Account Dayfile Buffer	6-24
B — Dump Brief Control Point Status	6-24
CBT — Dump Control Buffer Table	6-25
CCT — Dump Channel Control Table	6-25

CP — Dump Active Control Point Areas	6-26
CPO — Reset Default List Options	6-28
CT — Dump Channel Tables	6-28
DAYFILE — Dump System Dayfile Buffer	6-28
DB — Dump Disk Buffers	6-28
DBW — Dump Buffered Device/Buffer Statistics/PP-I/O Buffer Link Tables	6-29
DDB — Dump Dayfile Dump Buffer	6-29
DP — Dump Dayfile Buffer Pointers	6-29
EICB — Dump Environment Interface Communication Buffer	6-29
EJT — Dump Executing Job Table	6-30
EPB — Dump Extended Memory/PP Buffer	6-30
ERRLOG — Dump Error Log Buffer	6-30
EST — Dump Equipment Status Table	6-30
FNT — Dump System File Name/File Status Table	6-30
FOT — Dump Family Ordinal Table	6-31
HAT — Dump Hash Table	6-31
JC — Dump Job Control Area for Each Service Class	6-31
LC — Dump Low Central Memory	6-31
LDIS — Dump L-Display Buffer	6-32
LIDT — Dump Logical Identifier Table	6-32
MAINLOG — Dump Binary Maintenance Log Buffer	6-32
MCT — Dump Memory Control Table	6-32
MST — Dump Mass Storage/Track Reservation Tables	6-33
MTR — Dump CPU Monitor	6-33
MTRQUEUE — Dump CPUMTR Request and Recall Queues	6-34
ODIS — Dump Operator Display Buffer	6-34
PCP — Dump Pseudo-Control Point	6-34
PLD — Dump Peripheral Library Directory	6-35
PP — Dump PP Communication Areas	6-35
PROBE — Dump PROBE Data Tables	6-35
PST — Dump Program Status Table and Entry Point Directory	6-35
PUT — Dump Physical Unit Table	6-36
QFT — Dump Queued File Table	6-36
RCL — Dump Resident Central Library	6-36
RPL — Dump Resident Peripheral Library	6-37
SAB — Dump System Attribute Block	6-37
SDA — Dump Extended Statistical Data Area	6-37
SECDED — Dump SECDED Identifier Table	6-37
SST — Dump Subsystem Control Point/Subsystem Assignment Tables	6-38
TBDUMP — Dump Trace Buffer	6-38
Subsystem Dump/Analysis Directives	6-39
BATCHIO (BIO) — Dump Associated Memory for Analysis	6-39
IAF — Dump Associated Memory for Analysis	6-40
MAGNET (MAG) — Dump Associated Memory for Analysis	6-41
RHF — Dump Associated Memory for Analysis	6-42
Hardware Register Dump Directives	6-43
FMFREG — Dump PP Registers	6-43
IOUCR — Dump Concurrent PP Channel Registers	6-43
IOUMR — Dump IOU Maintenance Registers	6-43
LPVA — Load Central Memory PVA into Pseudo Register	6-44
MEMMR — Dump Memory Maintenance Registers	6-44
PROCA — Dump Processor Operand Cache	6-45
PROCW — Dump Processor Controlware Part Number and Revision Level	6-45

PROMR — Dump Processor Maintenance Registers	6-45
PROPM — Dump Processor Page Map	6-46
PRORF — Dump Processor Register File	6-46
PROSM — Dump Processor Segment Map	6-46
PROXP — Dump Processor Exchange Package	6-47
SC — Dump S/C Register	6-47
SETCPU — Set CPU Number	6-47
SETIOU — Set PP Dump Defaults	6-48
SETJPS — Change the JPS Register Value	6-48
SETRMA — Convert Address into RMA	6-49
SETVEP — Set Virtual Address Parameters	6-49
TRACEBK — Traceback Stack Frame Save Area	6-50
XP - Dump Deadstart Exchange Package	6-50
Buffer Controller Directive	6-51
Interactive Use of DSDI	6-52
Terminal Output Directives	6-54
C — Dump Memory in Instruction Parcel Format	6-54
CP — Dump Active Control Point Areas	6-55
D — Dump Memory in Byte Format	6-56
PP — Dump PP Communication Areas	6-57
Q — Dump PP Memory in Line Format	6-58
QOA, QOD, QXA, QXD — Dump 16-Bit PP Memory in Line Format	6-59
Example of DSDI Terminal Use	6-60
Printer Output Listing Examples	6-62

Express Deadstart Dump Interpreter (DSDI)

6

The deadstart dump interpreter (DSDI) is a utility program that converts selected portions of the binary information on an express deadstart dump (EDD) file into reports to be listed. The EDD file is generated on magnetic tape by the EDD utility, which may be run at deadstart time after a system malfunction has occurred. Refer to the CIP User's Handbook for complete information concerning the use of EDD.

Selection of data to be listed by DSDI is provided through input directives, either on an input file or on the DSDI command. Normal octal dumps of central memory, extended memory, and PP memory can be produced by these directives as well as specially formatted dumps of specific system tables and buffers.

EDD and DSDI offer these features:

- The EDD tape file contains a dump of memory, the executing exchange packages, the CYBER 170 status and control (S/C) registers, where applicable; maintenance registers for models 865 and 875 and CYBER 180-class machines, register file, exchange package, processor control store, processor operand cache, processor page map, processor segment map, where applicable; and all PPs except for PP0, which can be saved by either copying it to another PP or by biasing the PP switch on the deadstart panel. EDD can optionally dump all or part of extended memory and the selected buffer controllers. EDD and DSDI permit analysis of a system malfunction to be performed entirely offline.
- Because DSDI copies the EDD file to a word-addressable random file on mass storage, dump data can be accessed in any order.
- EDD requires only a small amount of time during deadstart because the data is transferred in binary form to magnetic tape.
- The tape file created by EDD can be retained on magnetic tape or mass storage until it is no longer needed. Thus, a selective listing can be generated at any time.
- General information from the EDD file appears in the title and subtitle line of each page of listed output. The first 50 columns of each input directive are included in the title line of the output list it produces. An input directive is provided that enables insertion of comments into the subtitle line.
- Use of DSDI is possible from an interactive terminal as well as from the batch environment. The output produced by several directives is formatted for terminal output (72 columns). From a batch environment, output is formatted for a 136-column printer.

DSDI can also be used to create a message file containing the contents of the maintenance registers (input/output unit, memory, and processor) that were dumped to the EDD file for CYBER 180-class machines. This message file, which is similar in format to the binary maintenance log (BML), can then be analyzed by an interpreter program, such as the Hardware Performance Analyzer (HPA), to produce maintenance register reports. Refer to the IOUMR, MEMMR, and PROMR DSDI input directives later in this section for additional information.

DSDI writes up to three messages to this message file: one message each for input/output unit errors, memory errors, and processor errors. If the status summary register for a particular element is zero (signifying no errors for that element), DSDI does not write the corresponding message to the message file.

In addition, the message file contains messages for any errors that were detected by the dedicated fault tolerance (DFT) processor but not yet written to the BML at the time the EDD file was created. These messages are identical in format to the ones that would have been written to the BML by the normal logging process.

Calling the Express Deadstart Dump Interpreter

Processing of the EDD file is initiated with the DSDI command. The format of the command is:

DSDI, p₁, p₂, . . . , p_n.

Each parameter p_i is either a keyword or a keyword equated to a value. All parameters are optional and order independent.

p_i	Description
B=bmlfile	The 1- to 7-character name of the message file on which the maintenance register messages and unlogged DFT error messages are to be written. The message file is similar in format to the BML and can be analyzed by HPA. This parameter is valid only if maintenance register messages or unlogged DFT error messages have been dumped to the EDD file.
B omitted	No message file is to be created.
D=fn	This parameter causes DSDI to write the random dump file to file fn. The original EDD file is not returned. The random file can then be used as the dump file on subsequent executions of DSDI, eliminating the need to read the entire dump tape on each call. This allows file fn to be predefined as a permanent mass storage file.
D	Creates a random dump file with the same name as the EDD file and returns the original EDD file.
D omitted	No random dump file is to be created.
DMB	A binary dump file produced by the DMB or LOADBC command is to be analyzed instead of an EDD file.
DMB omitted	EDD file is to be analyzed.

When the DMB parameter is specified for a DMB-created dump, DSDI uses directives as though the job were running at control point 1. The DMB command dumps the exchange package, central memory, and extended memory in binary (refer to the NOS Version 2 Reference Set, Volume 3). Use the following directives to dump selected portions of memory to be analyzed.

Directive	Description
CP,1	Dump exchange package.
RAC,1/C,fwa,lwa	Dump central memory in C format.
RAC,1/D,fwa,lwa	Dump central memory in D format.
RA,O/EC/C,fwa,lwa	Dump extended memory in C format.
RA,O/EC/D,fwa,lwa	Dump extended memory in D format.

P_i	Description
DMB omitted	<p><i>(Continued)</i></p> <p>fwa and lwa are the first and last word addresses of memory to be dumped. When this parameter is specified to analyze a LOADBC controlware dump, the only DSDI directive that can be used is:</p> <p style="padding-left: 40px;">BCDUMP/ops.</p> <p>Refer to the BCDUMP directive described later in this section for a description of the ops parameter.</p>
F=dumpfile	The 1- to 7-character name of the EDD file that DSDI will interpret.
F omitted	Same as F=DUMP.
FULL	Entire EDD file is read.
FULL omitted	When interpreting a dump for a CYBER 180-class system, the EDD file is read only up to the end of NOS memory. NOS/VE memory, external extended memory (STORNET or ESM) and buffer controller records are not read.
I=infile	The 1- to 7-character name of file on which input directives are written.
I omitted	Same as I=INPUT.
L=outfile	The 1- to 7-character name of file on which list output is to be written. The user must save or print the file.
L omitted	Same as L=OUTPUT, except that the file is automatically printed.
NR	EDD file is not to be rewound.
NR omitted	EDD file is to be rewound before processing.
P	<p>CMR pointers from the running system are to be used. Selecting this parameter causes the CMR pointers from the running system to be used to locate tables and buffer areas on the EDD file. This parameter is typically used when it is known that the CMR pointers on the EDD file were destroyed by system malfunction. Directives used to dump CMR will dump the pointers contained on the EDD file, not those from the running system.</p> <p>This parameter should be used only when the configuration of the running system is the same as the system in use when the EDD file was created.</p>
P omitted	CMR pointers from EDD file are to be used.
PD=n	Print density in number of lpi (3, 4, 6, or 8).

P_i	Description
PD	Same as PD=8.
PD omitted	Same as the job's default print density.
Z	Input directives are contained on the command after the terminator. The I parameter is ignored. This eliminates the need to use a separate input file for the directives when only a few directives are needed.
Z omitted	Input directives are not contained on the command. The system uses the I parameter.

When input directives appear on the DSDI command, the first character following the command terminator is the separator character for all directives on the command. Any display code character that is not used in any of the directives, including a space, can be used as the separator character. Each directive must be preceded by a separator and terminated by a period. The directives can extend to column 72 on the command. Continuation lines are not permitted.

For example, (slant used for separator):

```
DSDI, Z./SC./XP./P./D,0,20000./EC./D,0,10000.
```

If the directives are included in the input file, the following equivalent job would appear.

```
:
DSDI .
--EOR--
SC.          Dump CYBER 170 S/C register.
XP.          Dump executing exchange packages.
P.           Dump all PPs.
D,0,20000.   Dump the first 200008 locations of central memory.
EC.          Set memory type to extended memory.
D,0,10000.   Dump the first 100008 locations of extended memory.
--EOI--
```

A request for the EDD tape must precede the DSDI command. Since EDD writes information on a labeled, 7- or 9-track tape at a density of 800 characters per inch (cpi) for 7 track and 1600 or 6250 cpi for 9 track, the request should appear as follows:

```
LABEL,DUMP,D=density,tape,F=S,LB=KL,VSN=DUMP.
```

Parameter	Description
density	800 cpi for 7-track tape. 1600(PE) or 6250(GE) cpi for 9-track tape.
tape	MT for 7-track tape. NT for 9-track tape.

The request is presented in the resource mounting preview display and the job is rolled out until the tape is mounted and assigned. Although the default express dump file name (DUMP) is used in this example, a different file name can be specified, provided the same file name is also specified on the DSDI command (F parameter).

Input Directives

DSDI input directives provide the capability to selectively dump only those portions of the EDD file that are of interest. The input directives are grouped into the following categories.

Type of Directive	Description
List control	Enables you to control line printer page eject and print density, and to specify comments in subtitle line of the output listing.
File manipulation and control	Enables you to specify alternate files for DSDI input directives and listing output.
Central memory/extended memory dump	Provides octal dumps of specified portion of central memory or extended memory. Absolute or relative addresses can be printed (three print formats are available). Two additional dump formats are available to print 64-bit memory in hexadecimal for CYBER 180-class machines.
PP dump	Provides octal memory dumps of all or selected PPs (two print formats are available). PP analysis data can be included in the dumps. Six additional formats are available to print 16-bit PP memory for CYBER 180-class machines.
CMR dump	Provides specially formatted dumps of selected areas in central memory resident. These areas are specified by name rather than by address.
Subsystem dump/analysis	Provides specially formatted dumps of subsystem control points and associated tables and buffers.
Hardware register dump	Provides dumps of specified hardware registers.

Several of the DSDI input directives have interactive capabilities. The output produced by these directives is specially formatted for listing at an interactive terminal. Refer to Interactive Use of DSDI later in this section for information concerning directive entry and use of DSDI from an interactive terminal.

Refer to Printer Output Listing Examples, later in this section, for examples of listings produced by several of the DSDI input directives.

Directive Format

A directive has the following format.

```
directive, p1, p2, . . . , pn. comments
```

Field	Description
directive	The directive name starts in column 1. It is terminated by a separator or terminator character.
p _i	Parameter(s) for the directive. Depending on the requirements of the directive, the directive may have no parameters or a number of parameters.
comment	Any characters following the directive terminator are considered comments and are ignored by DSDI. However, the comments are included with the directive in the title line of each page of the output listing (a combined total of 50 characters appear in the title line).
separator	You can use any character, including a space, to separate the fields of a directive. The exceptions are: <ul style="list-style-type: none"> A to Z 0 to 9 + - * .

Refer to Interactive Use of DSDI later in this section for additional information concerning directive entry from an interactive terminal.

List Control Directives

The list control directives provide the capability to specify print density and page eject options, and to add comments in the subtitle line of each page listed.

EJ — Force Page Eject

The EJ directive forces to issue a page eject function before listing the output produced by the next directive processed. The EJ directive can also force a page eject upon reaching a specified point on the page being printed. The page eject function is performed automatically unless disabled by the EJOFF directive.

Format:

EJ, nn.

Parameter	Description
-----------	-------------

nn	Force page eject only if less than nn decimal lines remain on the current page. If nn is omitted, page eject is forced before listing the output from the next directive processed.
----	---

EJOFF — Turn Off Auto Page Eject

The EJOFF directive disables auto page eject. Until this directive is processed, DSDI automatically issues a page eject function before listing the output produced by each new directive.

Format:

EJOFF.

EJON — Turn On Auto Page Eject

The EJON directive enables auto page eject (default condition). DSDI automatically issues a page eject function before listing the output produced by each new directive processed. This directive has no effect unless auto page eject has been disabled by the EJOFF directive.

Format:

EJON.

PD — Reset Print Line Density

The PD directive resets the print line density to a value other than that specified on the DSDI command.

Format:

PD, n.

Parameter	Description
-----------	-------------

n	New print line density in number of lpi (3, 4, 6, or 8). If n is omitted or an incorrect value is specified, a diagnostic message is issued.
---	--

***. — Comment in Subtitle Line**

The *. directive specifies a comment that appears in the subtitle line of each page listed.

Format:

```
*.ccc...ccc
```

Parameter	Description
ccc...ccc	Comment; up to 36 characters are printed.

File Manipulation and Control Directives

File manipulation and control directives provide the capability to specify alternate files for DSDI input directives and listing output.

DISPOSE — Dispose Alternate List File to Print Queue

The DISPOSE directive causes the alternate list file specified by the OUTPUT directive to be disposed to the print queue. DSDI listing output then resumes on the original output file. This directive has no effect unless output has previously been assigned to an alternate list file (refer to the OUTPUT directive, next). Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

```
DISPOSE,username.
```

Parameter	Description
username	User name under which the remote batch terminal to receive the listing is logged in. If username is omitted, the listing is printed at a central site line printer.

OUTPUT — Assign Output to Alternate List File

The OUTPUT directive temporarily assigns DSDI listing output to a file other than that specified on the DSDI command. When the alternate file is disposed to the print queue (refer to the DISPOSE directive, above), output resumes on the original output file. If the alternate file is not disposed, both the original and the alternate output files remain at the job control point as local files. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

```
OUTPUT,filename.
```

Parameter	Description
filename	Name of alternate list file (from 1 to 7 characters). Only one alternate output file may be active at a time; filename cannot be the same name as the normal output file. If filename is omitted, the system assumes file name ALTRNT.

READ — Read Alternate Directives File

The READ directive causes DSDI to temporarily stop reading the current directives file and begin reading directives from the specified record on the named alternate file or from the current position if the record name is omitted. DSDI reads from the specified alternate directives file until an end-of-record is encountered (end-of-file or empty record if * is specified) and then resumes with the next directive on the original input directives file.

Format:

```
READ, filename, rec, *.
```

Parameter	Description
filename	Name of alternate directives file (local file).
rec	Optional record name. If rec is specified, file filename is searched for record rec from the current position to end-of-file or an empty record. If rec is not found, DSDI issues an error message. If rec is not specified, DSDI reads directives from the current position to end-of-record. Records must be in text format where the first word of the record is the record name, unless the file is assigned to an interactive terminal; then, directives may be entered directly.
*	Optional character specifying that DSDI is to read directives from all records until an end-of-file or an empty record is encountered.

REWIND — Rewind File

The REWIND directive repositions the specified file to beginning-of-information.

Format:

```
REWIND, filename.
```

Parameter	Description
filename	Name of file to be rewound.

Central Memory/Extended Memory Dump Directives

The central memory/extended memory dump directives provide the capability to dump any portion of central memory or extended memory in instruction parcel, byte, or word format. There are 6-bit display code character equivalents included with each format. Either absolute or relative memory locations may be dumped. Refer to CMR Dump Directives later in this section for directives used to dump specific portions of NOS central memory resident (CMR).

Dump Control Directives

Dump control directives select the type of memory to be dumped (central memory or extended memory) and the addressing mode to be used (absolute or relative).

ALLMEM — Extend Central Memory Dumps

The ALLMEM directive enables central memory dumps to extend past the central memory boundary on machines with central memory extension.

Format:

ALLMEM.

CM — Set Memory Type to Central Memory

The CM directive specifies that subsequent C, D, and E directives dump central memory locations. Unless the EC or UEC directive is specified, central memory locations are dumped by default.

Format:

CM.

EC — Set Memory Type to Extended Memory

The EC directive specifies that subsequent C, D, and E directives dump extended memory or unified extended memory (UEM) locations. Unless the EC or UEC directive is specified, central memory locations are dumped by default.

Format:

EC, emid.

or

EC, emid, fwa.

Parameter	Description
emid	Extended memory identification; emid can be UEM or ESM. If emid is omitted, UEM is assumed. This parameter is useful only on models 865, 875, and CYBER 180-class machines since all other models have only one type of extended memory available.
NOTE	
If ESM is specified when interpreting a dump of a CYBER 180-class machine, the FULL option must be specified on the DSDI command.	
fwa	First word address (fwa) divided by 1000. fwa is added to the fwa and last word address (lwa) parameters in subsequent C, D, and E directives when dumping UEM. If fwa is omitted, DSDI determines the beginning of unified extended memory from the pointer in CMR and adds it to the fwa and lwa parameters of the C, D, and E directives. This parameter is ignored on any machine other than models 865 and 875 and CYBER 180-class machines.

RA — Reset Reference Address

The RA directive specifies that subsequent C, D, and E directives dump memory locations relative to a specified reference address. Unless the RA or RAC directive is entered, absolute memory locations are dumped by default.

Format:

RA, nnnnnnn.

Parameter	Description
nnnnnnn	Reference address; addresses specified on subsequent C, D, and E directives are relative to this address.

Clearing the reference address specified on the most recent RA or RAC directive reenables absolute addressing. This is done by entering the RA directive in this format:

RA, 0.

RAC — *Reset Reference Address to RA of Control Point*

The RAC directive specifies that subsequent C, D, and E directives dump memory locations relative to the reference address of a specified control point. Unless the RA or RAC directive is entered, absolute memory locations are dumped by default.

Format:

RAC, nn.

Parameter Description	
nn	Control point number; addresses specified on subsequent C, D, and E directives are relative to the reference address of this control point.

Refer to the description of the RA directive to reenable absolute addressing.

UEC — *Set Memory Type to User-Access Extended Memory*

The UEC directive specifies that subsequent C, D, and E directives dump user-accessible extended memory locations. Unless the EC or UEC directive is specified, central memory locations are dumped by default.

Format:

UEC, fwa.

Parameter Description	
fwa	First word address divided by 1000 ₈ for user-accessible extended memory. fwa is added to the fwa and lwa parameters in subsequent C, D, and E directives. If fwa is omitted, DSDI determines the beginning of user-accessible extended memory from the pointer in CMR and adds it to the fwa and lwa parameters of the C, D, and E directives.

Memory Dump Directives

Memory dump directives specify the area of memory to be dumped and determine the format of the output listing (refer to Printer Output Listing Examples later in this section for a sample output listing). The CM and EC or UEC directives determine the type of memory to be dumped (default is central memory). Absolute memory locations are dumped unless relative addressing has been enabled (refer to RA and RAC directives earlier in this section).

C — Dump Memory in Instruction Parcel Format

The C directive causes the specified locations of central memory or extended memory to be dumped in four groups of five octal digits (three words per line) with 6-bit display code character equivalents. Repetitive data is suppressed.

Format:

C, fwa, lwa.

Parameter	Description
fwa	First word address to be dumped (mandatory).
lwa	Last word address, plus one location, to be dumped. If lwa is omitted, fwa+1 is assumed by default.

The output listing is read from top to bottom by column rather than across the page. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

D — Dump Memory in Byte Format

The D directive causes the specified locations of central memory or extended memory to be dumped in five groups of four octal digits (three words per line) with 6-bit display code character equivalents. Repetitive data is suppressed.

Format:

D, fwa, lwa.

Parameter	Description
fwa	First word address to be dumped (mandatory).
lwa	Last word address, plus one location, to be dumped. If lwa is omitted, fwa+1 is assumed by default.

The output listing is read from top to bottom by column rather than across the page. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

E — Dump Memory in Word Format

The E directive causes the specified locations in central memory or extended memory to be dumped in word format (four words per line) with 6-bit display code character equivalents.

Format:

```
E, fwa, lwa.
```

Parameter	Description
fwa	First word address to be dumped (mandatory).
lwa	Last word address, plus one location, to be dumped. If lwa is omitted, fwa+1 is assumed by default.

I — Dump 64-Bit Memory in Instruction Parcel Format

The I directive is valid on CYBER 180-class machines only. It prints the specified locations of central memory in four groups of four hexadecimal digits (three words per line) with 7-bit ASCII code character equivalents. The second format is for NOS/VE dumps.

Format:

```
I, fba, lba, ei.
```

or

```
I, fba, lba, asid.
```

Parameter	Description
fba	First byte address (fba), hexadecimal, to be printed (mandatory).
lba	Last byte address (lba), plus one location, to be printed. If lba is omitted, fba+8 is assumed by default.
ei	If ei is 1, fba and lba specify byte addresses relative to byte 0 of the environment interface. If ei is omitted, fba and lba specify absolute byte addresses.
asid	Specifies an actual segment identifier (ASID) for NOS/VE dumps. If you specify this parameter, the fba and lba parameters must represent byte offsets and have values in the range of from 0 to 7FFFFFFF hexadecimal. This parameter is optional.

When you call DSDI from a batch job the output from this directive is printed in three columns, read from top to bottom, one column at a time. When you call DSDI from an interactive job the words are displayed one word per line.

For NOS/VE dumps, the addresses given for the fba and lba parameters must be real memory addresses (RMAs) or, when translated to system virtual addresses (SVAs), the actual segment identifiers (ASIDs) must match. Also, the dump must contain all central memory (not just critical memory), and the FULL option must be specified on the DSDI command. Refer to appendix H for NOS/VE address formats to use with the DSDI utility.

W — Dump 64-Bit Memory in Word Format

The *W* directive is valid on CYBER 180-class machines only. It prints the specified locations of 64-bit memory in hexadecimal (4 words of 16 digits per line) with 7-bit ASCII code character equivalents. The second format is for NOS/VE dumps.

Format:

`w, fba, lba, ei.`

or

`w, fba, lba, asid.`

Parameter	Description
<code>fba</code>	First byte address, hexadecimal, to be printed (mandatory).
<code>lba</code>	Last byte address, plus one location, to be printed. If <code>lba</code> is omitted, <code>fba+8</code> is assumed by default.
<code>ei</code>	If <code>ei</code> is 1, <code>fba</code> and <code>lba</code> specify byte addresses relative to byte 0 of the environment interface. If <code>ei</code> is omitted, <code>fba</code> and <code>lba</code> specify absolute byte addresses.
<code>asid</code>	Specifies an actual segment identifier (ASID) for NOS/VE dumps. If you specify this parameter, the <code>fba</code> and <code>lba</code> parameters must represent byte offsets and have values in the range of from 0 to 7FFFFFFF hexadecimal. This parameter is optional.

For NOS/VE dumps, the addresses given for the `fba` and `lba` parameters must be real memory addresses (RMAs), or when translated to system virtual addresses (SVAs), the actual segment identifiers (ASIDs) must match. Also, the dump must contain all central memory (not just critical memory), and the `FULL` option must be specified on the `DSDI` command. Refer to appendix H for NOS/VE address formats to use with the `DSDI` utility.

PP Dump Directives

PP dump directives provide the capability to obtain a memory dump of all or selected PPs. On some CYBER 180-class machines, the values of the P, Q, K, and A registers are also displayed. Two dump formats are available: block format and line format.

AP — Dump Analysis of PP and PP Memory in Octal Line Format

The AP directive causes PP memory to be dumped in line format with 6-bit display code character equivalents (same format as the Q directive). Repetitive lines are suppressed and zero bytes are represented by hyphens (—). An analysis of the PP is printed before the memory dump. Analysis data includes the associated PP communications area, resident entry point call addresses, and read-only variables in direct cells. On some CYBER 180-class machines, the values of the P, Q, K, and A registers are also displayed. Certain direct cell variables are verified and those in error are indicated.

NOTE

Correct operation of this directive requires that the PP communication area on the EDD file be intact.

Format:

$P, n_1, n_2, \dots, n_m.$

Parameter	Description
n_i	Number of PP to be dumped or a program name. If a program name is specified, all PPs executing that program are dumped. A warning message is issued if an incorrect number is specified or the program name is not found in any PP. If n_i is omitted, all active PPs are dumped.

Refer to Printer Output Listing Examples for a sample of the printer output listing produced by this directive.

MPP — Move PP

The MPP directive causes the correct logical PP to be dumped if the logical position of PP00 has been changed prior to the full dump to tape. If the PPS-0/PPS-1 toggle switch has been toggled, the n parameter should not be specified. If PP00 has to be moved to another PP via a deadstart panel program, the n parameter should be specified.

Format:

MPP, n.

Parameter Description

n	Number of PP to which PP00 was moved. n cannot equal 0. If n is omitted, 10 is assumed (CYBER 170 and CYBER 70 Computer Systems with 20 PPs).
---	---

P — Dump PP Memory in Octal Block Format

The P directive causes PP memory to be dumped in block format, where each block represents 64 words of memory. The blocks are read by column (top to bottom), where each column contains eight 12-bit words in octal format numbered 0 through 7. There are eight columns in each block, numbered 0 through 7. Repetitive data is not suppressed and zero words are represented by hyphens (—). For models 865 and 875, bytes 7776₈ and 7777₈ contain the PP's R register.

Format:

P, n₁, n₂, . . . , n_m.

Parameter Description

n _i	Number of PP to be dumped. If omitted, all PPs are dumped.
----------------	--

PF — Dump FLPP Memory in Octal Block Format

The PF directive causes first-level peripheral processor (FLPP)¹ memory to be dumped in block format, where each block represents 64 words of memory. The blocks are read by column (top to bottom), where each column contains eight 12-bit words in octal format numbered 0 through 7. There are eight columns in each block, numbered 0 through 7. Repetitive data is not suppressed and zero words are represented by hyphens (—).

Format:

PF, n_1, n_2, \dots, n_m .

Parameter	Description
n_i	Number of FLPP to be dumped. If omitted, all FLPPs are dumped.

PMS — Read PP Select Switches

The PMS directive causes the dump of the S/C register (maintenance register for models 865 and 875) to be read to determine the current value of the PP memory select switches and the correct logical PP to be dumped, if the logical position of PP00 has been changed prior to the full dump to tape. If the PP memory select switches have been changed, this directive should be specified with the binary value of the switches prior to the change. This directive is meaningful only on a CYBER 170 Computer System with the exception of models 815, 825, 835, 845, and 855.

Format:

PMS, n .

Parameter	Description
n	Previous select switch setting; 0 through 11 ₈ . If n is the same value as that read from the S/C register, this directive is not meaningful.

1. Hardware manuals define peripheral processors making up a peripheral processor subsystem as PPs, and the first-level peripheral processors as peripheral processing units (PPUs). In this manual, first-level peripheral processors are referred to as FLPPs. FLPPs are available only on model 176.

PO — Dump 16-Bit PP Memory in Octal Block Format

The PO directive is valid on CYBER 180-class machines only. It prints PP memory in block format, where each block represents 64 words of memory. DSDI represents each 16-bit word by six octal digits. Words whose upper 4 bits are zero are represented by four octal digits. Blocks are read by columns (top to bottom), where each column contains eight 16-bit words numbered 0 through 7. There are eight columns in each block, numbered 0 through 7. Repetitive data is not suppressed and zero words are represented by hyphens (—). The PP's R register is displayed preceding the first line of the dump.

Format:

`PO, n1, n2, . . . , nm.`

Parameter Description

<code>n_i</code>	Number of PP to be dumped. If omitted, all PPs are printed.
----------------------------	---

PX — Dump 16-Bit PP Memory in Hexadecimal Block Format

The PX directive is valid on CYBER 180-class machines only. It prints 16-bit memory in block format, where each block represents 256 words of memory. Each 16-bit word is represented by four hexadecimal digits. Blocks are read by columns (top to bottom), where each column contains sixteen 16-bit words in hexadecimal format numbered 0 through F. There are 16 columns in each block, numbered 0 through F. Repetitive data is not suppressed and zero words are represented by hyphens (—). The PP's R register is displayed preceding the first line of the dump.

Format:

`PX, n1, n2, . . . , nm.`

Parameter Description

<code>n_i</code>	Number of PP to be printed. If omitted, all PPs are printed.
----------------------------	--

Q — Dump PP Memory in Octal Line Format

The Q directive causes PP memory to be dumped in line format with 6-bit display code character equivalents. Each line contains 16 bytes (PP words) printed in two sets of 8 bytes in octal format. Each set consists of an address, 8 bytes, and display code character equivalents. Repetitive lines are suppressed and zero bytes are represented by hyphens (—). Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

$Q, n_1, n_2, \dots, n_m.$

Parameter Description	
n_i	Number of PP to be dumped. If omitted, all PPs are dumped.

QF — Dump FLPP Memory in Octal Line Format

The QF directive causes FLPP² memory to be dumped in line format with 6-bit display code character equivalents. Each line contains 16 bytes printed in two sets of 8 bytes in octal format. Each set consists of an address, 8 data bytes, and 6-bit display code character equivalents. Repetitive lines are suppressed and zero bytes are represented by hyphens (—).

Format:

$QF, n_1, n_2, \dots, n_m.$

Parameter Description	
n_i	Number of FLPP to be dumped. If omitted, all FLPPs are dumped.

2. Hardware manuals define peripheral processors making up a peripheral processor subsystem as PPs, and the first-level peripheral processors as peripheral processing units (PPUs). In this manual, first-level peripheral processors are referred to as FLPPs. FLPPs are available only on model 176.

QOA — Dump 16-Bit PP Memory in Octal/ASCII Line Format

The QOA directive is valid on CYBER 180-class machines only. It prints PP memory in line format with 7-bit ASCII code character representation. Each line contains 16 words printed in two sets of eight. Each set consists of an octal address, eight words (six digits per word if the upper 4 bits are nonzero, four digits otherwise), and 7-bit ASCII code character equivalents. Repetitive lines are suppressed and zero words are represented by hyphens (—). The PP's R register is displayed preceding the first line of the dump. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

QOA, n_1, n_2, \dots, n_m .

Parameter Description

n_1 Number of PP to be printed. If omitted, all PPs are printed.

QOD — Dump 16-Bit PP Memory in Octal/Display Line Format

The QOD directive is valid on CYBER 180-class machines only. It prints PP memory in line format with 6-bit display code character representation. Each line contains 16 words printed in two sets of eight. Each set consists of an octal address, eight words (six digits per word if the upper 4 bits are nonzero, four digits otherwise), and 6-bit display code character equivalents. Repetitive lines are suppressed and zero words are represented by hyphens (—). The PP's R register is displayed preceding the first line of the dump. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

QOD, n_1, n_2, \dots, n_m .

Parameter Description

n_1 Number of PP to be printed. If omitted, all PPs are printed.

QXA — Dump 16-Bit PP Memory in Hexadecimal/ASCII Line Format

The QXA directive is valid on CYBER 180-class machines only. It prints PP memory in line format with 7-bit ASCII code character representation. Each line contains 16 words printed in two sets of eight. Each set consists of a hexadecimal address, eight words (four hexadecimal digits per word), and 7-bit ASCII code character equivalents. Repetitive lines are suppressed and zero words are represented by hyphens (—). The PP's R register is displayed preceding the first line of the dump. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

`QXA, n1, n2, . . . , nm.`

Parameter Description

`ni` Number of PP to be printed. If omitted, all PPs are printed.

QXD — Dump 16-Bit PP Memory in Hexadecimal/Display Line Format

The QXD directive is valid on CYBER 180-class machines only. It prints PP memory in line format with 6-bit display code character representation. Each line contains 16 words printed in two sets of eight. Each set consists of a hexadecimal address, eight words (four hexadecimal digits per word), and 6-bit display code character equivalents. Repetitive lines are suppressed and zero words are represented by hyphens (—). The PP's R register is displayed preceding the first line of the dump. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

`QXD, n1, n2, . . . , nm.`

Parameter Description

`ni` Number of PP to be printed. If omitted, all PPs are printed.

CMR Dump Directives

The CMR dump directives provide the capability to selectively dump specified areas of central memory resident. Refer to the NOS Version 2 Systems Programmer's Instant for detailed illustrations of central memory resident.

Successful use of the CMR dump directives is dependent upon the integrity of central memory resident at the time EDD was performed. Most important is the integrity of the CMR pointers on the EDD file. If these pointers are not intact, the dump produced by DSDI may prove meaningless. Thus, if it is suspected that the CMR pointers are not intact, specifying the P keyword on the DSDI command allows DSDI to use the CMR pointers from the running system. This option should be used only when the configuration of the running system is the same as the system in use at the time the EDD file was created. If the CMR pointers on the EDD file are not intact, the integrity of the other areas of central memory is also questionable. In this case, the output produced by the CMR dump directives may be unpredictable.

ACCOUNT — Dump Account Dayfile Buffer

The ACCOUNT directive causes the account dayfile pointers and buffer to be dumped in word format (four words per line) with 6-bit display code character equivalents. This format is the same format as that for the E memory dump directive. This directive also dumps the buffer in a line-by-line format, as on the DSD A display.

Format:

```
ACCOUNT.
```

B — Dump Brief Control Point Status

The B directive dumps all the control points in a format similar to the DSD B,A display.

Format:

```
B.
```

CBT — Dump Control Buffer Table

The CBT directive causes all control buffer table entries to be dumped unless ordinal numbers are specified, then only the specified control buffer table entries are dumped. The dump is formatted to reflect the appropriate parameter fields for each word. Control buffer table entries that are output to a printer are listed in two columns. The format is the same format as that for the MST directive. The control buffer table entries can also be output to a terminal, but only when ordinal numbers are specified. Control buffer table entries that are output to a terminal are listed in a single column format.

Format:

`CBT, n1, n2, . . . , nm.`

Parameter Description

n_i	Ordinal number of the control buffer table entry to be dumped. If n_i is omitted, all control buffer table entries are dumped.
-------	--

CCT — Dump Channel Control Table

The CCT directive causes all nonzero channel control table entries to be dumped unless ordinal numbers are specified, then only the specified channel control table entries are dumped. The dump is formatted to reflect the appropriate parameter fields for each word. Channel control table entries that are output to a printer are listed in two columns. The format is the same format as that for the MST directive. The channel control table entries can also be output to a terminal, but only when ordinal numbers are specified. Channel control table entries that are output to a terminal are listed in a single column format.

Format:

`CCT, n1, n2, . . . , nm.`

Parameter Description

n_i	Ordinal number of the channel control table entry to be dumped. If n_i is omitted, all nonzero channel control table entries are dumped.
-------	--

CP — Dump Active Control Point Areas

The CP directive causes all active control point areas, or a selected subset, to be dumped. List options provide the capability to dump only desired portions of the control point area and other control point related data. Specified without any parameters, the CP directive dumps all active control point areas and pseudo-control point areas. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

CP.

or

CP, n_1 /ops₁, n_2 /ops₂, . . . , n_m /ops_m.

Parameter Description

n_i	Control point number.
ops _i	List options; a string of up to 10 characters indicating the portion of the control point area, or control point related data, to be dumped. If list options are specified, they apply only to the control point number with which they are associated. Valid options are:

ops _i	Description
X	Dumps the exchange package, parameter summary, control point's entry in the executing job table, and contents of the memory locations indicated by the values in each exchange package register. These register values are treated as relative addresses within the field length; invalid addresses are represented with contents of zero. The current command and special entry point parameters are not printed for the system control point. Refer to Printer Output Listing Examples later in this section for a sample of the printer output listing produced.
T	Provides a detailed dump of the control point area with a description of each parameter field and NOSTEXT symbol for each word.
A ³	Dumps the job dayfile pointers and buffer in word format with 6-bit display code character equivalents. This format is the same format as the E memory dump directive. This option also dumps the buffer in a line-by-line format, as on the DSD A display.
F ³	Provides a dump of the 3-word FNT entry, EST, and mass storage track chain, if one exists, for all files attached to the specified control point.

3. This option does not dump information for the system control point.

Parameter Description

ops;_i	<i>(Continued)</i>
ops;_i	Description
C ⁴	Dumps the field length of a specified control point in instruction parcel format with 6-bit display code character equivalents. This format is the same format as the C memory dump directive. Repetitive data is suppressed.
D ⁴	Dumps the field length of a specified control point in byte format with 6-bit display code character equivalents. This format is the same format as the D memory dump directive. Repetitive data is suppressed.
E ⁴	Dumps the field length of a specified control point in word format with 6-bit display code character equivalents. This format is the same format as the E memory dump directive.
G	Dumps the control point area in instruction parcel format with 6-bit display code character equivalents. This format is the same format as the C memory dump directive.
H	Dumps the control point area in byte format with 6-bit display code character equivalents. This format is the same format as the D memory dump directive.
I	Dumps the control point area in word format with 6-bit display code character equivalents. This format is the same format as the E memory dump directive.
M ⁴	Dumps the user-accessible extended memory of a specified control point in byte format with 6-bit display code character equivalents. This format is the same format as the D memory dump directive.
N ⁴	Dumps the negative field length of a specified control point. This format is the same format as the D memory dump directive.
P	Provides a dump and analysis of all active PPs associated with control point n.
default	If n is specified with no corresponding list options, options XTAF are selected automatically. Refer to Printer Output Listing Examples later in this section for a sample of the printer output listing produced. The default options selected can be changed with the CPO directive.

4. This option does not dump information for the system control point.

All of the list options specified for a particular control point area (up to 10) are processed. For example, if the C and D options are both specified, the control point field length is dumped twice, once in instruction parcel (C) format and again in byte (D) format.

CPO — Reset Default List Options

The CPO directive selects a new string of default list options for the CP directive.

Format:

CPO, ops .

Parameter	Description
ops	New default list options for CP directive, a string of up to 10 characters. Refer to the description of the CP directive for a list of valid option characters.

CT — Dump Channel Tables

The CT directive causes the channel status table, EJT assignment table, and channel controlware table to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

CT.

DAYFILE — Dump System Dayfile Buffer

The DAYFILE directive causes the system dayfile pointers and buffer to be dumped in word format (four words per line) with 6-bit display code character equivalents. This format is the same format as that for the E memory dump directive. This directive also dumps the buffer in a line-by-line format, as on the DSD A display.

Format:

DAYFILE .

DB — Dump Disk Buffers

The DB directive causes all disk buffers, or a selected subset, to be dumped to a printer in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

DB, id, n₁, n₂, . . . , n_m .

Parameter	Description
-----------	-------------

id	Specifies the machine identifier of the mainframe from which to dump the disk buffers. If id is omitted, the default machine identifier is taken from CMR of the mainframe that is dumping the disk buffers.
n _i	Number assigned to a disk buffer to be dumped. If n _i is omitted, all active disk buffers are dumped.

DBW — Dump Buffered Device/Buffer Statistics/PP-I/O Buffer Link Tables

The DBW directive causes the list control words from the buffered device table, the buffer statistics table, and the PP-I/O buffer link table to be dumped. The dump is formatted to reflect the appropriate parameter fields for each word. Output to a printer is listed in two columns. Output to a terminal is listed in one column. The format is the same format as that for the MST directive.

Format:

DBW.

DDB — Dump Dayfile Dump Buffer

The DDB directive causes the dayfile dump buffer to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

DDB.

DP — Dump Dayfile Buffer Pointers

The DP directive causes the dayfile buffer pointers to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

DP.

EICB — Dump Environment Interface Communication Buffer

The EICB directive causes the environment interface (EI) communication buffer to be dumped. The dump is formatted to reflect the appropriate parameter fields and COMSDFT symbols for each word. This directive also dumps the dedicated fault tolerance/operating system (DFT/OS) buffer in instruction parcel format (four groups of four hexadecimal digits with 7-bit ASCII code character equivalents).

EJT — Dump Executing Job Table

The EJT directive provides a detailed dump of each entry in the executing job table, including a description of each parameter field and the NOSTEXT symbol for each word.

Format:

EJT.

EPB — Dump Extended Memory/PP Buffer

The EPB directive causes the extended memory/PP buffer to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

EPB.

ERRLOG — Dump Error Log Buffer

The ERRLOG directive causes the error log pointers and buffer to be dumped in word format (four words per line) with 6-bit display code character equivalents. This format is the same format as that for the E memory dump directive. This directive also dumps the buffer in a line-by-line format, as on the DSD A display.

Format:

ERRLOG.

EST — Dump Equipment Status Table

The EST directive provides a detailed dump of each entry in the equipment status table, with a description of each entry and the NOSTEXT symbol for each word.

Format:

EST.

FNT — Dump System File Name/File Status Table

The FNT directive provides a detailed dump of each entry in the system file name/file status table, with a description of each entry and the NOSTEXT symbol for each word.

Format:

FNT.

FOT — Dump Family Ordinal Table

The FOT directive causes the family ordinal table to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

FOT.

HAT — Dump Hash Table

The HAT directive causes all nonzero hash table entries to be dumped unless ordinal numbers are specified, in which case only the specified hash table entries are dumped. The dump is formatted to reflect the appropriate parameter fields for each word. Hash table entries that are output to a printer are listed in two columns. The format is the same format as that for the MST directive. The hash table entries can also be output to a terminal, but only when ordinal numbers are specified. Hash table entries that are output to a terminal are listed in a single-column format.

Format:

HAT, n_1, n_2, \dots, n_m .

Parameter	Description
n_i	Ordinal number of the hash table entry to be dumped. If n_i is omitted, all nonzero hash table entries are dumped.

JC — Dump Job Control Area for Each Service Class

The JC directive causes the job control area for each job service class to be dumped. The dump is formatted to reflect the appropriate parameter fields and NOSTEXT symbol for each word. This directive also dumps the service class control table (SCT) in a line-by-line format.

Format:

JC.

LC — Dump Low Central Memory

The LC directive causes DSDI to dump the contents of low central memory (that is, central memory locations 0 through 177₈). Each word is divided into the appropriate parameter fields. Each field is listed on a separate line with a description of the parameter. The absolute address and 6-bit display code character equivalents are also listed for each word.

Format:

LC.

LDIS — Dump L-Display Buffer

The LDIS directive causes the L-display buffer to be dumped in instruction parcel format with 6-bit display code character equivalents. This format is the same format as that for the C memory dump directive.

Format:

```
LDIS.
```

LIDT — Dump Logical Identifier Table

The LIDT directive causes the logical identifier table to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

```
LIDT.
```

MAINLOG — Dump Binary Maintenance Log Buffer

The MAINLOG directive dumps the binary maintenance log pointers and buffer in word format, four words per line, with 6-bit display code character equivalents.

Format:

```
MAINLOG.
```

MCT — Dump Memory Control Table

The MCT directive dumps the contents of the memory control table (MCT). For each control point and pseudo-control point, there are two lines. The first line displays the control point (or pseudo-control point) having the next lower reference address, the one having the next higher reference address, the unassigned CM FI/100₈ between the control point (or pseudo-control point) having the next higher reference address, and the job's reference address, field length, negative field length, and one byte of recovery information. The second line displays the corresponding information for extended memory.

Format:

```
MCT.
```

MST — Dump Mass Storage/Track Reservation Tables

The MST directive dumps the MST (Mass Storage Table) and the TRT (Track Reservation Table) for the specified device(s). If no devices are specified, the MST directive dumps all mass storage devices. If the specified device is a shared device, the MRT (Mass Storage Recovery Table) is also dumped. If the device is a shared buffered device, the BST (Buffer Status Table) is also dumped.

The dump format for the mass storage tables reflects the appropriate parameter fields and NOSTEXT symbol of each word. The portion of the dump describing the track reservation tables is presented in byte format with 6-bit display code character equivalents. The track link byte ordinal and status bits (three groups of 4 bits) are indicated for each word. The MRT and the BST, if present, are dumped in a format similar to the TRT. Refer to Printer Output Listing Examples later in this section for a sample of the printer output listing produced by this directive.

Format:

```
MST,est1,est2,...,estn.
```

Parameter Description	
est _i	EST ordinal of equipment whose mass storage table is to be dumped. If est _i is omitted, all mass storage and track reservation tables are dumped.

MTR — Dump CPU Monitor

The MTR directive causes the CPU monitor to be dumped. Exchange packages are dumped in exchange package format. The program area is dumped, using relative addressing in instruction parcel format with 6-bit display code character equivalents. This format is the same format as that for the C memory dump directive.

Format:

```
MTR.
```

MTRQUEUE — Dump CPUMTR Request and Recall Queues

The MTRQUEUE directive causes all CPUMTR request and recall queues, or a selected subset, to be dumped.

Format:

```
MTRQUEUE,idents.
```

Parameter	Description
idents	A string of up to four characters indicating the queues to be dumped. The queue identifiers can be specified in any order and are dumped in the order specified. Default is to dump all four queues. Valid queue identifiers are:
idents	Queue Name
C	CPUCIO request.
P	PP request.
R	Recall.
W	CPU request.

ODIS — Dump Operator Display Buffer

The ODIS directive causes the operator display buffer (used for communication between an operator and a remote diagnostic user) to be dumped in instruction parcel format with 6-bit display code character equivalents. This format is the same format as that for the C memory dump directive.

Format:

```
ODIS.
```

PCP — Dump Pseudo-Control Point

The PCP directive causes all active pseudo-control point areas, or a selected subset, to be dumped. The format is the same as that for the CP control point dump directive.

PLD — Dump Peripheral Library Directory

The PLD directive causes the PP library directory to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

PLD.

PP — Dump PP Communication Areas

The PP directive causes the contents of all PP communication areas to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive. The control point assignment, channel assignment, and monitor function are listed with each communication area. Refer to Printer Output Listing Examples later in this section for a sample of the printer output listing produced by this directive. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

PP.

PROBE — Dump PROBE Data Tables

The PROBE directive causes the PROBE data tables to be dumped in byte format with 6-bit display code character equivalents.

Format:

PROBE.

PST — Dump Program Status Table and Entry Point Directory

The PST directive causes the program status table and the entry point directory to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

PST.

PUT — Dump Physical Unit Table

The PUT directive causes all nonzero physical unit table entries to be dumped unless ordinal numbers are specified, then only the specified physical unit table entries are dumped. The dump is formatted to reflect the appropriate parameter fields for each word. Physical unit table entries that are output to a printer are listed in two columns. The format is the same format as that for the MST directive. The physical unit table entries can also be output to a terminal, but only when ordinal numbers are specified. Physical unit table entries that are output to a terminal are listed in a single column format.

Format:

```
PUT, n1, n2, . . . , nm.
```

Parameter Description	
n _i	Ordinal number of the physical unit table entry to be dumped. If n _i is omitted, all nonzero physical unit table entries are dumped.

QFT — Dump Queued File Table

The QFT directive provides a detailed dump of each entry in the queued file table, with a description of each parameter field and the NOSTEXT symbol for each word.

Format:

```
QFT.
```

RCL — Dump Resident Central Library

The RCL directive causes the resident CPU library to be dumped in instruction parcel format with 6-bit display code character equivalents. This format is the same format as that for the C memory dump directive. A warning is printed if the resident central library is empty.

Format:

```
RCL.
```


RPL — Dump Resident Peripheral Library

The RPL directive causes the resident PP library to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive. A header line for each PP program dumped indicates the name of the program and its length in bytes. Each succeeding line contains 10 bytes (two central memory words) of the PP program. The PP address of the first byte in each line, relative to address zero of the PP, is also listed.

Format:

RPL.

SAB — Dump System Attribute Block

The SAB directive causes the system attribute block to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

SDA — Dump Extended Statistical Data Area

The SDA directive causes the extended statistical data area to be dumped. Each word is divided into the appropriate parameter fields and is listed with a description of the parameter. The absolute address and 6-bit display code character equivalents are also listed for each word.

Format:

SDA.

SECDED — Dump SECDED Identifier Table

The SECDED directive causes the single error correction double error detection (SECDED) identifier table to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

SECDED.

SST — Dump Subsystem Control Point/Subsystem Assignment Tables

The SST directive causes the subsystem control point table and subsystem assignment table to be dumped in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

Format:

SST.

TBDUMP — Dump Trace Buffer

The TBDUMP directive dumps the specified trace buffer. The functions traced are listed, followed by the absolute and control point area addresses that were saved, and the buffer IN pointer. The trace buffer data is displayed in the same format as that for the D memory dump directive.

Format:

TBDUMP, MTR.

or

TBDUMP, PPU.

Parameter	Description
MTR	MTR to CPUMTR function trace buffer.
PPU	PP to CPUMTR function trace buffer.

Subsystem Dump/Analysis Directives

The subsystem dump/analysis directives provide the capability to selectively dump portions of central and/or PP memory associated with a specific subsystem (BIO, IAF, MAG, and RHF). Although many other directives already described in this section can be used to dump the same areas of memory, these directives dump those areas most frequently analyzed when subsystem-related malfunctions occur. In addition, many of the dumps are specially formatted to provide a detailed description of the area being dumped. If the specified subsystem was not active at the time the EDD file was created, an error message is issued.

BATCHIO (BIO) — Dump Associated Memory for Analysis

The BATCHIO directive causes the areas of central and/or PP memory that are most frequently analyzed when BATCHIO malfunctions are indicated to be dumped. The areas and type of memory dumped are determined by the list options specified.

Formats:

BATCHIO, ops .

or

BIO, ops .

Parameter	Description
ops	List options; a string of characters indicating the areas of memory to be dumped. If no options are specified, all three options (PBN) are selected by default and are processed in order as listed.
ops	Description
P	Provides analysis and full memory dump of PPs having resident copies of 1CD, 1IO, QAP, QAC, or DSP. The output listing generated is the same (in format and content) as that produced by entering the AP directive in this format: AP, 1CD, 1IO, QAP, QAC, DSP. Refer to the description of the AP directive earlier in this section for additional information.
B	Provides specially formatted dumps of each active BATCHIO buffer point. Included with the dump of each buffer point is the associated equipment type and FET, as well as EST and FNT/FST entries.
N	Provides a dump of the negative field length associated with the BATCHIO control point in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.

IAF — Dump Associated Memory for Analysis

The IAF directive causes the areas of central and/or PP memory that are most frequently analyzed when IAF malfunctions to be dumped. The areas and type of memory are determined by the list options specified. The IAF current entry word (SSPA) is always printed at the beginning of the listing, in byte format, regardless of which list options are specified.

Format:

IAF, ops .

Parameter	Description
ops	List options; a string of up to five characters indicating the areas of memory to be dumped. If no options are specified, four options (ETLP) are selected by default and processed in order as listed.
ops	Description
C	Provides analysis of the IAF command table.
E	Dumps the IAF reentry table in byte format (two words per line) with 6-bit display code character equivalents. The first word in each line is preceded by its ordinal within the table.
T	Provides a specially formatted dump of the IAF terminal table in which each word reflects the appropriate parameter fields. The message status table entry is included when dumping network terminal tables. In addition, each word is preceded by a description of the parameter fields and its COMSREM symbol. Terminal table entries that are empty, except for having status bit 0 in word 3 (VROT entry) set, are not printed.
L	Provides a dump of pot link table and all pots. The pot link table is dumped in byte format with pot link byte ordinals indicated for each word, but no 6-bit display code character equivalents. Repetitive pot link table entries are suppressed. The pots are dumped in word format, three lines per pot, with the first line containing only the pot number.
P	Provides analysis and full memory dump of all PPs having resident copies of TLX, 1TA, ITN, and 1TO. This option also provides an analysis and dump of all PPs having resident copies of 1RO and 1RI that are associated with control points of interactive origin. The output listing generated is the same (in format and content) as that produced by entering the AP directive in this format:

AP, TLX, 1TA, 1TN, 1TO, 1RO, 1RI .

The exception is that the AP directive also dumps all PPs having copies of 1RO and 1RI rather than only those associated with control points of interactive origin. Refer to the description of the AP directive, earlier in this section, for additional information.

MAGNET (MAG) — Dump Associated Memory for Analysis

The MAGNET directive causes the areas of central and/or PP memory that are most frequently analyzed when a malfunction within the Magnetic Tape Subsystem is indicated to be dumped. The areas and type of memory dumped are determined by the list options specified.

Format:

MAGNET, ops .

or

MAG, ops .

Parameter Description

ops List options; a string of characters indicating the areas of memory to be dumped. If no options are specified, all four options (UQSP) are selected by default and are processed in order as listed.

ops Description

- U** Provides a specially formatted dump of the Magnetic Tape Subsystem unit descriptor tables with associated FET, EST, and FNT/FST. Each word of a unit descriptor table is formatted to reflect appropriate parameter fields. In addition, each word is preceded by a description of the parameter fields and its COMSMTX symbol. If extended labels are present, they appear with the FET in the output listing. The FET also indicates the address and control point number of the user.
- Q** Provides a dump of the Magnetic Tape Subsystem queue table in byte format (two words per line) with 6-bit display code character equivalents. The first word in each line is preceded by its ordinal within the table.
- S** Provides a dump of the Magnetic Tape Subsystem active stage job table, staging tape VSN list, and stage request table, if they are defined. The tables are dumped in byte format (two words per line) with 6-bit display code character equivalents. The first word in each line is preceded by its ordinal within the table.
- P** Provides analysis and full memory dump of all PPs having resident copies of 1MT. The output listing generated is the same (in format and content) as that produced by entering the AP directive in this format:

AP, 1MT.

Refer to the description of the AP directive earlier in this section for additional information.

RHF — Dump Associated Memory for Analysis

The RHF directive causes the areas of central and/or PP memory that are most frequently analyzed when RHF malfunctions are indicated to be dumped. The areas and type of memory dumped are determined by the list options specified.

Format:

RHF, ops.

Parameter	Description
ops	List options; a string of characters indicating the areas of memory to be dumped. If no options are specified, all three options (ACP) are selected by default and are processed in order as listed.
ops	Description
A	Provides a dump of the RHF dayfile buffer in word format with 6-bit display code character equivalents. This format is the same format as that for the E memory dump directive. This option also dumps a standard dayfile.
C	Provides a dump of the RHF field length in byte format with 6-bit display code character equivalents. This format is the same format as that for the D memory dump directive.
P	Provides analysis and full memory dump of all active PPs associated with the control point.

Hardware Register Dump Directives

The hardware register dump directives provide the capability to dump specified hardware registers.

FMFREG — Dump PP Registers

The FMFREG directive is valid on CYBER 180-class machines only. For each PP register (P, Q, K, A) stored while processing a fatal mainframe error, DSDI prints the PP register contents.

Format:

```
FMFREG.
```

IOUCR — Dump Concurrent PP Channel Registers

The IOUCR directive is valid only on CYBER machines that support concurrent channels and concurrent PPs. For each concurrent PP channel specified, DSDI prints the channel number, channel type, IOU status register contents, T register contents, fault status mask, and test mode operand generator. The channel number and channel type are octal; all of the other fields are hexadecimal.

Format:

```
IOUCR, ch1, ch2, . . . , chn.
```

Parameter	Description
ch _i	Channel number of the concurrent PP channel register to be printed. If omitted, all concurrent PP channel registers are printed.

IOUMR — Dump IOU Maintenance Registers

The IOUMR directive is valid on CYBER 180-class machines only. For each input/output unit (IOU) maintenance register specified, DSDI prints the hexadecimal register number, the register contents, and the register description. For registers containing error indicators, DSDI prints a description of each error.

Format:

```
IOUMR, first, last.
```

Parameter	Description
first	First register (hexadecimal) to be printed. If omitted, printing begins with register 00.
last	Last register +1 (hexadecimal) to be printed. If omitted, printing ends with register first+1.

If no parameters are specified, DSDI prints all maintenance registers.

LPVA — Load Central Memory PVA into Pseudo Register

The LPVA directive is valid on CYBER 180-class machines only. You use this directive when you are doing an analysis of a NOS/VE dump tape. The dump must contain all memory (not just critical memory), and the FULL option must be specified on the DSDI command. The LPVA directive retrieves a copy of 48 bits starting at the specified address and loads them into the pseudo register named PVA (Process Virtual Address). This value is printed at your terminal and saved for use by subsequent directives.

Format:

```
LPVA, address, offset.
```

Parameter Description

address	Specifies the beginning byte address of the 48 bits loaded into PVA. You can enter this parameter in any of the acceptable NOS/VE address formats. Refer to appendix H for NOS/VE address formats to use with the DSDI utility. This parameter is required.
offset	Specifies the hexadecimal number of bytes from the start-address to begin loading the 48-bit address. This parameter is optional.

MEMMR — Dump Memory Maintenance Registers

The MEMMR directive is valid on CYBER 180-class machines only. For each memory maintenance register specified, DSDI prints the hexadecimal register number, the register contents, and the register description. For registers containing error indicators, DSDI prints a description of each error.

Format:

```
MEMMR, first, last.
```

Parameter Description

first	First register (hexadecimal) to be printed. If omitted, printing begins with register 00.
last	Last register+1 (hexadecimal) to be printed. If omitted, printing ends with register first+1.

If no parameters are specified, DSDI prints all maintenance registers.

PROCA — Dump Processor Operand Cache

The PROCA directive is valid on CYBER 180 model 990 and 995 machines only. This directive prints the contents of the processor operand cache (POC) record from the EDD file. The dump format has three columns. The first column contains the specified word numbers, the second column contains the contents of the corresponding words from the control memory part of the POC record in hexadecimal, and the third column contains the contents of the corresponding words from the data memory part of the POC record in hexadecimal.

Format:

```
PROCA, first, last.
```

Parameter	Description
first	First word number (hexadecimal) to be printed. If omitted, printing begins with word 0.
last	Last word number (hexadecimal) to be printed. If you specify first but omit last, printing ends with word first+1. If you omit both first and last, printing ends with word FFF ₁₆ .

PROCW — Dump Processor Controlware Part Number and Revision Level

The PROCW directive is valid on CYBER 180-class machines only. It prints the processor controlware part number and the revision level. For model 990 and 995 machines, DSDI extracts the controlware revision level and date from the last 256-bit word of the processor control store (PCS) record from the EDD file.

Format:

```
PROCW.
```

PROMR — Dump Processor Maintenance Registers

The PROMR directive is valid on CYBER 180-class machines only. For each processor maintenance register specified, DSDI prints the hexadecimal register number, the register contents, and the register description. For registers containing error indicators, descriptions of each error are printed.

Format:

```
PROMR, first, last.
```

Parameter	Description
first	First register (hexadecimal) to be printed. If omitted, printing begins with register 00.
last	Last register+1 (hexadecimal) to be printed. If omitted, printing ends with register first+1.

If no parameters are specified, all maintenance registers for both processors are printed. If first and last are specified, then the specified registers are dumped for all processors.

PROPM — Dump Processor Page Map

The PROPM directive is valid on CYBER 180 model 990 and 995 machines only. The directive prints the contents of the processor page map (PPM) record from the EDD file. The listing consists of 128 data lines plus column headings. Each line contains a set number (0 to 3), the entry number (0 to $1F_{16}$) within its set, and the contents of that entry in hexadecimal.

Format:

```
PROPM.
```

PRORF — Dump Processor Register File

The PRORF directive is valid on CYBER 180-class machines only. For all CYBER 180-class machines except models 990 and 995, it prints the maintenance channel interface port number followed by the hexadecimal register number and contents of each processor register specified.

Format:

```
PRORF, first, last.
```

Parameter Description

first	First register to be printed. If omitted, printing begins with register 00.
last	Last register+1 to be printed. If omitted, printing ends with register first+1.

If no parameters are specified, the entire register file for both processors is printed. If first and last are specified, then the specified registers are dumped for all processors.

For CYBER 180 model 990 and 995 machines, enter the PRORF directive without parameters. DSDI produces a listing of the entire contents of each of the following records from the EDD file: PRF, PIS, PRG, and PRH. Refer to Printer Output Listing Examples later in this section for a sample of the listing DSDI produces for model 990 and 995 machines.

PROSM — Dump Processor Segment Map

The PROSM directive is valid on CYBER 180 model 990 and 995 machines only. The directive prints the contents of the processor segment map (PSM) record from the EDD file. The listing consists of 64 data lines plus column headings. Each line contains a set number (0 or 1), the entry number (0 to $1F_{16}$) within its set, and the contents of that entry in hexadecimal.

Format:

```
PROSM.
```

PROXP — Dump Processor Exchange Package

The PROXP directive is valid on CYBER 180-class machines only. It prints the maintenance channel interface port number followed by the hexadecimal register number, the contents of each processor register, and (for most registers) a description of the register's contents. The exchange packages for all processors are printed.

Format:

PROXP.

For NOS/VE dumps, use the following format:

PROXP, address.

Parameter	Description
address	Specifies the first byte address of the exchange package. You can enter this parameter in any of the acceptable NOS/VE address formats. Refer to appendix H for NOS/VE address formats to use with the DSDI utility. If you omit this parameter, the current contents of the central processor registers are printed. If the memory address you specify is for the stack frame save area, the only registers printed are the registers recorded in the save area.

SC — Dump S/C Register

The SC directive is valid only on CYBER 170 Computer Systems with the exception of models 815, 825, 835, 845, and 855; it causes the S/C registers (maintenance registers for models 865 and 875) to be dumped.

Format:

SC.

SETCPU — Set CPU Number

The SETCPU directive is valid on any CYBER model that supports 170 mode in both CPUs. It sets up processor pointers for the CPU specified in the directive.

Format:

SETCPU, number.

Parameter	Description
number	Specifies the number of the CPU (0 or 1). Default is 0.

SETIOU — Set PP Dump Defaults

The SETIOU directive presets the parameter defaults for processing PP memory dumps to allow printing of concurrent PP memory.

Format:

SETIOU, option.

Parameter Description

option	Specifies whether to process concurrent PPs or nonconcurrent PPs. If no option is specified, option N is assumed.
--------	---

option Description

C	Process concurrent PPs to allow printing 8192 words of concurrent PP memory. This option is valid only on CYBER machines that support concurrent channels and concurrent PPs.
N	Process nonconcurrent PPs to allow printing 4096 words of PP memory.

SETJPS — Change the JPS Register Value

The SETJPS directive is valid on CYBER 180-class machines only. You use this directive when you are doing an analysis of a NOS/VE dump tape. The dump must contain all memory (not just critical memory), and the FULL option must be specified on the DSDI command. The SETJPS directive sets the pseudo register named JPS (Job Process State) to a value different than the value of the JPS register found in the dump.

Format:

SETJPS, address.

Parameter Description

address	Specifies the location of, or the actual new value for, the JPS register. You can enter this parameter in any of the acceptable NOS/VE address formats. Refer to appendix H for NOS/VE address formats to use with the DSDI utility. If you specify address as a process virtual address (PVA) or system virtual address (SVA), the system loads the value stored at that location into the pseudo JPS register. If you specify address as a real memory address (RMA), the system loads that value into the pseudo JPS register. This parameter is required.
---------	---

SETRMA — Convert Address into RMA

The SETRMA directive is valid on CYBER 180-class machines only. You use this directive when you are doing an analysis of a NOS/VE dump tape. The dump must contain all memory (not just critical memory), and the FULL option must be specified on the DSDI command. The SETRMA directive converts the specified address into a real memory address (RMA) and saves that value in a pseudo register (called RMA) for use in subsequent directives.

Format:

SETRMA, address.

Parameter	Description
address	Specifies the address to be converted to an RMA. You can enter this parameter in any of the acceptable NOS/VE address formats. Refer to appendix H for NOS/VE address formats to use with the DSDI utility. This parameter is required.

SETVEP — Set Virtual Address Parameters

The SETVEP directive is valid on CYBER 180-class machines only. You can use this directive when you are doing an analysis of a NOS/VE dump tape. The dump must contain all memory (not just critical memory), and the FULL option must be specified on the DSDI command. The SETVEP directive changes the virtual memory parameters used for the NOS/VE DSDI memory display directives (I and W).

Format:

SETVEP, address, length, mask, mps.

Parameter	Description
address	Specifies the beginning address of the page table as a hexadecimal number of from 1 to 6 digits. This parameter is required.
length	Specifies the length of the page table as a hexadecimal number of from 1 to 6 digits. This parameter is required.
mask	Specifies the value of the page size mask as a hexadecimal number of from 1 to 6 digits. This parameter is required.
mps	Specifies the contents of the pseudo register named monitor process state (MPS) as a hexadecimal number of from 1 to 6 digits. This parameter is required.

TRACEBK — Traceback Stack Frame Save Area

The TRACEBK directive is valid on CYBER 180-class machines only. You use this directive when you are doing an analysis of a NOS/VE dump tape. The dump must contain all memory (not just critical memory), and the FULL option must be specified on the DSDI command. The TRACEBK directive displays the exchange package found at the specified address; and, for each stack frame, the module name and contents of the registers saved.

Format:

```
TRACEBK, address.
```

Parameter	Description
address	Specifies the address of the exchange package where the traceback starts. You can enter this parameter in any of the acceptable NOS/VE address formats. Refer to appendix H for NOS/VE address formats to use with the DSDI utility. If you omit this parameter, the system uses the exchange package currently in the central processor.

If you enter this directive interactively, the output includes just the minimum save area (the P register and address registers A0, A1, and A2). If you enter this directive in a batch job with line printer output, the output contains the entire stack frame as saved by the task. The actual contents depends on the information saved by the task.

XP - Dump Deadstart Exchange Package

The XP directive causes the CPU exchange package executing at the time of deadstart to be dumped. If there are two CPUs in the system, both exchange packages in execution at the time of deadstart are dumped.

Format:

```
XP.
```

Buffer Controller Directive

The buffer controller (BCDUMP) directive provides the capability to dump specified buffer controllers and list verification information for controlware dump files. The BCDUMP directive can interpret either a dump file produced by the LOADBC utility or a dump file produced by EDD. If the dump file was produced by EDD, the dump buffer controller option must have been selected at the time the dump was taken, and the FULL option must be specified on the DSDI command. The BCDUMP directive causes the selected buffer controllers to be dumped.

Format:

BCDUMP/ops .

or

BCDUMP, cc₁, cc₂, . . . , cc_n .

or

BCDUMP, cc₁/ops₁, cc₂/ops₂, . . . , cc_n/ops_n .

Parameter Description

cc _i	Channel number of buffer controller to be dumped. If no channels are specified, all buffer controllers are dumped.
ops _i	List options; a character string indicating the line format, interpretation, and verification of the selected buffer controllers to be dumped. If no options are specified, options H and D are selected by default.
ops_i	Line Format
H	Hexadecimal line format (default).
O	Octal line format.
Options A and D indicate interpretation.	
ops_i	Interpretation
A	7-bit ASCII code interpretation.
D	6-bit display code interpretation (default).
The following options indicate verification.	
ops_i	Verification
V	List verification information (used with the controlware dump file produced by the LOADBC utility).
V omitted	Do not list verification information.

Interactive Use of DSDI

Incorporated within DSDI is an interactive facility that allows several of the directives already described in this section to be entered interactively from a terminal. This interactive facility is designed to provide the following additional capabilities.

- Allows preliminary examination of the EDD file to determine which areas should be listed in detail at a line printer.
- Allows examination of certain areas of the EDD file not listed during normal operational procedures following a system malfunction. Typically, predefined portions of the EDD file are listed following a system malfunction.
- Allows online examination of the EDD file from a remote location.

Refer to Example of DSDI Terminal Use later in this section for an example showing interactive use of DSDI.

When the DSDI command is entered from an interactive terminal, there will be a 10- to 60-second delay before input directives can be entered. During this delay, DSDI is copying the EDD file to a random-access mass storage file. The length of the delay depends on device speed and current system activity. When DSDI is able to accept input directives, it will issue the following prompt to the terminal.

```
ENTER DIRECTIVES- -  
?
```

Directives are entered following the question mark prompt. Only one directive can be entered at a time, and each directive is restricted to one line. The format is the same as described for batch input (refer to Directive Format earlier in this section).

Generally, any of the DSDI input directives can be entered at an interactive terminal. However, the output produced by many of the directives is formatted for listing only at a line printer (136 columns) and cannot be listed at the terminal (72 columns). The L parameter on the DSDI command initially determines the disposition of the list output. If a file name is not specified, list output is assigned to file OUTPUT by default (that is, the terminal). In this case, entry of directives that produce output that cannot be listed at the terminal results in the message:

```
DIRECTIVE RESTRICTED TO PRINTER OUTPUT.
```

If a list output file name is specified on the DSDI command, all input directives can be entered at the terminal. All list output (including error messages) is written to the specified file.

These two input directives are provided to further control the disposition of list output:

Directive	Description
OUTPUT,filename.	Assigns output to alternate file filename (file name OUTPUT is not allowed; that is, alternate list output cannot be assigned to the terminal). If filename is omitted, the system assumes file name ALTRNT. While this directive is active, all input directives can be entered at the terminal. All list output (except error messages) is written to file filename and is formatted for transmission to a line printer. Error messages are written directly to the terminal.
DISPOSE.	Disposes the alternate list file (specified in the OUTPUT directive) to the print queue. Output will be printed at the central site line printer. All subsequent list output resumes on the original output file specified on the DSDI command.

Refer to File Manipulation and Control Directives earlier in this section for additional information concerning use of these directives.

Terminal Output Directives

The following directives produce output formatted for listing at an interactive terminal.

C — Dump Memory in Instruction Parcel Format

The C directive causes the specified locations of central memory or extended memory to be dumped in four groups of five octal digits (one word per line) with 6-bit display code character equivalents. No pagination is processed for terminal output. The CM, EC, and UEC directives (refer to Central Memory/Extended Memory Dump Directives earlier in this section) determine the type of memory to be dumped; default is central memory. The RA or RAC directive (refer to Central Memory/Extended Memory Dump Directives earlier in this section) must be entered to dump relative addresses; default is absolute addressing.

Format:

C, fwa, lwa.

Parameter	Description
fwa	First word address to be dumped (mandatory).
lwa	Last word address, plus one location, to be dumped. If omitted, fwa+1 is assumed by default.

Figure 6-1 shows the terminal output produced by the C directive.

```
? C, 6230, 6240.
0006230 34240 10100 00012 50036 1TAA AU 3
0006231 00764 70000 00000 10113 * AAK
0006232 04154 70000 00000 10113 DM* AAK
0006233 00004 67446 74000 10005 - - A E
0006234 05153 05700 00000 00000 EMX.
0006235 00000 00000 00000 00000
0006236 00000 00000 00000 00000
0006237 00000 00005 05111 14422 EEII9R
```

Figure 6-1. C Directive Output

CP — Dump Active Control Point Areas

The CP directive causes the job sequence name and control point area address for each control point to be dumped.

Format:

CP.

Figure 6-2 shows the terminal output produced by the CP directive.

CP 01	CP 02	CP 03	CP 04	CP 05	CP 06	CP 07	CP 10
IAF	AALT	NAM		AAKW	AAAG	AAAF	AAAE
0200	0400	0600	1000	1200	1400	1600	2000
CP 11	CP 12	CP 13	CP 14	CP 15	CP 16	CP 17	CP 20
AAAD	AAAC	AAAB	AALN		AALU		
2200	2400	2600	3000	3200	3400	3600	4000
CP 21	CP 22	CP 23	CP 24	CP 25	CP 26	CP 27	CP 30
							RBF
4200	4400	4600	5000	5200	5400	5600	6000
CP 31	CP 32	CP 33	CP 34 ← Control Point Number				
MAG	BID		SYS ← Job Sequence Name at Control Point				
6200	6400	6600	7000 ← Control Point Area Address				

Figure 6-2. CP Directive Output

D — Dump Memory in Byte Format

The D directive causes specified locations of central memory or extended memory to be dumped in five groups of four octal digits (one word per line) with 6-bit display code character equivalents. No pagination is processed for terminal output. The CM, EC, and UEC directives (refer to Central Memory/Extended Memory Dump Directives earlier in this section) determine the type of memory to be dumped; default is central memory. The RA or RAC directive (refer to Central Memory/Extended Memory Dump Directives earlier in this section) must be entered to dump relative addresses; default is absolute addressing.

Format:

D, fwa, lwa.

Parameter	Description
fwa	First word address to be dumped (mandatory).
lwa	Last word address, plus one location, to be dumped. If omitted, fwa+1 is assumed by default.

Figure 6-3 shows the terminal output produced by the D directive.

```
? D,6230,6240.
0006230 3424 0101 0000 0125 0036 1TAA AU 3
0006231 0076 4700 0000 0001 0113 * AAK
0006232 0415 4700 0000 0001 0113 DM* AAK
0006233 0000 4674 4674 0001 0005 - - A E
0006234 0515 3057 0000 0000 0000 EMX.
0006235 0000 0000 0000 0000 0000
0006236 0000 0000 0000 0000 0000
0006237 0000 0000 0505 1111 4422 EEII9R
```

Figure 6-3. D Directive Output

PP — Dump PP Communication Areas

The PP directive causes the PP number, executing program name, control point assignment, and input register address for each PP communication area to be dumped.

Format:

PP.

Figure 6-4 shows the terminal output produced by the PP directive.

										PP Number
Program in Execution						Control Point Assignment				
PP00	PP01	PP02	PP03	PP04	PP05	PP06	PP07	PP10	PP11	
MTR-34	DSD-34	QAC-16	----	----	1R0-05	1SJ-34	1MT-31	----	QAC-14	
7200	7210	7220	7230	7240	7250	7260	7270	7300	7310	
PP20	PP21	PP22	PP23	PP24	PP25	PP26	PP27	PP30	PP31	
PIP-03	CPD-34	QAC-02	----	1I0-32	----	----	----	1MA-01	----	
7320	7330	7340	7350	7360	7370	7400	7410	7420	7430	
PP00										Input Register Address

7440										

Figure 6-4. PP Directive Output

Q — Dump PP Memory in Line Format

The Q directive causes the specified locations of PP memory to be dumped in line format. Each line contains 8 bytes (PP words) with 6-bit display code character equivalents. Repetitive lines are suppressed and zero bytes are represented by hyphens (—).

Format:

Q,n,fwa,lwa.

Parameter	Description
n	Number of PP to be dumped.
fwa	First word address to be dumped.
lwa	Last word address, plus one location, to be dumped.

NOTE

fwa and lwa are automatically adjusted so that the dump limits fall within a multiple of 10₈ words.

This format is valid only for terminal output. If attempted from a job of batch origin or while an alternate list file is active, the fwa and lwa parameter will be interpreted as PP numbers.

Figure 6-5 shows the terminal output produced by the Q directive.

? Q,5,0,100.												
0000	0003	2020	3340	----	0614	----	4334	0117	CPPo5	FL	81A0	
0010	0064	0001	7772	0100	0006	1073	1401	6072	A	A	FH	LA
0020	2250	3225	----	0027	0012	4402	5747	5751	R/ZU	W	J93.*.(
0030	0011	7646	----	0001	0141	0600	----	----	I	-	AA6F	
0040	1501	1116	2014	0074	0203	----	----	----	MAINPL	BC		
0050	3404	2330	0035	6213	1707	----	4000	6675	1DSX 2	KOG	5	
0060	----	4521	----	6101	0001	0153	0001	0532	+Q	A	AAS	AEZ
0070	0001	0100	1000	0003	6000	6250	6251	6252	AA H	C	/	()

Figure 6-5. Q Directive Output

QOA, QOD, QXA, QXD — Dump 16-Bit PP Memory in Line Format

These four directives are valid on CYBER 180-class machines only. Each directive prints specified locations of 16-bit PP memory in line format. Each line contains eight PP words in octal or hexadecimal with 6-bit display code or 7-bit ASCII code character representations. Repetitive lines are suppressed and zero bytes appear as hyphens (—).

Directive Format	Numeric Representation	Character Representation
QOA,n,fwa,lwa,R.	Octal	7-bit ASCII code
QOD,n,fwa,lwa,R.	Octal	6-bit display code
QXA,n,fwa,lwa,R.	Hexadecimal	7-bit ASCII code
QXD,n,fwa,lwa,R.	Hexadecimal	6-bit display code

Parameter Description

n	Number of PP to be printed.
fwa	First word address to be printed.
lwa	Last word address, plus one location, to be printed.
R	If specified, the R register is printed.

This directive format is valid only for terminal output. If it is used in a job of batch origin or while an alternate list file is active, DSDI interprets the fwa and lwa parameters as PP numbers.

Example of DSDI Terminal Use

This example illustrates how DSDI might be used following a system malfunction to analyze portions of the EDD file from an interactive terminal. It is assumed that an EDD file was created during normal system recovery procedures. Vertical spacing has been expanded to permit commentary. The example begins after the login sequence has been completed.

batch	Enter Batch subsystem.
RFL,0.	
/label,dump,vsn=dump,lb=k1,f=s,mt,d=800	The LABEL command is entered to assign the EDD dump tape to this job. Use of the vsn parameter allows the job to be rolled out while the tape is mounted and assigned.
/get,altdir	Retrieve alternate directives file ALTDIR (refer to the example in figure 6-6).
/dsdi.	Calls DSDI, which copies EDD tape to a random mass storage file.
ENTER DIRECTIVES - -	Enter terminal output directives (refer to Terminal Output Directives earlier in this section) to list any portion of the EDD file at the terminal. DSDI issues the prompt (?) when it is ready to accept a new directive.
.	
.	
.	
? output,altout.	List output produced by subsequent directives is written to local file ALTOUT. This allows entry of directives that produce line printer formatted output.
.	
.	
.	
? read,altdir.	All input directives in alternate directives file ALTDIR are read and processed. List output is written to local file ALTOUT. DSDI does not request terminal input until the last directive on ALTDIR is processed.
.	
.	
.	
? dispose.	Dispose local file ALTOUT to the print queue for listing at the central site line printer. Output produced by subsequent directives is listed at the terminal.
EXPRESS DUMP COMPLETE (FL USED	DSDI is terminated by entering a carriage return in response to the ? prompt.
xxxxxxB)	
/	

In summary, the following operations were performed by DSDI. First, small areas of the dump file were listed at the terminal for preliminary examination. This was done both to analyze the cause and effect of the system failure, as well as determine the extent of line printer listings required. An appropriate comment may be placed in the list file subtitle at this time via the *.ccc...ccc directive.

Next, directives were entered to generate the necessary line printer listings. These listings are generally extensive, or contain specially formatted output that cannot be listed at the terminal. Thus, output was written to an alternate list file named ALTOUT.

After all necessary directives had been entered from the terminal, an alternate directives file (ALTDIR) was read. ALTDIR is a permanent file containing input directives necessary to obtain a printer listing of specific areas in the dump file that are frequently examined following a system failure (for example, CMR tables and buffers, PP memory, and so forth). Figure 6-6 illustrates a typical alternate directives file. Comments describing areas of the dump file to be listed appear, along with the directive, in the title line of the output listing.⁵

When DSDI finished processing the last directive in file ALTDIR, it again issued the ? prompt to the terminal requesting further directive input from the keyboard. At this time, the DISPOSE directive was entered causing file ALTOUT to be printed at the central site line printer. Refer to Printer Output Listing Examples later in this section for examples of printer output listings.

```
DUMP
LC.LOW CENTRAL MEMORY POINTERS
PP.PP COMMUNICATION AREA
CPO,H.
EJOFF.
EJ.
CP.CONTROL POINT AREAS
CPO,XTAF.
DP.DAYFILE BUFFER POINTERS
EJ.
EST.EQUIPMENT STATUS TABLE
FNT.SYSTEM FILE NAME TABLE
EJ.
MST.MASS STORAGE TABLES
JC.JOB CONTROL PARAMETERS
CP,30/T.SYSTEM CONTROL POINT AREA
ACCOUNT.ACCOUNT FILE BUFFER
ERRLOG.ERROR LOG BUFFER
EJON.
AP.ACTIVE PPS
```

Figure 6-6. Sample Alternate Directives File

5. All records in an alternate directives file must be in text format: the first word in each record contains only the record name.

Printer Output Listing Examples

The listings illustrated in this section are samples of the line printer output listings produced by several of the input directives described under Input Directives earlier in this section.

Each page of output listings begins with two header lines: a title line and a subtitle line. The header lines are formatted as follows.

```
RA =      current input directive                DUMP nn DSDI   yy/mm/dd. hh.mm.ss.  PAGE  n
0000000 dump type hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM.  NOS 2  comments
```

Header	Description
RA=0000000	Indicates absolute addressing is in effect. If relative addresses were being dumped (RA or RAC directive entered), a nonzero reference address would appear in this field.
current input directive	Input directive currently being processed, including comments (50-character field).
DUMP nn	Reflects the EDD tape number currently being used (nn is a 2-digit dump identifier assigned during EDD).
yy/mm/dd. hh.mm.ss.	This field reflects the date and time when DSDI was executed.
dump type	Indicates type of memory currently being dumped (central memory, extended memory, or a PP number). If the AP directive is processed, the PP number is followed by the name of the program currently being dumped. If the CP or PCP directive is processed, CP or PCP is displayed, followed by the CP/PCP number, followed by CM or EM. If the EJT, EST, FNT, MST or QFT directive is processed, the table name and current entry ordinal is displayed.
h.mm.ss. yy/mm/dd.	Time of day, date, system title line, and system version name taken from CMR.
comments	Up to 36 characters of comments specified on a *.ccc...ccc input directive.

Example 1:

Figure 6-7 illustrates the printer output listing produced by the PP directive (dump PP communication areas).




Figure 6-7. Sample PP Directive Output

Example 2:

Figure 6-8 illustrates the printer output listing produced by the AP directive (dump analysis of PP and PP memory in line format).

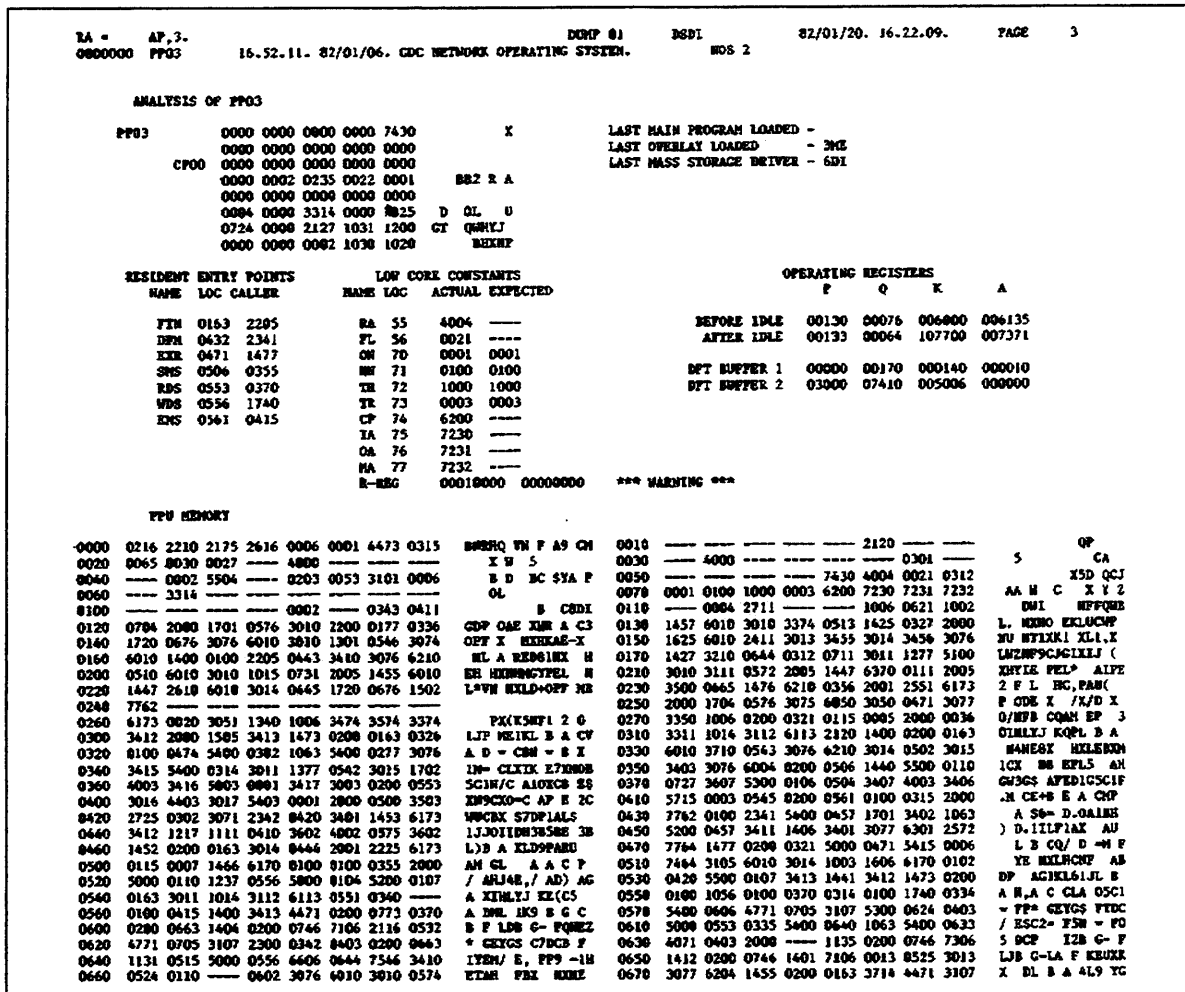


Figure 6-8. Sample AP Directive Output

Example 3:

Figure 6-9 illustrates the printer output listing produced by the CP directive (dump active control points). The default list options (XTAF) are used to dump the control point. This example consists of six pages. Also, notice that the columns cross page boundaries; that is, the left column is read continuously, from the top of the second page to the middle of the fourth page. The sequence then continues at the top of the right column on the second page.

```

RA * CP,2. DUMP 00 DSDI 83/06/02, 22.36.17. PAGE 1
000000 CH 22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM. NOS 2

0400 - CONTROL POINT 02

CPO2 EXCHANGE PACKAGE

P 33557 A0 22577 B0 0 (A0)=0000 0000 0000 0002 2575 BU (B0)=0000 0000 0000 0000 0000
RA 120300 A1 1 B1 3 (A1)=0000 0000 0000 0000 0000 (B1)=3444 4400 0000 0000 0000 199
FL 444500 A2 26250 B2 31053 (A2)=0000 0000 0000 0000 0024 T (B2)=0400 0320 2500 0000 0000 D CPU
EN 70070000 A3 14424 B3 0 (A3)=0000 0000 0000 0004 3521 B20 (B3)=0000 0000 0000 0000 0000
RAE 0 A4 33070 B4 43 (A4)=0000 0000 0000 0003 3244 C19 (B4)=0000 0000 0000 0000 0000
FILE 0 A5 33072 B5 31347 (A5)=0000 0000 0000 0003 3273 C2 (B5)=0400 0320 5100 0000 0000 D CP(
RA 400 A6 1 B6 33646 (A6)=0000 0000 0000 0000 0000 (B6)=5110 0336 3001 0003 3274 CHC3XA C2
0 A7 33070 B7 43515 (A7)=0000 0000 0000 0003 3244 C29 (B7)=0000 0004 0435 0204 3715 002B04H

XD 7777 7777 7700 0000 0000 (X0)=0000 0000 0000 0000 0000
X1 0000 0000 0000 0000 0000 (X1)=0000 0000 0000 0000 0000
X2 0000 0000 0000 0000 0024 T (X2)=0000 0000 0000 0000 0000
X3 0000 0000 0000 0004 3521 B20 (X3)=1000 0150 0000 0000 0000 H A/
X4 0000 0000 0000 0003 3243 C28 (X4)=0400 0321 6630 2205 0501 D CO XRECA
X5 0000 0000 0000 0003 3273 C2 (X5)=0000 0000 0000 0000 0000
X6 2205 1400 0000 0000 0000 RCL (X6)=0000 0000 0000 0000 0000
X7 0000 0000 0000 0003 3244 C29 (X7)=0400 0320 1130 0530 0505 D CP1XEXEC

PARAMETER SUMMARY

USER NAME NETOPS SENSE SWITCHES 00
PRIMARY FILE CCL R1 000000 X List Option
ORIGIN TYPE SYOT CCL R2 000000
TIME USED 0000000222 CCL R5 000000
TIME REMAINING 1014223327 CCL EF 00
SUBSYSTEM ID 7770 CPU PRIORITY 0077
CPU STATUS X SERVICE CLASS XXXX SSSC

MESSAGE 1 - MP/SW102 22.05.09.LL1NK L003103,EN,RL=3,H=N,S
MESSAGE 2 - SREQUEST PR= DISPLAY
CURRENT COMMAND - NIP(MIN=199,ISTP=YES,FSTP=YES,MC=500)
LAST DAYFILE MESSAGE - 22.05.09.MP/SW102 22.05.09.LL1NK L003103,EN,RL=3,H=N,SW102 /103,SW102 /003
SPECIAL ENTRY POINTS -

EJT ENTRY
025000 1601 1555 0137 7777 0002 MAR A4 B
025001 0000 0000 4800 0002 7776 5 B
025002 0000 0000 0000 0000 0000
025003 0700 0000 0000 0000 0036 6 3
    
```

Figure 6-9. Sample CP Directive Output (Continued)

```

BA = CP,2.        DUMP 00   DSDI         83/06/02. 22.36.17.   PAGE   2
000000 CM        22.75.04. 83/06/02. CDC NETWORK OPERATING SYSTEM NOS 2

          Beginning of T List Option
CONTROL POINT AREA

CPU STATUS          STSN 000420 5 /A A | RESERVED            BFCM 000475 00000000
CPU SUB-CP ACTIVE  0 01 | CONTROL STATEMENT ADDRESS 000000
NUMBER PPS ASSIGNED | DEMAND FILE RANDOM INDEX 000000
ERROR FLAGS        0 0000 | RESERVED            ALRM 000476 0077 000000
RESERVED           0 0 | MAX MAGNETIC TAPES 7 7 7 7 7 7 7 7
SUSPEND ON ROLLOUT FLAG | MAX REMOVABLE PACKS 7 7 7 7 7 7 7 7
ROLLOUT CONTROL   0 00 | MAX DEFER BATCH JOBS 7 7 7 7 7 7 7 7
RESERVED           0 0000 | RESERVED            7 7 7 7 7 7 7 7
RECALL CONTROL    JCIW 000421 0176 0001 A JA | MAX TIME LIMIT 7 7 7 7 7 7 7 7
RESERVED           0 0 | MAX SBU LIMIT 7 7 7 7 7 7 7 7
CPU SELECTION     0 7770 | MAX FIELD LENGTH 7 7 7 7 7 7 7 7
SUBSYSTEM ID      0 0 | MAX ECS FIELD LENGTH 7 7 7 7 7 7 7 7
JOB CONTROL/BDIS* FLAGS | MAX LINES PRINTED 7 7 7 7 7 7 7 7
USER PRIVACY, PRESERVE ECS | MAX CARDS PUNCHED 7 7 7 7 7 7 7 7
HMB* LINKAGE      0 001221 C C44B R | RESERVED            ACLM 000477 0000 7777
CPU RECALL CRITERION | DAYFILE MESSAGES COUNT 7777
JOB TERMINATION OPTION | CONTROL STATEMENTS CNT 7777
QUEUE FILE DISPOSITION | RESERVED            77
RESERVED           0 00000 | RS PRUS COMM 777777
RESV/1 STATEMENT FLAG | ACCOUNT ACCESS WORD AACH 000500 00000000000083777777
CHANGE MEN OP/CCL LEVEL | RSP IJL NAT. WPM 000501 0000002 P C
RESERVED           0 0000 | RESERVED            000000
JOB ORIGIN TYPE   JOTM 0000 0000 | MAP IJL ACCUMULATOR MPAN 000502 00000000000000000000
OPER ASSIGNED EQUIP OAEW 0000 0000 | AUC1/2/3 ACCUMULATORS AUTM 000502 00000000000000000000
NEGATIVE FL/1000 FLSW 000424 0004 0000 B JCB | AUC4/5 ACCUMULATORS AM2W 000503 00000000000000000000
ASETH RECOVERY MFL 0000 00001203 | RESERVED            TCAV 000504 0000 000000
BA/1000           00001203 | LENGTH BUFFER 0 00000000
FL/1000           00001203 0466 | ADDRESS BUFFER 0 000000 0000
ZERO             00001177 4004 | LENGTH BUFFER 1 00000000
BA/1000 - MFL/1000 | ADDRESS BUFFER 1 SEPM 000505 0 000000
4000B * MFL/1000 00001177 4004 | SPEC ENTRY POINT FLAG 000505 0
FL/1000          00001177 4004 | RSVB/LRB*/CLB* EN PT FLAGS 0
RESERVED          0 00000225 BU | ARG*/MFL*/SSM* EN PT FLAGS 0
ECS BA/1000      00000225 | SSJA*/MFL*/SSM* EN PT FLAGS 0
ECS FL/1000      00000225 0000 | RESERVED            0
PUNCH MODE/OVERRIDE | RESTART/SUPPRESS BWP* FLAGS 0
RESERVED          0 00000000 | DR* FILE FLAGS 0
RESERVED/INSTALLATION | DR* FILE FLAGS 0
PP PAUSE FLAGS   0 0000 | DR* FILE FLAG (0 = ENTIRE FL) 0
SENSE SWITCHES  0 00 | SSJA* PARAMETER BLK ADDRESS 000000
RESERVED          0 00 | SYSTEM PROC CALL WORD SPCU 000506 00000000000000000000 A
MESSAGE 1 AREA   W51W 000430 16205025103433355555 NP/SNT02 | CCL - RFG JCDW 000507 00 000000
000431 5535357334057334457 22.05.09. | CCL - DATA 0001
000432 14141116130025163333 LLINK L00 | RESERVED            00000000
000433 36343334560514562234 3103,EN,ML | CCL - EF JCRV 000510 00 000000
000434 56345610461456230800 =3,M-N,S | CCL - R5 000000
MESSAGE 2 AREA   W62W 000435 53220521250523245547 $REQUEST * | CCL - R2 000000
000436 13475504112320140131 K* DISPLAY | CCL - R1 000000
000437 00000000000000000000 | SEC NER/K BIS FLAGS W64W 000511 0 AS645=AS=
INSTALLATION WORD 0 I88W 000440 00000000000000000000 | ECS AND CH CHN MORE FLAGS 0
INSTALLATION WORD 1 IN7W 000441 00000000000000000000 | INPUT BUFFER ADDRESS 014043
INSTALLATION WORD 2 IN2W 000442 00000000000000000000 | RIGHT SCREEN BFR AND 014054
INSTALLATION WORD 3 IN5W 000443 00000000000000000000 | LEFT SCREEN BFR AND 014054
MAP OPTIONS/LIB FLAG LB1W 000512 00
    
```

Figure 6-9. Sample CP Directive Output (Continued)

RA = CP,2.		22.15.04. 03/06/02. C0C NETWORK OPERATING SYSTEM		DUMP 00	DSDI	03/06/02. 22.36.17.	PAGE 3
00000000 ON					NOS 2		
INSTALLATION WORD 4	INAM 000444	000000000000000000				REMOVE/LOADER FLAG	00
INSTALLATION WORD 5	INSM 000445	000000000000000000				RESERVED FOR LOADER	0000
INSTALLATION WORD 6	INSH 000446	000000000000000000				INTERACTIVE DEMON CONTROL	0000
INSTALLATION WORD 7	INSH 000447	000000000000000000				GLOBAL LIB INDICATORS	00000000
LIMIT FLAGS	ACTM 000450	00		WWR		SECOND LIB/GLOBAL END	LB2M 000513 000000000000000000
OVERFLOW FLAG	SRM	0000				FIRST LIB/GLOBAL END	LB3M 000514 000000000000000000
SRU ACCUMULATOR/OVERFLOW		00001651271722				PRINT ADDR LAST FLD EXEC	EDCM 000515 0000
CP ACCUM (NUMPOINTS/4)	CPTM 000451	00000001074853556000		AGS		ECS FL FOR SMP= CALL	PPM
EO ACCUMULATORS	IOAM 000452	00314540000000000416		Y+S	BN	FL FOR SMP= CALL	0000
R13 = R1 + R3	RP1V 000453	000036		3	3	B	
R14 = R1 + R4		000036					
RESERVED							
ADDER ACCUMULATOR	ADAM	0000002					
R1 = 1000	RP2M 000454	023420		STP	O/		
R12 = R1 + R2		001750					
RESERVED		00000000					
DISABLE SRU ACCUM, CPN	RP3M 000455	0000025700		B.	7X		
EOH		0000004230					
SRU ACCOUNT BLOCK LMT	STLM 000456	777777					
COMPUTED SRU JOB STEP LIMIT		777777777777					
RESERVED	SRJM 000457	0000		DCO			
SRU JOB STEP LIMIT		777777					
SRU ACCUM AT JOB STEP START		0000040304					
RESERVED	CPJM 000460	0000		H			
CP TIME JOB STEP LIMIT		777777					
CP ACCUM AT JOB STEP START		0000000016					
CHANGE/ACCOUNTING FLAG	PPM 000461	00					
RESERVED		00					
SRU VALIDATION LIMIT		7777					
FW ORIGINAL OF PROFILE FILE		0000					
LEVEL-3 BLOCK TRACK		0000					
LEVEL-3 BLOCK SECTOR		0000					
RESERVED	000462	000000000000000000					
MAX FL FOR JOB STEP	FLCM 000463	3765		4	F	4	
LAST CARD FL (GNL)		0600					
MAXIMUM FL FOR ENTIRE JOB		3765					
RESERVED		00000000					
JOB STEP MAX ECS FL	ELCM 000464	0200		B	B		
LAST CARD ECS FL (GNL)		0000					
JOB MAX ECS FL		0200					
RESERVED		00000000					
EQUIP ASSIGNED COUNT	BACH 000465	0001		A			
RESERVED	FLTM	000					
SCHEDULER STATUS		0					
NEGATIVE FL FOR ROLLIN		0000					
ECS FL REQUEST		0000					
ON FL REQUEST		0000					
TRM SUBSYSTEM	TRSM 000466	00					
RESERVED		00					
TERMINAL NUMBER	TTM	0000					
TERMINAL INT ADDR	TIAM	000000					
OUTPUT POINTER	TIOM	000000					
RESERVED	PPCM 000467	000000000000		H			
EST END OF FAMILY DEVICE		0010					
LIMIT FOR SIZE OF BAF		7					
LIMIT FOR NUMBER OF PF		7					
LIMIT COWN SIZE IAPP		7					

Figure 6-9. Sample CP Directive Output (Continued)

```

BA = CP,2.                DUMP 00          DSDI          83/06/02. 22.36.17.    PAGE 4
00000000 CH              22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM      NOS 2

LIMIT FOR SIZE OF IAPP          7          RECALL REQUESTS IR-S          000553 00000000000000000000
USER NAME          UJWH 000470 16052417202300          METOPS 4          RECALL REQUESTS IR-S          000554 00000000000000000000
CHARGE FLAG, USER INDEX          377772          RECALL REQUESTS IR-S          000555 00000000000000000000
NO EXIT FLAG          EECM 000471 2          P          B50          RECALL REQUESTS IR-S          000556 00000000000000000000
RESERVED          000          RECALL REQUESTS IR-S          000557 00000000000000000000
RETRIEVE DATA          0077          RECALL REQUESTS IR-S          000560 00000000000000000000
TERMINAL INPUT POINTER TIME          000000          RECALL REQUESTS IR-S          000561 00000000000000000000
RETRIEVE DATA          024070          RECALL REQUESTS IR-S          000562 00000000000000000000
EJT ORDINAL OF JOB          TFSM 000472 0017          0          RECALL REQUESTS IR-S          000563 00000000000000000000
PRIMARY FILE FMT OFFSET          0000          RECALL REQUESTS IR-S          000564 00000000000000000000
RESERVED          00          RECALL REQUESTS IR-S          000565 00000000000000000000
ROLLOUT TIME          TERM          000          RECALL REQUESTS IR-S          000566 00000000000000000000
EVENTS DESCRIP TOR          CSFN 000475 00          00000000          B          IAT          RECALL REQUESTS IR-S          000567 00000000000000000000
LOGOUT EP/LOG, EP/LOG REA          0          RECALL REQUESTS IR-S          000570 00000000000000000000
PR CHARGE/USER/DIMBIT DEC          2          RECALL REQUESTS IR-S          000571 00000000000000000000
GCR FLAG/CS COUNT          00000011          RECALL REQUESTS IR-S          000572 00000000000000000000
NEXT STATEMENT INDEX          0134          RECALL REQUESTS IR-S          000573 00000000000000000000
LIMIT INDEX          0065          RECALL REQUESTS IR-S          000574 00000000000000000000
INPUT/SKIP FLAGS          CSSM 000474 0          0065          F7A7A B          RECALL REQUESTS IR-S          000575 00000000000000000000
EST ORDINAL          006          RECALL REQUESTS IR-S          000576 00000000000000000000
FIRST TRACK          4201          RECALL REQUESTS IR-S          000577 00000000000000000000
COMMENT TRACK          4201          T List Option resumes at top of
CURRENT SECTION          0002          right column on sheet 2 of 6.
OVERLAP WORD COUNT          0000

----- DAYFILE POINTERS AND BUFFER ----- A List Option

00120275 0011 7777 0000 0000 0046 I          -
00120276 0000 0000 0114 0000 0104          NL AD
00120277 0007 4200 4200 0010 0003 67 7 M C

00117777 5535557334057354157 55232520052226112311 1716550701116050400 00000000000000000000 22.05.06. SUPERVISION GAINED
00120000 5535 3567334057334157 16205023163433355555 5535357334057334157 5503020552605222311 22.05.06.NP/SMT02 22.05.06. CCP VERSI
00120007 17160555353636565514 05260514005540433756 55260122110116240055 4044370000000000000000 00 33, LEVEL 594, VARIANT 594
00120015 5535557334057334157 16205023163433355555 5535357334057334157 55202205261117252355 22.05.06.NP/SMT02 22.05.06. PREVIOUS
00120017 08235516170405009533 33335455202205261117 25235574235516170405 00553333300000000000 00 CS MODE 000, PREVIOUS NS MODE 000
00120025 5535557334057334157 16205023163433355555 5535357334057334157 14141116130055163333 22.05.09.NP/SMT02 22.05.09.LINK L00
00120027 3634333640516562214 36365610461656231634 3335555503333365623 1634333555550343336 5403,EN,NS,M,N,SMT02 /003,SMT02 /103
00120033 00000000000000000000 5535557334057334157 16205023163433355555 5535357334057334157 5535357334057334157 22.05.09.NP/SMT02 22.05.09.
00120037 14141116130055163333 36343336560516562214 54365610461656231634 3335555503433365623 55353573337373657 SMT02 /003 NS/ 22.04.43.
00120043 163433355555555555 000000000000000000 44550317152014052405 04570000000000000000 SMT02, DUMP NP00199 COMPLETED.
00120047 23163433355555555555 25152055162033333444 44550317152014052405 23163433355555555555 22.04.43.NS/ 22.04.43.SMT02 , L
00120053 55353573337373657 162350555555555555 55353573337373657 23163433355555555555 22.04.43.NS/ 22.04.46.NS/
00120057 17010411160755511511 055555555555555555 55353573337373657 162350555555555555 000000000000000000 0000000000000000 22.05.09.NP/SMT02 22.05.09.
00120063 55353573337373657 23163433355555555555 17010411160755511511 07161355525252500000 22.04.46.SMT02 , LOADING (BTGLN ),
00120067 5535357333737404157 162350555555555555 5535357333737404157 23163433355555555555 22.04.46.NS/ 22.04.36.SMT02 , L
00120073 17010411160755511511 025555555555555555 5535357333737404157 162350555555555555 0000000000000000 0000000000000000 22.04.47.NS/
00120077 5535357333737404157 23163433355555555555 17010411160755511511 24050457000000000000 22.04.57.SMT02 , LOAD COMPLETED.
00120105 5535357334057334157 16205023163433355555 5535357334057334157 162050231634333555 22.05.06.NP/SMT02 22.05.06.NPU SMT02
00120107 55355601035634333600 00000000000000000000 5535357334057334157 16205023163433355555 ,AC,NLS 22.05.06.NP/SMT02

DAYFILE LINES IN BUFFER

NS/ 22.04.43.SMT02 , DUMP NP00199 COMPLETED.
22.04.43.NS/ 22.04.43.SMT02 , LOADING (MIC ) .
22.04.46.NS/ 22.04.46.SMT02 , LOADING (BTGLN ) .
    
```

Figure 6-9. Sample CP Directive Output (Continued)

RA #	CP,2.	DUMP 00	DSDI	83/06/02. 22.36.18.	PAGE	5
0000000	CR	22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM	MOS 2			
22.04.56.NS/		22.04.56.SN102 , LOADING UNCB)				
22.04.57.NS/		22.04.57.SN102 , LOAD COMPLETED.				
22.05.06.NP/SN102		22.05.06.NP SN102 ,AC,103				
22.05.06.NP/SN102		22.05.06. SUPERVISION GAINED				
22.05.06.NP/SN102		22.05.06. CCP VERSION 33, LEVEL 594, VARIANT 594				
22.05.06.NP/SN102		22.05.06. PREVIOUS CS MODE 000, PREVIOUS NS MODE 000				
22.05.09.NP/SN102		22.05.09.LLJNK L003103,EM,RL=3,H=N,SN102 /003,SN102 /103				
22.05.09.NP/SN102		22.05.09.LLJNK L003103,EM,RL=3,H=N,SN102 /103,SN102 /003				
ATTACHED FILES ← F List Option						
0304	INPUT	IN	FMT - 1116 2025 2400 0001 1700 INPUT AO FST - 0006 4175 4175 0002 0005 F6 6 B E TRACK CHAIN - 4175 0002	FUT 0000 0000 0000 0000 0000 EST 7700 6210 0013 1377 0421 1277 H KK D&J		
0312	NMR	LO	FMT - 1601 1500 0000 0000 1500 NMR N FST - 0006 4176 4176 0004 0005 F6 6 D E TRACK CHAIN - 4176 0005	FUT 0000 0000 0000 0000 0000 EST 7700 6210 0013 1377 0421 1277 H KK D&J		
0315	OUTFIL	LO	FMT - 1725 2406 1114 0000 1500 OUTFIL N FST - 0007 4205 4205 0001 0105 67E7E AAE TRACK CHAIN - 4205 0003	FUT 0000 0000 0000 0000 0000 EST 7702 6210 0013 1377 0421 1302 H KK D&KB		
0320	ZZZZZC0	LI	FMT - 3232 3232 3203 3301 1007 ZZZZZC0M6 FST - 0006 4175 4175 0001 0005 F6 6 A E TRACK CHAIN - 4175 0002	FUT 0000 0000 0000 0000 0000 EST 7700 6210 0013 1377 0421 1277 H KK D&J		
0323	ZZZZZC2	LO	FMT - 3232 3232 3203 3500 1907 ZZZZZC2 M6 FST - 0006 4177 4177 0002 0307 F6 6 BCG TRACK CHAIN - 4177 0002	FUT 0000 0000 0000 0000 0000 EST 7700 6210 0013 1377 0421 1277 H KK D&J		
0326	OUTPUT	PN	FMT - 1725 2420 2524 0000 1200 OUTPUT J FST - 0010 6742 6742 0001 0005 H 7 7 A E TRACK CHAIN - 6742 0001	FUT 0000 0000 0700 0000 0000 6 EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE		
0331	NRF1	LO	FMT - 1622 0634 0000 0000 1500 NRF1 N FST - 0007 4210 4210 0002 0705 67N7N B6E TRACK CHAIN - 4210 0003	FUT 0000 0000 0000 0000 0000 EST 7702 6210 0013 1377 0421 1302 H KK D&KB		
0334	NRF2	LO	FMT - 1622 0635 0000 0000 1500 NRF2 N FST - 0006 4200 4200 0001 0705 F7 7 AGE TRACK CHAIN - 4200 0003	FUT 0000 0000 0000 0000 0000 EST 7700 6210 0013 1377 0421 1277 H KK D&J		
0337	LIST	PN	FMT - 1411 2324 0000 0000 1200 LIST J FST - 0010 7052 7052 0001 0005 H J J A E TRACK CHAIN - 7052 0001	FUT 0000 0000 0700 0000 0000 6 EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE		
0342	ZZZZZPP	PN	FMT - 3232 3232 3220 2000 1200 ZZZZZPP J	FUT 0000 0000 0700 0000 0000 6		

RA #	CP,2.	DUMP 00	DSDI	83/06/02. 22.36.18.	PAGE	6
0000000	CR	22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM	MOS 2			
		FST - 0010 7053 7053 0001 0005 H S B A E TRACK CHAIN - 7053 0001		EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE		
0345	ZZZZMB	PN	FMT - 3232 3232 0415 0200 1200 ZZZZMB J FST - 0010 7054 7054 0001 0005 H = = A E TRACK CHAIN - 7054 0001	FUT 0000 0000 0700 0000 0000 6 EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE		
0350	ZZZZZM	LO	FMT - 3232 3232 3204 1600 1500 ZZZZZM N FST - 0007 4216 4216 0022 0703 67N7N B6C TRACK CHAIN - 4216 0022	FUT 0000 0000 0000 0000 0000 EST 7702 6210 0013 1377 0421 1302 H KK D&KB		

Figure 6-9. Sample CP Directive Output

Example 4:

Figure 6-10 illustrates the printer output listing produced by the MST directive (dump mass storage/track reservation table). The MST is listed in two columns. The left column is read from top to bottom, perhaps across page boundaries, and continues at the top of the right column. The track reservation table is listed in single column following the MST.

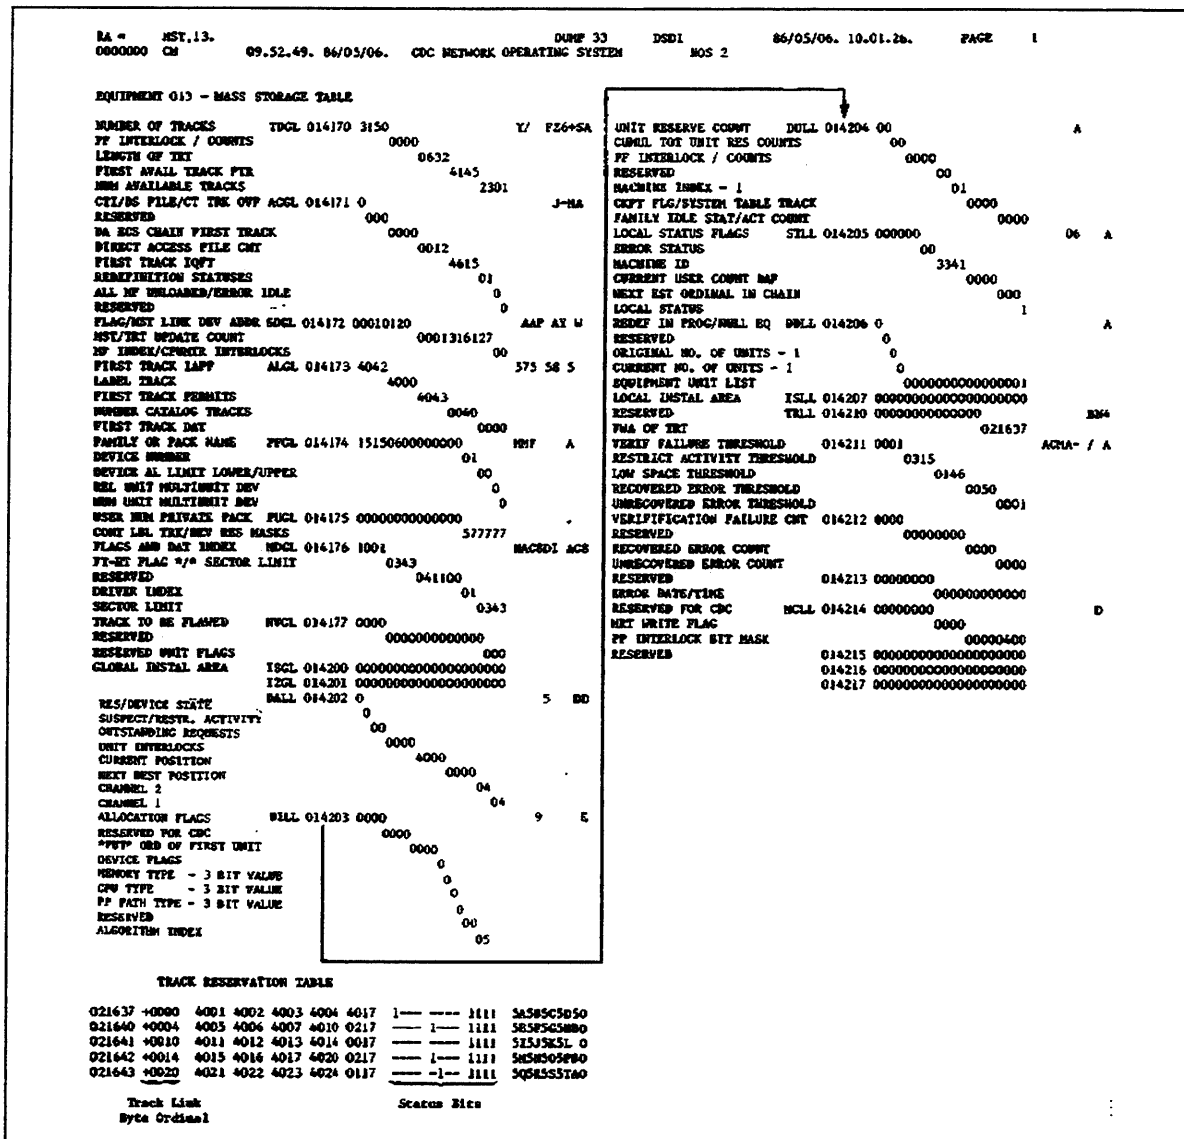


Figure 6-10. Sample MST Directive Output

Example 5:

Figure 6-11 illustrates the printer output listing produced by the C, D, and E memory dump directives (instruction parcel, byte, and word format, respectively). The same portions of central memory are dumped in each format. Auto page eject has been disabled using the EJOFF directive to allow listing the output from all three memory dump directives on one page.

```

RA = C,50,110. C - FORMAT DUMP. DUMP 01 DSDI 82/01/20. 16,22,27. PAGE 17
000000 CH 16,32,11. 82/01/06. CDC NETWORK OPERATING SYSTEM. NOS 2

000050 0000 0000 0000 0000 0000 0000061 7777 0000 0000 0000 0000 0000072 74647 56475 12000 00000 J
----- 0000062 0000 0000 01065 34137 AFS64 0000073 00033 36401 00007 30005 CO A E
000052 0000 0000 0000 07240 5 0000063 0000 0000 0000 0000 0000 0000074 00033 56404 00035 00025 C2 D C/ U
000053 30766 01030 10057 43001 X HXHE XA 0000064 0000 0000 0000 51426 ELV 0000075 00035 56406 00057 50002 C F E B
000054 34137 44234 12147 30200 1KL71JL B 0000065 00042 71100 0000 0000 0000 0000076 00000 00000 00000 00000
000055 01633 01401 00034 17761 A XLA C6 0000066 0000 0000 0000 0000 0000 0000077 00007 75400 00400 07200 = 5
000056 00000 00000 00000 07164 ----- 0000100 00000 00000 00000 00000
000057 77775 12230 00320 00000 (RX Z 0000070 0000 03300 24040 40000 0 TDP
000060 00751 10000 00000 00000 I 0000071 74500 00122 25000 07421 / ARU Q 0000107 00054 15700 01065 34117 E64 AFS60

CH B,50,110. D - FORMAT DUMP.

000050 0000 0000 0000 0000 0000 0000061 7777 0000 0000 0000 0000 0000072 7464 7564 7512 0000 0000 J
----- 0000062 0000 0000 0001 0653 4137 AFS64 0000073 0003 3564 0100 0073 0005 CO A E
000052 0000 0000 0000 0000 7240 5 0000063 0000 0000 0000 0000 0000 0000074 0003 3564 0400 0350 0025 C2 D C/ U
000053 3076 6070 3070 0574 3001 X HXHE XA 0000064 0000 0000 0000 0005 1426 ELV 0000075 0003 3564 0600 0575 0002 C F E B
000054 3413 1442 3412 1473 0200 1KL71JL B 0000065 0004 2711 0000 0000 0000 0000076 0000 0000 0000 0000
000055 0163 3074 0100 0341 7761 A XLA C6 0000066 0000 0000 0000 0000 0000 0000077 0000 7754 0000 4000 7200 = 5
000056 0000 0000 0000 0000 7164 ----- 0000100 0000 0000 0000 0000 0000
000057 7777 5122 3000 3200 0000 (RX Z 0000070 0000 0333 0024 0404 0000 0 TDP
000060 0075 1100 0000 0000 0000 I 0000071 7450 0001 2225 0000 7421 / ARU Q 0000107 0005 4137 0001 0653 4117 E64 AFS60

CH E,50,110. E - FORMAT DUMP.

000050 000000000000000000 000000000000000000 0000000000000007240 30766010301005743001 5X HXHE XA
000054 34131442341214730200 01433014010003417761 0000000000000007164 77775122300032000000 1KL71JL B A XLA C6 CRX 2
000060 00751100000000000000 77770000000000000000 00000000000106534137 00000000000000000000 I AFS64
000064 00000000000000051426 00042711000000000000 00000000000000000000 00000000000000000000 ELV DWI
000070 00000330024040400000 74500001222500007421 74647564751200000000 0003364010000730005 0 TDP / ARU Q J CO A E
000074 00033564040003500025 00035564040005750002 00000000000000000000 00007754000040007200 C2 D C/ U C F E B = 5
000100 00000000000000000000 00000000000000000000 00000000000000000000 00000000000000000000
000104 00000000000000000000 00000000000000000000 00000000000000000000 00054137000106534117 E64 AFS60
    
```

Figure 6-11. Sample C Directive Output

Example 6:

Figure 6-12 illustrates the printer output listing of the system file name table produced by the FNT dump directive. This table is printed in the same format as that produced by the D memory dump option (refer to example 5).

```

RA = FNT. DUMP 00 DS01 89/02/09. 15.13.46. PAGE 13
0000000 FNT 0000 14.36.03. 88/06/21.1641 CY170-730 SN964 NOS DEVELOPMENT. NOS 71F2/20R4/5GT.

SYSTEM FILENAME TABLE

FNT 0000
FILE NAME FNTG 025374 23312324051500 SYSTEM H EST ORDINAL FSTG 025375 0006 F5E
RES/FAT INDEX 00 FIRST TRACK 4005
FILE TYPE 10 USER COUNT FOR MOD/APP/RD 00000000
RESERVED/INTERLOCK 00 ATTACH MODE/UPDATE COUNT 0000

FNT 0001
FILE NAME FNTG 025376 22233004413700 R5MD64 K EST ORDINAL FSTG 025377 0077 5M
RES/FAT INDEX 00 FIRST TRACK 4016
FILE TYPE 13 USER COUNT FOR MOD/APP/RD 00000000
RESERVED/INTERLOCK 00 ATTACH MODE/UPDATE COUNT 0000

FNT 0002
FILE NAME FNTG 025400 22233026413700 R5XV64 K EST ORDINAL FSTG 025401 0077 5M
RES/FAT INDEX 00 FIRST TRACK 4015
FILE TYPE 13 USER COUNT FOR MOD/APP/RD 00000000
RESERVED/INTERLOCK 00 ATTACH MODE/UPDATE COUNT 0000

FNT 0003
FILE NAME FNTG 025402 20221706111403 PRDFILC K EST ORDINAL FSTG 025403 0077
RES/FAT INDEX 00 FIRST TRACK 6445
FILE TYPE 13 USER COUNT FOR MOD/APP/RD 00000000
RESERVED/INTERLOCK 00 ATTACH MODE/UPDATE COUNT 0000

FNT 0004
FILE NAME FNTG 025404 26011411042523 VALIDUS K EST ORDINAL FSTG 025405 0077
RES/FAT INDEX 00 FIRST TRACK 6674
FILE TYPE 13 USER COUNT FOR MOD/APP/RD 00000000
RESERVED/INTERLOCK 00 ATTACH MODE/UPDATE COUNT 0000
    
```

Figure 6-12. Sample FNT Directive Output

Example 7:

Figure 6-13 illustrates the printer output listing produced by the PRORF directive for CYBER 180 model 990 and 995 machines.

```

18A = PRORF. DUMP 00 DS01 85/06/07. 17.23.56. PAGE 1
0000000 CH 09.06.46. 85/07/29. CDC NETWORK OPERATING SYSTEM NOS 2.

RADIAL NCI = 1

CPU CURRENT INSTRUCTION REGISTER          INSTRUCTION BUFFER STACK
0 0000 0800 0800 2000                      0 C700 9322 0800 AAD4
1 0000 2000 0000 0078                      1 FB10 28A6 AF32 F250
2 0000 0000 0000 0000                      2 9812 9812 1000 0004
3 0000 0000 0000 0078                      .
4 FFFF 0000 F0FF E70F                      .
5 FFFF F3BC 10FF E803                      .
6 0003 E000 00FF E049                      62 0042 225A 8C40 4224
7 3CFF 0000 00FD E000                      63 44C3 E8CA BE66 CE10

EXECUTING WORDS OF SOFT CONTROL MEMORIES    REGISTER UNIT
ACW R2 0000 2000 0000 0000                  A0 0000 1084 0000 6700
ACW R3 0048 0000 0000 0000                  A1 0000 FFFF 880D 0000
ACW R4 0000 0000 0000 0000                  A2 0000 FFFF 8000 0000
BBP 0000 0000 0000 0000                      .
0000 0000 0000 0000                          .
EPM SCR 0000 4900 0001 C200                  SKE 07FF 03F2 4E41 5500
EPM ETM 0000 0000 0100 F200                  SIF 07FF 03F2 4E41 5500
I00 CS 0000 0000 0000 0000
I00 R4P 0000 0000 0000 0000
I00 IDS FFFF FFFF FFFF FFFF                    HISTORY FILE
LSW 0000 0000 0000 0000                      0 X-REG 0000 0000 0000 0000
SWA 00 0000 0000 8FFF F019                    0 A-REG 1000 0000 003E
0000

ERROR INFORMATION TABLE
0 0000 0000 0000 F200
1 0000 0000 0101 F210
2 0000 0000 0202 F220
.
.
.
14 0000 0000 DE0E F2ED
15 0000 0000 DF0F F2FD
F X-REG 0000 0000 0000 0000
F A-REG 1000 0000 0000
F P, H 0002 4680 8900
    
```

Figure 6-13. Sample PRORF Directive Output

NOS allows you to deadstart from a disk rather than from a tape. To deadstart from a disk, you must first use the **INSTALL** command to install a deadstart file that you specify to a deadstart disk. The disk that you select must be available on a single-unit device such as an 844 or 885-11/12,¹ and must have the Common Test and Initialization (CTI) software. (For more information about CTI, refer to the CYBER Initialization Package User's Handbook.) If the deadstart disk is a shared device in a multimainframe environment, a full **INITIALIZE** must have been previously done. The calling job must be of system origin or you must be validated for system origin privileges and the system must be in debug mode. In addition, in a secured system, **INSTALL** is allowed only from jobs with security administrator privileges.

You must re-execute the **INSTALL** command if a full **INITIALIZE** becomes necessary, because **PFDUMP** and **PFLoad** do not dump or load the system deadstart file (SDF).

The format of the command is:

```
INSTALL, filename, EQest.
```

Parameter Description

filename	Disk file (assigned to the job) to be installed as an SDF. Default is SYSTEM , which must be assigned to the job. The name SDF is a reserved file name and cannot be specified for filename.
----------	---

INSTALL stops reading the system deadstart file at the first EOF; therefore, the new running system does not include any changes made through **SYSEDIT** after a previous deadstart.

If the deadstart file to be installed is on tape, it must first be copied to disk before **INSTALL** can be used.

NOTE

The deadstart file to be installed must contain a valid OPLD directory. Note that when you use **LIBEDIT** to update a deadstart file and write the new deadstart file directly to tape, the OPLD record on the tape will not contain valid disk addresses. Therefore, if you want the deadstart file on tape and you intend to later use the **INSTALL** command on the deadstart file, you should have **LIBEDIT** write the new deadstart file to disk, and then copy the file to tape using the **COPY** command (refer to the NOS 2 Reference Set, Volume 3).

est	EST ordinal of the deadstart disk on which file filename is to be installed.
-----	--

1. The system deadstart file cannot be installed to an 819, 885-42, 887, or 9853 disk.

FLAW K Display	8-3
INITIALIZE K Display	8-8
Machine Recovery (MREC) Utility K Display	8-16
MREC Procedures	8-17
Using MREC with the K Display	8-17
Using MREC without the K Display	8-22
MREC Unit and Controller Reservations	8-22
Mass Storage Extended Subsystem (MSE) K Display	8-24
Network Access Method (NAM) K Display	8-26
Display Control Characters	8-28
NAM K-Display Operation	8-29
NAM Mode Maintenance Commands	8-31
DB Command	8-31
DE Command	8-32
DU Command	8-32
FL Command	8-33
LB Command	8-33
LE Command	8-34
LR Command	8-34
RS Command	8-35
ST Command	8-36
Queue File Transfer Facility (QTF) K Display	8-40
QTF Selection Classes	8-40
File Selection Process	8-40
Visualizing the File Selection Process	8-41
Queued File Size Index	8-43
Queued File Size and Transfer Times	8-43
QTF K-Display Usage	8-44
QTF Operation Under RHF	8-44
QTF Operation Under NAM	8-45
QTF K Display Commands	8-46
CLASS Command	8-47
DISABLE Command	8-48
ENABLE Command	8-48
HELP Command	8-49
IDLE Command	8-49
INCLUDE Command	8-50
SCHED Command	8-51
STATUS Command	8-52
QTF Physical Identifier K Display	8-54
QTF Selection Class K Display	8-56
QTF Transfer K Display	8-57
STOP Command	8-58
REDEFINE K Display	8-59

Remote Batch Facility (RBF) K Display	8-70
Remote Host Facility (RHF) K Display	8-73
RHF Initiation	8-74
Operator Interface	8-74
RHF Commands Under K display	8-76
Application Table Display	8-77
Network Path Table Display	8-78
RHF Commands Available Under Application or Network Path Display	8-81
RHF Termination	8-81
SCOPE 2 Station Facility (SSF) K Displays	8-82
Operator Interface	8-82
Station Login	8-82
Enabling and Disabling File Transfers	8-82
Station Recovery	8-83
Station Disconnection and Logout	8-83
File Transfer Limit Commands	8-84
Staged File Transfer Commands	8-84
Spooled File Transfer Commands	8-85
Transaction Facility (TAF) K Displays	8-86
TAF Initialization K Display	8-86
Restarting TAF K Display	8-90
TAF Status K Display	8-90
TAF K-Display Commands	8-92
TAF/CRM Status K Displays and Commands	8-96

This section documents the K displays listed next, along with the utilities used to present them. Other K-display utilities described elsewhere in this manual include the permanent file utilities, queued file utilities, and the status/control register simulator. K-display utilities described in other manuals include those for the CDCNET network (refer to the CDCNET network manuals) and the printer support utility (refer to the NOS Version 2 Operations Handbook).

K Display	Description
FLAW	Disk flaw mapping.
INITIALIZE	Disk initialization.
Machine Recovery (MREC)	Multimainframe device recovery.
Mass Storage Extended Subsystem (MSE)	K display of the MSE.
Network Access Method (NAM)	Operator interface to NAM.
Queue File Transfer Facility (QTF)	K display of the QTF application.
REDEFINE	Online reconfiguration display.
Remote Batch Facility (RBF)	K display of the RBF Subsystem.
Remote Host Facility (RHF)	K display of the RHF Subsystem.
SCOPE 2 Station Facility (SSF)	K display of the SSF Subsystem.
Transaction Facility (TAF)	K display of the TAF Subsystem.

By using the K display, a job can place information on the console screen and receive information from the keyboard. The information is passed to the job by DSD. Normally, these displays are used for utility programs.

The job first issues a request message on the B,0 display, asking you to bring up the K display.

When this happens, enter:

K, jsn.

jsn	Description
jsn	Job sequence name of the requesting job.

Once the display is assigned to a job, you can enter data by typing K. followed by data. The data is transferred to a specified area of the job's field length when you terminate the entry. If more than 50 characters are entered as data, the message

LINE TOO LONG.

appears on the screen. DSD does not accept the entry until the data string is shortened.

K displays are job oriented. The job sequence name associated with each K display appears at the top of the screen next to the display designator and name.

All parameter entries must be prefixed by K period (K.). However, after pressing CR or NEXT for the first parameter entry, everything but the K. is erased. This allows another parameter to be entered without entering K. first. All examples in this section show K. although you may not have to enter it. If you have to enter a DSD command during parameter entry, backspace to erase the K., enter the command, and then continue parameter entry by entering K. and the parameter.

FLAW K Display

The FLAW utility reserves (flaws) tracks on any mass storage device during normal system operation. Each entry identifies an area of mass storage that is unusable (flawed area) and prevents the system from accessing it. Since 881 and 883 disk packs normally contain flaw information in the utility sector, the FLAW utility should be used on an 881 or 883 only to specify additional areas not currently in the utility sector. Obtain flaw addresses from the customer engineer.

Flawing tracks on mass storage devices is accomplished using the FLAW K displays (figures 8-1 and 8-2).

```

K. FLAW.                AAAM

                          CURRENT SLF, CLF, SPF, AND CPF ENTRIES

EQ = 0                   EST ORDINAL OF DEVICE.
NUM VALUES ENTERED     TRT OCTAL VALUES * = DUPLICATE

1
2
3
4
5
6
7
8
9
10
11
12
13
14

ENTER EST ORDINAL.

```

Figure 8-1. FLAW K Display (Left Screen)

```

K.  FLAW.                AAAM

MASS STORAGE DEVICE FLAWING

SLF, CLF, SPF AND CPF ENTRIES

SLF = NNNN      SET LOGICAL TRACK FLAW.
CLF = NNNN      CLEAR LOGICAL TRACK FLAW.
SPF = XX        SET PHYSICAL TRACK FLAW.
CPF = XX        CLEAR PHYSICAL TRACK FLAW.

WHERE
  NNNN = VALID LOGICAL TRACK NUMBER.
  XX   = C000,T00,S00 (CYLINDER,TRACK,SECTOR.)
        = A0000000 (ADDR/10B, FOR DE/DP ONLY.)

NUMERIC VALUES ARE CONSIDERED OCTAL UNLESS THE
DIGITS 8 OR 9 APPEAR IN THE NUMBER, OR A *D* IS
APPENDED TO THE NUMBER.

SEE LEFT SCREEN FOR CURRENT ENTRIES,
UP TO 14 ENTRIES ALLOWED
    
```

Figure 8-2. FLAW K Display (Right Screen)

All console entry is performed under DSD control. The sequence is:

1. Call the FLAW utility by entering:

```
X.FLAW.
```

2. Bring the K display to the left console screen by entering:

```
K,jsn.
```

jsn	Description
jsn	Job sequence name of the job requesting the K display (the message REQUEST *K* DISPLAY appears on the B,0 display).

3. You can bring the flawing instructions up on the right screen by entering KK.

4. Specify the mass storage device on which tracks are to be flawed. Enter:

```
K.EQ=est.
```

est	Description
est	EST ordinal of the mass storage device.

5. Enter flaws. A maximum of 14 flaw entries is allowed for each call to the FLAW utility. In addition, there are four types of flaw entries that may be specified. The general format for the K display entry is:

```
K.xtk=ta.
```

Variable	Description										
xtk	Specifies one of the following types of flaw entries.										
	<table border="1"> <thead> <tr> <th>xtk</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>SPF</td> <td>Sets the track reservation table (TRT) entry for the specified physical block (track) address(es) in extended memory or disk to indicate that the block is unavailable for use.</td> </tr> <tr> <td>CPF</td> <td>Clears that TRT flaw entry for the specified physical block address(es) in extended memory or disk to indicate that the block is available for use.</td> </tr> <tr> <td>SLF</td> <td>Sets the TRT flaw entry for the specified logical track to indicate that the track is unavailable for use.</td> </tr> <tr> <td>CLF</td> <td>Clears the TRT flaw entry for the specified logical track to indicate that the track is available for use.</td> </tr> </tbody> </table>	xtk	Description	SPF	Sets the track reservation table (TRT) entry for the specified physical block (track) address(es) in extended memory or disk to indicate that the block is unavailable for use.	CPF	Clears that TRT flaw entry for the specified physical block address(es) in extended memory or disk to indicate that the block is available for use.	SLF	Sets the TRT flaw entry for the specified logical track to indicate that the track is unavailable for use.	CLF	Clears the TRT flaw entry for the specified logical track to indicate that the track is available for use.
xtk	Description										
SPF	Sets the track reservation table (TRT) entry for the specified physical block (track) address(es) in extended memory or disk to indicate that the block is unavailable for use.										
CPF	Clears that TRT flaw entry for the specified physical block address(es) in extended memory or disk to indicate that the block is available for use.										
SLF	Sets the TRT flaw entry for the specified logical track to indicate that the track is unavailable for use.										
CLF	Clears the TRT flaw entry for the specified logical track to indicate that the track is available for use.										
ta	Specifies the track address to be reserved. (Refer to table 8-1.) ta can be one of the following:										
	<table border="1"> <thead> <tr> <th>ta</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>trk</td> <td>Logical track address for a disk or extended memory. (Use with SLF and CLF.) The variable trk can be any octal number in the specified range for the particular type of disk drive or extended memory.</td> </tr> <tr> <td>Aaddr</td> <td>Physical block (track) address for extended memory. (Use with SPF or CPF.) addr is the extended memory address divided by 10₈.</td> </tr> <tr> <td>Aaddr– Aaddr</td> <td>Range of physical block addresses for extended memory. (Use with SPF or CPF.)</td> </tr> <tr> <td>Ccyl, Ttrk, Ssec</td> <td>Physical track address for the disk. (Use with SPF or CPF.) The variables cyl, trk, and sec can be any octal number in the specified range for the particular type of disk drive.</td> </tr> </tbody> </table>	ta	Description	trk	Logical track address for a disk or extended memory. (Use with SLF and CLF.) The variable trk can be any octal number in the specified range for the particular type of disk drive or extended memory.	Aaddr	Physical block (track) address for extended memory. (Use with SPF or CPF.) addr is the extended memory address divided by 10 ₈ .	Aaddr– Aaddr	Range of physical block addresses for extended memory. (Use with SPF or CPF.)	Ccyl, Ttrk, Ssec	Physical track address for the disk. (Use with SPF or CPF.) The variables cyl, trk, and sec can be any octal number in the specified range for the particular type of disk drive.
ta	Description										
trk	Logical track address for a disk or extended memory. (Use with SLF and CLF.) The variable trk can be any octal number in the specified range for the particular type of disk drive or extended memory.										
Aaddr	Physical block (track) address for extended memory. (Use with SPF or CPF.) addr is the extended memory address divided by 10 ₈ .										
Aaddr– Aaddr	Range of physical block addresses for extended memory. (Use with SPF or CPF.)										
Ccyl, Ttrk, Ssec	Physical track address for the disk. (Use with SPF or CPF.) The variables cyl, trk, and sec can be any octal number in the specified range for the particular type of disk drive.										

6. Initiate flawing of the specified device by entering:

K.GO.

The FLAW utility provides two messages in the system dayfile that indicate the results of the flawing operation. The first message is:

```
nn TRACKS FLAWED.
```

nn	Description
nn	Octal number of tracks that were successfully flawed.

The second message appears only if some of the flaws specified were not processed. This occurs when the track specified for flawing is already reserved by the system (but not as a flawed track). In this case, the following message also appears in the system dayfile.

```
nn FLAWS NOT PROCESSED,list.
```

Variable	Description
nn	Octal number of flaws not processed.
list	List of the logical tracks that were not flawed.

The entries described here are similar to those entered in the APRDECK (refer to section 3, Deadstart Decks) for flawing a device at deadstart time. However, the flaw entries specified using the FLAW utility or DSD command INITIALIZE are not recovered if the device is initialized at deadstart time. Only the flaw entries specified in the APRDECK will be recovered. If a device is initialized during normal system operation (INITIALIZE command), all flaws specified in the devices TRT, including those entered using the FLAW utility or INITIALIZE command, will be recovered providing the device has a good label at the time of initialization. If the label is bad, or cannot be recognized, all current flaws are cleared. The left console screen (figure 8-3) shows all flaw entries made through the FLAW utility and INITIALIZE command. It lists the flaw entry and its logical track equivalent. If the same logical track is referenced by more than one flaw entry, an asterisk appears to the right of those entries, so that you are aware that only the last entry takes effect.


```

K.  FLAW                AAAM
      CURRENT SLF, CLF, SPF, AND CPF ENTRIES

EQ = 14                EST ORDINAL OF DEVICE.
NUM  VALUES ENTERED   TRT OCTAL VALUES  * = DUPLICATE

  1  SLF=4000.          4000                *
  2  SLF=4000B.         4000                *
  3  SLF=4001.          4001
  .
  .
  .
14

SLF=4001.

```

Figure 8-3. FLAW Utility K Display (Left Screen)

Table 8-1. Flawing Information for Disks and Extended Memory

Device	Device Type	Logical Track Range ¹	Cylinders/ Device ¹	Tracks/ Cylinder ¹	Sectors/ Track ¹
819, single density	DV	4000-5465	633	12	24
819, double density	DW	4000-7153	1466	12	24
834	DD	4000-7135	1457	12	40
836	DG	4000-6565	1273	30	57
844-21, half track	DI	4000-7137	630	22	30
844-21, full track	DK	4000-7137	630	23	30
844-41/44, half track	DJ	4000-7147	1464	23	30
844-41/44, full track	DL	4000-7147	1464	23	30
885-11/12, half track	DM	4000-7221	1511	50	40
885-11/12, full track	DQ	4000-7221	1511	50	40
885-42	DB	4000-7221	1511	12	40
887, 4K sector	DF	4000-7343	1562	4	46
887, 16K sector	DH	4000-7343	1562	4	13
895	DC	4000-7351	1565	17	3
9853	DN	4000-7726	2601	23	25
Extended memory	DE/DP	4000-7620			
5832, 1X SSD	EA	4000-7745	1514	4	3
5832, 2X SSD	EB	4000-7745	1514	4	6
5833, 1X Sabre	EC	4000-7775	3135	7	6
5833, 1XP Sabre	ED	4000-7775	3135	7	6
5833, 2X Sabre	EE	4000-7761	3135	7	13
5833, 2XP Sabre	EF	4000-7761	3135	7	13

1. Numbers are in octal.

INITIALIZE K Display

The INITIALIZE command can be used to reconfigure certain removable devices (844-21 and 844-41/44) to suit your needs. For example, if you currently have two single-unit 844 packs (DK-1s), both packs can be initialized and linked together to form a multispindle device (DK-2). However, this can be done only if the devices to be linked meet the following requirements:

- Same device type.
- Same channels.
- Same share status (shared or nonshared).
- Removable.
- Not currently in use.

The INITIALIZE command must be entered to set the initialize status for each device to be chained. Current multispindle devices can also be initialized providing all packs that form the device are mounted in a logical order as determined by the unit numbers list on the E,C display. It is only necessary to enter the INITIALIZE command for the first unit of a current multispindle device. The format of the INITIALIZE command is described in section 5, DSD Commands.

NOTE

Examine the FAMC and DAFC fields in the family status display (E,F) before entering the INITIALIZE command. The user count for the device must be zero before this command is valid.

The following steps describe the procedures to initialize and (if necessary) flaw tracks on a mass storage device.

1. Enter the INITIALIZE command for the device(s) to be initialized. Examine the B,0 display for the following message:

```
REQUEST *K* DISPLAY.
```

Note the job sequence name of the job displaying the message is CMS.

2. Activate the K display for that job by entering:

```
K, CMS.
```

The INITIALIZE K display (figure 8-4) appears on the left console screen. All parameters required to initialize and (if necessary) flaw the specified device are entered through the K display.

```

K. INITIALIZE.      CMS

MASS STORAGE DEVICE INITIALIZATION                PAGE 1 OF 2

INITIALIZE EQUIPMENT
  014

INITIALIZE REQUEST          *TOTAL*

OPTION                      DESCRIPTION

EQ = 014      EST ORDINAL OF DEVICE (5-777).
FM = 0        FAMILY NAME OR PACK NAME IF TY=X (1-7 CHAR).
DN = 0        DEVICE NUMBER (1-77).
UN = 0        PRIVATE AUXILIARY DEVICE USER NAME (1-7 CHAR).
TY = 0        ACCESS TYPE (F OR X ).
LA = 0        LOWER ACCESS LEVEL (1-7 CHAR).
UA = 0        UPPER ACCESS LEVEL (1-7 CHAR).
DM = 0        DEVICE MASK (0-377).
SM = 0        SECONDARY MASK (0-377).
NC = 0        CATALOG TRACKS (0-200).
NP = 0        NUMBER OF PACKS (1-7).

ENTER OLD FAMILY NAME.
```

Figure 8-4. INITIALIZE Command K Display

The K display lists the options used to initialize a device (under the OPTION column as shown in figure 8-4). Refer to table 8-2 for a description of each option. Refer to table 8-3 for additional information. The EST ordinal of the device to be initialized is listed under the INITIALIZE EQUIPMENT heading (see figure 8-4).

Table 8-2. Device Initialization Options

Option¹	Description
EQ	EST ordinal of the device to be initialized. For multispindle devices, this ordinal must be the first of the consecutive units in the multispindle chain.
FM	The 1- to 7-character family name. Specifies the permanent file family in which the initialized device is to be included. All devices must have a family name or pack name. The name 0 (single character zero) is reserved and cannot be used. This option cannot be used to change the family name of the link device in a multimainframe environment. If you specify TY=X, this option specifies a 1- to 7-character packname to be associated with an auxiliary device. To clear an existing entry, enter FM=NULL.
DN	The 2-digit (octal) logical device number (from 1 to 77 ₈) that uniquely identifies the device in its permanent file family. This option cannot be entered if you specify TY=X.
UN	The 1- to 7-character user name. This option is specified only when initializing an auxiliary device (TY=X). If specified, the device is considered to be a private auxiliary device. Only the user name specified will be allowed to create files on the device (use SAVE, REPLACE, or DEFINE commands). To clear an existing entry, enter UN=NULL.
TY=F	Initialized device may contain direct and indirect access permanent files. However, if you specify DM=0, only direct access files can reside on the device. If you specify SM=0 and DM=0, the device can contain only special system permanent files. Indirect access files can reside only on a master device (that is, DM≠0).
TY=X	Initialized device is an auxiliary device. This is a mass storage device that is not part of a permanent file family. An auxiliary device is a supplementary permanent file storage device that may be privately owned (UN option specified) or can be shared by many users (UN not specified). Auxiliary devices can contain direct or indirect access permanent files.
LA	Lower limit for the access level of the device. This establishes the lowest access level of files that can be stored on the device (must not be less than the lower limit for the access level of the device as given in its EST entry).
UA	Upper limit for the access level of the device. This establishes the highest access level of files that can be stored on the device (must not be greater than the upper limit for the access level of the device as given in its EST entry).

1. Device initialization options may be changed only if the total initialization level (AL) is specified in the INITIALIZE command.

(Continued on next page)

Table 8-2. Device Initialization Options

Option¹	Description
<i>(Continued from previous page)</i>	
DM	The 3-digit (octal) device mask (from 0 to 377 ₈). This option is required whenever a permanent file master device is being initialized. It defines which users will have this device as their master device. This option cannot be entered if you specify TY=X.
SM	The 3-digit (octal) secondary device mask (from 0 to 377 ₈) used to control the residence of direct access files. This option cannot be entered if you specify TY=X.
NC	Octal number of catalog tracks (from 0 to 200 ₈ ; must be a power of 2). This option is used only if the number of catalog tracks specified as the system default for the device type is not satisfactory.
NP	Number of physical units to be included in a multispindle device. The default value is 1. Each unit to be included in the multispindle chain must be defined as removable in the current disk status display (E,M).

2. Device initialization options may be changed only if the total initialization level (AL) is specified in the INITIALIZE command.

Table 8-3. Track Flawing Options

Option¹	Description²
SPF	Sets the track reservation table (TRT) entry for the specified physical block (track) address(es) in extended memory or disk to indicate that the block is unavailable for use.
CPF	Clears the TRT flaw entry for the specified physical address(es) in extended memory or disk to indicate that the block is available for use.
SLF	Sets the TRT flaw entry for the specified logical track to indicate that the track is unavailable for use.
CLF	Clears the TRT flaw entry for the specified logical track to indicate that the track is available for use.

1. Flawing of 881 and 883 disk packs is automatic; only flaws additional to the current flaw information in the utility sector should be entered.

2. Refer to the APRDECK description in section 3, Deadstart Decks, for entry formats.

3. Enter the INITIALIZE command for each additional device to be initialized. This can also be done before activating the K display. In either case, only the first device specified will be listed (by EST ordinal) in the K display. Thus, to update the K display to show additional devices, enter the following command:

K.RERUN.

If more than one device is listed, they are initialized one at a time as they appear in the list from left to right. Multispindle devices (more than one EST ordinal) are considered one device.

4. The system has already checked the label on each mass storage device. If the label was found to be good, it is necessary to enter device initialization options (requested in messages that appear in the K display) to identify the device. This is to ensure that the device selected is indeed the correct device to be initialized. The messages appear automatically and are displayed until the correct options are entered. If an incorrect option is entered, it is ignored. Refer to table 8-2 for a description of each option.

Examine the family status display (E,F) to determine the current option values.

The following messages may be displayed and the appropriate response should be entered. If none of these messages appears, the device label was not recognized or was found to be bad. In this event, proceed to step 5.

Message	Description
ENTER OLD DEVICE NUMBER	This message appears if the device to be initialized is a permanent file family device. Enter the following response: K.DN=devicenumber.
ENTER OLD FAMILY NAME	This message appears only if more than one family of permanent file devices is currently active in the system. Enter the following response: K.FM=familyname.
ENTER OLD PACK NAME	This message appears only if the device to be initialized is an auxiliary device. Enter the following response: K.PN=packname.
ENTER OLD USER NAME	This message appears only if the auxiliary device to be initialized is a private auxiliary device (associated with a specific user name). Enter the following response: K.UN=username. The user name is written to the account dayfile when the device is mounted. If you discover that the wrong device was specified in the INITIALIZE command, you can clear the initialize status for that device by entering: K.CLEAR. The device to be processed by the clear entry must be a valid device. That is, the device cannot have a device number that conflicts with another device in its family name or a pack name that duplicates one already in the system. Its mask bits must meet standard requirements. The leftmost device in the list of devices to be initialized is cleared. One of the preceding messages will then be displayed for the next device to be initialized (if any) providing the label on that device is good.

- When the following message appears on the K display, enter the options that specify the new characteristics to be associated with the device when it is initialized (refer to tables 8-2 and 8-3).

ENTER PARAMETERS

The new options can be entered as a string or one at a time. For example:

K.option₁=value₁, . . . , option_n=value_n.

or

K.option₁=value₁.

K.option₂=value₂.

K.option_n=value_n.

If flaw entries are to be specified (refer to table 8-3), they must be entered singly, as shown in the last example. If the label on the device being initialized was good, all current flaws on that device are normally recovered. However, if the label was not recognized or was bad, the flaw entries cannot be recovered and must be entered (if necessary) using this mechanism. A maximum of 14 flaw entries are permitted. In addition to the SLF, CLF, SPF, and CPF entries, the flaw information recorded in the utility sector on an 881, 883, or 885 disk pack is read during initialization of 844 equipment, and the appropriate areas are reserved by the system automatically.

If the NP option is specified (NP > 1), the device is to be initialized as a multispindle device. In this case, the number of packs specified by NP indicates the number of spindles to be linked. This is the next n number of devices waiting to be initialized. Each device must be defined as removable and mounted on consecutive physical unit numbers. To determine if a device is defined as removable, examine the disk status display (E,M). If the units are configured correctly, the label on each unit is checked. If any label is not recognized or is bad, that unit is free for initialization and chaining. However, if the label is good, the message

ENTER IDENTITY OF EQest

appears in the K display (est is the EST ordinal of the device). Enter one of the following responses.

K.DN=devicenumber.

K.FM=familyname, DN=devicenumber.

K.PN=packname.

K.PN=packname, UN=username.

This is a precautionary measure to ensure that the devices specified are the correct devices to be chained.

- After all the necessary options have been entered for a specific device, enter the following command to proceed with the initialization.

K.GO.

If more devices are waiting to be initialized, repeat steps 4 through 6 of this procedure for each device.

INITIALIZE K Display

Example 1:

Suppose you want to create and initialize a multispindle device by linking spindles from two separate units. Assume your EQPDECK has the following entries:

```
EQ40=DL,ST=OFF,UN=4,CH=27/30.  
EQ41=DL,ST=OFF,UN=5,CH=27/30.
```

and you want to initialize multispindle device EQ40, which consists of units 4 and 5. Enter the following commands:

```
ON,40.  
ON,41.  
INITIALIZE,AL,40,41.  
MOUNT,40.  
MOUNT,41.  
K,CMS.
```

Enter device initialization options as described in steps 4 and 5.

```
K.NP=2.
```

When the message ENTER IDENTITY OF EQest appears in the K display, respond as described in step 5.

```
K.GO.
```


Example 2:

Assume that you want to initialize a multispindle device and your EQPDECK already defines the device as a multispindle device:

```
EQ40=DL, ST=OFF, UN=4/5, CH=27/30.
```

In this case, enter the following commands:

```
ON, 40.
INITIALIZE, AL, 40.
MOUNT, 40.
K, CMS.
```

Enter device initialization options as described in steps 4 and 5.

```
K.GO.
```

Example 3:

Suppose you want to initialize a single-spindle device and your EQPDECK defines the device as a multispindle device:

```
EQ40=DL, ST=OFF, UN=4/5, CH=27/30.
```

Mount the pack you want to initialize on unit 4 and enter the following commands:

```
ON, 40.
INITIALIZE, AL, 40.
MOUNT, 40.
K, CMS.
```

Enter device initialization options as described in steps 4 and 5.

```
K.PN=1.
K.GO.
```

Machine Recovery (MREC) Utility K Display

When a machine that has been in a multimainframe configuration has a malfunction that requires a level 0 deadstart, before you can begin the deadstart, the machine recovery utility program (MREC) must be run on every machine that shares devices with the machine that is not working properly. MREC releases local mass storage space on the shared devices and clears interlocks set before the malfunction occurred.

For example, assume that machines A and B share disk unit 12 and machines B and C share disk unit 13 as shown in figure 8-5.




Figure 8-5. Machine Configurations

If machine A must be deadstarted using a level 0 deadstart, MREC must be run on machine B to recover disk unit 12. Machine C need not be involved since it is not aware of the existence of disk unit 12. However, if machine B must be deadstarted using a level 0 deadstart, MREC must be run on machine A to recover disk unit 12 and on machine C to recover disk unit 13.

NOTE

Once MREC has been run for an inoperative machine, any level of deadstart on the machine other than 0 is not possible.

MREC Procedures

The MREC utility can be used through keyboard entry to a K display or by direct keyboard entry under DSD or DIS control without using the K display.

Using MREC with the K Display

The following procedure describes K-display usage for MREC operations.

1. Call the MREC utility by entering this DSD command:¹

```
X.MREC.
```

2. Examine the DSD system status display (B,0). When MREC is scheduled to a control point, it is indicated on the B,0 display. The message

```
REQUEST *K* DISPLAY
```

appears in the message field for that control point.

3. Activate the K display for that control point by entering:

```
K, jsn.
```

jsn	Description
jsn	MREC job sequence name.

The K display for MREC (figure 8-6) appears on the left console screen.

The MREC left screen K display lists all of the devices that are shared by the machine on which MREC is being run. The machine identifier (MID) of this machine is also given. Information describing the shared devices is given in the following format.

```
eq type un dn fm/pn mids sharing device
```

Header	Description
eq	EST ordinal of the shared device.
type	Device type.
un	Unit number of the device.
dn	Device number.
fm/pn	Family name/pack name.
mids sharing device	Machine identifiers of other machines that are currently accessing the device. If there is an * by the machine identifier, the machine is determined to be down.

¹ Under DIS control, the command MREC calls the MREC utility.

```

K. MREC.                jsn
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
LVL0   - LVL7          MID=AA   NOS version

MACHINE RECOVERY UTILITY                                PAGE 1 OF 1

  EQ   TYPE   UN   DN   FM/PN   MIDS SHARING DEVICE
  ---  ---   --  --  ---     -----
  10   DJ     0   0   PACKV2   73  72  71  70
  11   DJ     1
  12   DJ     2
  13   DJ     3   30  MAINTV2   73  72  71  70
  14   DL     4   40  SYS606   73  72  71  70
  15   DI     1   3   SYST06   73  72  71  70
  16   DI     2
  17   DI     3   0   R4IAE    73  72  71  70
  20   DJ     4

ID OF DOWNED MACHINE = 72
EQ(S) TO RECOVER = ALL
  
```

Figure 8-6. MREC K Display (Left Screen)

4. The MREC right screen K display contains the parameters and commands that may be selected. Refer to table 8-4 (MREC parameters) and table 8-5 (MREC commands) for more information. To activate the MREC right screen K display (figures 8-7 and 8-8), enter:

KK.

5. Enter the MREC parameters in the following format:

K.param₁=value₁,param₂=value₂, . . . ,param_n=value_n.

The parameters entered (and error messages, if any) are displayed on the lower portion of the MREC left screen K display.

6. If you want to reset the parameters to their default values or refresh the device descriptions on the MREC left screen K display, enter:

K.RERUN.

The parameters can then be reentered.

7. Enter the desired parameters and initiate MREC processing by entering:

K.GO.

When processing is complete, the message

PROCESSING COMPLETE

appears at the bottom of the left screen.

8. After all MREC operations are complete, end the MREC utility by entering:

K.STOP.

Table 8-4. MREC Parameters

Parameter	Description
ID=id	The 1- or 2-character machine identifier of the inoperative machine that is to be processed. This parameter must be entered before processing can take place.
EQ=est	EST ordinals of the devices to be processed. Only devices shared between the machine that is down and the machine on which MREC is running are processed. The form of the entry is: EQ=est ₁ ,est ₂ , . . . ,est _n . If you omit this parameter, EQ=ALL is assumed. In this case, all devices shared between this machine and the inoperative machine are processed.
OP=x	You can enter this parameter only when using the K display (that is, it cannot be used if the MREC utility is called by a command) and then, use it only if a unit or controller cannot be accessed by MREC due to physical hardware reservations. Its use is invalid if a unit reservation is not in effect.
x	Description
R	Directs the MREC utility to release all unit reservations (using the GRENADE function; refer to section 2, Deadstart) for 844-41/44 equipment.
I	Directs the MREC utility to ignore certain functions on the equipment for which the reservation message was issued. Functions that do not require the unit to be accessed are still performed.

Table 8-5. MREC Commands

Command	Description
GO.	Directs the MREC utility to proceed with processing the entered parameters.
RERUN.	Reinitializes the device descriptions and parameters on the K display and reruns the MREC utility.
STOP.	Terminates the MREC utility and ends the K-display interaction.
+	Pages the left screen K display forward to the next screen.
-	Pages the left screen K display backward to the first screen.
)	Pages the right screen K display forward to the next screen.
(Pages the right screen K display backward to the previous screen.

```

K.  MREC.          jsn

MREC                                     PAGE 1 OF 2

PARAMETER DESCRIPTIONS

ID      MID OF DOWNED MACHINE TO PROCESS.
EQ      EST ORDINALS OF DEVICES SHARED BETWEEN THIS MACHINE
        AND DOWNED MACHINE TO PROCESS.  ENTER AS
        EQ=XXX,XXX,...,XXX.  THE DEFAULT IS EQ=ALL.
OP      MAY BE ENTERED IF 844 RESERVE SITUATIONS OCCUR.
        THE RESPONSE APPLIES ONLY TO THE CURRENT SITUATION.
        OP=I      IGNORE CURRENT DEVICE.
        OP=R      CLEAR ALL UNIT RESERVES ON THAT CONTROLLER.

K DISPLAY COMMANDS

GO      INITIATE PROCESSING OF DEVICES SPECIFIED.
RERUN   REINITIALIZE K-DISPLAY AND RERUN PROGRAM.
STOP    TERMINATE PROGRAM.
+       PAGE LEFT DISPLAY FORWARD.
-       PAGE LEFT DISPLAY BACKWARD.
(       PAGE RIGHT DISPLAY FORWARD.
)       PAGE RIGHT DISPLAY BACKWARD.

```

Figure 8-7. MREC K Display (Right Screen) (Page 1 of 2)

```

MREC                                     PAGE 2 OF 2
DESCRIPTION OF TABLE ENTRIES.

EQ      EST ORDINAL OF EQUIPMENT.
UN      UNIT NUMBER.
DN      DEVICE NUMBER.
FM/PN   FAMILY OR PACK NAME.
MIDS SHARING DEVICE  MACHINE ID-S OF OTHER MACHINES
                    CURRENTLY SHARING THE DEVICE.
* (BESIDE MID)      THE MACHINE IS DOWN. IT CANNOT BE
                    DETERMINED IF MACHINES RUNNING IN
                    SHARED RMS MODE ARE DOWN.

```

Figure 8-8. MREC K Display (Right Screen) (Page 2 of 2)

Using MREC without the K Display

You can also call the MREC utility and specify the appropriate parameters without using the K display. You can do this by making a single keyboard entry (under DSD or DIS control) in the following format:

```
X.MREC(param1=value1,param2=value2,...,paramn=valuen)
```

The parameters (param_i=value_i) are described in table 8-4. The ID=id parameter must be entered.

When the MREC command is entered with parameters, the K display is not activated and processing occurs automatically. If an error occurs using this procedure, the message

```
REQUEST *K* DISPLAY
```

appears on the DSD B,0 display. Activate the K display and continue as described under Using MREC With the K Display earlier in this section.

MREC Unit and Controller Reservations

When attempting to access a device, MREC may find the controller access or unit reserved by another machine. When this occurs, the following message appears on the left screen K display.

```
EQest,CHcc, CONTROLLER RESERVED.
```

or

```
EQest,UNuu, UNIT RESERVED.
```

Variable	Description
est	EST ordinal of the device.
cc	Channel number.
uu	Physical unit number (from 0 to 77 ₈).

Assuming the inoperative machine is the machine holding the reservation, you can clear the reservation or direct MREC to clear it by using the following procedures:

To clear a controller reservation, initiate deadstart on the machine that is down.

To clear a unit reservation, perform one of the following procedures:

- On a device that is connected to a 7155 controller, initiate deadstart on the machine that is down.
- On a device that is not connected to a 7155 controller, toggle the OFF LINE/ON LINE switch on the back of the drive to OFF LINE and then back to ON LINE.
- If none of the preceding procedures can be performed, select the OP=R parameter to clear an 844 device reservation.

NOTE

Do not select the OP=R parameter unless the other reservation clearing procedures cannot be performed. It is recommended that all machines in a multimainframe environment be put in IDLE status or put in STEP mode when you select the OP=R parameter.

Once the correct action has been taken, enter

K.GO.

to continue processing. If the reservation is still not cleared, the appropriate reserved message reappears on the left screen K display. Repeat one of the reservation clearing procedures described above or enter:

K.OP=I.

This directs the MREC utility to ignore certain functions on the device. Processing may then continue.

If a device or controller in an independent shared device multimainframe environment is reserved by a down machine, the previously mentioned messages appear on the system status display (B,0) instead of on the K display. Use the procedures already described to clear the reservation.

Mass Storage Extended Subsystem (MSE) K Display

MSE uses the K display to present messages that require your action. You can use the MSE K display to reply to these messages and to open and close staging/destaging operations. You can also use the MSE K display to display a specific storage module's ON/OFF status, the number of DRDs within the storage module for staging/destaging, and other information including the job sequence name, file name, user index, and PRU count for the files being processed. In addition, you can set the maximum number of DRDs to be used for staging/destaging operations within a specified storage module.

When MSE enters messages into the K display, a request for the K display appears on the B,0 display. Enter K,MSE to activate the display for MSE. After you have responded to all K-display messages, the request for the K display terminates.

The MSE K display provides space for two messages with up to three lines per message. If you request information about a specific storage module, the message area is replaced by the information concerning the storage module. Figures 8-9 and 8-10 show the message and storage module formats of the MSE K display.

Valid MSE commands are:

Command	Description
K.m.GO.	Enters a GO response to the message at message ordinal m.
K.STAGE.	Toggles staging operations between open and close. Staging is normally open (permitted).
K.DESTAGE.	Toggles destaging operations between open and close. Destaging is normally open (permitted).
K.SMi.	Displays information for storage module i (i is an SM identifier A through H). The display shows SM ON/OFF/MAINTENANCE status, the number of DRDs for staging/destaging, and DRD activity including the job sequence name, file name, user index, and PRU count for the files being processed. The display also shows if a DRD is being used for staging or destaging, a DRD is under control of a utility, or the state of a DRD is maintenance mode, off, or idle.
K.SMi,DRD=ST=s.	Sets the maximum number of DRDs to be used for staging within storage module i (i is an SM identifier A through H). s is an integer 0 through 2; an entry of 0 turns off staging for SMi.
K.SMi,DRD=DS=d.	Sets the maximum number of DRDs to be used for destaging within storage module i (i is an SM identifier A through H). d is an integer 0 through 2; an entry of 0 turns off destaging for SMi.

MSE clears a message when an acceptable action is taken. Usually this action is entering a GO response to the message at message ordinal m. However, for actions such as emptying the exit tray on a storage module, the message is cleared when the hardware status indicates the action has occurred and you enter the GO response. Refer to appendix B of the NOS Version 2 Operations Handbook for the appropriate action for each message.

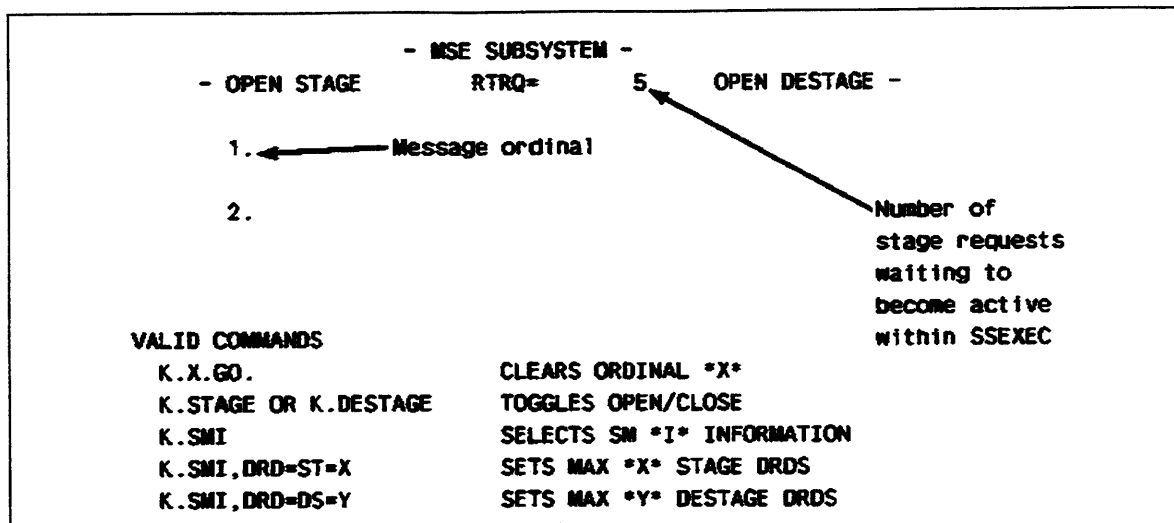


Figure 8-9. MSE K Display (Message Format)




Figure 8-10. MSE K Display (Storage Module Format)

Network Access Method (NAM) K Display

The NAM K display provides a common, centralized interface by which network applications (CS, NS, NVF, or NETOU) can communicate with the host operator (HOP).

Figure 8-11 shows the format of the NAM K display. Table 8-6 explains the fields that appear on the NAM K display.

		data area				
	message					
Alert Line	ALERTS	app1	app2	app3	app4	app5
HOP Message Line	HOP	hop message line				
Response Line	appname	operator entry		repeat message		

Figure 8-11. NAM K-Display Format

Table 8-6. NAM K-Display Fields

Field	Description
data area	The portion of the screen that receives the output or diagnostic message as a result of a command that you entered. This portion of the NAM K display is operated as a scrollable paged device. Each line of data enters the display at the bottom of the data area and forces the previous lines to shift up by one line. The previous top line is lost. However, whenever you turn the page-wait status on and more than a screen full of information is written to the data area, the prompt MORE DATA.. is displayed on the bottom line of the data area. The display remains fixed until you enter the + character to display the next page of information (refer to Display Control Characters later in this section).
message	The system prompt READY.. indicates that you can make additional keyboard entries. The system prompt MORE DATA.. indicates that you can enter + to see more data.
appl...app5	The alert line is a list of applications requesting your attention.
hop message line	This field contains a message received from another host.
appname	The name of the application (CS, NS, NAM, NVF, or NETOU) with which you are currently interacting.
operator entry	This field contains up to 40 characters of the last command that you entered. Commands of more than 40 characters are truncated.
repeat message	Whenever a command cannot be accepted because the system was not done processing a previous command, you get the system prompt REPEAT.. in this field. Repeat the command.

NAM supports both the right and left screen K displays. The right K display provides help information.

The NAM K display is available at all times during NAM execution, the message REQUEST K DISPLAY appears on the B,0 display when the K display is not assigned to NAM and some application has requested operator intervention. You must assign the K display to NAM to interact with NAM or an application.

You assign the K display to NAM by entering:

K, NAM.

Figure 8-12 shows the NAM K display as it appears when you initially assign it.

```
READY . .  
ALERTS  
NAM
```

Figure 8-12. NAM K Display

The NAM K display has two operating modes:

- NAM mode.
- Application mode.

Initially, the display is in NAM mode. The appearance of NAM indicates that you are interacting directly with NAM.

Display Control Characters

You can use the following four characters to control the NAM K display. They are valid in both NAM mode and application mode. Each of the following characters must be entered in the first character position after K.

Character	Description
*	When in application mode, the * character causes the display to revert to NAM mode. The last command is aborted and any information generated by the application, after the asterisk is entered, is discarded by NAM. The page-wait status is automatically turned on. The * character is ignored if the display is already in NAM mode.
/	When in application mode, the / character aborts the last command without altering the assignment of the NAM K display. The / character always results in a READY.. prompt. When in NAM mode, the / character is ignored.
+	When in application mode, the + character turns the page-wait status on, if the current page-wait status is off. If the current page-wait status is on, then the command K.+ displays the next available page of information. The page-wait status is always on in NAM mode. The page-wait status is initialized to off each time you enter application mode. Therefore, whenever you select an application, the page-wait status automatically changes from on to off and you can set it as desired.
-	The - character turns off the page-wait status. The - character is ignored if entered in application mode while the page-wait status is off. Any attempt to turn the page-wait status off while in NAM mode is ignored without sending any diagnostic message.

NAM K-Display Operation

Initially, the display is in NAM mode. The application mode is invoked by a NAM mode command that indicates the application with which you wish to interact. The application mode is indicated by the appearance of the name of that application in the appname portion of the response line (refer to figure 8-10). In application mode, all entries except the asterisk are passed by NAM to a previously designated application. An asterisk character returns the display to NAM mode.

Entering any command (processed either by NAM or by an application) immediately causes the command to appear in the operator entry field. When processing of the command completes, the prompt READY.. appears in the bottom line of the data area. If you key something in, other than one of the display control characters *, /, +, or - before the prompt READY.. appears, the character is displayed in your entry field but the command is not acted upon and the message REPEAT.. appears to the right of your entry.

Table 8-7 describes the commands for assigning the K display to one of the applications CS, NS, NVF, or NETOU, and other frequently used commands. If appname or the remote host number are invalid or the application is not currently active, you will receive a diagnostic message. These commands are available in NAM mode only.

Table 8-7. NAM Mode Commands

Command	Description
K.AP=appname	Assigns the NAM K display to the specified application. The application name appears in the appname portion of the response line (refer to figure 8-10). appname is one of the supervisory programs CS, NS, NVF, or NETOU. The command satisfies any alert request posted by the application. When you type the application name, that application name is removed from the alert line, the page-wait status is turned off, and the last page of the application's recent history buffer appears in the data area of the display (refer to Recent History Command in section 15, Network Operations).
K.AP	Assigns the NAM K display to the application whose name appears in the leftmost position of the alert line. The page-wait status is turned off and the last page of the application's recent history buffer appears in the data area of the display.
K.HELP	Displays information about the NAM mode commands on the right K display.
K.IG=appname	Causes NAM to ignore all alert requests from the specified application. If the application had an alert request pending, its name is removed from the alert list. NAM informs the application that its alert request was acknowledged and ignored. The rest of the display remains unaltered. This command is cancelled automatically when you enter the AP=appname command. The assignment of the NAM K display is not altered by the IG=appname command.

(Continued on next page)

Table 8-7. NAM Mode Commands

Command	Description														
<i>(Continued from previous page)</i>															
K.IG	Causes NAM to ignore all alert requests from the application whose name appears in the leftmost position of the alert line. NAM informs the application that its alert request was acknowledged and ignored. The rest of the display remains unaltered. This command is cancelled automatically when you enter the AP=apname command.														
K.END	Clears the NAM K display. It causes the data area of the display to go blank but leaves the alert line and response line unaltered. The END command is used to stop the generation of the status display (figure 8-12). Otherwise, if the K display is dropped, NAM automatically ends after 60 to 90 seconds.														
K.SEND,rrr,string.	Causes NAM to send the message text if a logical link exists between the specified remote host and this host. This command applies only to hosts that are connected using 255x NPUs, since no logical link exists between hosts that are connected through either a public data network or a CDCNET network. Use the NAM ST command to determine which hosts are connected.														
<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>rrr</td> <td>Host node number of the remote host (from 1 to 255).</td> </tr> <tr> <td>string</td> <td>Message text of no more than 40 characters.</td> </tr> </tbody> </table> <p>The following message is displayed on the specified remote host's NAM K display after the alert line.</p> <pre>hh.mm.ss. FROM sss string</pre> <table border="1"> <thead> <tr> <th>Variable</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>hh.mm.ss</td> <td>Time the message is received.</td> </tr> <tr> <td>sss</td> <td>Host node number of the sending host (from 1 to 255).</td> </tr> <tr> <td>string</td> <td>Message text of no more than 40 characters.</td> </tr> </tbody> </table>		Parameter	Description	rrr	Host node number of the remote host (from 1 to 255).	string	Message text of no more than 40 characters.	Variable	Description	hh.mm.ss	Time the message is received.	sss	Host node number of the sending host (from 1 to 255).	string	Message text of no more than 40 characters.
Parameter	Description														
rrr	Host node number of the remote host (from 1 to 255).														
string	Message text of no more than 40 characters.														
Variable	Description														
hh.mm.ss	Time the message is received.														
sss	Host node number of the sending host (from 1 to 255).														
string	Message text of no more than 40 characters.														
K.ST	Causes the NAM status display to appear on the K-display screen.														

NAM Mode Maintenance Commands

While in NAM mode, the HOP can enter the following network maintenance commands. The commands are sent to NAM or to applications to determine the status of the network and aid in the debugging of network problems. The HOP may select NAM or a single network supervisory application or NAM and all network supervisory applications using the following commands:

Command	Description
DB ²	Activates the online debugging code.
DE ²	Deactivates the online debugging code.
DU ²	Causes NAM to dump its field length and/or request an application to dump its field length.
FL	Changes NAM's maximum field length.
LB ²	Begins the logging of network traffic in the debug log file.
LE ²	Ends the logging of network traffic in the debug log file.
LR ²	Releases the debug log file.
RS ²	Dumps the statistics data to a permanent file, resets the statistics counters to 0 (zero), and continues gathering statistics.
ST	Causes the NAM status display to appear on the K-display screen (refer to figure 8-13).

DB Command

The debug begin (DB) command causes NAM to turn on its online debugging code and/or request an application to turn on its online debugging code.

The command format is:

DB=mode .

mode	Description
ALL	NAM turns on its online debugging code and requests all applications to turn on their online debugging code.
appname	NAM requests the specified application to turn on its online debugging code. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM turns on its online debugging code.

2. These commands are supported by the network supervisory applications CS, NS, and NVF. They are ignored by IAF, TAF, and many other network applications.

DE Command

The debug end (DE) command causes NAM to turn off its online debugging code and/or request an application to turn off its online debugging code.

The command format is:

DE=mode .

mode	Description
ALL	NAM turns off its online debugging code and requests all applications to turn off their online debugging code.
appname	NAM requests the specified application to turn off its online debugging code. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM turns off its online debugging code.

DU Command

The dump (DU) command causes NAM to dump its field length and/or request an application to dump its field length.

The command format is:

DU=mode .

mode	Description
ALL	NAM dumps its field length and requests all applications to dump their field length.
appname	NAM requests the specified application to dump its field length. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM dumps its field length.

FL Command

The field length (FL) command changes NAM's maximum field length to the specified value. If the field length specified is less than NAM's current maximum field length or NAM's running field length plus 1000₈, you will receive a diagnostic message.

If NAM reaches its maximum field length, use the FL command to increase NAM's maximum field length. If the K display is in application mode when NAM reaches its maximum field length, NAM automatically switches the K display to the NAM mode.

Until you enter the FL command, all other NAM commands are rejected and NAM stops servicing the network. Frequent need for this command indicates that the initial field length for NAM is inadequate and should be increased by changing the MAXFL parameter of the NIP command described in the NOS Version 2 Installation Handbook.

The command format is:

```
FL=fl.
```

Parameter	Description
fl	Specified maximum field length in octal. fl must not exceed 360000 ₈ . If it does, NAM issues a diagnostic message to the K display.

LB Command

The log begin (LB) command causes NAM to begin logging network traffic in its debug log file and/or request an application to begin logging network traffic in its debug log file (refer to the NAM Version 1 Host Application Programming Reference Manual for information on creating the application's log file).

The command format is:

```
LB=mode.
```

mode	Description
ALL	NAM begins logging network traffic in its debug log file and requests all applications to begin logging network traffic in their debug log files.
appname	NAM requests the specified application to begin logging network traffic in its debug log file. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM begins logging network traffic in its debug log file. NAM must be installed with the network trace option or this command is ignored. Refer to the NOS Version 2 Installation Handbook for more information.

To release the debug log file, enter the log release (LR) command.

LE Command

The log end (LE) command causes NAM to end logging network traffic in its debug log file and/or request an application to end logging network traffic in its debug log file.

The command format is:

LE=mode .

mode	Description
ALL	NAM ends logging network traffic in its debug log file and requests all applications to end logging network traffic in their debug log files.
appname	NAM requests the specified application to end logging network traffic in its debug log file. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM ends logging network traffic in its debug log file.

To release the debug log file, enter the log release (LR) command.

LR Command

The log release (LR) command causes NAM to release its debug log file (if one exists) and/or request an application to release its debug log file. When the local debug log file is released, its contents are copied to a permanent file. Logging continues on a new local debug log file.

The command format is:

LR=mode .

mode	Description
ALL	NAM releases its debug log file and/or requests all applications to release their debug log files.
appname	NAM requests the specified application to release its log debug file. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM releases its debug log file.

RS Command

The reset statistics (RS) command causes NAM to dump its statistics to a permanent file and restart gathering statistics and/or request an application to dump its statistics to a permanent file and restart gathering statistics (refer to the NAM Version 1 Host Application Programming Reference Manual for descriptions of the network statistics). You can use this command to help analyze NAM and/or application performance.

The command format is:

```
RS=mode.
```

mode	Description
ALL	NAM dumps its statistics to a permanent file, resets the statistics counters to 0 (zero), and continues gathering statistics and/or requests all applications to dump their statistics to permanent files, reset their statistics counters to 0 (zero), and continue gathering statistics.
appname	NAM requests the specified application to dump its statistics to a permanent file, reset its statistics counters to 0 (zero), and continue gathering statistics. The specified application must be accessing the network or you will receive a diagnostic message.
NAM	NAM dumps its statistics to a permanent file, resets the statistics counters to 0 (zero), and continues gathering statistics.

ST Command

The status (ST) command causes the NAM status display (figure 8-13) to appear on the K-display screen.

The command format is:

ST.

The NAM status display provides status information regarding all applications, couplers, and logical links currently active in the host. This display is refreshed periodically to display the current conditions of the applications and logical links.

The first line of the NAM status display contains the network invocation number (NIN), host regulation level (REG LVL), the number of applications accessing the network, and the maximum field length of NAM.

NIN = 014													REG LVL = 3			NO OF APPLS = 5				MAXFL = 100000			
APP		JSN	STATUS	I	NCN	AC	NSM	NDM	TIME UP														
IAF		IAF	000000		30		5	2	07.54.35														
RBF		RBF	000000		3		8	4	08.10.00														
TAF		TAF	000000		10		1	0	08.25.46														
TVF		AABA	001000		1		0	0	08.04.00														
NVF		AAAM	000000	N	2		5	2	09.09.45														
EST													EM	HN	NSM	NHM	NLM	IVTSTAT	PRUST	NPUREJ			
054														01	5			0000	46125	5			
LOG-LINK													HN	TN	H	N	S	T	NCN	AC	NHDQ	NLDQ	TIME UP
													01	01	0	0	S		44				08.01.30

Figure 8-13. NAM Status Display

The NAM status display shows the application status entries in the following format:

```
app      jsn          status  i          ncn ac  nsm ndm timeup
```

Header	Description
app	Name of the application. The application name may be one of the following CDC-written applications or a site-written application.
app	Description
ATF	Automated Tape Facility.
CS	Communications Supervisor.
CYBIS	CYBIS-NAM Interface.
I AF	Interactive Facility.
ITF	Interactive Transfer Facility.
MCS	Message Control System.
NJF	Network Job Entry Facility.
NETFS	Network File Server.
NETLS	Network Log Server.
NLTERM	Network Log Terminator.
NETOU	Network Operator Utility.
NS	Network Supervisor.
NVF	Network Validation Facility.
PSU	Printer Support Utility.
PTF	Permanent File Transfer Facility.
PTFS	Permanent File Transfer Facility Server.
QTF	Queued File Transfer Facility.
QTFS	Queued File Transfer Facility Server.
RBF	Remote Batch Facility.
TAF	Transaction Facility.
TCF	Terminal Cluster Facility.
TLF	Tie Line Facility.
TVF	Terminal Verification Facility.
VEIAF	NOS/VE Interactive Facility.

Header	Description																														
jsn	Job sequence name of the executing job table entry.																														
status	Status of the application in octal. Each of the 18 bits (numbered left to right) represents the following condition:																														
	<table border="1"> <thead> <tr> <th>Bit Number</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td></td> <td>Force flag.</td> </tr> <tr> <td>1</td> <td>Wait flag.</td> </tr> <tr> <td>2</td> <td>Rollout flag.</td> </tr> <tr> <td>3</td> <td>ON flag.</td> </tr> <tr> <td>4–5</td> <td>Swap flag.</td> </tr> <tr> <td>6</td> <td>IN flag.</td> </tr> <tr> <td>7–14</td> <td>Reserved.</td> </tr> <tr> <td>15–17</td> <td>NVF response flags.</td> </tr> <tr> <td>i</td> <td>Ignore alert flag. If set to Y, NAM ignores an alert request from the application. If set to N, NAM accepts an alert request from the application. This field is blank if the application is not allowed to use the NAM K display.</td> </tr> <tr> <td>ncn</td> <td>Number of current connections to the application.</td> </tr> <tr> <td>ac</td> <td>Reserved.</td> </tr> <tr> <td>nsm</td> <td>Number of asynchronous supervisory messages currently queued for the application.</td> </tr> <tr> <td>ndm</td> <td>Number of data messages and synchronous supervisory messages queued for the application.</td> </tr> <tr> <td>timeup</td> <td>Time, in the format hh.m.ss, that the application became active.</td> </tr> </tbody> </table>	Bit Number	Description		Force flag.	1	Wait flag.	2	Rollout flag.	3	ON flag.	4–5	Swap flag.	6	IN flag.	7–14	Reserved.	15–17	NVF response flags.	i	Ignore alert flag. If set to Y, NAM ignores an alert request from the application. If set to N, NAM accepts an alert request from the application. This field is blank if the application is not allowed to use the NAM K display.	ncn	Number of current connections to the application.	ac	Reserved.	nsm	Number of asynchronous supervisory messages currently queued for the application.	ndm	Number of data messages and synchronous supervisory messages queued for the application.	timeup	Time, in the format hh.m.ss, that the application became active.
Bit Number	Description																														
	Force flag.																														
1	Wait flag.																														
2	Rollout flag.																														
3	ON flag.																														
4–5	Swap flag.																														
6	IN flag.																														
7–14	Reserved.																														
15–17	NVF response flags.																														
i	Ignore alert flag. If set to Y, NAM ignores an alert request from the application. If set to N, NAM accepts an alert request from the application. This field is blank if the application is not allowed to use the NAM K display.																														
ncn	Number of current connections to the application.																														
ac	Reserved.																														
nsm	Number of asynchronous supervisory messages currently queued for the application.																														
ndm	Number of data messages and synchronous supervisory messages queued for the application.																														
timeup	Time, in the format hh.m.ss, that the application became active.																														

The NAM status display shows the coupler status for all NPU and MDI/MTI coupler entries in the EST currently being serviced by NAM. These entries have the following format:

```
est  em  hn  nsm  nhm  nlm  ivtstat  prust  npurej
```

Header	Description
est	Equipment status table ordinal of the front end NPU or MDI/MTI.
em	Equipment mnemonic NP for the NPU or ND for the MDI/MTI.
hn	Host node number of the coupler.
nsm	Number of asynchronous supervisory messages currently queued for the coupler.
nhm	Reserved.
nlm	Reserved.
ivtstat	Number of characters transferred downline on interactive connections in the last 30 seconds.

Header	Description
prust	Number of characters transferred downline on PRU connections in the last 30 seconds.
npurej	Number of times data was rejected by the NPU.

The NAM status display shows the logical link entries in the following format.

```
hn tn h n s t ncn ac nhdq nldq timeup
```

Header	Description
hn	Node number of the NPU or MDI/MTI coupler.
tn	Terminal node number of the NPU or MDI/MTI on host-to-network logical links, or host node number of the remote host on host-to-host logical links.
h	Logical link regulation level as reported by the host.
n	Regulation level as reported by the NPU or MDI/MTI. n can have the following values:
n Description	
0	Logical link is down or disabled for data.
1	Only asynchronous supervisory messages are allowed.
2	Only asynchronous supervisory messages and high priority connections are allowed.
3	All network connections are allowed.
s	Supervision indicator. If CS is using this logical link to send supervisory messages, the field is set to S; otherwise, it is blank.
t	Host-to-host logical link indicator. If this logical link is between two hosts, the field is set to H; otherwise, it is blank.
ncn	Number of current connections on the logical link.
ac	Reserved.
nhdq	Reserved.
nldq	Reserved.
timeup	Time, in the format hh.mm.ss., that the logical link became active.

Queue File Transfer Facility (QTF) K Display

The QTF utility provides a multistreaming capability for transferring queued files between mainframes using connections through RHF or NAM. QTF allows you to specify which files are to be transferred. You do this by defining a set of selection classes. You can define or modify values for selection classes using the CLASS command described later in this subsection. You can specify a range of file sizes and the maximum number of simultaneous connections allowed. You can also specify a number of reserved connections that cannot be used by lower-priority selection classes.

QTF Selection Classes

Selection classes are designated by a single letter in the range from A to L. Selection class A is the highest priority, and L is the lowest priority.

QTF uses selection classes to find files to assign to the PIDs (physical identifiers) for remote hosts. Associated with each PID is one or more LIDs (logical identifiers). The system PID/LID table is created by system program CLDT from the LID configuration file LIDCMid (refer to LID Configuration File in section 10) and is viewed and modified using the L-display utility LIDOU (refer to LIDOU L Display in section 9). QTF reads the system PID/LID table to determine which PIDs have connections through RHF or NAM and which LIDs are associated with each PID.

File Selection Process

Each copy of QTF has a fixed number of connection slots available. The goal of the QTF file scheduling algorithm is to find files for as many of these connection slots as possible within the constraints of the selection class definitions. On each pass through its file selection process, QTF proceeds through all defined selection classes from the highest priority A to the lowest priority L. For each selection class, QTF attempts to find one file for each available PID that does not already have a file assigned. The file selection process continues until one of the following occurs:

- All connection slots are in use, or
- All PID/selection class combinations have been checked and
 - Are in use, or
 - Are unavailable (disabled or the maximum number of connections per selection class has been reached), or
 - No files were found.
- The sum of the reserved connections for higher-priority selection classes is equal to or exceeds the number of free connection slots.

Visualizing the File Selection Process

To visualize the file selection process, imagine that selection classes and PIDs form a two-dimensional matrix of jacks on a telephone switchboard with PIDs assigned across the columns and selection classes assigned down the rows. Above each jack is a light. When a light is flashing, a queued file is available for the corresponding PID/selection class combination. The plug cables represent the connection slots. A jack can have only one plug inserted at a time. Inserting a plug into a jack establishes a connection and causes the light to stop flashing and a file transfer to begin. A plug is removed from a jack after one or more files have been transferred or a connection error occurs.

Figure 8-14 shows the connection slot array (plugs) and the PID/selection class matrix (switchboard).




Figure 8-14. Connection Slot Array and PID/Selection Class Matrix

On each file selection cycle, QTF looks for a free plug cable (an unused connection slot) and, starting at the upper left corner of the PID/selection class switchboard, searches across each row in turn looking for a flashing light (queued file matching the PID/selection class combination). When the plug is inserted into the jack, that connection is established. If another plug cable is free, the search continues from the current jack until QTF runs out of plugs or the bottom right corner of the switchboard is reached.

For each selection class row for which a reserved number of connections is defined, QTF checks the number of plugs inserted in that row when the rightmost jack for the row is reached. If the number of plugs inserted is less than the number of reserved connections, QTF places enough plugs in unused jacks in that row until the number of reserved jacks is filled, only one free plug remains, or the row is filled; whichever occurs first.

For each selection class row for which a maximum number of connections is defined and the maximum is reached, QTF continues searching at the leftmost jack of the next row down.

You can minimize the number of different operating system calls QTF must make to cover all of the switchboard jacks. Do this by either minimizing the number of selection classes (number of rows) or by defining overlapping selection criteria for selection classes. For example, if selection classes A, B, and C all have the same file size index range, and QTF does not find a file for a PID in selection class row A, QTF can consider selection class rows B and C checked as well. If selection class D has a different, nonoverlapping file size index range, QTF must call the operating system to check the jacks in selection class D regardless of the results of the checks for selection classes A, B, and C (assuming that the file selection process has not yet terminated).

Queued File Size Index

QTF file size selection criteria are specified in terms of a queued file size index. The NOSTEXT symbols FSI1 through FSI6, defined in NOS common deck PPCOM (refer to the NOS Version 2 Installation Handbook), allow you to completely and uniquely partition the range of queued file sizes into seven subranges. Each subrange is identified by an integer value in the range from 1 to 7, known as a file size index. The FSI n values define end points for each file size subrange in units of PRUs (640 6-bit characters).

FSI File size subrange

1	$0 \leq \text{size} < \text{FSI1}$
2	$\text{FSI1} \leq \text{size} < \text{FSI2}$
3	$\text{FSI2} \leq \text{size} < \text{FSI3}$
4	$\text{FSI3} \leq \text{size} < \text{FSI4}$
5	$\text{FSI4} \leq \text{size} < \text{FSI5}$
6	$\text{FSI5} \leq \text{size} < \text{FSI6}$
7	$\text{FSI6} \leq \text{size}$

Each NOS queued file is assigned a file size index value based on its size. This value is displayed on the DSD Q displays (refer to the NOS Version 2 Operations Handbook).

You can examine the NOS common deck PPCOM to determine the default values for the FSI n symbols. However, to modify the current values, you need to rebuild NOSTEXT, the NOS operating system, and QTF under RHP. For each FSI n symbol, the value of FSI n must be less than the value of FSI $n+1$.

The file size ranges are displayed on the QTF help display (figure 8-15).

Queued File Size and Transfer Times

Given a typical set of file size index values, the following shows the approximate file transfer time in minutes for various line speeds. The time is based on typical line utilization and only one transfer active through the line. NAM line speed is shown in minutes.

FSI	File Size (PRUs)	1200 bps	9600 bps	19.2k bps	56k bps	RHF/ LCN
1	319	21	1.5	<1.0	<1.0	<0.5
2	511	34	4.2	2.5	<1.0	<0.5
3	1023	68	8.5	4.2	1.5	<0.5
4	2047	136	17.0	8.5	3.0	<0.5
5	4095	272	34.0	17.0	6.0	<0.5
6	8191	544	68.0	34.0	12.0	<0.5
7	999999	>544	>68.0	>34.0	>12.0	>0.5

QTF K-Display Usage

The following procedures describe DSD K-display usage for QTF operations. QTF is used by both the RHF and NAM subsystems.

QTF Operation Under RHF

Normally, QTF is defined as an autostart application in the RHF configuration file (refer to RHF Configuration Files in section 10). That means a QTF job starts automatically when the RHF subsystem is initiated. The RHF application table K display shows how many copies of QTF are defined and how many copies are currently active.

QTF terminates when the K.IDLE command terminates the RHF subsystem. To manually terminate a copy of QTF, bring up the RHF application table display by entering the following commands:

```
K, RHF .  
K .APPL .
```

When the display appears, enter:

```
DISABLE, ord.
```

where ord is the ordinal of an active QTF in the RHF application table display. The message QTF, NETWORK IDLEDOWN IN PROGRESS should appear in the QTF job dayfile. QTF completes any files in progress and then ends. If QTF does not terminate within a reasonable amount of time, use the DSD DROP command to terminate QTF.

If a fault condition occurs, QTF normally restarts itself. To manually restart QTF, bring up the RHF application table display by entering the following commands:

```
K, RHF .  
K .APPL .
```

When the display appears, enter:

```
ENABLE, ord.
```

where ord is the QTF ordinal in the RHF application table display.

You can determine the jsn of QTF by using the DSD system status display (B,O) and the RHF application table display. When QTF is active at a control point, it is indicated on the B,O display.

Activate the K display for that QTF by entering:

```
K, jsn.
```

where jsn is the job sequence name for QTF. The QTF transfer display appears on the left screen.

QTF Operation Under NAM

Normally, QTF is started by NAMI during NAM subsystem initiation. The NAM status K display shows how many copies of QTF are currently active.

QTF terminates automatically when the NAM subsystem is terminated. To manually terminate a copy of QTF, use the NVF IDLE or DISABLE commands (refer to NVF Control Commands in section 15) to either idle QTF or immediately terminate QTF. Enter the following commands:

```
K, NAM.
K. * .
K. AP=NVF.
K. IDLE, AP=QTF.   or   K. DISABLE, AP=QTF.
```

The message QTF, NETWORK IDLEDOWN IN PROGRESS or QTF, NETWORK SHUTDOWN should appear in the QTF job dayfile. If the IDLE,AP=QTF command is used, QTF completes any files in progress and then ends. If QTF does not terminate within a reasonable amount of time, use the DISABLE,AP=QTF command or DSD DROP command to terminate QTF.

If a fault condition occurs, QTF normally restarts itself. To manually restart QTF, enter the following command (refer to the NAMI command in section 16):

```
X. NAMI (RS=QT)
```

You can determine the jsn of QTF by using the NAM status K display. Bring up the NAM status K display by entering the following commands:

```
K, NAM.
K. * .
K. ST.
```

Activate the K display for that QTF by entering the following command:

```
K, jsn.
```

where jsn is the job sequence name for QTF. The QTF transfer display appears on the left screen.

QTF K Display Commands

QTF commands have the same general format whether they are entered through the K display or read from an input file using the INCLUDE command.

- Each QTF command must be on one line and only the first 80 columns of a command line are significant.
- A command consists of a command verb followed by 0 to 10 parameters.
- A command verb and parameters are separated by one or more spaces or a comma.
- A command is terminated by a single period.
- Most command verbs and parameters can be abbreviated. The exceptions are IDLE and STOP.

QTF accepts the following K display commands:

Command	Description
CLASS	Define or modify QTF selection class values.
DISABLE	Change the status of QTF elements from enabled to disabled.
ENABLE	Change the status of QTF elements from disabled to enabled.
HELP	Displays information about the available QTF commands.
IDLE	Begins the idle-down process of QTF.
INCLUDE	Causes QTF commands to be read from an input file.
SCHED	Modify QTF scheduling parameters.
STATUS	Displays information about the status of file transfers in progress, PIDs, or QTF selection classes.
STOP	Causes an immediate termination of QTF.
+	Page forward to the next screen of a multipage display.
-	Page backward to the previous screen of a multipage display.
.	Comment line (anything following a period is ignored).

CLASS Command

The CLASS command allows you to define or modify QTF selection class values.

The command format is:

```
K.CLASS, SC=selclass, FSI=fsi, MAXIMUM=max, RESERVE=res.
```

Parameter Description

selclass	Selection class name; specify one character in the range from A to L. Selection class A is the highest priority class and L is the lowest priority class. This parameter is required.								
fsi	File size index; specify the file size index or range of indexes using one of the following formats: <table border="1" data-bbox="380 646 1343 835"> <thead> <tr> <th>Format</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>fsi</td> <td>File size index in the range from 1 to 7.</td> </tr> <tr> <td>fsil..fsiu</td> <td>File size index range, where fsil ≤ fsiu.</td> </tr> <tr> <td>*</td> <td>An * specifies all file sizes. It is equivalent to entering 1..7. If FSI is omitted, the default is FSI=*</td> </tr> </tbody> </table>	Format	Description	fsi	File size index in the range from 1 to 7.	fsil..fsiu	File size index range, where fsil ≤ fsiu.	*	An * specifies all file sizes. It is equivalent to entering 1..7. If FSI is omitted, the default is FSI=*
Format	Description								
fsi	File size index in the range from 1 to 7.								
fsil..fsiu	File size index range, where fsil ≤ fsiu.								
*	An * specifies all file sizes. It is equivalent to entering 1..7. If FSI is omitted, the default is FSI=*								
max	Maximum connections; specify the maximum number of simultaneous connections for files in this selection class in the range from 0 to the maximum available connections (specified by the MAXCONS parameter in the QTF SCHED command) or specify *. If 0 or * is specified, all connections can be used for files in this selection class. However, only one connection per PID per selection class is permitted. Default is 0 or *.								
res	Reserved connections; specify the number of connections reserved for files in this selection class in the range from 0 to the maximum available connections (specified by the MAXCONS parameter in the QTF SCHED command). If zero is specified, no connections are reserved for this selection class. For each selection class, the number of available connections is reduced by the sum of reserved connections for all higher-priority selection classes. To prevent a possible overcommitment of connections, QTF allows only mxcons-1 connections to be reserved. Default is zero.								

The following default selection class is defined at QTF initiation:

```
CLASS, SC=A, FSI=*, MAXIMUM=*, RESERVE=0.
```

which allows one file transfer to any one remote host with no restrictions on file size and no reserved connections. To modify the default selection class or define other classes, use additional CLASS commands.

DISABLE Command

The **DISABLE** command changes the status of the selected QTF elements from enabled to disabled.

The command format is:

K.DISABLE, element.

Parameter	Description						
element	One of the following: PID=pid SC=selclass SC=selclass,PID=pid						
<table border="1"> <thead> <tr> <th>Option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>pid</td> <td>Physical identifier of a given mainframe. Specify PID=* to disable all PIDs.</td> </tr> <tr> <td>selclass</td> <td>Selection class in the range from A to L. Specify SC=* to disable all selection classes.</td> </tr> </tbody> </table>		Option	Description	pid	Physical identifier of a given mainframe. Specify PID=* to disable all PIDs.	selclass	Selection class in the range from A to L. Specify SC=* to disable all selection classes.
Option	Description						
pid	Physical identifier of a given mainframe. Specify PID=* to disable all PIDs.						
selclass	Selection class in the range from A to L. Specify SC=* to disable all selection classes.						

ENABLE Command

The **ENABLE** command changes the status of the selected QTF elements from disabled to enabled.

The command format is:

K.ENABLE, element.

Parameter	Description						
element	One of the following: PID=pid SC=selclass SC=selclass,PID=pid						
<table border="1"> <thead> <tr> <th>Option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>pid</td> <td>Physical identifier of a given mainframe. Specify PID=* to enable all PIDs.</td> </tr> <tr> <td>selclass</td> <td>Selection class in the range from A to L. Specify SC=* to enable all selection classes.</td> </tr> </tbody> </table>		Option	Description	pid	Physical identifier of a given mainframe. Specify PID=* to enable all PIDs.	selclass	Selection class in the range from A to L. Specify SC=* to enable all selection classes.
Option	Description						
pid	Physical identifier of a given mainframe. Specify PID=* to enable all PIDs.						
selclass	Selection class in the range from A to L. Specify SC=* to enable all selection classes.						

HELP Command

The HELP command displays the available QTF K display commands, file size ranges in decimal PRUs, and the parameters for the QTF SCHED command.

The command format is:

K.HELP.

The HELP K display (figure 8-15) appears on the right screen.

HELP	FSI	SIZE RANGE
COMMANDS		
CLASS, SC=A, FSI=N. .N, MAXIMUM=N, RESERVE=N.	1	0.. 319
DISABLE, ELEMENT.	2	320.. 511
ENABLE, ELEMENT.	3	512.. 1023
IDLE.	4	1024.. 2047
INCLUDE, FILE=PFN.	5	2048.. 4095
SCHED, PARAMETERS.	6	4096.. 8191
STATUS, TYPE. (TYPE=PID, SC, OR TRANSFER)	7	8192..999999
STOP.		
+ OR -	PARAMETERS =	(FILE SIZE RANGES IN DECIMAL PRUS)
	SELSECS=N	
ELEMENT = ONE OF-	MAXCONS=N	
SC=X OR *	TIMEOUT=N	
PID=XXX OR *	DISABLE=N	
SC=X, PID=XXX	RETRY=N	
X=A, B, C, . . . , L.	LDTRFSH=N	

Figure 8-15. QTF HELP K Display (Right Screen)

IDLE Command

The IDLE command causes QTF to idle down and terminate. The message QTF, NETWORK IDLEDOWN IN PROGRESS appears in the QTF job dayfile and IDLEDOWN appears on the B,0 display. QTF completes any file transfers in progress and then ends.

The command format is:

K.IDLE.

QTF does not accept an abbreviation for this command.

To restart QTF, enter the following command:

X.NAMI (RS=QT)

INCLUDE Command

The INCLUDE command causes QTF to rewind the specified input file and begin processing QTF commands from that file.

The command format is:

```
K.INCLUDE,FILE=infile.
```

Parameter Description

infile	File name from which QTF commands are to be processed. QTF searches for a local file with name infile. If the local file is not found, QTF attempts to GET or ATTACH permanent file infile from user name SYSTEMX. If QTF performs a GET or ATTACH, it returns the file after processing all of the commands. This parameter is required.
--------	---

A file of QTF commands processed by using the INCLUDE command can itself contain another INCLUDE command. Upon executing the second INCLUDE command, QTF switches to the new file and does not read anything following the INCLUDE command in the old file.

SCHEM Command

The SCHEM command allows you to modify QTF scheduling parameters.

The command format is:

```
K.SCHED,SELSECS=selsecs,MAXCONS=mxcons,TIMEOUT=tmomins,
  DISABLE=dismins,RETRY=maxrtry,LDTRFSH=ldtmins.
```

Parameter Description

Parameter	Description
selsecs	File selection recall period in the range from 1 to 4095 seconds; specifies the minimum amount of time between attempts to acquire new queued files. Whenever QTF is unable to fill all connection slots with files, it waits selsecs seconds before again attempting to acquire files for unused slots. Default is 20 seconds.
mxcons	Maximum simultaneous connection slots (to all hosts) in the range from 1 to 4. Default is 4 connections.
tmomins	Timeout period in the range from 1 to 60 minutes; specifies the maximum amount of time QTF waits for a reply from a remote host or access method. For the RHF variant of QTF, tmomins also specifies how long QTF waits for NAD code conversion resources when applicable. Default is 2 minutes.
dismins	Disable period in the range from 1 to 60 minutes; specifies the amount of time for which QTF will not attempt transfers to a PID after an error occurs. You can use the ENABLE command to clear the error disabled condition. Default is 2 minutes.
maxrtry	Maximum retries in the range from 1 to 50 retries; specifies the maximum number of immediate retries that QTF attempts to complete a file transfer before error disabling the PID. When QTF cannot complete a connection request (the connection is rejected by the local or remote access method or QTFS), QTF disables the PID without any immediate retrys. Once the PID is disabled, QTF returns the file to the local queue and does not attempt to acquire any files for LIDs associated with the PID until either dismins minutes have elapsed or you use the ENABLE,PID command to enable the PID. Default is 2 retries.
ldtmins	PID/LID table refresh period in the range from 1 to 60 minutes; specifies how often QTF updates its internal tables from the system PID/LID table. Default is 2 minutes.

STATUS Command

The STATUS command displays the status of selected QTF elements or file transfers.

The command format is:

```
K.STATUS,display.
```

display	Description
PID	Selects the QTF physical identifier K display which shows the status of PIDs, selection classes for each PID, and the number of queued files transferred.
SC	Selects the QTF selection class K display which shows the selection class definitions and the number of files transferred.
TRANSFER	Selects the QTF transfer K display which shows transfers in progress and the number of files already transferred.

Each QTF status display consists of a two-line display header and a variable number of display items. The lower portion of the screen contains a QTF command and message area. Figure 8-16 shows the general format for the QTF status displays. The display fields are described following the figure.

```

STATUS,type          nnnnnn FILES SENT    xxx PAGE n OF m
display header
display item 1
display item 2
.
.
.
display item k
*
COMMAND- <----- last cmd (1-50) ----->
idledown  <----- last cmd (51-80) ----->
<----- opr msg ----->
    
```

Figure 8-16. General Format of QTF Status K Display

Field	Description
type	Display type; PID, SC, or TRANSFER.
nnnnnn	Number of queued files transferred.
xxx	Access method; RHF or NAM.
n	Current page number.
m	Maximum page number.
last cmd	The last QTF command processed. If the command is incorrect, an asterisk (*) appears in the line above the command to show the approximate position of the character or parameter in error. If the error is in the second part of the command (positions 51 to 80), a plus sign (+) appears instead of an asterisk.
idledown	This field contains the message *IDLEDOWN* if QTF is in idle-down mode; otherwise, this field is blank.
opr msg	This field contains an operator message about the last QTF command processed. Refer to the NOS Version 2 Operations Handbook for message descriptions.

QTF Physical Identifier K Display

Figure 8-17 shows a sample QTF STATUS,PID K display.

STATUS, PID		nnnnnn	FILES SENT	RHF	PAGE 1 OF 1
PID	STATUS	SC DISABLED	SC IN-USE		
M01	ENABLED	-----	AB-----		
M02	DISABLED	-----	-----		
M03	NO LIDS	A-----L	-----		
M04	NO PATH	-----	-----		
M05	REJECTING	-----	-----		

Figure 8-17. QTF STATUS,PID K Display

Each entry in the QTF physical identifier display appears in the following format:

```
pid  status  sc disabled  sc in-use
```

Header	Description
pid	Physical identifier of the remote host specified in the NOS system LID table.
status	Status of the PID; status can be one of the following:
status	Description
DISABLED	The PID was disabled by the QTF DISABLE command. Use the QTF ENABLE command to change the PID status to ENABLED.
ENABLED	The PID is available for file selection for enabled selection classes.
NO LIDS	No enabled LIDs are associated with this PID in the system LID table or the PID is disabled. Use the LIDOU L-display utility to display and modify the system LID table.
NO PATH	The system LID table does not define a path in service to this PID through the access method (RHF or NAM) that is available to this copy of QTF. Use the LIDOU L-display utility to determine whether an RHF or NAM path to the PID is defined and its status. If the path status is OUT OF SERVICE, use the RHF network path table display or the NAM status display to display the status of the network elements linking the two hosts.

Header	Description								
status	<i>(Continued)</i>								
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>REJECTING</td> <td> <p>QTF has temporarily disabled file transfers to the PID because QTF was unable to establish a connection to the PID or was unable to successfully transfer a file after the maximum number of retries. QTF returns the PID to ENABLED status after the specified disable period has elapsed. (Refer to the QTF SCHED command for descriptions of the RETRY and DISABLE parameters.) You can use the QTF ENABLE command to immediately return the PID to ENABLED status. If the PID returns to REJECTING status, examine the QTF dayfile to determine whether a connection reject or a particular file is causing the problem.</p> <p>Currently, for RHF paths and for some types of NAM paths, the path status in the system LID table does not reflect the status of the remote PID or access method but rather the status of the local host network elements used to access the PID. Thus, a remote host may be out-of-service (due to scheduled maintenance, for example), yet the local system LID table shows the path status as IN SERVICE. Each time the disable period elapses, QTF attempts to send files to the PID. If the access method rejects the connection, QTF returns the PID to REJECTING status.</p> </td> </tr> <tr> <td>sc disabled</td> <td>Selection classes that are unavailable for this PID. They were disabled by the QTF DISABLE command. Use the QTF ENABLE command to enable a particular PID and selection class combination.</td> </tr> <tr> <td>sc in-use</td> <td>Selection classes currently in use for this PID. The QTF transfer display shows the ordinals for connection slots and the selection class associated with each slot.</td> </tr> </tbody> </table>	status	Description	REJECTING	<p>QTF has temporarily disabled file transfers to the PID because QTF was unable to establish a connection to the PID or was unable to successfully transfer a file after the maximum number of retries. QTF returns the PID to ENABLED status after the specified disable period has elapsed. (Refer to the QTF SCHED command for descriptions of the RETRY and DISABLE parameters.) You can use the QTF ENABLE command to immediately return the PID to ENABLED status. If the PID returns to REJECTING status, examine the QTF dayfile to determine whether a connection reject or a particular file is causing the problem.</p> <p>Currently, for RHF paths and for some types of NAM paths, the path status in the system LID table does not reflect the status of the remote PID or access method but rather the status of the local host network elements used to access the PID. Thus, a remote host may be out-of-service (due to scheduled maintenance, for example), yet the local system LID table shows the path status as IN SERVICE. Each time the disable period elapses, QTF attempts to send files to the PID. If the access method rejects the connection, QTF returns the PID to REJECTING status.</p>	sc disabled	Selection classes that are unavailable for this PID. They were disabled by the QTF DISABLE command. Use the QTF ENABLE command to enable a particular PID and selection class combination.	sc in-use	Selection classes currently in use for this PID. The QTF transfer display shows the ordinals for connection slots and the selection class associated with each slot.
status	Description								
REJECTING	<p>QTF has temporarily disabled file transfers to the PID because QTF was unable to establish a connection to the PID or was unable to successfully transfer a file after the maximum number of retries. QTF returns the PID to ENABLED status after the specified disable period has elapsed. (Refer to the QTF SCHED command for descriptions of the RETRY and DISABLE parameters.) You can use the QTF ENABLE command to immediately return the PID to ENABLED status. If the PID returns to REJECTING status, examine the QTF dayfile to determine whether a connection reject or a particular file is causing the problem.</p> <p>Currently, for RHF paths and for some types of NAM paths, the path status in the system LID table does not reflect the status of the remote PID or access method but rather the status of the local host network elements used to access the PID. Thus, a remote host may be out-of-service (due to scheduled maintenance, for example), yet the local system LID table shows the path status as IN SERVICE. Each time the disable period elapses, QTF attempts to send files to the PID. If the access method rejects the connection, QTF returns the PID to REJECTING status.</p>								
sc disabled	Selection classes that are unavailable for this PID. They were disabled by the QTF DISABLE command. Use the QTF ENABLE command to enable a particular PID and selection class combination.								
sc in-use	Selection classes currently in use for this PID. The QTF transfer display shows the ordinals for connection slots and the selection class associated with each slot.								

QTF Selection Class K Display

Figure 8-18 shows a sample QTF STATUS,SC K display.

STATUS, SC	nnnnnn	FILES SENT	NAM	PAGE 1 OF 1		
SC	STATUS	FSI RANGE	MAXIMUM	RESERVED	IN-USE	AVAILABLE
A	ENABLED	1 .. 6	----	----	1	2
B	ENABLED	1 .. 6	----	----	1	2
C	DISABLED	7 .. 7	1	----	----	----
D	UNDEFINED					
E	UNDEFINED					
F	UNDEFINED					
G	UNDEFINED					
H	UNDEFINED					
I	UNDEFINED					
J	UNDEFINED					
K	UNDEFINED					
L	UNDEFINED					

Figure 8-18. QTF STATUS,SC K Display

Each entry in the QTF selection class display appears in the following format:

sc status fsi range maximum reserved in-use available

Header	Description								
sc	Selection class.								
status	Status of the selection class; status can be one of the following: <table border="1" data-bbox="402 1045 1360 1325"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DISABLED</td> <td>The selection class was disabled by the QTF DISABLE command. Use the QTF ENABLE command to change the selection class status to ENABLED.</td> </tr> <tr> <td>ENABLED</td> <td>The selection class is available for file selection for enabled PIDs.</td> </tr> <tr> <td>UNDEFINED</td> <td>The selection class has not been defined. Use the QTF CLASS command to define or modify a selection class.</td> </tr> </tbody> </table>	status	Description	DISABLED	The selection class was disabled by the QTF DISABLE command. Use the QTF ENABLE command to change the selection class status to ENABLED.	ENABLED	The selection class is available for file selection for enabled PIDs.	UNDEFINED	The selection class has not been defined. Use the QTF CLASS command to define or modify a selection class.
status	Description								
DISABLED	The selection class was disabled by the QTF DISABLE command. Use the QTF ENABLE command to change the selection class status to ENABLED.								
ENABLED	The selection class is available for file selection for enabled PIDs.								
UNDEFINED	The selection class has not been defined. Use the QTF CLASS command to define or modify a selection class.								
fsi range	File size index range. This field shows the inclusive range of file size index values that are allowed for this selection class.								
maximum	Maximum number of connection slots allowed for files in this selection class. This field contains dashes (—) if no maximum is specified.								
reserved	Number of connection slots reserved for files in this selection class. This field contains dashes (—) if no connection slots are reserved.								

Header	Description
in-use	Number of connection slots that have been assigned files in this selection class. This field contains dashes (—) if no connection slots are in use.
available	Number of free connection slots available for assigning files in this selection class. This field contains dashes (—) if no connection slots are available.

QTF Transfer K Display

Figure 8-19 shows a sample QTF STATUS,TRANSFER K display.

STATUS,TRANSFER		nnnnnn	FILES SENT	RHF	PAGE 1 OF 1			
ACN	SC	FILE	LID/PID	DC	DD	PRUS	STARTED	STATUS
1	A	ABCD	SV6 M03	PR	C8	1234	11.03.45	*START XFR CONNECTION ESTABLISHED TO M03.
2	A	AAED	MFF MFF	IN	US	22	11.10.00	WAIT REM ACQUIRED DC=IN, ST=MFF, DO=MA2.
3								
4								

Figure 8-19. QTF STATUS,TRANSFER K Display

Each entry in the QTF transfer display appears in the following format:

```
acn  sc  file  lid/pid  dc  dd  prus  started  status
      message
```

Header	Description
acn	Application connection number; this is the ordinal of the connection slot.
sc	Selection class associated with this connection slot.
file	Job sequence name associated with the queued file assigned to this connection slot.
lid	Logical identifier of the host associated with this connection slot.
pid	Physical identifier of the remote host to which this slot is connected or is being connected.
dc	Disposition code of the queued file. Refer to the ROUTE command in the NOS Version 2 Reference Set, Volume 3 for a list of disposition codes and their meanings.
dd	Data declaration. This field defines the format of the file data during transfer. Values of C6 and C8 indicate that the file is to be sent in code-conversion mode. Values of UH, US, and UU indicate that a binary transfer is to be used.

Header	Description																						
prus	Size of the file in physical record units (PRUs) in decimal. One PRU contains the equivalent of 640 6-bit characters.																						
started	This field shows the time (hh.mm.ss) at which the connection to the remote host was initiated. The time is updated when the connection is established and at the beginning of each file transfer sequence.																						
status	Status of the connection or file transfer; status can be one of the following: <table border="1" data-bbox="402 464 1356 825"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ACQ FILE</td> <td>File has been acquired.</td> </tr> <tr> <td>CONNECT</td> <td>Issuing connect request.</td> </tr> <tr> <td>END CONN</td> <td>Connection terminating.</td> </tr> <tr> <td>FILE XFR</td> <td>File transfer in progress.</td> </tr> <tr> <td>IDLE</td> <td>Connection established.</td> </tr> <tr> <td>QUEUED</td> <td>File transfer successful.</td> </tr> <tr> <td>START XFR</td> <td>Initiating file transfer.</td> </tr> <tr> <td>STOP XFR</td> <td>File transfer completing.</td> </tr> <tr> <td>WAIT ACQ</td> <td>Waiting for acquire queued file response.</td> </tr> <tr> <td>WAIT REM</td> <td>Waiting for response from remote host or access method.</td> </tr> </tbody> </table>	status	Description	ACQ FILE	File has been acquired.	CONNECT	Issuing connect request.	END CONN	Connection terminating.	FILE XFR	File transfer in progress.	IDLE	Connection established.	QUEUED	File transfer successful.	START XFR	Initiating file transfer.	STOP XFR	File transfer completing.	WAIT ACQ	Waiting for acquire queued file response.	WAIT REM	Waiting for response from remote host or access method.
status	Description																						
ACQ FILE	File has been acquired.																						
CONNECT	Issuing connect request.																						
END CONN	Connection terminating.																						
FILE XFR	File transfer in progress.																						
IDLE	Connection established.																						
QUEUED	File transfer successful.																						
START XFR	Initiating file transfer.																						
STOP XFR	File transfer completing.																						
WAIT ACQ	Waiting for acquire queued file response.																						
WAIT REM	Waiting for response from remote host or access method.																						
message	This field contains the last dayfile message issued by QTF for this connection slot. Refer to the NOS Version 2 Operations Handbook for message descriptions.																						

STOP Command

The STOP command causes QTF to terminate immediately. The message QTF, NETWORK SHUTDOWN appears in the QTF job dayfile. QTF abandons any file transfers in progress, returns files to the system I/O queues, and ends.

The command format is:

```
K.STOP.
```

QTF does not accept an abbreviation for this command.

REDEFINE K Display

Use the REDEFINE utility to reconfigure 844 and 885-11/12³ disk drives online and thereby logically eliminate a unit that is malfunctioning without performing a level 0 initial deadstart. A single unit or multiunit device⁴ that fails can be replaced with an unused unit. Unused units must be available or made available across channels or on the same channel as the failing unit by physically moving the disk pack from the failing unit to the replacement unit. Table 8-8 gives the equipment requirements for reconfiguration.

Table 8-8. Equipment Requirements for Reconfiguration

Operation To Be Performed	Current Equipment	Replacement	Equipment	Special Considerations
	Must be unloaded, removable	Must be in EST	Must be unloaded, removable	
Add or Return a Unit	Yes	No	Yes	Equipment must have less than eight units.
Delete a Unit	Must be removable; if not unloaded, the unit must be in the EST	NA	NA	Equipment must have at least one unit.
Replace a Unit	No	Not necessary	Yes, if in EST	The system does not perform label verification on packs that are unloaded before the failing unit is replaced. If the unloaded pack is not moved from the failing unit or the wrong pack is moved, the system issues an error message when the replacement unit is accessed.
Recable a Unit	No	NA	No	None.

3. When reconfiguring an 885-11/12 disk drive, a customer engineer must be present.

4. When reconfiguring a multiunit device, all units of the device must be on the same controller(s).

NOTE

There are special cases when you cannot perform an online reconfiguration. In these cases, you must perform a level 0 deadstart and define the failing device to another drive. The special cases are:

- Reconfiguring an 885-11/12 disk drive that is the only system device defined.
 - Reconfiguring an 844 or 885-11/12 disk drive that is the only temporary device defined.
 - Reconfiguring a device designated as an independent shared device.
 - Reconfiguring an 819 or 885-42 disk drive.
-

The reconfiguration sequence follows:

1. Request the reconfiguration of the disk device defined by EST ordinal est by entering:

```
REDEFINE, est .
```

The message REQUEST *K* DISPLAY appears at the appropriate control point on the system status display (B,0). You can reenter the REDEFINE command as many times as there are devices to be reconfigured. Multispindle devices that include two or more units are considered to be one device.

If an INITIALIZE command is being processed, the REDEFINE entry is not processed until the initialization is complete.

2. Bring the initial REDEFINE K display (see figure 8-20) to the left console screen by entering:

```
K, jsn .
```

jsn	Description
jsn	Job sequence name of the requesting job.

To bring up the right console screen K display, enter:

```
KK .
```

As shown in the left screen of figure 8-20, the current EST description of the device being reconfigured appears under the heading CURRENT EQUIPMENT CONFIGURATION. The EST description of the device as changes are made appears under the heading REQUESTED EQUIPMENT CONFIGURATION. Any devices listed under IDLED EQUIPMENTS are devices that have been previously selected by your REDEFINE command but have yet to be processed in the reconfiguration run.

The left screen K display may not list all relevant devices. If more than one device is listed, they are processed one at a time as they appear in the list with one exception: all shared devices are processed prior to nonshared devices.

You are guided through the reconfiguration process in two ways. First, by the list of commands and parameters shown on the right screen K display of figure 8-21, and second, by the system responses and error messages.

```

K.                CMS
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
LVL0   - LVL7      MID=AA   NOS version

                MASS STORAGE FAILURE RECOVERY

CURRENT EQUIPMENT CONFIGURATION
  011 DJ-1  CH24,01  UL=01  -- -- -- -- -- ST = I---

REQUESTED EQUIPMENT CONFIGURATION
  011 DJ-1  CH24,01  UL=01  -- -- -- -- -- ST = I---
      CHANGED UNITS

IDLED EQUIPMENTS
  010 DJ-1  CH01,24  UL=00  -- -- -- -- -- ST = I---

```

Figure 8-20. REDEFINE K Display (Left Screen)

```

K DISPLAY COMMANDS
CLEAR          CLEAR IDLE AND SUSPEND ON CURRENT EQUIPMENT.
END            CLEAR ALL DEVICE IDLES AND END *CONFIG*.
GO             INITIATE PROCESSING OF ENTERED CONFIGURATION.
RERUN         RESTART *CONFIG* UTILITY PROCESSING.
RESET         RESET CURRENT EQUIPMENT TO DEFAULT PARAMETERS.
SUSPEND       SUSPEND SYSTEM OPERATION ON CURRENT EQUIPMENT.
CH=C1,C2     C1 AND OPTIONALLY C2 ARE NEW ACCESS CHANNELS.
EQ=EEE       SET EST ORDINAL EEE TO BE PROCESSED.
UL=U1,U2,... SET UNIT LIST AS SPECIFIED.
UR=UU        UNIT UU IS TO BE RECABLED WITH A NEW DRIVE.

```

Figure 8-21. REDEFINE K Display (Right Screen)

3. Enter all of the valid parameters (refer to table 8-9) that define the reconfiguration characteristics for the specified device. Parameters are entered singly. Parameters are processed left to right; no terminator is necessary.

If you enter an incorrect parameter (for example, the wrong channel number), the error can be corrected by entering the correct parameter.

4. Enter

K.GO

to initiate processing of the parameters when all parameters and commands (refer to table 8-10) have been entered for the specific device. Enter this command after each set of parameters to signal the system to go ahead with the reconfiguration you have defined. If more devices remain to be reconfigured, repeat steps 3 and 4.

5. Enter

K.END

to end a reconfiguration run when there are no more devices to be processed.

Table 8-9. Reconfiguration Parameters

Parameter	Description
CH=c ₁ ,c ₂	<p>Specifies the channel numbers (c₁ and, optionally, c₂) to be used under the new device definition.</p> <p>Use this parameter when the entire device is to be redefined to alternate channel number(s). You can also use it to add or delete channels from an equipment definition. New channels can be added by specifying new channel numbers for a defined equipment. Channels can be deleted by specifying CH=. when deleting a unit. Acceptable values for c₁ and c₂ are from 0 to 13₈ for systems having 10 or less PPs; from 0 to 13₈ and from 20₈ to 33₈ for systems having more than 10 PPs. Leading zeros can be omitted.</p>
EQ=est	<p>Specifies the equipment with EST ordinal est is to be processed. Enter this parameter when the processing of equipments is order dependent; for example, when a device must first be made available before it can replace a failing device. The system ignores all parameters not processed before you enter the EQ=est parameter. Parameters entered after EQ=est refer to the specified equipment until you enter another EQ=est or a command that causes the next equipment in the list to be selected for processing.</p>
UL=u ₁ ,u ₂ ,...,u _n	<p>Specifies the unit list for the new configuration.</p> <p>When a unit is to be replaced, added, or deleted, the entire unit configuration must be entered with this parameter. The equipment must be unloaded (multimainframe mode) or otherwise have an unavailable status if a unit is to be added or deleted. Any number of units can be changed. By specifying UL=. the current unit configuration is deleted.</p>
NOTE	
<p>If the unit number specified in the UL= parameter represents a unit that is not defined in the EST, the unit number is accepted without validation. Ensure that the unit number entered represents a valid device.</p>	
UR=un	<p>Specifies that the device with unit number un is to be recabled. This parameter is used when a device is to be physically replaced by a new device with the same unit number. More than one device can be specified for recabling at the same time.</p>

Table 8-10. Reconfiguration Commands

Command	Description
ABORT	Discontinues processing of the current command. It is entered only in response to a detected error condition and is used to initiate error recovery procedures. This command cannot be followed immediately by an END command.
CLEAR	Clears the suspend and redefinition request status for the current equipment. Label and read/write verification of the device is performed. This command should be entered only when no further processing is desired for the current equipment.
END	Terminates the reconfiguration processing. The K-display parameters are set to the default values as control is returned. This command cannot be entered immediately after an ABORT command.
GO	Initiates the processing of specified reconfiguration parameters previously entered.
IGNORE	Informs the system on which the command was entered to ignore processing on this device (multimainframe mode only). This command should be entered during an add or delete unit reconfiguration run on the machine(s) within the multimainframe complex which, for control reasons, cannot add to or delete from the specified equipment. Also, you can use this command to ignore an error message pertaining to a marginally unacceptable servo timing check, which is performed automatically when reconfiguring an 885-11/12 disk drive.
NEXT	Enter this command in response to an 885-11/12 disk drive servo timing check that meets requirements. This command causes processing to continue with the next device or the next step of processing. This command is valid only when reconfiguring 885-11/12 disk drives.
RECHECK	Retries a verification/diagnostic process that previously gave an error. Only the commands RECHECK, ABORT, and IGNORE are accepted by the system when an error message is issued.
RERUN	Sets the parameters to default values and updates the list of equipment to be reconfigured.
RESET	Resets the parameters to default values.
SUSPEND	Causes system processing on the specified device to be suspended indefinitely while the device is in a not ready state. Only diagnostic access to the device is allowed. All other jobs accessing the device will be unable to continue until the device is returned to a ready state. More than one equipment can be suspended at the same time.

You can stop the reconfiguration procedure by entering either the **K.ABORT** command or the **K.CLEAR** command (depending on which command is posted in the list of valid commands on the K display). In either case, processing advances to the next device in the list of devices under **IDLED EQUIPMENTS** (see figure 8-20).

If the device specified is a shared device in a multimainframe environment and reconfiguration is not desirable on one or more of the mainframes, use the following procedure:

1. Enter at the console of each mainframe for which reconfiguration of the shared device is not desired.

```
REDEFINE,est.
```

est	Description
est	EST ordinal of the shared device.

2. Enter

```
K,jsn.
```

jsn	Description
jsn	Job sequence name of the job requesting the REDEFINE K display.

3. Enter

```
K.IGNORE
```

and processing on the shared device in the list is ignored by that mainframe. The machine must wait until the shared device is done with its processing.

4. Enter

```
K.END
```

to end the reconfiguration procedure.

Enter either the **RERUN** or **RESET** command to clear the **IGNORE** command.

Figure 8-22 shows the output for a sample reconfiguration run.

ORD	TYPE	CHANNELS	UNITS	STATUS
06	DJ-1	CH26,32	UL=06	ST = ----
07	DJ-1	CH26,32	UL=07	ST = ----
10	DJ-1	CH32	UL=01	ST = ----
11	DJ-1	CH13	UL=02	ST = ----
12	DI-1	CH26	UL=03	ST = ----
13	DI-1	CH32	UL=11	ST = ----
14	DI-1	CH26,32	UL=04	ST = I---

Figure 8-22. Reconfiguration Run Output

Following are examples of the reconfiguration of disk devices and how their status changes in the **REDEFINE K** display.

Example 1:

This example illustrates returning a unit to the system.

Assume the unit of est ordinal 7, the system's spare disk drive, was used to replace a failing disk drive. Now that the defective unit has been repaired, it is to be returned as the unit of the spare disk drive.

Disk configuration before the REDEFINE:

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
7	DJ-0	00	-	I—

Enter the following commands:

Command	Description
REDEFINE,7.	7 is the EST ordinal of the shared device.
K.jsn.	jsn is the job sequence name of the job requesting the REDEFINE K display.
K.CH=32,26	32 and 26 are the channel numbers to be used by est ordinal 7.
K.UL=7	7 is the unit number of the device being returned.
K.GO	Initiates the processing of all the parameters and commands entered for EST ordinal 7.

The system responds with the message:

EQ 7 REDEFINITION COMPLETE.

To end the reconfiguration, enter:

K.END

Disk configuration after the REDEFINE:

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
7	DJ-1	32, 26	7	—

Example 2:

This example illustrates reconfiguring a failing unit.

Assume EST ordinal 10 and EST ordinal 11 are defined in the EST display. Equipment 10 is a spare unit currently not being used. Equipment 11 has a pack mounted and is the failing device. Before the failing device can be reconfigured, the spare unit must be removed from the EST. The following stipulations apply when deleting a unit from the EST:

- If the EST ordinal is defined in a single mainframe environment or if it is not shared in a multimainframe environment, then having device unavailable status (U status in the E,M display) for that EST ordinal is sufficient.
- If the EST ordinal is shared in a linked shared device multimainframe environment, then that EST ordinal must have device unavailable status and be globally unloaded (U and N status in the E,M display) before a reconfiguration can be performed.

Disk configuration before the REDEFINE:

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
10	DJ-1	32	1	I—
11	DJ-1	13	2	—

To perform the reconfiguration, enter the following commands:

Command	Description
REDEFINE,10.	10 is the EST ordinal of the spare device.
REDEFINE,11.	11 is the EST ordinal of the failing device.
K.jsn.	jsn is the job sequence name of the job requesting the REDEFINE K display.
K.CH=.	Deletes the channel number(s) assigned to EST ordinal 10 from the EST.
K.UL=.	Deletes the unit number of EST ordinal 10 from the EST.
K.GO	Initiates the processing of all the parameters and commands entered for EST ordinal 10.

The system responds with:

EQ 10 REDEFINITION COMPLETE.

Enter the reconfiguration commands for the failing device as follows:

Command	Description
K.CH=32	32 is the channel number that was assigned to EST ordinal 10.
K.UL=1	1 is the unit number that was assigned to EST ordinal 10.
K.GO	Initiates the processing of all the parameters and commands entered for EST ordinal 11.

REDEFINE K Display

The system responds with:

```
SPIN DOWN UNIT 02.
```

02 is the unit number of EST ordinal 11.

After the unit is spun down, the system responds with:

```
MOVE PACK FROM UNIT 02 TO UNIT 01 AND SPIN UP.
```

After the pack is moved and the new unit is spun up, the system responds with:

```
EQ 11 REDEFINITION COMPLETE.
```

To end the reconfiguration, enter:

```
K.END
```

Disk configuration after the REDEFINE:

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
10	DJ-0	00	—	—
11	DJ-1	32	1	—

Example 3:

This example illustrates reconfiguring devices across channels.

Units may be reconfigured across channels with the REDEFINE command. Assume that EST ordinal 12 is the failing unit (or possibly the failing channel) and is on channel 26. EST ordinal 13 is the spare unit on channel 32. The following commands illustrate a reconfiguration across channels.

Disk configuration before the REDEFINE:

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
12	DI-1	26	3	—
13	DI-1	32	11	—

Enter the following commands:

Command	Description
REDEFINE,12.	12 is the EST ordinal of the failing device.
REDEFINE,13.	13 is the EST ordinal of the spare device.
K.jsn.	jsn is the job sequence name of the job requesting the REDEFINE K display.
K.EQ=13	Selects the spare unit to be processed first.
K.SUSPEND	Suspends the spare unit to clear it from the EST.

The system responds with:

SYSTEM USAGE OF EQ 13 SUSPENDED.

To continue with the reconfiguration on EST ordinal 12, enter:

- K.EQ=12 Selects the failing unit for processing.
- K.CH=32 Assigns the old channel number of EST ordinal 13 to EST ordinal 12.
- K.UL=11 Assigns the old unit number of EST ordinal 13 to EST ordinal 12.
- K.GO Initiates the processing of all the parameters and commands entered for EST ordinal 12.

The system responds with:

EQ 12 REDEFINITION COMPLETE.

Disk configuration at this point in example 3.

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
12	DI-1	32	11	I—
13	DI-1	32	11	IS—

All SUSPEND status flags must be cleared before a reconfiguration run can be ended. The system therefore automatically selects EST ordinal 13 again for the next EST ordinal to be processed.

To continue with the reconfiguration, enter:

Command	Description
K.CH=26	Assigns the old channel number of EST ordinal 12 to EST ordinal 13.
K.UL=3	Assigns the old unit number of EST ordinal 12 to EST ordinal 13.
K.GO	Initiates the processing of all the parameters and commands entered for EST ordinal 13.

The system responds with:

EQ 13 REDEFINITION COMPLETE.

To end the reconfiguration, enter:

K.END

Disk configuration after the REDEFINE:

EST Ordinal	Equipment Type	Channel Number(s)	Unit Number	Status
12	DI-1	32	11	—
13	DI-1	26	3	—

Remote Batch Facility (RBF) K Display

Whenever RBF is in operation, user connections and activity are shown in the RBF K display (refer to figure 8-23).

RBF Version Number	PSR Level of RBF Software	Number of Terminals Active (Decimal)	Number of Batch Devices Active (Decimal)			
REMOTE BATCH TERMINAL STATUS. RBF VER 1.2- 525 TERMINALS = 1 DEVICES = 2						
TERM. NAME.	DEV. TYP.	TCLASS. /STATE.	USR/JSN. .	FAMILY/ FILESIZE.	ERRORS.	FORMS. CODE.
LOET1	CO.	200UT.	AADC .	NZDGYR.	.	.
	CR1.	STOP.
	LP1.	GD .	AJZA .	40 .	.	.

Figure 8-23. RBF K Display

Select the RBF K display with this DSD command:

K, RBF.

Data on the RBF display is updated periodically to reflect the current status of batch terminals connected to RBF.

Terminals are identified by termname, a unique name defined by the site. Terminals are displayed in alphabetical order. There are multiple lines on the K display per terminal; one line for the terminal console, and one line for each batch device (card reader, line printer, card punch, or plotter).

If there are more terminals and devices active than can be displayed on one screen, the message

MORE LINES FOLLOW.

appears at the lower left corner of the display. You can display additional screens by entering:

K.+

The K.+ command advances the display page-by-page and end-around from the last page to the first. All screens are displayed in a forward direction. Use of the console input K.- does not move the screen back to the previous display.

Entries in the display have the following format:

term	dev	tclass	usr/jsn	family/		forms
name	typ	/state		filesize	errors	code

Header	Description																												
term name	Name of the terminal. A unique name, termname, assigned by the installation during network startup. This field is used for console devices only.																												
dev typ	Device type. The device type code is one of the following: <table border="1"> <thead> <tr> <th>Code</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>CO</td> <td>Console device.</td> </tr> <tr> <td>CPi</td> <td>Card punch.</td> </tr> <tr> <td>CRi</td> <td>Card reader.</td> </tr> <tr> <td>LPi</td> <td>Line printer.</td> </tr> <tr> <td>PLi</td> <td>Plotter.</td> </tr> </tbody> </table> <p>The ordinal number, i, of the device can range from 1 to 7.</p>	Code	Description	CO	Console device.	CPi	Card punch.	CRi	Card reader.	LPi	Line printer.	PLi	Plotter.																
Code	Description																												
CO	Console device.																												
CPi	Card punch.																												
CRi	Card reader.																												
LPi	Line printer.																												
PLi	Plotter.																												
tclass/state	If this field is in the same line as the termname, the contents of the field is the terminal class mnemonic (refer to the Remote Batch Facility Reference Manual). If this field occurs in a line not containing a termname, the device status code is one of the following: <table border="1"> <thead> <tr> <th>Code</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ABRT</td> <td>File in transmission is to be discarded.</td> </tr> <tr> <td>CONN</td> <td>Device is connected (initial state).</td> </tr> <tr> <td>END</td> <td>Device stops transmission at EOI.</td> </tr> <tr> <td>ENDA</td> <td>File being aborted; device will stop at EOI.</td> </tr> <tr> <td>ENDC</td> <td>End connection.</td> </tr> <tr> <td>ENDI</td> <td>Device will stop at EOI; idle down requested.</td> </tr> <tr> <td>GO</td> <td>Device is ready for input.</td> </tr> <tr> <td>NULL</td> <td>Console is not connected, but RBF devices are. This code appears only when the device type is a console (device type code CO).</td> </tr> <tr> <td>PREC</td> <td>Preconnect status (device connection not complete).</td> </tr> <tr> <td>STOP</td> <td>Device is not ready for transmission of data.</td> </tr> <tr> <td>STPA</td> <td>Device is stopped; current file to be aborted.</td> </tr> <tr> <td>STPE</td> <td>Device is stopped.</td> </tr> <tr> <td>STPI</td> <td>Device is stopped due to idle down request.</td> </tr> </tbody> </table>	Code	Description	ABRT	File in transmission is to be discarded.	CONN	Device is connected (initial state).	END	Device stops transmission at EOI.	ENDA	File being aborted; device will stop at EOI.	ENDC	End connection.	ENDI	Device will stop at EOI; idle down requested.	GO	Device is ready for input.	NULL	Console is not connected, but RBF devices are. This code appears only when the device type is a console (device type code CO).	PREC	Preconnect status (device connection not complete).	STOP	Device is not ready for transmission of data.	STPA	Device is stopped; current file to be aborted.	STPE	Device is stopped.	STPI	Device is stopped due to idle down request.
Code	Description																												
ABRT	File in transmission is to be discarded.																												
CONN	Device is connected (initial state).																												
END	Device stops transmission at EOI.																												
ENDA	File being aborted; device will stop at EOI.																												
ENDC	End connection.																												
ENDI	Device will stop at EOI; idle down requested.																												
GO	Device is ready for input.																												
NULL	Console is not connected, but RBF devices are. This code appears only when the device type is a console (device type code CO).																												
PREC	Preconnect status (device connection not complete).																												
STOP	Device is not ready for transmission of data.																												
STPA	Device is stopped; current file to be aborted.																												
STPE	Device is stopped.																												
STPI	Device is stopped due to idle down request.																												

Header	Description										
usr/jsn	If a file transfer is in progress, jsn is displayed in this field. If no file is being transferred, then the user name appears in this field.										
family/filesize	If this field is in the same line as the termname, the content of this field is the family associated with the username. The user index and family combination determine the terminal identifier used by the system for routing jobs. Otherwise, the contents of this field is the size of the output file in PRUs.										
errors	Contents of this field, if present, is one of the following error messages: <table border="1"> <thead> <tr> <th>Message</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DISK ERROR</td> <td>Data has been lost due to an unrecoverable disk failure.</td> </tr> <tr> <td>DISK FULL</td> <td>A disk-full indication was received while RBF was attempting to write to disk.</td> </tr> <tr> <td>NOT READY</td> <td>The device the user specified is not available (for example, the line printer is out of paper or the card reader has a card jam).</td> </tr> <tr> <td>QUEUE FULL</td> <td>The system input queue has reached its limit of jobs waiting to begin execution.</td> </tr> </tbody> </table>	Message	Description	DISK ERROR	Data has been lost due to an unrecoverable disk failure.	DISK FULL	A disk-full indication was received while RBF was attempting to write to disk.	NOT READY	The device the user specified is not available (for example, the line printer is out of paper or the card reader has a card jam).	QUEUE FULL	The system input queue has reached its limit of jobs waiting to begin execution.
Message	Description										
DISK ERROR	Data has been lost due to an unrecoverable disk failure.										
DISK FULL	A disk-full indication was received while RBF was attempting to write to disk.										
NOT READY	The device the user specified is not available (for example, the line printer is out of paper or the card reader has a card jam).										
QUEUE FULL	The system input queue has reached its limit of jobs waiting to begin execution.										
forms code	Forms code for output devices as specified by the user and defined by the site.										

Remote Host Facility (RHF) K Display

The Remote Host Facility (RHF) links NOS to a loosely coupled network (LCN) providing transfer of permanent files, queued files, and maintenance facilities for the LCN hardware. Each system in an LCN configuration is connected to one or more LCN trunks by network access devices (NADs). Several types of NADs are available, allowing the connection of various types of computer systems to an LCN. Each system has an RHF that provides some or all of the following capabilities.

RHF runs at a control point and contains an operator interface package, control tables, network application code, and the PP routines that drive the RHF network hardware.

The following applications are also available as part of RHF:

Application	Description
PTF, PTFS	Permanent File Transfer Facility (PTF) and Permanent File Transfer Facility Servicer (PTFS). PTF and PTFS provide users access to remote permanent files. A local user activates PTF with the MFLINK command (refer to the NOS Version 2 Reference Set, Volume 3). When a remote user enters the MFLINK command, RHF activates a PTFS application on the local host to service the remote request.
MHF	Maintenance Host Facility (MHF). MHF automatically loads controlware of local NADs during RHF initiation and performs maintenance logging. If a local NAD fails while RHF is active, MHF automatically dumps the NAD's memory to a permanent file and reloads the controlware, for a description of LNAD statement parameters that control automatic NAD dumping and loading (refer to section 10, LID/RHF Configuration Files). MHF also periodically checks all local NADs and any remote NADs enabled for logging (refer to the network path table display described later in this section) and copies new NAD error log entries to the binary maintenance log (BML). MHF starts automatically when RHF is initiated and remains active until RHF terminates.
QTF, QTFS	Queue File Transfer Facility (QTF) and Queue File Transfer Facility Servicer (QTFS). QTF and QTFS allow you to transfer input and output files to a remote system. When RHF is initiated, the system automatically activates QTF. When a remote host QTF application has a file to transfer, RHF initiates QTFS on the local host to service the remote request.

The following subsections describe the initiation, operation, control, and termination of the RHF.

RHF Initiation

Before network operation can begin, you must initiate RHF using the DSD RHFffff command (refer to section 5, DSD Commands).

When RHF is initiated, it starts all enabled applications defined by the network configuration as autostart applications (refer to RHF Configuration Files in section 10). These applications are typically MHF and QTF.

Application MHF loads controlware into all local NADs that are defined in RHF's configuration with AUTOLOAD=YES and have an EST status of ON.

Operator Interface

The operator interface for RHF consists of using the following displays:

Display	Description
Application table	Lists all active applications.
Network path table	Shows how the remote hosts are connected to the local host.

An example of each display is shown later in this section. To understand how to use the information provided in the displays it is helpful to understand the sample LCN network shown in figure 8-24.

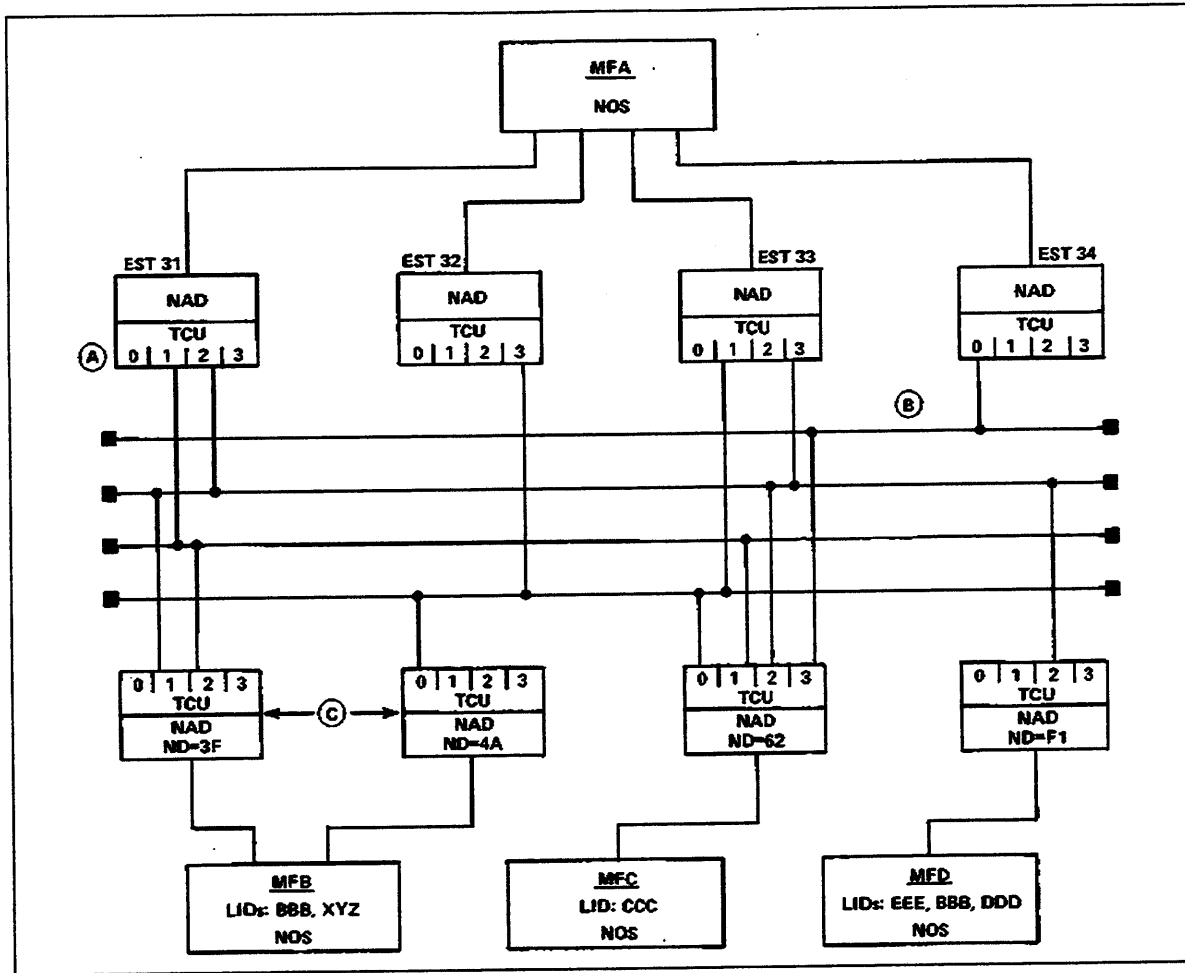


Figure 8-24. Sample LCN

MFA is the local NOS host and MFB, MFC, and MFD are the remote NOS hosts. MFA has four EST ordinals set up during installation (EST 31, 32, 33, and 34) to be used with the three remote hosts. Each NAD can connect to a maximum of four channels that can be used to communicate with NADs. These trunk control units (TCUs) are represented by the numbers 0, 1, 2, and 3 in the diagram (refer to A in figure 8-24).

The horizontal lines between the NADs of the local mainframe (MFA) and the NADs of the remote mainframes (MFB, MFC, and MFD) depict the connections between the NADs. For example, the top horizontal line shows the connections between the channels starting at TCU 0 of the local host and TCU 3 of the remote host (refer to B in figure 8-24).

Also during the installation process, each remote NAD is given a remote NAD address to uniquely identify that particular NAD. For example, the remote host MFB has two NADs associated with it. Their remote NAD addresses are 3F and 4A (refer to C in figure 8-24).

RHF Commands Under K display

Use the RHF K display to communicate with RHF. Bring up this display with this DSD command:

K, RHF.

Two K displays are available under RHF, the application table display (see figure 8-25) and the network path table display (see figure 8-26). When you first enter K,RHF, the application table display appears. The next time and subsequent times you enter K,RHF, the display that appears will be either the application table or the network path table, whichever was displayed last.

The following commands are available under the RHF K display:

Command	Function
APPL	Displays the application table (see figure 8-25).
IDLE	Begins the idle-down process of RHF and all its associated applications.
PATH	Displays the network path table (see figure 8-26).
ord,ND=rna,AC=rnac, DD=dd,RT=rteb, LT=lteb,LOG=status	Modifies entries associated with the network path table display ordinal ord.
+	Pages the current display forward.
-	Pages the current display backward.

Application Table Display

The application table display (see figure 8-25) lists all active applications. In this figure, USRAP represents an application written by the site.

ORD	APPL	ENABLED	MAX COPIES	ACTIVE COPIES	MAX CONNECTS
1	QTF	YES	1	1	4
2	QTFS	YES	4	3	1
3	PTF	YES	4	0	1
4	PTFS	YES	4	0	1
5	USRAP	YES	1	1	6
6	MHF	YES	1	1	1

ACTIVE APPLICATIONS (NETON PERFORMED)

JOB NAME	JOBORD	CONNECTS	JOB NAME	JOBORD	CONNECTS
ABCA	26	3	AARM	22	1
AAQT	31	1	ACAL	33	2
ACXQ	24	0	ACXY	27	0

Figure 8-25. Application Table Display

Each entry in the display appears in the following format:

```

ord  appl  enabled  maxcopies  activecopies  maxconnects
jobname  jobord  connects

```

Header	Description
ord	Application table display ordinal.
appl	Name of the application.
enabled	Specifies whether the application communicates with RHF.
maxcopies	Maximum number of copies of the application that can simultaneously communicate with RHF.
activecopies	Number of copies of the application that are currently communicating with RHF.
maxconnects	Maximum number of network connections allowed by each copy of the application.
jobname	Job sequence name of the application.
jobord	Executing job table ordinal of the application.
connects	Number of network connections currently used by this copy of the application.

Network Path Table Display

The network path table display, illustrated in figure 8-26, shows how the remote hosts are connected to the local host. In this figure, A 1 means that a TCU is enabled. A 0 means that a TCU is not enabled. The leftmost digit of these entries corresponds with TCU 0, the next with TCU 1, and so forth. For example, if a NAD has TCUs 1 and 2 enabled, the corresponding entry should be 0110. Refer to figure 8-24.

ORD	PID	EST	CH	ENABLED	LOCAL TRUNK ENABLES	REMOTE TRUNK ENABLES	REMOTE NAD ADDR	DEST	ACCESS CODE	CON COUNT	LOG ERRS
1	MFB	031	6	YES	0110	0110	3F	0	FOFO	2	NO
2		032	7	YES	0001	1000	4A	0	FOFO	1	NO
3		033	10	NO	0001	0100	3F	0	FOFO	0	NO
4		033	10	YES	0100	1000	4A	0	FOFO	0	NO
5	MFC	031	6	NO	0110	0110	62	0	FOFO	0	NO
6		032	7	NO	0001	1000	62	0	FOFO	0	NO
7		033	10	YES	0101	1010	62	0	FOFO	2	NO
8		034	11	YES	1000	0001	62	0	FOFO	0	NO
9	MFD	031	6	YES	0010	0010	F1	0	FOFO	0	NO
10		033	10	YES	0001	0010	F1	0	FOFO	0	NO

Figure 8-26. Network Path Table Display

Each entry in the display appears in the following format:

```

ord  pid  est  ch  enabled  local  remote  remote  dest  access  con  log
      trunk trunk nad      dest  code   count  errs
      enables enables addr

```

Header	Description
ord	Network path table display ordinal.
pid	Physical identifier of the remote mainframe for which the path is defined.
est	EST ordinal of the NAD that defines the local end of the path.
ch	Channel number of the NAD that defines the local end of the path.
enabled	Specifies whether RHF uses the path for starting new connections. For example, if a customer engineer wants to run diagnostics on a trunk connecting two NADs (local and remote), you would disable the appropriate path. When the connection count on that path falls to zero, the customer engineer can run concurrent diagnostics on that trunk without disturbing the operation of RHF.
local trunk enables	Bit pattern specifying which TCUs to use on the local NAD for communications with the remote NAD.
remote trunk enables	Bit pattern specifying which TCUs the remote NAD uses in returning communications to the local NAD.
remote nad addr	Hardware address (hexadecimal) of the remote NAD.
dest	Destination device address (hexadecimal).
access code	Access code of the remote NAD; the software access code is the two leftmost characters and the hardware access code is the two rightmost characters.
con count	Number of connections currently using this path.
log errs	Specifies whether the MHF application logs trunk errors (detected by the remote NAD) into the binary maintenance log (BML).

When the network path table is displayed, you can enter the following command to make changes.

```
ord,ND=rna,AC=rnac,DD=dd,RT=rteb,LT=lteb,LOG=status.
```

Parameter	Description
ord	Network path table display ordinal.
rna	New remote NAD address in hexadecimal.
rnac	New remote NAD access code in hexadecimal.
dd	New remote NAD destination device address in hexadecimal.
rteb	New remote trunk enable bit pattern (nonzero 4-bit binary number with the leftmost bit representing TCU 0).
lteb	New local trunk enable bit pattern (nonzero 4-bit binary number with the leftmost bit representing TCU 0). The same number of bits must be set for both RT and LT or the command will be rejected.
status	Status of error logging for the specified remote NAD. Enter YES to have MHF periodically copy the trunk errors occurring on the specified remote NAD to the binary maintenance log (BML). MHF does not copy trunk errors when you enter NO.

ord is the only required parameter and it must come first. All other parameters are optional and can appear in any order, but at least one, besides ord, must be present.

NOTE

Any change to rna, dd, or status for a given path results in an equivalent change for all paths using that remote NAD. This is because rna, dd, and status are associated with the remote NAD rather than the path to the remote NAD.

When RHF finds a faulty local NAD, it turns off the NAD's EST entry, activates the MHF application, and issues the following job dayfile message:

```
RHF, NAD ON EST est HAS BEEN TURNED OFF
```

MHF dumps the NAD's memory, reloads the NAD's controlware, and turns on the NAD's EST entry.

When RHF finds a path configured to use a faulty or nonexistent local NAD trunk control unit (TCU), it disables the path and notifies you with the following message on the B,O display:

```
RHF,BAD TCU ON PATH xxx, PATH TURNED OFF.
```

Use the GO,RHF command to acknowledge the message.

RHF Commands Available Under Application or Network Path Display

When the application table or network path table is displayed, you can enter the following commands to enable or disable the application or network path.

ENABLE, ord.

or

DISABLE, ord.

Parameter Description

ord	Ordinal of the application or network path on the current display.
-----	--

For example, if the application table display is up, ENABLE,ord enables the application specified by the display ordinal ord.

If an application is being enabled (whether currently enabled or disabled) and that application is defined as an autostart application, RHF initiates a copy of that application. This feature may be used, for example, if the MHF application is normally disabled and you want to start logging NAD errors. It may also be used if the copy of the QTF application has been dropped accidentally and a new copy of the QTF application must be started.

RHF Termination

You can terminate an RHF operation using the IDLE command.

The command format is:

IDLE.

When you enter the IDLE command the system waits for applications to stop and the message IDLE-DOWN IN PROGRESS appears on the system status display. When activity stops, the system drops RHF and the message RHF ENDED appears. The IDLE command allows a gradual shutdown of RHF activities. No new connections are allowed.

SCOPE 2 Station Facility (SSF) K Displays

The following commands control operation of the SSF Subsystem.⁵ Initiate the SSF Subsystem before issuing any commands. SSF can be initiated automatically when the system is brought up, either manually by entering an SSFffff command, or by entering the ENABLE,SSF command.

Operator Interface

The SSF uses the K display for operator communications; therefore, the K display must be assigned to SSF before any other SSF Subsystem commands can be entered. Use the following command to assign the K display to SSF.

```
K, SSF.
```

Station Login

After SSF is assigned to a control point, SSF searches the equipment status table (EST) for a SCOPE 2 communication coupler. If SSF finds an available coupler, it asks the SCOPE 2 mainframe to establish communications. If no coupler is found, a message appears at the console indicating that SSF is ready to log in to a SCOPE 2 mainframe. To log in, first enter an ON command (refer to section 5, DSD Commands, for a description of the ON command), then enter the following command:

```
K.LOGIN.
```

Enabling and Disabling File Transfers

File transfers between NOS and SCOPE 2 must be enabled before actual file transfers can begin; however, this occurs automatically when the SSF Subsystem is logged in. You can disable file transfers by entering the following command:

```
K.OFFSTAT,x.
```

Parameter Description

x	Specifies the SCOPE 2 mainframe PID.
---	--------------------------------------

When you enter a K.OFFSTAT command, file transfers in progress are completed, but no new file transfers are initiated. This command does not affect command, display, or message processing.

File transfers can be reenabled using the following command:

```
K.ONSTAT,x.
```

Parameter Description

x	Specifies the SCOPE 2 mainframe PID.
---	--------------------------------------

⁵ Refer to appendix C for more information on the SSF Subsystem.

Station Recovery

If NOS fails, follow the SSF initiation procedure as described in section 5, DSD Commands, to recover the station. SSF does not require reinitiation if SCOPE 2 fails or if there is an FLPP failure that results in a reload of the FLPP.

Station Disconnection and Logout

Station disconnection refers to the termination of communications between SSF and SCOPE 2 as a result of a hardware or software error or as a result of an operator command. When the disconnection occurs, all pending requests to the SCOPE 2 mainframe (such as screen refreshes, messages, and acknowledgments) are dropped.

Before entering an operator command to disconnect SCOPE 2, you should perform the following steps:

1. Idle down the file transfers by entering the K.OFFSTAT command. The station is idle when all staging SPOT jobs are finished and the spooling SPOT job has been swapped out.
2. If the system is in restricted or privileged mode, enter the CLROP command to relinquish operator control of SCOPE 2. Refer to the SCOPE 2.1 Operator's Guide for a description of the CLROP command.

To disconnect the SCOPE 2 mainframe, you can either drop the station control point by issuing a STOP or IDLE Subsystem command as described in section 5, DSD Commands, or you can log out of a specific SCOPE 2 mainframe by using the following command:

```
K.LOGOUT.
```

or

```
K.LOGOUT,pid.
```

Parameter Description

pid	PID of the mainframe you want to disconnect (required only in a multiframe environment).
-----	--

File Transfer Limit Commands

The following commands allow you to set limits on the number of file transfers that can be in progress concurrently for various types of file operations. The system default value for each of the following commands is 2; however, the default values can be changed, as an installation option, when SSF is installed. Refer to the NOS Version 2 Installation Handbook for more information on the default values for file transfer limit commands.

Staged File Transfer Commands

The following commands control the number of concurrent SPOT jobs performing staged file operations:

Command	Description
K.PURGE,pid,n.	Specifies the maximum number of concurrent SPOT jobs allowed to purge files from a NOS permanent file device. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of purge operations ($0 \leq n \leq 7$).
K.READTP,pid,n.	Specifies the maximum number of concurrent SPOT jobs allowed to read tape files from NOS to SCOPE 2. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of read operations ($0 \leq n \leq 7$).
K.WRITETP,pid,n.	Specifies the maximum number of concurrent SPOT jobs allowed to write tape files from SCOPE 2 to NOS. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of write operations ($0 \leq n \leq 7$).
K.GETPF,pid,n.	Specifies the maximum number of concurrent SPOT jobs allowed to read permanent files from NOS to SCOPE 2. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of permanent file read operations ($0 \leq n \leq 7$).
K.SAVEPF,pid,n.	Specifies the maximum number of concurrent SPOT jobs allowed to write files from SCOPE 2 to a NOS permanent file device. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of permanent file write operations ($0 \leq n \leq 7$).

Spooled File Transfer Commands

The following commands control the number of I/O files that the spooling SPOT job xSTA can transfer concurrently. Separate limits can be defined for both input and output files; however, the system limits the combined total of input and output files to four concurrent transfers.

Command	Description
K.INPUT,pid,n.	Specifies the maximum number of concurrent input file transfers from NOS terminals to SCOPE 2. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of input files ($0 \leq n \leq 4$).
K.OUTPUT,pid,n.	Specifies the maximum number of concurrent output file transfers from SCOPE 2 to NOS terminals. pid is the PID of the SCOPE 2 mainframe. n is the maximum number of output files ($0 \leq n \leq 4$).

Transaction Facility (TAF) K Displays

The following commands control operation of the TAF Subsystem. Initiate the TAF Subsystem by using the TAFfff command before issuing these commands.

TAF Initialization K Display

When the transaction executive is brought to a control point, the message REQUEST *K* DISPLAY appears at the control point if a DISPLAY,ON command is specified in the TAF configuration file. Respond with the entry:

K, TAF.

The TAF initialization K display appears on the system console as shown in figure 8-27.

K.	TAF
hh.mm.ss. yy/mm/dd.	CDC NETWORK OPERATING SYSTEM
	MID=AA NOS version
TAF INITIALIZATION OPTIONS	
OPTION	DESCRIPTION
SCP = 31	SUB CONTROL POINTS (2 - 31).
CMB = 40	COMMUNICATION BLOCKS (19 - 40).
REC = NO	RECOVERY MODE (YES OR NO).
MFL = 376500B	MAXIMUM FIELD LENGTH (40000B - 376500B).
ECS = 0K	EXTENDED MEMORY FIELD LENGTH (0K - 400K).
TLF = TASKLIB	TASK LIBRARY FILE (1-7 CHARACTERS).
INT = CRF,NONE	INITIALIZE CRF RECOVERY FILES (NONE,ALL OR 1-7 DIGITS).
INT = CRM,NONE	INITIALIZE CRM RECOVERY FILES (NONE OR ALL).
ERO = CRF,NO	ERROR OVERRIDE (YES OR NO).
BFL = 70000B	CMM BASE FL (20000B - 200000B).
EFL = 0B	CMM EXPANDABLE FL (0 - 100000B).
TFL = 30000B	CMM TARGET FL (10000B - 100000B).

Figure 8-27. TAF Initialization K Display (Left Screen)

Any of the following initialization commands can then be entered. If no values are to be changed, enter the command:

K.END

Values are decimal unless otherwise indicated.

Command	Description
K.BFL= <i>n</i>	Changes the starting and minimum field length allocated by TAF to CRM ($20000_8 \leq n \leq 200000_8$). Default is 70000_8 .
K.CMB= <i>n</i>	Changes the maximum number of communication blocks allowed to the TAF Subsystem ($19 \leq n \leq 40$). Default is 40.
K.ECS= <i>n</i>	Sets the extended memory field length to be used by the transaction executive; <i>n</i> is octal thousands of words. Default is 0.
K.EFL= <i>n</i>	Changes the maximum additional central memory field length made available to CRM for buffers and capsules ($0 \leq n \leq 100000_8$). Default is 0.
K.END	Ends input of the transaction executive initialization parameters. Initialization is completed when the TAF status K display appears (figure 8-28).
K.ERO=CRF, <i>op</i>	Specifies whether to override certain I/O and logic errors when processing the communication recovery files (CRF).
<i>op</i>	Description
NO	Aborts if I/O or logic errors are encountered while processing the communication recovery files. This is the default setting.
YES	I/O or logic errors encountered on a run unit header record result in the loss of that run unit with no indication to the terminal user. The loss is noted on the recovery report.
	I/O or logic errors encountered on a message record within a run unit result in a loss of that run unit. A status field is set in the run unit header, allowing TAF to inform the terminal user of the run unit loss.
	I/O or logic errors encountered on the CRF header record result in an unconditional abort of the TAF Subsystem.

Command	Description														
K.GO	Ends input of the transaction executive initialization parameters. Initialization is complete when the TAF status K display appears (figure 8-28).														
K.INT=typ,fileid	<p>Specifies which communication recovery files (CRF) are to be initialized. This is the only way to initialize a CRF. Files specified on a RECOVER command in the TAF configuration file and specified in this command are initialized. Files specified on a RECOVER command but not specified in this command are used for recovery.</p> <p>This command also specifies whether CYBER record manager (CRM) data base recovery files are initialized or recovered. For CRM recovery files, this command is valid regardless of TAF assembly parameters.</p> <table border="1"> <thead> <tr> <th>typ</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>CRF</td> <td>Communication recovery files. This parameter is valid only if the installation parameter IPTAR equals 1.</td> </tr> <tr> <td>CRM</td> <td>CYBER Record Manager after-image and before-image recovery files.</td> </tr> </tbody> </table> <table border="1"> <thead> <tr> <th>fileid</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>n</td> <td>A digit from 1 to 7 that defines a CRF to be initialized. The digit corresponds to the ID parameter on the RECOVER command in the TAF configuration file. This parameter is not valid for CRM recovery files.</td> </tr> <tr> <td>ALL</td> <td> <p>If typ is CRF, all communication recovery files defined by RECOVER commands in the TAF configuration file are initialized. If typ is CRM, all CRM recovery files are initialized.</p> <p>This parameter must be used with caution when typ is CRM since the CRM update history currently on the after-image recovery files is lost.</p> </td> </tr> <tr> <td>NONE</td> <td>If typ is CRF, communication recovery files are initialized; all communication recovery files specified in RECOVER commands in the TAF configuration file are recovered. If typ is CRM, all CRM data bases are recovered based on information in the existing recovery files. No CRM recovery files are initialized. This fileid is the default for both typ=CRF and typ=CRM.</td> </tr> </tbody> </table>	typ	Description	CRF	Communication recovery files. This parameter is valid only if the installation parameter IPTAR equals 1.	CRM	CYBER Record Manager after-image and before-image recovery files.	fileid	Description	n	A digit from 1 to 7 that defines a CRF to be initialized. The digit corresponds to the ID parameter on the RECOVER command in the TAF configuration file. This parameter is not valid for CRM recovery files.	ALL	<p>If typ is CRF, all communication recovery files defined by RECOVER commands in the TAF configuration file are initialized. If typ is CRM, all CRM recovery files are initialized.</p> <p>This parameter must be used with caution when typ is CRM since the CRM update history currently on the after-image recovery files is lost.</p>	NONE	If typ is CRF, communication recovery files are initialized; all communication recovery files specified in RECOVER commands in the TAF configuration file are recovered. If typ is CRM, all CRM data bases are recovered based on information in the existing recovery files. No CRM recovery files are initialized. This fileid is the default for both typ=CRF and typ=CRM.
typ	Description														
CRF	Communication recovery files. This parameter is valid only if the installation parameter IPTAR equals 1.														
CRM	CYBER Record Manager after-image and before-image recovery files.														
fileid	Description														
n	A digit from 1 to 7 that defines a CRF to be initialized. The digit corresponds to the ID parameter on the RECOVER command in the TAF configuration file. This parameter is not valid for CRM recovery files.														
ALL	<p>If typ is CRF, all communication recovery files defined by RECOVER commands in the TAF configuration file are initialized. If typ is CRM, all CRM recovery files are initialized.</p> <p>This parameter must be used with caution when typ is CRM since the CRM update history currently on the after-image recovery files is lost.</p>														
NONE	If typ is CRF, communication recovery files are initialized; all communication recovery files specified in RECOVER commands in the TAF configuration file are recovered. If typ is CRM, all CRM data bases are recovered based on information in the existing recovery files. No CRM recovery files are initialized. This fileid is the default for both typ=CRF and typ=CRM.														

Command	Description
K.MFL=n	Sets the maximum field length to be used by the transaction executive ($40000_8 \leq n \leq 376500_8$). Default is 376500_8 .
K.REC=a	Specifies whether to set the recovery bit in the user area of each terminal status table entry (YES or NO). If YES, the user recovery bit is set. If NO, the value of the user recovery bit is not changed from what it was before the command was issued. Default is NO.
K.SCP=n	Changes the number of subcontrol points ($2 \leq n < 31$). Default is 31.
K.STOP	Aborts the TAF Subsystem initialization unconditionally.
K.TFL=n	Changes the value used as the upper bound for TARGET. This is the amount of memory CRM uses for data and index blocks ($10000_8 \leq n \leq 100000_8$). Default is 30000_8 . For more information, refer to the CYBER Record Manager Advanced Access Methods Version 2 Reference Manual.
K.TLF=filename	Changes the name of the system task library file (any valid file name). Default is TASKLIB.

Restarting TAF K Display

On a level 3 recovery deadstart, or if the TAF Subsystem aborts, the TAF procedure file will automatically restart TAF by transferring control to the TAF automatic recovery program, if sense switch 4 is set. The TAF automatic recovery program recovers the central memory pointers and variables defined during TAF initialization.

If a DISPLAY,ON command is in the TAF configuration file, the automatic recovery program brings up the K display on the left console screen. This display is identical to the TAF initialization K display, except that values specified by the TAF initialization commands replace any default values that were in the TAF initialization K display.

TAF Status K Display

When the TAF Subsystem is executing, the TAF status K display (figure 8-28) indicates the:

- Latest transaction sequence number.
- Number of words of unused memory.
- Maximum field length.
- Global task dump limit.
- Subsystem default values for memory dump parameters.

The subsystem default values are used to control memory dumps when parameters are not included in the CMDUMP or DSDUMP command. Any of these default values can be changed by specifying the corresponding parameter in the K.DSDUMP command.

K.	TAF
hh.mm.ss. yy/mm/dd.	CDC NETWORK OPERATING SYSTEM
	MID=AA NOS version
TAF STATUS	
TRANSACTION SEQUENCE NUMBER	5
UNUSED FIELD LENGTH	7700
MAXIMUM FIELD LENGTH	376000
GLOBAL TASK DUMP LIMIT	0
CMDUMP/DSDUMP DEFAULT VALUES	
OPTION	DESCRIPTION
FW = 0	FWA OF TASK MEMORY TO BE DUMPED.
LW = 100000	LWA OF TASK MEMORY TO BE DUMPED.
EP = 1	IF EP=1, DUMP EXCHANGE PACKAGE.
	IF EP=0, DO NOT DUMP EXCHANGE PACKAGE.
OQ = BC	OUTPUT QUEUE.
QD = 0	QUEUE DESTINATION.

Figure 8-28. TAF Status K Display (Left Screen)

The default parameter values for the CMDUMP and DSDUMP commands are given on the TAF status K display shown in figure 8-28.

Parameter	Description								
EP	Exchange package: <table border="1"> <thead> <tr> <th>EP</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Do not dump the exchange package.</td> </tr> <tr> <td>1</td> <td>Dump the exchange package.</td> </tr> </tbody> </table>	EP	Description	0	Do not dump the exchange package.	1	Dump the exchange package.		
EP	Description								
0	Do not dump the exchange package.								
1	Dump the exchange package.								
FW	First word address of the task memory to be dumped.								
LW	Last word address of the task memory to be dumped.								
OQ ⁶	Output queue: <table border="1"> <thead> <tr> <th>OQ</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>BC</td> <td>Local batch.</td> </tr> <tr> <td>RB</td> <td>Remote batch.</td> </tr> <tr> <td>PF</td> <td>Permanent file.</td> </tr> </tbody> </table>	OQ	Description	BC	Local batch.	RB	Remote batch.	PF	Permanent file.
OQ	Description								
BC	Local batch.								
RB	Remote batch.								
PF	Permanent file.								
QD ⁶	Queue destination: Printer identifier (if OQ=BC). User name (if OQ=RB). Permanent file name (if OQ=PF).								

⁶ The functioning of this parameter is described in the TAF Version 1 Reference Manual.

TAF K-Display Commands

When the transaction executive is at its control point, the following commands can be entered from the system console or submitted from tasks using the KPOINT request. Any task can issue the K.DUMP command. Only tasks that reside on the system task library can issue the other K-display commands. (Refer to the TAF Version 1 Reference Manual for additional information on the KPOINT request and the system task library.)

Command	Description
K.ASSIGN,est or K.ASSIGN,est,db,n	<p>Assigns a magnetic tape unit to be used for a journal file. est is the EST ordinal of the tape unit. The first form of the command makes unit est available for the transaction executive to assign to the next tape journal file that encounters end of reel. Two tape units may be preassigned. If a tape has not been preassigned in this manner, an end-of-reel on a journal file causes subsequent entries for that file to be placed on disk.</p> <p>The second form of the command forces journal file n (n=1, 2, or 3) for data base db, defined as a tape file, from disk to tape. The transaction executive copies the data from the disk journal file to tape est and places all subsequent entries for that file on the tape. This command is necessary after the transaction executive initialization to assign tape units to the tape journal files or after an end-of-reel on a tape journal file when no tape had been preassigned to the transaction executive. All data residing on the disk for the tape journal file must fit on the tape assigned by this command or the transaction executive will unload the tape and issue the message *UNABLE TO USE TAPE*.</p>
K.DEBUG	Turns on the application interface program (AIP) debug option, which logs all messages on trace file ZZZZZDN. Use this command only when TAF is installed with the DEBUG option.
K.DROP,n	Drops an executing task at subcontrol point n.
K.DSDUMP,FW=addr, LW=addr,EP=pkg, OQ=outq,QD=qdest, DB=db	Allows you to modify the standard system default parameters controlling memory dumps. The command does not directly cause a dump. Rather, it sets default values to be used when a subsequent CMDUMP request is received or when abort conditions occur. Refer to TAF Status K Display, earlier in this section, for explanations of the parameters.

Command	Description
K.DUMP,fwa,lwa	<p data-bbox="542 237 1351 426">Dumps all or part of the field length of the Transaction Facility from the first word address (fwa) to the last word address (lwa) of the area to be dumped. The default value for fwa is 0 and for lwa is 377777₈. The default base is octal. If no parameters are specified, the entire field length is dumped. The output is routed to a printer that has an ID of 0.</p> <p data-bbox="542 447 1351 573">Unlike other K-display commands, the K.DUMP command can be issued from any task. Other K-display commands can be issued by tasks only if they are on the system task library (refer to the TAF Version 1 Reference Manual).</p> <p data-bbox="542 594 1351 751">Since secure information may be contained in a dump of the Transaction Facility, the following safeguards have been set up to protect dumped information; however, the installation must take the ultimate responsibility for the protection of dumped information.</p> <ul data-bbox="542 772 1351 1257" style="list-style-type: none"> <li data-bbox="542 772 1351 930">• The global task dump limit (GTDL) can be set by the K.DUMPLIM command to limit the number of times the K.DUMP command can be issued from tasks. The initial value of the GTDL is 0 (zero), so the K.DUMP command is disabled from use by a task by default. <li data-bbox="542 951 1351 1108">• For all dumps of the Transaction Facility, whether you initiated it or a task did, a one-page header precedes the dump. This header page indicates the output is secure and should be given only to the TAF central site systems analyst. <li data-bbox="542 1129 1351 1257">• When the Transaction Facility is dumped, the message TAF FIELD LENGTH DUMP RELEASED is issued to the system dayfile, the Transaction Facility dayfile, and line one of the control point.

Command	Description
K.DUMPLIM,n	<p>Sets GTDL to value n ($0 \leq n \leq 9999999$). If n is not specified, the GTDL is set to 0.</p> <p>The GTDL is the number of times the K.DUMP command can be issued from tasks. This value is displayed on the TAF status K display shown in figure 8-28. The initial value of the GTDL is 0. When the GTDL is 0, no dumps of the Transaction Facility can occur from tasks. Thus, the K.DUMP command is disabled from tasks by default.</p> <p>To enable the K.DUMP command for tasks, issue the K.DUMPLIM command to set the GTDL to a value greater than zero. Each time a task issues a K.DUMP command, the GTDL is decreased by one until it equals zero. When the first K.DUMP command is issued from a task with the GTDL equal to zero, the message</p> <p style="text-align: center;">GLOBAL TASK DUMP LIMIT EXHAUSTED</p> <p>is issued to the Transaction Facility dayfile, the system dayfile, and line one of the control point. Also, the message</p> <p style="text-align: center;">DUMPS LOST</p> <p>is displayed on the K display in place of the value of the GTDL. This message remains until the value of GTDL is set to a value greater than or equal to zero. The K.DUMPLIM command should be used with care in system tasks, since this might allow unauthorized users to alter the GTDL.</p>
K.IDLE	<p>Idles down the transaction control point. Once idle down has been initiated, no new transactions will be permitted but currently executing transactions will be allowed to finish.</p>
K.JEND,db,n	<p>Forces end-of-reel processing (writes an EOI and rewinds the file) on tape journal file n of data base db. If n is not a tape journal file, the command is ignored.</p>
K.MAXFL,n	<p>Alters the transaction executive maximum field length. The transaction executive does not attempt to obtain more than n words of storage. This command is rejected if the value for n is more than 376500₈ or less than the field length currently required for TAF.</p>
K.MESSAGE,TN=b. message	<p>Directs the transaction executive to send message to a terminal specified by terminal name b.</p>

Command	Description
K.NODEBUG	Turns off the application interface program (AIP) debug option, which logs all messages on trace file ZZZZZDN. Use this command only when TAF is installed with the DEBUG option.
K.OFFTASK,a,db	Disables the use of task a, where a is the task name in the data base db task library directory (dbTASKL). The data base name db is not specified for tasks in the system task library.
K.ONTASK.a,db	Reverses the effect of a previous OFFTASK command for the specified task a in the data base db task library directory (dbTASKL). The data base name db is not specified for tasks in the system task library (TASKLIB).
K.ROLLTIM,nnnnnn	Changes the amount of time that TAF will retain its field length between communication input messages; nnnnnn is specified in units of milliseconds. Refer to the installation parameter ITRTL in the NOS Version 2 Installation Handbook.
K.SWITCH	Causes the console K display to change to a display listing all allowable TAF K display commands. When K.SWITCH is entered a second time, the previous K display returns. This command activates task KDIS and forces TAF to remain rolled in.
K.TBCON,nn	Changes the number of TAF/CRM batch concurrency users; nn is less than or equal to the value specified on the TBCON command in TAF's configuration file (TCF).
K.TST,TN=a,DB=db, UU=nnnn,UL=mmmm, NN=b	Changes entries in the terminal status table for terminal a. The following entries can be changed: data base name db, user area upper 12 bits (nnnn), user area lower 12 bits (mmmm), and new terminal name b. The changes do not affect the network and simulation files. Do not use this command if the terminal is logged in.

TAF/CRM Status K Displays and Commands

You can use the K display to monitor the status of CRM, CRM data bases, or CRM data base permanent files while TAF/CRM is running if CRMTASK is present on the system task library. To get the CRMTASK K display enter:

K.DIS,CRMTASK.

After the K display is assigned to the task, the display in figure 8-29 is brought to the left screen.

```

K.                                TAF
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                                MID=AA   NOS version

                                CRMTASK COMMANDS

CRMSTAT                          SEE ALL CRM DATA BASES.
CRMSTAT,DB                       SEE CRM DATA BASE DB.
CRMSTAT,DBPFN                    SEE CRM DATA BASE FILE DBPFN.

DEDOWN,DB                        DOWN DATA BASE DB.
DEDOWN,DBPFN                     DOWN DATA BASE FILE DBPFN.
DBUP,DB                          UP DATA BASE DB.
DBUP,DBPFN                       UP DATA BASE FILE DBPFN.

END                              END CRMTASK PROCESSING.
MENU                             SEE THIS DISPLAY.
+                                PAGE LEFT SCREEN FORWARD.
-                                PAGE LEFT SCREEN BACKWARD.

```

Figure 8-29. CRMTASK Commands K Display (Left Screen)

The following commands are available under CRMTASK.

Command	Description
K.CRMSTAT	Selects the CRM status K display (figure 8-30). This display shows the status of all CRM data bases.
K.CRMSTAT,db	Selects the CRM data base status K display (figure 8-31). This display shows the status of the specific CRM data base with identifier db. It also shows the file name and status of the before-image files.
K.CRMSTAT,dbfile	Selects the CRM file status K display (figure 8-32). This display shows the status of the specific permanent file in a CRM data base with file name dbfile.
K.DBDOWN,db	Makes a specific CRM data base with identifier db unavailable for processing. When this command executes, the CRMTASK command directory appears on the K display.
K.DBDOWN,dbfile	Makes a specific permanent file in a CRM data base with file name dbfile unavailable for processing. When this command executes, the CRMTASK command directory appears on the K display.
K.DBUP,db	Makes a specific CRM data base with identifier db available for processing. When this command executes, the CRMTASK command directory appears on the K display.
K.DBUP,dbfile	Makes a specific permanent file in a CRM data base with file name dbfile available for processing. When this command executes, the CRMTASK command directory appears on the K display.
K.END	Ends CRMTASK K-display processing.
K.MENU	Returns to the CRMTASK commands K display (figure 8-29).
K.+	Pages the left screen forward to the next screen.
K.-	Pages the left screen backward to the first screen.

```

K.                TAF
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                    MID=AA   NOS version

                    CRM STATUS

TRANSACTIONS IN INPUT QUEUE =      0
ACTIVE TRANSACTIONS =              1
TRANSACTIONS IN OUTPUT QUEUE =     0

DATA                DATA                DATA
BASE  STATUS        BASE  STATUS        BASE  STATUS

TT    UP            BA    UP            BB    UP
BC    UP            BD    UP            BE    UP
BF    IDLE          BG    UP            BH    UP
BI    UP            BJ    IDLE          CA    UP
CB    UP            CC    UP            CD    UP
CE    DOWN          CF    UP            CG    DOWN
CH    UP            CI    DOWN          CJ    UP
    
```

Figure 8-30. CRM Status K Display (Left Screen)

```

K.                TAF
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                    MID=AA   NOS version

----- CRM DATA BASE STATUS -----

DATA BASE NAME =  AB           AFTER IMAGE FILE = ZZABA01
DATA BASE STATUS = UP         PRU-S REMAINING = 20

----- BEFORE IMAGE FILES STATUS -----
FILE    STATUS        FILE    STATUS        FILE    STATUS

ZZABB01  UP           ZZABB02  UP           ZZABB03  UP
ZZABB04  DOWN         ZZABB05  UP           ZZABB06  DOWN
ZZABB07  UP

----- DATA BASE FILES STATUS -----
FILE    STATUS        FILE    STATUS        FILE    STATUS

ABPFN01  UP           ABPFN02  UP           ABPFN03  IDLE
ABPFN04  UP           ABPFN05  DOWN         ABPFN06  UP
ABPFN07  UP           ABPFN08  UP           ABPFN09  UP
ABPFN10  DOWN         ABPFN11  UP           ABPFN12  UP
    
```

Figure 8-31. CRM Data Base Status K Display (Left Screen)

```
K.                TAF
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                    MID=AA    NOS version

                    CRM FILE STATUS

FILE NAME =    TTRNFDM
FILE STATUS = UP
RECOVERABLE = YES

PACK NAME =    PACKNAM
DEVICE TYPE = DEFAULT
ATTACH MODE = M

SIZE OF PRIMARY KEY =      30
NUMBER OF ALTERNATE KEYS = 6

ACTIVE USERS = 4
ACTIVE LOCKS = 8
```

Figure 8-32. CRM File Status K Display (Left Screen)

L-Display Utilities

9

FOTD L Display	9-2
LIDOU L Display	9-3
QDSPLAY L Display	9-6
SCTD L Display	9-9
SDSPLAY L Display	9-11
SDSPLAY Commands	9-16
Changing System Attributes	9-17
L.keyword=value Commands	9-17
SUBSYST L Display	9-19

This section documents the following L displays and the utilities used to present them.

Utility	Description
FOTD	Displays family ordinal table (FOT).
LIDOU	Displays logical identifier table (LID).
QDSPLAY	Displays the contents of a file in the queued file table (QFT).
SCTD	Displays the contents of the service class control table, which shows the valid service classes for each origin type.
SDSPLAY	Displays system values used to control job flow.
SUBSYST	Displays subsystem information.

The L display is like the K display in that you must start the program by entering the name of the utility. When the L display is ready for use, DSD assigns it to the left screen automatically. Only one L-display program is active at a time. The program need not occupy a control point or memory while you are looking at the display.

Using the L display, you can run any utility program you have created. The L display is an interface between your program and DSD. This interface allows the site analyst to write CPU programs that show displays on the system console.

All entries must be prefixed by L period (L.). However, when pressing CR or NEXT after the first entry, everything but the L. is erased. This allows another command to be entered without entering L. first. All examples in this section show L., although you may not have to type it. If it becomes necessary to enter a DSD command during parameter entry, simply backspace to erase the L., enter the command, and then continue by typing L. and the entry.

After you call a specific utility, enter the first command in the following format:

```
L.commandstring.
```

Parameter	Description
commandstring	Any input (command, data, or parameter) that is defined by the job as valid input.

Execute your own L-display utility by entering the following command:

```
LDISopt.
```

Parameter	Description
opt	One to three characters defining your utility.

FOTD L Display

The FOTD L-display utility displays all the family names known to the system and the corresponding family ordinal. To initiate the FOTD utility enter one of the following commands:

```
FOTD,L=outfile,LO=option.  
or
```

```
FOTD,outfile,option.
```

Parameter	Description						
outfile	Output file name. This parameter is valid only if a list option is specified. The default outfile is file OUTPUT.						
option	List option. Enter one or more of the following: <table border="1" data-bbox="462 674 1369 854"> <thead> <tr> <th>option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.</td> </tr> <tr> <td>L</td> <td>Formats the data for a line printer.</td> </tr> </tbody> </table>	option	Description	D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.	L	Formats the data for a line printer.
option	Description						
D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.						
L	Formats the data for a line printer.						

If you do not specify outfile and option, the data is written to the L-display buffer once and the utility ends.

Figure 9-1 shows the FOTD L display.

```
L. FOTD  
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM  
MID=AA NOS version  
  
FAMILY ORDINAL TABLE DISPLAY  
  
DEFAULT FAMILY = FAMILY1 AVAILABLE ENTRIES = 4B  
  
ORD. FAMILY ORD. FAMILY ORD. FAMILY ORD. FAMILY  
1. FAMILY1 2. FAMILY2 3. FAMILY3 4. FAMILY4  
5. FAMILY5 6. FAMILY6 7. FAMILY7 10. FAMILY8  
11. FAMILY9 12. FAMIL10 13. FAMIL11 14. FAMIL12
```

Figure 9-1. FOTD L Display

The FOTD L display contains the default family name and the octal number of FOT entries that remain available for assignment. The display also shows each family name and its corresponding family ordinal.

LIDOU L Display

The LIDOU utility displays the LID table configuration and allows you to add or delete LIDs, LID attributes, and PID attributes in the LID table configuration. This modified configuration is retained across a level 1, 2, or 3 recovery deadstart. A level 0 deadstart causes a return to the configuration defined by the LID configuration file.

To initiate the LIDOU utility, enter one of the following commands:

```
LIDOU.  
or
```

```
LIDOU,I=infile,L=outfile.
```

The first command format can be entered only from the console under DSD control; also, all input and output is processed using the LIDOU L display. Figure 9-2 shows the initial LIDOU L display (LID display).

NOTE

The information shown in figure 9-2 corresponds to the sample LCN shown in section 8.

From this display you can select the HELP display shown in figures 9-3 and 9-4. You can toggle between these two L displays as described in the HELP display.

The second command format can be entered under DIS control or by a system origin batch job. In either case, the I and L parameters must be specified.

Parameter	Description
I=infile	LIDOU reads input directives from file infile. Valid input directives are the same as the commands shown in the HELP L display (figure 9-3) without the L. prefix and without the BACK, HELP, +, or – commands.
L=outfile	LIDOU writes output to file outfile.

```

L. LIDOU
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
MID=AA NOS version

PID LID ENABLED ATTR MFTYPE NETWORK STATUS

MFB --- YES H-V- NOS RHF IN SERVICE
MBF YES ----
BBB YES S---
XYZ YES S---
---

MFC --- YES -L-- NOS RHF IN SERVICE
MFC YES ----
CCC YES S--N
---

MFD --- YES -LVN NOS SSF IN SERVICE
MFD YES ----
DDD NO S---
EEE NO S---
BBB YES S---
---
```

Figure 9-2. LIDOU L Display (LID)

```

L. LIDOU
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
MID=AA NOS version

LID OPERATOR UTILITY HELP DISPLAY PAGE 1 OF 2

COMMAND DESCRIPTION

L.END TERMINATE LIDOU
L.BACK FORMAT LID DISPLAY
L.HELP FORMAT HELP DISPLAY
L.+ PAGES THE L DISPLAY FORWARD
L.- PAGES THE L DISPLAY BACKWARD
L.OUT ROUTES A LISTING OF THE LID TABLE
L.PID,PA=XXZ SET ALL APPLICABLE PID ATTRIBUTES
L.PID,NT=NNN,NA=Z ENABLE/DISABLE NETWORK STATUS
L.PID,LD=LID,LA=YYYZ ADD OR SET ATTRIBUTES FOR LID
L.PID,LD=LID,DELETE DELETE LID FROM LID TABLE
L.PID,MF=MAINFRM SET MAINFRAME DESCRIPTION

X CAN BE V OR N. Y CAN BE S, B, V OR N. Z CAN BE E OR D.
ALL COMMANDS CAN BE ENTERED FROM THE HELP OR LID DISPLAY
ENTER L.+ FOR MORE HELP (COMPLETE ATTRIBUTE DESCRIPTIONS)
```

Figure 9-3. LIDOU L Display (HELP) (Page 1 of 2)

```

L. LIDOU
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
MID=AA NOS version

LID OPERATOR UTILITY HELP DISPLAY PAGE 2 OF 2

ATTRIBUTE DESCRIPTIONS

S - STORE AND FORWARD LIDS FOR LINKED PIDS ONLY.
B - LOOPBACK LIDS FOR HOST PIDS ONLY.
V - VALIDATE NOTE THAT S AND B CANNOT
D - DISABLE BE SPECIFIED TOGETHER AND
E - ENABLE ALSO THAT E AND D CANNOT
N - NOLIST BE SPECIFIED TOGETHER.

PARAMETER DESCRIPTIONS

PID IS THE THREE CHARACTER PHYSICAL IDENTIFIER
LID IS AN ASSOCIATED THREE CHARACTER LOGICAL IDENTIFIER
X CAN BE V OR N. Y CAN BE S, B, V OR N. Z CAN BE E OR D.
MAINFRM IS A (ONE TO SEVEN CHARACTER) MAINFRAME DESCRIPTOR
NNN IS THE NETWORK (*RHF*, *SSF*, *NAM*, OR *ALL*)
ENTER L.- TO SEE THE AVAILABLE COMMANDS

```

Figure 9-4. LIDOU L Display (HELP) (Page 2 of 2)

Although not shown in the HELP L display, some or all PID characteristics can be changed with one command by including all of the appropriate parameters on that command. For example, all PID characteristics for a given PID can be changed with the following command:

```
L.pid,PA=xxz,MF=mainfrm,NT=nnn,NA=z.
```

QDSPLAY L Display

The DSD utility QDSPLAY displays the content of a queued file listed in the queued file table (QFT). ASCII8 files are automatically translated into display code before they are displayed.

NOTE

On a secured system, the console must be in security unlock status to accept this command (refer to the UNLOCK,username,password command in section 5).

To bring the QDSPLAY L display to the left console screen, enter the following command:

```
QDSPLAY, jsn.
```

Parameter	Description
jsn	Job sequence name of the queued file you want to examine.

Figures 9-5 and 9-6 show the initial QDSPLAY L display.

```

L. QDSPLAY
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                        MID=AA   NOS version

QDSPLAY COMMANDS                                PAGE 1 OF 2

COMMAND          DESCRIPTION

BACK             RETURN FROM HELP/SS TO PREVIOUS FILE DISPLAY.
BOI              POSITION FILE TO BOI AND DISPLAY DATA.
DAY              POSITION PRINT FILE TO DAYFILE AND DISPLAY DATA.
DROP            DROP FILE FROM QUEUE AND TERMINATE.

END              RETURN FILE TO QUEUE AND TERMINATE.
HELP            LIST L DISPLAY COMMANDS FOR QDSPLAY UTILITY.
LINE            SET FILE DISPLAY MODE TO LINE FORMAT.
OCTAL           SET FILE DISPLAY MODE TO OCTAL FORMAT.

SS              DISPLAY IN OCTAL FROM BEGINNING OF SYSTEM SECTOR.
+              PAGE DISPLAY FORWARD (CIRCULAR FOR SS).
-              PAGE DISPLAY BACKWARD.

```

Figure 9-5. Initial QDSPLAY L Display (Page 1 of 2)

L. QDSPLAY	
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM	
MID=AA NOS version	
QDSPLAY COMMANDS	PAGE 2 OF 2
COMMAND	DESCRIPTION
SL	SKIP 1 LINE FORWARD IN FILE AND DISPLAY DATA.
SL,N	SKIP N LINES FORWARD IN FILE AND DISPLAY DATA.
SP	SKIP 1 PAGE FORWARD IN FILE AND DISPLAY DATA.
SP,N	SKIP N PAGES FORWARD IN FILE AND DISPLAY DATA.
SR	SKIP 1 RECORD FORWARD IN FILE AND DISPLAY DATA.
SR,N	SKIP N RECORDS FORWARD IN FILE AND DISPLAY DATA.
SR,*	SKIP TO EOI IN FILE.
BR	SKIP 1 RECORD BACKWARD IN FILE AND DISPLAY DATA.
BR,N	SKIP N RECORDS BACKWARD IN FILE AND DISPLAY DATA.

Figure 9-6. Initial QDSPLAY L Display (Page 2 of 2)

Valid QDSPLAY commands are:

Command	Description
L.BACK	Returns from the L.SS or L.HELP command displays and displays the contents of the file at the current position. This command is ignored if you are already displaying the contents of the file.
L.BOI	Positions the file to the beginning-of-information (BOI) and then displays the data from that point.
L.BR,r	Skips backward r records in the file and displays the data from the beginning of that record. If r is not specified, the default is one record. The record count, r, is decimal unless a B postradix is specified to make the number octal. The maximum skip count is 377777 ₈ . A BOI terminates the skip regardless of the number of records specified or skipped. This command is not allowed when the HELP display or system sector is being displayed.
L.DAY	Positions the print file to the beginning of the dayfile record and displays the data from that point.
L.DROP	Drops the file from the queue and terminates the QDSPLAY utility.
L.END	Returns the file to the queue and terminates the QDSPLAY utility.
L.HELP	Lists the L display commands for the QDSPLAY utility.
L.LINE	Sets the file display mode to line format (not allowed when displaying the system sector).

Command	Description
L.OCTAL	Sets the file display mode to octal format.
L.SL,n	Skips n lines forward in the file and displays the data from the beginning of that line. If n is not specified, the default is one line. The line count, n, is decimal unless a B postradix is specified to make the number octal. The maximum skip count allowed is 377777 ₈ . A line can be up to 14 words. An EOR, EOF, or EOI ends the skip regardless of the number of lines specified or skipped. This command is allowed only when displaying data in line format and is not allowed when the HELP display or system sector is being displayed.
L.SP,p	Skips forward p pages in the file and displays the data from the beginning of that page. If p is not specified, the default is one page. The page count, p, is decimal unless a B postradix is specified to make the number octal. The maximum skip count allowed is 377777 ₈ . A page is based on a carriage control 1 as the first character in a line. An EOR, EOF, or EOI ends the skip regardless of the number of pages specified or skipped. The command is allowed only when displaying data in line format and is not allowed when the HELP display or system sector is being displayed.
L.SR,r	Skips forward r records in the file and displays the data from the beginning of that record. If r is not specified, the default is one record. If an asterisk (*) is specified for the record count, r, a skip forward to the EOI occurs. The record count, r, is decimal unless a B postradix is specified making the number octal. The maximum skip count allowed is 377777 ₈ . An EOI ends the skip regardless of the number of records specified or skipped. This command is not allowed when the HELP display or system sector is being displayed.
L.SS	Displays the data from the beginning of the system sector in octal display code format.
L.+	Pages the QDSPLAY, HELP, or SS display forward.
L.-	Pages the HELP display backward.

SCTD L Display

The service class control table L-display (SCTD) utility displays the contents of the service class control table (SCT). To begin the SCTD utility, enter one of the following commands:

```
SCTD,L=outfile,LO=option.
```

or

```
SCTD,outfile,option.
```

Parameter	Description						
outfile	Output file name. This parameter is valid only if a list option is specified. The default outfile is file OUTPUT.						
option	List option. Enter one or more of the following: <table border="1" data-bbox="451 674 1351 852"> <thead> <tr> <th>option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.</td> </tr> <tr> <td>L</td> <td>Formats the data for a line printer.</td> </tr> </tbody> </table>	option	Description	D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.	L	Formats the data for a line printer.
option	Description						
D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.						
L	Formats the data for a line printer.						

Figures 9-7 and 9-8 show the SCTD L display. This display shows the various origin types along with the service classes valid for each origin type.

At the bottom of figure 9-8, the display shows how the service classes are prioritized. The priority levels are P0 through P7.

When entering the DSD CLASS or PCLASS command, use this utility to see the service classes for each origin type or the priority level. Refer to the DSD CLASS or PCLASS command in section 5 for further information.

Valid SCTD L display commands are:

Command	Description
L.+	Pages the display forward.
L.-	Pages the display backward.
L.END	Terminates the SCTD utility.

```
L. SCTD
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
MID=AA NOS version

SERVICE CLASS CONTROL TABLE DISPLAY PAGE 1 OF 2

SYSTEM ORIGIN SERVICE CLASSES

SY BC RB TS DI NS SS MA CT IO
I1 I2 I3

LOCAL BATCH ORIGIN SERVICE CLASSES

BC

PRIORITY ORDERED SERVICE CLASSES

P0 P1 P2 P3 P4 P5 P6 P7
BC IO I1 I2 I3
```

Figure 9-7. SCTD L Display (Page 1 of 2)

```
L. SCTD
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
MID=AA NOS version

SERVICE CLASS CONTROL TABLE DISPLAY PAGE 2 OF 2

REMOTE BATCH ORIGIN TYPE SERVICE CLASSES

RB

INTERACTIVE ORIGIN SERVICE CLASSES

TS DI

PRIORITY ORDERED SERVICE CLASSES

P0 P1 P2 P3 P4 P5 P6 P7
BC IO I1 I2 I3
```

Figure 9-8. SCTD L Display (Page 2 of 2)

SDSPLAY L Display

The SDSPLAY utility displays system values used to control job flow. This utility allows you to alter the queue priorities and weighting factors associated with the input, output, and executing queues for each service class. Also, you can alter the service limits associated with each service class. System values that you change using the SDSPLAY utility are valid only until the next deadstart.

To initiate the SDSPLAY utility, enter one of the following commands:

```
SDSPLAY, L=outfile, LO=option, I=infile.
```

or

```
SDSPLAY, outfile, option, infile.
```

Parameter	Description						
outfile	Output file name. This parameter is valid only if LO=L or I=infile is specified. The default outfile is file OUTPUT.						
option	List option. Enter one of the following: <table border="1"> <thead> <tr> <th>option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.</td> </tr> <tr> <td>L</td> <td>Formats the data for a line printer.</td> </tr> </tbody> </table>	option	Description	D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.	L	Formats the data for a line printer.
option	Description						
D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.						
L	Formats the data for a line printer.						
infile	Input file name. If infile is specified, the L display is not activated and SDSPLAY reads input directives from file infile rather than from the system console via the L display. This option allows you to run an SYOT job that changes the system scheduling parameters without direct operator intervention.						

Figures 9-9 through 9-11 show the SDSPLAY-HELP display. This display consists of three screens. The first screen (figure 9-9) shows the display commands used to control the utility. The other two screens (figures 9-10 and 9-11) show the service class definitions.

```

L. SDSPLAY - HELP.
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version

SDSPLAY COMMANDS                                PAGE 1 OF 3

COMMAND                                DESCRIPTION

CLASS=XX                                DISPLAY THE DATA FOR SERVICE CLASS XX.
END                                      END THE PROCESSING OF SDSPLAY.
GO                                       SET ALL ENTERED VALUES INTO THE SYSTEM.
HELP                                    SHOW L DISPLAY COMMANDS FOR SDSPLAY.
RESET                                   IGNORE ALL VALUES ENTERED SINCE THE
                                         LAST GO.

+                                       PAGE THE DISPLAY FORWARD.
-                                       PAGE THE DISPLAY BACKWARD.

```

Figure 9-9. SDSPLAY-HELP Display (Page 1 of 3)

L.	SDSPLAY - HELP	
	SERVICE CLASS DEFINITIONS	PAGE 2 OF 3
	SY	SYSTEM SERVICE CLASS.
	BC	BATCH SERVICE CLASS.
	RB	REMOTE BATCH SERVICE CLASS.
	TS	INTERACTIVE SERVICE CLASS.
	DI	DETACHED INTERACTIVE SERVICE CLASS.
	NS	NETWORK SUPERVISOR SERVICE CLASS.
	SS	SUBSYSTEM SERVICE CLASS.
	MA	MAINTENANCE SERVICE CLASS.
	CT	COMMUNICATION TASK SERVICE CLASS.

Figure 9-10. SDSPLAY-HELP Display (Page 2 of 3)

L.	SDSPLAY - HELP	
	SERVICE CLASS DEFINITIONS	PAGE 3 OF 3
	I0	INSTALLATION CLASS 0.
	I1	INSTALLATION CLASS 1.
	I2	INSTALLATION CLASS 2.
	I3	INSTALLATION CLASS 3.
	THE FOLLOWING SERVICE CLASSES MAY BE DEFINED AT YOUR SITE.	
	I4	INSTALLATION CLASS 4.
	I5	INSTALLATION CLASS 5.
	I6	INSTALLATION CLASS 6.
	I7	INSTALLATION CLASS 7.
	I8	INSTALLATION CLASS 8.
	I9	INSTALLATION CLASS 9.
	II	INSTALLATION CLASS I.
	IJ	INSTALLATION CLASS J.
	IK	INSTALLATION CLASS K.
	IL	INSTALLATION CLASS L.

Figure 9-11. SDSPLAY-HELP Display (Page 3 of 3)

Figures 9-12 through 9-14 show the SDSPLAY-CLASS display. This display consists of three screens. The first screen shows queue priorities and weighting factors associated with the specified service class. The other two screens show service limits associated with the specified service class. For additional information concerning queue priorities and service limits, refer to the QUEUE and SERVICE commands in section 5, DSD Commands.

NOTE

The current number of jobs for the specified service class is displayed at the bottom of the second screen. However, this value may not reflect the actual count of jobs running under the displayed service class since the displayed value is updated only at the initial processing of the SDSPLAY utility or after issuing an L.RESET or L.GO command.

```

L. SDSPLAY - CLASS.
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version
ATTRIBUTES FOR SERVICE CLASS SY                PAGE 1 OF 3

      OPTION                                DESCRIPTION

INLP = 7770B    INPUT QUEUE LOWER PRIORITY.
INUP = 7776B    INPUT QUEUE UPPER PRIORITY.
OTLP = 7000B    OUTPUT QUEUE LOWER PRIORITY.
OTUP = 7776B    OUTPUT QUEUE UPPER PRIORITY.
EXLP = 2000B    EXECUTION QUEUE LOWER PRIORITY.
EXUP = 7000B    EXECUTION QUEUE UPPER PRIORITY.
EXIL = 4000B    EXECUTION QUEUE INITIAL LOWER PRIORITY.
EXIP = 7000B    EXECUTION QUEUE INITIAL PRIORITY.
EXTP = 4004B    INITIAL TERMINAL PRIORITY AFTER I/O.

INWF = 1B       INPUT QUEUE WEIGHTING FACTOR.
EXWF = 1B       EXECUTION QUEUE WEIGHTING FACTOR.
OTWF = 1B       OUTPUT QUEUE WEIGHTING FACTOR.

      (ALL PRIORITY VALUES ARE 0-7777B).
      (WEIGHTING FACTORS MUST BE A POWER OF 2 FROM 1 - 4000B).
    
```

Figure 9-12. SDSPLAY-CLASS Display (Page 1 of 3)

L. SDSPLAY - CLASS.
 ATTRIBUTES FOR SERVICE CLASS SY PAGE 2 OF 3

OPTION	DESCRIPTION
CT = 100B	CP TIME SLICE (1-7777B MILLISECONDS*64).
CM = 20B	CM TIME SLICE (0-7777B SECONDS).
NJ = 7777B	MAXIMUM NUMBER OF JOBS (0-7777B).
CP = 6770B	CONTROL POINT SLICE PRIORITY (0-7777B).
FL = 3777B	MAXIMUM CM FL/100B, ANY JOB (0-3765B).
AM = 77777B	MAXIMUM CM FL/100B, ALL JOBS (0-77777B).
EC = 3777B	MAXIMUM EM FL/1000B, ANY JOB (0-7777B).
EM = 7777B	MAXIMUM EM LF/1000B, ALL JOBS (0-7777B).
PR = 30B	CPU PRIORITY (2-77B).
SE = 20B	CPU SLICE EXTENSION (0-7777B MS).
RS = 4B	CPU RECALL SLICE (1-7777B MS).
US = 20B	UNEXTENDED CPU SLICE (1-7777B MS).

THE CURRENT JOB COUNT FOR THIS SERVICE CLASS IS 2B.

Figure 9-13. SDSPLAY-CLASS Display (Page 2 of 3)

L. SDSPLAY - CLASS.
 ATTRIBUTES FOR SERVICE CLASS SY PAGE 3 OF 3

INDEX	VALUE	DESCRIPTION
DS = 0	(UNLIMITED)	MAXIMUM DIRECT ACCESS FILE SIZE (0-7).
FC = 0	(UNLIMITED)	MAXIMUM NUMBER PERMANENT FILES (0-7).
CS = 0	(UNLIMITED)	MAXIMUM PRUS, ALL INDIRECT FILES (0-7).
FS = 0	(UNLIMITED)	MAXIMUM PRUS, ONE INDIRECT FILE (0-7).

OPTION	DESCRIPTION
DT = DI	SERVICE CLASS ON DETACH (NOT SS OR DS).
TD = 341B	SUSPENSION TIMEOUT DELAY (0-7777B SECONDS*10B).

NOTE - EXTP, DT AND TD ARE USED ONLY FOR INTERACTIVE JOBS.

Figure 9-14. SDSPLAY-CLASS Display (Page 3 of 3)

SDSPLAY Commands

You can enter the following SDSPLAY commands from either the HELP or CLASS displays and from any screen.

Command	Description
L.CLASS=xx	Selects the CLASS display and displays the system attributes for the service class specified by xx, where xx is any service class defined on the HELP display.
L.END	Terminates the SDSPLAY utility. If you made changes to system attributes, the changes will not be saved unless you issue a L.GO command before issuing the L.END command.
L.HELP	Selects the HELP display and displays the first screen.
L.GO	Saves any changes you made to system attributes using the L.keyword=value commands by entering the new values into the system tables. The current display remains on the screen.
L.RESET	Clears any changes you made to system attributes using the L.keyword=value commands since you issued the last L.GO command. The current display remains on the screen. Following a L.RESET command, the CLASS displays will display values from the system tables.
L.+	Pages the HELP or CLASS display forward to the next screen. When the last screen is being displayed, the next L.+ command advances the display to the first screen. You cannot page from one display (CLASS or HELP) to the other display; nor can you page from one service class to another service class within the CLASS displays.
L.-	Pages the HELP or CLASS display backward to the previous screen. When the first screen is being displayed, this command has no effect on the display.

Changing System Attributes

To alter system attributes using the SDSPLAY utility, you select the CLASS display and specify the service class using the L.CLASS command. Next, you make changes using L.keyword=value commands. The values you enter apply only to the service class being displayed. You can enter multiple keyword=value pairs separated by commas in one command line. The only limit is the L display input line length restriction.

When you change system values, the updated values are displayed immediately if the corresponding screen is being displayed. If your changes affect a different screen, the updated values will be displayed when you page to that screen. Your changes will not be entered into the system tables until you enter a L.GO command. You can cancel your changes and reset the system values by entering a L.RESET command, but it must be entered before you enter the L.GO command.

You can change system values for a different service class by entering another L.CLASS command followed by more L.keyword=value commands.

L.keyword=value Commands

You can enter the following L.keyword=value commands from a CLASS display. Although you can enter numeric values in either octal or decimal, most of the values are displayed in octal. Octal is the input default; you specify decimal values with a radix of D following the number. Also, numbers containing an 8 or 9 are interpreted as decimal.

The keywords are listed in the same order as they appear in the CLASS display.

Keyword	Description
INLP	Priority assigned to a file entering the input queue. The value of INLP can range from 0 to 7777 ₈ , but must be less than the value of INUP.
INUP	Highest priority a file can reach in the input queue; aging stops when this priority is reached. The value of INUP can range from 0 to 7777 ₈ , but must be greater than the value of INLP.
OTLP	Priority assigned to a file entering the output queue. The value of OTLP can range from 0 to 7777 ₈ , but must be less than the value of OTUP.
OTUP	Highest priority a file can reach in the output queue; aging stops when this priority is reached. The value of OTUP can range from 0 to 7777 ₈ , but must be greater than the value of OTLP.
EXLP	Lowest priority for a job in the execution queue. The value of EXLP can range from 0 to 7777 ₈ , but must be less than the value of EXUP.
EXUP	Priority assigned to a job when initially assigned to a control point. However, online interactive jobs with terminal I/O available are scheduled at EXTP priority. The value of EXUP can range from 0 to 7777 ₈ , but must be greater than the value of EXLP.
EXIL	Priority a job receives when it initially exhausts its CM or CPU time slice. Thereafter, whenever the time slice is exhausted the job's priority will be set to the value of EXLP. The value of EXIL can range from 0 to 7777 ₈ , but must be in the range of values for EXLP and EXUP.

Keyword	Description
EXIP	Initial scheduling priority for an executing job. The value of EXIP can range from 0 to 7777 ₈ , but must be in the range of values for EXLP and EXUP.
EXTP	Terminal job priority assigned to a terminal job that has terminal input or output available. The value of EXTP can range from 0 to 7777 ₈ , but must be in the range of values for EXLP and EXUP.
INWF	Weighting factor for input queue priority calculations. The value of INWF must be a power of 2 in the range 1 to 4000 ₈ .
EXWF	Weighting factor for execution queue priority calculations. The value of EXWF must be a power of 2 in the range 1 to 4000 ₈ .
OTWF	Weighting factor for output queue priority calculations. The value of OTWF must be a power of 2 in the range 1 to 4000 ₈ .
AM	Maximum field length divided by 100 ₈ for all jobs of the specified service class. Refer to the SERVICE command in section 3.
CM	Central memory time slice in seconds. Refer to the SERVICE command in section 3.
CP	Control point slice priority. Refer to the SERVICE command in section 3.
CS	Cumulative size in PRUs allowed for all indirect access permanent files. Refer to the SERVICE command in section 3.
CT	Control point time slice (seconds). Refer to the SERVICE command in section 3.
DS	Size in PRUs allowed for individual direct access permanent files. Refer to the SERVICE command in section 3.
DT	Service class to which a detached job is assigned. Refer to the SERVICE command in section 3.
EC	Maximum user-accessible extended memory field length in words divided by UEBS for any job of the specified service class. Refer to the SERVICE command in section 3.
EM	Maximum extended memory length in words divided by UEBS for all jobs of the specified service class. Refer to the SERVICE command in section 3.
FC	Number of permanent files allowed. Refer to the SERVICE command in section 3.
FL	Maximum field length divided by 100 ₈ for any job of the specified service class. Refer to the SERVICE command in section 3.
FS	Size in PRUs allowed for individual indirect access permanent files. Refer to the SERVICE command in section 3.
NJ	Maximum number of jobs. Refer to the SERVICE command in section 3.
PR	CPU priority. Refer to the SERVICE command in section 3.
RS	Recall CPU slice (milliseconds). Refer to the SERVICE command in section 3.

Keyword	Description
SE	CPU slice extension (milliseconds). Refer to the SERVICE command in section 3.
TD	Suspension timeout delay. Refer to the SERVICE command in section 3.
US	Unextended CPU slice (milliseconds). Refer to the SERVICE command in section 3.

SUBSYST L Display

The SUBSYST L-display utility displays information about all the subsystems supported by NOS. To begin the SUBSYST utility enter one of the following commands:

```
SUBSYST, L=outfile, LO=option.  
or
```

```
SUBSYST, outfile, option.
```

Parameter	Description						
outfile	Output file name. This parameter is valid only if a list option is specified. The default outfile is file OUTPUT.						
option	List option. Enter one of the following: <table border="1" data-bbox="446 966 1351 1148"> <thead> <tr> <th>option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.</td> </tr> <tr> <td>L</td> <td>Formats the data for a line printer.</td> </tr> </tbody> </table>	option	Description	D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.	L	Formats the data for a line printer.
option	Description						
D	Formats the data for the DSD L display. This is the default list option if the parameters outfile and option are not specified.						
L	Formats the data for a line printer.						

If you do not specify outfile and option, the data is written to the L display.

Refer to the NOS Version 2 Operations Handbook for further information regarding the SUBSYST L display.

LID Configuration File	10-1
Creating the LIDCMid File	10-1
NPID Statement	10-2
NLID Statement	10-3
Equipment Configuration Example	10-4
Specifying the Size of the LID Table	10-6
Building or Rebuilding the LID Table	10-6
Listing the LID Table	10-7
Changing LIDS Using the LIDOU Display	10-7
Specifying Requirements for Running NOS/VE in Dual-State Mode	10-7
RHF Configuration Files	10-8
RCFGEN Utility	10-8
RCFGEN Input	10-10
Statement Order	10-10
Network Configuration Statements	10-11
LNDR Statement	10-11
APPL Statement	10-11
NPID Statement	10-13
PATH Statement	10-13
LNAD Statement	10-14
RNAD Statement	10-16
DEBUG Statement	10-16
CHARGE Statement	10-17
LCN Examples	10-18
QTF Configuration Requirements	10-20
RHF Configuration File Directives	10-20
NAM Configuration File Directives	10-21
QTF Procedure File	10-22
Procedure File Example	10-23

This section contains information about LID configuration files, RHF configuration files, and QTF configuration requirements.

LID Configuration File

You create a LID configuration file to define the physical and logical machines available for your system. After a level 0 deadstart, the system program CLDT searches for the LID configuration file on user index 377777g. It will use the information in this file to generate a logical identifier table (LDT) in central memory.

This section describes how to create and maintain the LID configuration file (LIDCMid). Specifically, how to:

- Create the LIDCMid file using NPID and NLID configuration statements.
- Specify the size of the LID table using the CMRDECK LDT entry.
- Build or rebuild the LID table using the X.CLDT command.
- List the current LID table using the LISTLID command.
- Change LIDs and LID attributes using the LIDOU display.
- Specify the requirements for running NOS/VE in dual-state mode.

Creating the LIDCMid File

The first line of the LIDCMid file should be a 1- to 7-character name for the LID configuration file. Use this format for naming the file:

```
LIDCMid
```

Value	Description
id	2-character machine identifier for the mainframe.

You can place comment lines in the file by putting an asterisk in column 1 or by putting the word COMMENT beginning in column one.

Example:

```
LIDCM64  
*LIDCMID FILE FOR MAINFRAME 64.
```

You should define each PID (physical identifier) and all the LIDs (logical identifiers) associated with the PID by using NPID and NLID statements.

NPID Statement

Use one NPID statement to define each physical mainframe available to your system. Use this format for the NPID statement:

```
NPID, PID=pid, ENABLED=status, MFTYPE=type, AT=at, NETDIS=sub.
```

Parameter	Description								
PID=pid	Physical identifier; pid is the 3-character physical identifier for the mainframe. For NOS, pid should be Mid, where id is the 2-character machine identifier defined by the MID=id CMRDECK entry. Required.								
ENABLED=status	Specifies whether the mainframe identified by PID is available. You can enter YES or NO. Default is YES.								
MFTYPE=type	Mainframe type; type is any 1- to 7-character string specifying the mainframe type or mainframe identifier (for example, NOSBE, NOS, CY200, NOSMF2, CLSH1). Required.								
AT=at	Specifies the attribute associated with the mainframe identified by PID. You can specify multiple attributes by using a slant separator (AT=VALID/NLIST) or by repeating the parameter (AT=VALID,AT=NLIST).								
	<table border="1"> <thead> <tr> <th>at</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>VALID</td> <td>Specifies that USER command prevalidation is required. This parameter causes the local host to verify that a user is authorized. This prevalidation is performed in addition to the normal user validation performed at the remote host.</td> </tr> <tr> <td>NVALID</td> <td>Specifies that USER command prevalidation is not required.</td> </tr> <tr> <td>NLIST</td> <td>Specifies that this PID is not available to the end user through the LISTLID command.</td> </tr> </tbody> </table> <p>The default attributes are: USER command prevalidation is not required and PID is listable using the LISTLID command.</p>	at	Description	VALID	Specifies that USER command prevalidation is required. This parameter causes the local host to verify that a user is authorized. This prevalidation is performed in addition to the normal user validation performed at the remote host.	NVALID	Specifies that USER command prevalidation is not required.	NLIST	Specifies that this PID is not available to the end user through the LISTLID command.
at	Description								
VALID	Specifies that USER command prevalidation is required. This parameter causes the local host to verify that a user is authorized. This prevalidation is performed in addition to the normal user validation performed at the remote host.								
NVALID	Specifies that USER command prevalidation is not required.								
NLIST	Specifies that this PID is not available to the end user through the LISTLID command.								
NETDIS=sub	Specifies which network accesses to the mainframe identified by PID are disabled. By default, all network accesses are enabled. You should not specify NETDIS for the host PID. Any combination of network accesses may be used (for example, NETDIS=RHF/SSF).								
	<table border="1"> <thead> <tr> <th>sub</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>NAM</td> <td>Network Access Method.</td> </tr> <tr> <td>RHF</td> <td>Remote Host Facility.</td> </tr> <tr> <td>SSF</td> <td>SCOPE Station Facility.</td> </tr> </tbody> </table>	sub	Description	NAM	Network Access Method.	RHF	Remote Host Facility.	SSF	SCOPE Station Facility.
sub	Description								
NAM	Network Access Method.								
RHF	Remote Host Facility.								
SSF	SCOPE Station Facility.								

NLID Statement

Use NLID statements after each NPID statement to define the relationship between each PID and its associated LIDs and the attributes of the LIDs. Use this format for the NLID statement:

```
NLID,LID=lid,ENABLED=status,AT=at.
```

Parameter	Description
LID=lid	Logical identifier; lid is the 3-character logical identifier for the mainframe identified by the last NPID statement. One lid entry must be the same as the last pid. Required.
ENABLED=status	Specifies whether the mainframe identified by LID is available. You can enter YES or NO. Default is YES.
AT=at	Specifies the attribute associated with the mainframe identified by LID. You can specify multiple attributes by using a slant separator (AT=STOREF/VALID) or by repeating the parameter (AT=STOREF,AT=VALID).
at	Description
LOOPB	Loopback capability for RHF testing. This testing can be performed on one mainframe with one NAD. The attribute of LOOPB is valid only for LIDs defined for the host (local) mainframe. You should not specify the LOOPB attribute if the STOREF attribute is specified.
STOREF	Store and forward capability indicating that the specified mainframe will act as part of the network in store and forward mode. Data will pass through the machine and into the network. You should not specify the STOREF attribute if the LOOPB attribute is specified.
VALID	Specifies that USER command prevalidation is required and can be specified only if the STOREF attribute is specified. This parameter causes the mainframe defined as being in store and forward mode (AT=STOREF) to verify that a user is authorized. This prevalidation is performed in addition to the normal user validation performed at the remote host that processes the job.
NVALID	Specifies that USER command prevalidation is not required and can be specified only if the STOREF attribute is specified.
NLIST	Specifies that this LID is not available to the end user through the LISTLID command. LIDs with this attribute might be reserved for onsite maintenance purposes.

Equipment Configuration Example

Figure 10-1 shows an equipment configuration consisting of six mainframes. Figure 10-2 shows the LIDCMid files for the three NOS mainframes in the configuration: LIDCM64, LIDCM42, and LIDCM05.

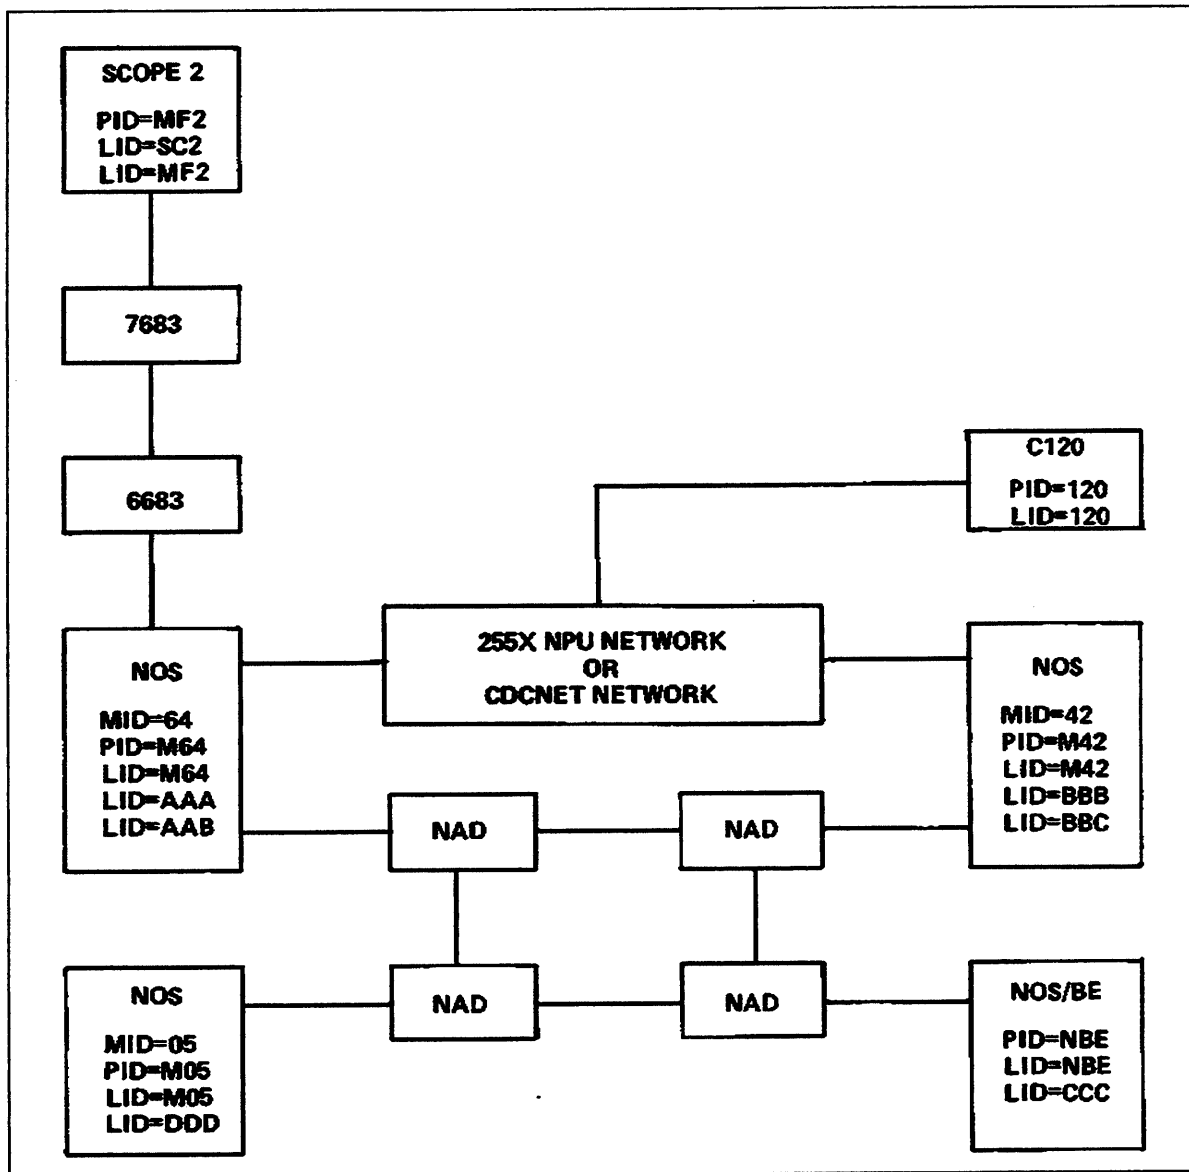


Figure 10-1. Equipment Configuration

```

LIDCM64
*LIDCM64      LIDCMID FILE FOR MAINFRAME 64.  NOTE THAT LID=LBK IS
*             DEFINED AS A LOOPBACK LID AND THAT THE NAM NETWORK IS
*             DISABLED TO NOS PID=M42.
NPID,PID=M64,MFTYPE=NOSHOST,AT=VALID.
NLID,LID=M64.
NLID,LID=AAA.
NLID,LID=AAB.
NLID,LID=LBK,AT=LOOPB.
NPID,PID=M42,MFTYPE=NOS42,NETDIS=NAM.
NLID,LID=M42.
NLID,LID=BBB.
NLID,LID=BBC.
NLID,LID=NBE,AT=STOREF/NLIST.
NLID,LID=CCC,AT=STOREF.
NPID,PID=MF2,MFTYPE=SCOPE2,AT=NVALID.
NLID,LID=MF2.
NLID,LID=SC2.
NPID,PID=120,MFTYPE=C120,AT=NVALID.
NLID,LID=120.
NPID,PID=M05,MFTYPE=NOS05.
NLID,LID=M05.
NLID,LID=DDD.
NLID,LID=NBE,AT=STOREF.
NLID,LID=CCC,AT=STOREF.

LIDCM42
*LIDCM42      LIDCMID FILE FOR MAINFRAME 42.  NOTE THAT LID=NBE IS
*             DEFINED AS A STORE AND FORWARD LID FOR THE NOS/BE MAINFRAME.
NPID,PID=M42,MFTYPE=NOSHOST.
NLID,LID=M42.
NLID,LID=BBB.
NLID,LID=BBC.
NLID,LID=M64,AT=STOREF.
NLID,LID=CCC,AT=STOREF.
NLID,LID=NBE,AT=STOREF.
NLID,LID=LBK,AT=LOOPE.
NPID,PID=M64,MFTYPE=NOS64,NETDIS=NAM.
NLID,LID=M64.
NLID,LID=AAA.
NLID,LID=AAB.
NPID,PID=NBE,MFTYPE=NOSBE,AT=NVALID.
NLID,LID=NBE.
NLID,LID=CCC.
NPID,PID=120,MFTYPE=C120,AT=NVALID.
NLID,LID=120.

LIDCM05
*LIDCM05      LIDCMID FILE FOR MAINFRAME 05.  AGAIN, NOTE THAT
*             LID=NBE IS DEFINED AS A STORE AND FORWARD LID FOR THE
*             NOS/BE MAINFRAME.
NPID,PID=M05,MFTYPE=M05.
NLID,LID=M05.
NLID,LID=LBK,AT=LOOPE.
NLID,LID=M64,AT=STOREF.
NLID,LID=CCC,AT=STOREF.
NLID,LID=NBE,AT=STOREF.
NPID,PID=NBE,MFTYPE=NOSBE,AT=NVALID.
NLID,LID=CCC.
NLID,LID=NBE.
NPID,PID=M64,MFTYPE=NOS64.
NLID,LID=M64.

```

Figure 10-2. LIDCMid Files for Example Equipment Configuration

Specifying the Size of the LID Table

You should specify the size of the LID table (LDT) in central memory with this CMRDECK entry:

```
LDT=nnnn.
```

Value	Description
nnnn	Number of central memory words allocated to the LDT (default=1, minimum=0, maximum=1100g).

You can calculate the value for nnnn by using this formula:

$$(3+lid) * pid$$

Value	Description
lid	Maximum number of LIDs allowed per PID.
pid	Total number of PIDs in all networks.

Building or Rebuilding the LID Table

When you deadstart the system, the system program CDLT automatically looks for the LIDCMid file under user index 377777g. It will use the file to automatically create the LID table in central memory.

If the system cannot find the file, or if there are errors in the file, an error message will be displayed in the dayfile. You can then create or correct your LIDCMid file and execute the X.CLDT command to build or rebuild the LID table in central memory.

To build or rebuild the LID table in central memory, follow these steps:

1. Create or correct your LIDCMid file and ensure that the file is stored as an indirect access file on user index 377777g.
2. Ensure that all subsystems that access the LID table (NAM, RHF, and SSF) are not active.
3. Enter this command:

```
X.CLDT
```

The system will build or rebuild the LID table in central memory.

Listing the LID Table

You can use the LISTLID command to list the configuration of the LID table. The LISTLID command can be executed through an interactive or batch environment. Use this format for the command:

LISTLID

```
LISTLID, LID=lid, PID=pid, L=list.
```

Parameter	Description
LID=lid	List the attributes of PIDs in which LID lid exists. LID lid must be enabled and it must be a listable LID. If the LID keyword is specified, a LID value lid must also be specified. If the LID keyword is omitted, the default is to list all listable LIDs.
PID or PID=pid	List the attributes of a given LID (LID=lid specified) or all LIDs (LID omitted) under a given PID (PID=pid specified) or all PIDs (PID specified). If the PID keyword is omitted, the default is to list all listable LIDs specified by the LID parameter.
L=list	Specifies the local file to which all listable output is to be written. If the L keyword is omitted, the default is file OUTPUT.

Changing LIDS Using the LIDOU Display

You can change LIDs by using the LIDOU display. Refer to section 9, L-Display Utilities, for information about the LIDOU display.

Specifying Requirements for Running NOS/VE in Dual-State Mode

You must create the appropriate configuration files before users can route batch jobs from the NOS side of a dual-state system to the NOS/VE side.

You must define the LIDs for NOS/VE batch jobs in the LIDCMid file. You can define a maximum of 30 NOS/VE LIDs.

Create the indirect access file LIDVEid (id is a 2-character machine identifier for the mainframe) on user index 377777g. The LIDVEid file is used by the dual-state IRHF job. This file contains the LIDs to be used for NOS/VE batch jobs, in the following format:

```
nvelid 1
nvelid 2
:
nvelid 30
```

Each LID must be on a separate line.

In the LIDCMid file used by CLDT, you should use an NPID statement to define a PID for the NOS/VE side of your dual-state system and include NLID statements for each of the LIDs in the LIDVEid file. The PID value used is not significant but must be unique. In a QTF network environment, you should also define the NOS/VE LIDs as store and forward (AT=STOREF) under the NOS host PID to allow storing and forwarding of batch jobs from other hosts through the NOS side of a dual-state system to the NOS/VE side.

RHF Configuration Files

You must define the RHF configurations including all NADs, applications, and physical identifiers (PIDs) to be used by or accessible to RHF. You may define multiple configurations. Each RHF in the network must have its own definition of the configuration; each definition will be different from the definition of the configuration for other RHF's in the network. Use the RCFGEN utility to define these configurations.

RCFGEN Utility

The RCFGEN utility reads network configuration statements to create the RHF configuration record, which RHF uses to define the network and determine proper access.

Each configuration record has a name of one to seven letters or digits, starting with a letter. The default name set by RCFGEN when creating a record (and used by RHF when searching for a record) is RCFpid, where pid is the three-character physical identifier of the host. You can use the RN=recname parameter on the RCFGEN and RHF commands to specify a different configuration record name.

When RHF is initiated, it searches for the specified configuration record first on local file RCFILE, then on permanent file RCFpid (user name SYSTEMX, direct or indirect access), and finally on the system library.

By default, RHF searches for configuration record name RCFpid. If you provide a subsystem procedure file for RHF, you can use the command RHF,RN=recname to specify a different name. Refer to the NOS Version 2 Installation Handbook for information about creating a procedure file for RHF.

The following command calls the RCFGEN utility.

```
RCFGEN, I=input, L=list, O=output, RN=recname, LO=option.comments
```

The parameters are order independent and optional; if omitted, the defaults are used.

Parameter	Description
I=input	Specifies the local file from which the network configuration statements are to be read. The default is file INPUT. If only the keyword I is specified, file COMPILE is used.
L=list	Specifies the local file to which all listable output is to be written. The default is file OUTPUT.
O=output	Specifies the local file to which the configuration tables are to be written. The default is file RCFILE.
RN=recname	Specifies the name given to the configuration record being generated. The default name, RCFpid, is used if RN is omitted, if only the keyword RN is specified, or if recname is 0.
LO=option	Specifies the list options used when generating listable output. If LO is omitted or option is 0, only the network configuration statements and diagnostics are listed. If only the keyword LO is specified, macro definitions and table-generating definitions are also listed. LO=ALL specifies all list options. (The non-default list options may help analysis of RCFGEN or RHF problems.)
comments	Specifies an optional 1- to 70-character string that is placed in the prefix table of the configuration record being generated. (During initiation, RHF displays the string in the job dayfile.)

RCFGEN Input

The input to RCFGEN consists of network configuration statements. The syntax of these statements conforms to COMPASS statements. Configuration statements (except LNAD and RNAD statements) must not start before character position 3.

Use the following network configuration statements.

Statement	Description
LNDR	Defines the maximum number of local NAD drivers (NDRs) allowed to execute at one time.
APPL	Defines application programs that are allowed to access RHF.
NPID	Defines the physical identifier of a remote mainframe.
PATH	Defines the paths to a remote mainframe through the LCN network.
LNAD	Defines information necessary to address local NADs.
RNAD	Defines the addressing information necessary to access a remote NAD.
DEBUG	Defines debug parameters.
CHARGE	Defines the charge that is transferred to a user control point for each RHF call.

Statement Order

Network configuration statements must be specified in the following order to properly define a network.

1. LNDR statement (if necessary).
2. APPL statement to define application programs.
3. Sets of NPID and PATH statements that define all portions of the network. The PATH statement must be associated with a given physical mainframe (NPID statement). The following structure is required of these statements when defining a network:
 - a. NPID statement.
 - b. All paths (PATH statement) associated with the preceding NPID.
4. LNAD statements to define local NADs.
5. RNAD statements to define hardware addressing of remote NADs.
6. DEBUG and CHARGE statements (order independent).

At least one of each configuration statement is required in the configuration file with the exceptions of the LNDR, DEBUG, and CHARGE statements. Defaults are specified in the individual statement descriptions.

Network Configuration Statements

The network configuration statements are described in the order that they must be specified to define a network.

LNDR Statement

The maximum number of NAD drivers is the maximum number of PPs that may contain a NAD driver at one time, regardless of the number of NADs and the number of drivers allowed per NAD.

To define the maximum NDRs allowed for all NADs, enter:

```
LNDR    MAXNDRS=nn
```

Parameter	Description
MAXNDRS=nn	Maximum number of PPs that may contain NDRs at any one time. Default is 4. The maximum value is the number of PPs available for NAD drivers. nn must not be less than the largest value specified for MAXNDRS on any LNAD statement.

APPL Statement

Each application definition uses additional RHF field length. This additional field length is equal to $\text{mxcopy} \times (5 + 3 \times \text{mxcons})$ central memory words (mxcopy and mxcons are defined below). To define an RHF application, enter:

```
APPL    NAME=name, ENABLED=status, MXCONS=mxcons, MXCOPYS=mxcopy, SVR=status,
        ASTART=status, SYSORG=status
```

Parameter	Description
NAME=name	Application name; name is 1- to 7-alphanumeric characters, where the first character must be alphabetic. Required.
ENABLED=status	Availability of the application when RHF is initiated. status may be YES or NO. Default is YES.
MXCONS=mxcons	Maximum simultaneous connections that are allowed for this application. The maximum value for mxcons is 127. Default is 1.
MXCOPYS=mxcopy	Maximum number of simultaneously active copies of this application that are allowed. The maximum value for mxcopy is 127. Default is 1.

Parameter	Description
SVR=status	Servicer program status; specifies if the application is a servicer, that is, an application program that is started automatically by RHF upon request of an application program on another host. status may be YES or NO. Default is NO.
ASTART=status	Application startup status; specifies whether the application is started when RHF is initiated or when the operator enables the application. status may be YES or NO. If NO is specified, the application is started by the user, or for a servicer application, by a request from a remote application. Default is NO.
SYSORG=status	System origin status; specifies whether the application must have privileged application status to gain access to RHF. status may be YES or NO. If status is YES, the application must have an SSJ= entry point and be loaded from the system library. If status is NO, the application need not reside on a system library; however, the user must have CUCP (system control point) validation. Default is YES.

In determining the number of allowed connections and copies of an application, note that each NAD has a maximum of 127 active connections. This number is restricted to 35 during NAD controlware loading but may be increased by modifying the appropriate NAD controlware load parameters in LOADBC.

When defining the APPL statements you must follow certain restrictions for system-supplied applications QTF, QTFS, PTF, PTFS, and ITF. The maximum simultaneous connections should be set to one (default) for QTFS, PTF, and PTFS. The maximum number of simultaneously active copies should be set to one (default) for QTF and ITF. The maximum simultaneous connections for QTF should be set to four unless installation parameters ACNMAXC for QTF and MAXFILEXFR for FIP are changed. The SVR=YES parameter must be specified for QTFS and PTFS and must not be specified for QTF, PTF, or ITF. ASTART=YES should be specified for QTF.

For ITF, the maximum simultaneous connections should be equal to the value of the PI parameter on the ITF command in the JOBITF record on the NAMI startup master file. The released default is 2.

The MHF application is defined by default in RHF's configuration file; no separate definition is needed. If you choose to define MHF, you must specify ASTART=YES to allow initial loading of controlware into local NADs.

NPID Statement

Each remote mainframe definition requires three words in RHF's field length. To define a remote mainframe, enter:

```
NPID  PID=pid,ENABLED=status,MFTYPE=type
```

Parameter	Description
PID=pid	Physical identifier of the remote mainframe; pid must be a unique 3-character physical identifier. Required.
ENABLED=status	Availability of the mainframe identified by PID. status may be YES or NO. Default is YES.
MFTYPE=type	Mainframe type; type is any 1- to 7-character string specifying the mainframe type or mainframe identifier. For example, you could use NOS1, NOS2, NOS/BE, or CY200. Required.

PATH Statement

Every path defined for a remote mainframe requires two words in RHF's field length. To define a path for a remote mainframe, enter:

```
PATH  ENABLED=status,LT=tttt,RT=rrrr,RNAD=raddr,LNAD=laddr,AC=aaaa
```

Parameter	Description
ENABLED=status	Availability of the path when RHF is initialized. status may be YES or NO. Default is YES.
LT=tttt	Local trunk control units (TCUs) enabled; tttt is a 4-digit nonzero binary number indicating the network trunk connections for the local NAD. Required.
RT=rrrr	Remote TCUs enabled; rrrr is a 4-digit nonzero binary number indicating the network trunk connections for the remote NAD. Required.
RNAD=raddr	Symbolic address of the remote NAD entry for this path referenced in the RNAD statement. Required.
LNAD=laddr	Symbolic address of the local NAD entry for this path referenced in the LNAD statement. Required.
AC=aaaa	Access code; aaaa is a 4-digit hexadecimal access code for the remote NAD. Default is 0.

At least one PATH statement is required for each PID defined.

LNAD Statement

Each local NAD definition requires five words in RHF's field length. To define a local NAD, enter:

```
laddr      LNAD  CH=ch, MAXNDRS=n, DEDICATE=status, CMPATHS=nn, CMBUFFS=nn,
            AUTODUMP=status, AUTOLOAD=status, TRACE=trw, LOG=status
```

Parameter	Description
laddr	Symbolic address referenced in a preceding PATH statement. This parameter is required and must begin in character position 1 or 2.
CH=ch	Channel number; ch is a 2-digit octal number of the channel to which the NAD is connected.
MAXNDRS=n	Maximum number of NAD drivers that may be assigned at one time to this NAD. n can be from 1 to 3. Default is 1.
DEDICATE=status	Dedicated channel status; specifies whether the driver will always hold the NAD channel reservation between consecutive blocks of one I/O request. status may be YES or NO. Default is YES. YES should be specified unless some non-CDC driver requires high-performance access to the NAD channel.
CMPATHS=nn	Maximum number of convert mode paths. You can specify a number from 0 to 63. The default is 0. If nn is 0, then code conversion is done in the CPU. Refer to table 10-1 for additional information.
CMBUFFS=nn	Maximum number of convert mode buffers. You can specify a number from 0 to 63. The default is 0. Refer to table 10-1 for additional information.
AUTODUMP=status	Automatic dump status; specifies whether RHF will automatically dump the NAD's memory if the NAD fails. status may be YES or NO. Default is NO. (YES is intended for fault analysis along with TRACE=YES/FULL).
AUTOLOAD=status	Automatic load status; specifies whether RHF will automatically load the NAD's controlware during initialization and reload the NAD if it fails. status may be YES or NO. Default is YES.

Parameter	Description
TRACE=trw	<p>NAD trace word; controls the value of trace word in the NAD microcode initialization parameters used by MHF when loading the NAD (refer to the 380-170 Network Access Device Reference Manual for a complete description). trw may be NO (no trace), YES (standard trace buffers), or FULL (maximum trace buffers). Default is NO.</p> <p>Because TRACE decreases NAD performance, TRACE should be specified only for fault analysis. AUTODUMP=YES should be specified along with TRACE=YES or FULL to capture the NAD memory dump. TRACE=YES or FULL produces a trace word value of 2954 or 529F hexadecimal, respectively.</p>
LOG=status	<p>Error logging status; specifies if local NAD trunk errors are to be transmitted by MHF to the mainframe's binary maintenance log. status may be YES or NO. Default is YES.</p>

The software relies on the NAD to determine the actual number of paths and buffers reserved for code conversion. The algorithm used by the NAD is:

Paths reserved = the smaller of (paths requested) or (buffers reserved * 2/3).
 Buffers reserved = the smaller of (buffers requested) or (total buffers - 2).

The code conversion parameters should be adjusted according to the average number of concurrently active connections doing code conversion and the average size of the files being converted. Table 10-1 gives some suggested values for the code conversion parameters. In this table, the term average connection implies four conversion connections. The term average file size is a file approximately 500 PRUs long.

Table 10-1. Suggested Code Conversion Parameters

Connections	96K-Byte NAD	128K-Byte NAD
Less than average connections, average file size	2 paths 3 buffers	2 paths 3 buffers
Less than average connections, large file sizes	2 paths 5 buffers	2 paths 5 buffers
Average connections, average file sizes	4 paths 5 buffers	6 paths 8 buffers
Greater than average connections, average file sizes	6 paths 9 buffers	7 paths 11 buffers
Greater than average connections, large file sizes	5 paths 9 buffers	6 paths 13 buffers

RNAD Statement

Every remote NAD defined requires four words in RHF's field length. To define a remote NAD, enter:

```
raddr    RNAD    ND=nn, DD=d, LOG=status
```

Parameter	Description
raddr	Symbolic address referenced in a preceding PATH statement. This parameter is required and must begin in character position 1 or 2.
ND=nn	Remote NAD address; nn is a 2-digit hexadecimal address of the remote NAD. Default is 0.
DD=d	Remote NAD exit port; d is a 1-digit hexadecimal address of the exit port of the remote NAD. Default is 0.
LOG=status	Error logging status; specifies if remote NAD trunk errors are to be recorded by MHF in the mainframe's binary maintenance log. status may be YES or NO. Default is NO. Do not specify LOG=YES if the defined NAD is a local NAD (as in the case of a loopback path). Instead, use the LOG parameter on the LNAD statement to control error logging.

DEBUG Statement

The DEBUG statement controls the manner in which RHF uses queue entries.

```
DEBUG    TRACE=status
```

Parameter	Description
TRACE=status	RHF trace status; specifies whether the RHF trace is on or off. status may be YES or NO. Default is NO. NO specifies that RHF trace is off and queue entries freed by RHF are placed at the top of the empty queue and reused immediately. RHF uses a slightly smaller amount of processing time when the RHF trace is off. Default is NO. YES specifies that RHF trace is on and a queue entry is reused only after all queue entries ahead of it have been used. Also, when the RHF trace is on, you can use it for analysis of an RHF dump and resolution of the associated RHF problem.

CHARGE Statement

The CHARGE statement specifies the amount of system resources a user control point is charged for an RHF call. RHF distinguishes two different types of calls; those that require a large amount of processing time and those that require a small amount of processing time.

CHARGE TYPE=type, CPA=cpa, CPB=cpb, IO=io, CMFL=cm, PP=pp

Parameter	Description						
TYPE=type	Type of call for which the charge is being specified.						
	<table border="1"> <thead> <tr> <th>Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Requires small amount of RHF processing time.</td> </tr> <tr> <td>2</td> <td>Requires large amount of RHF processing time.</td> </tr> </tbody> </table>	Type	Description	1	Requires small amount of RHF processing time.	2	Requires large amount of RHF processing time.
Type	Description						
1	Requires small amount of RHF processing time.						
2	Requires large amount of RHF processing time.						
CPA=cpa	Time to be charged for central processor cpa (decimal milliseconds). Default is 2 milliseconds for a type 1 call; 10 milliseconds for a type 2 call.						
CPB=cpb	Time to be charged for central processor cpb (decimal milliseconds). Default is 0.						
IO=io	Input/output time to be charged (decimal milliseconds). Default is 0.						
CMFL=cm	Central memory field length to be charged. cm is the field length in octal divided by 100 ₈ . Default is 10 ₈ .						
PP=pp	PP time to be charged (decimal milliseconds). Default is 0.						

A CHARGE statement is not required. You may, however, enter two CHARGE statements: one for type 1 calls, a second for type 2 calls.

LCN Examples

Figure 10-3 illustrates a sample LCN network and is the basis of the network configuration statements shown in figures 10-4 and 10-5. Figure 10-4 shows the network configuration statements for the RHF on mainframe MFA. Figure 10-5 shows the network configuration statements for the RHF on mainframe MFD.

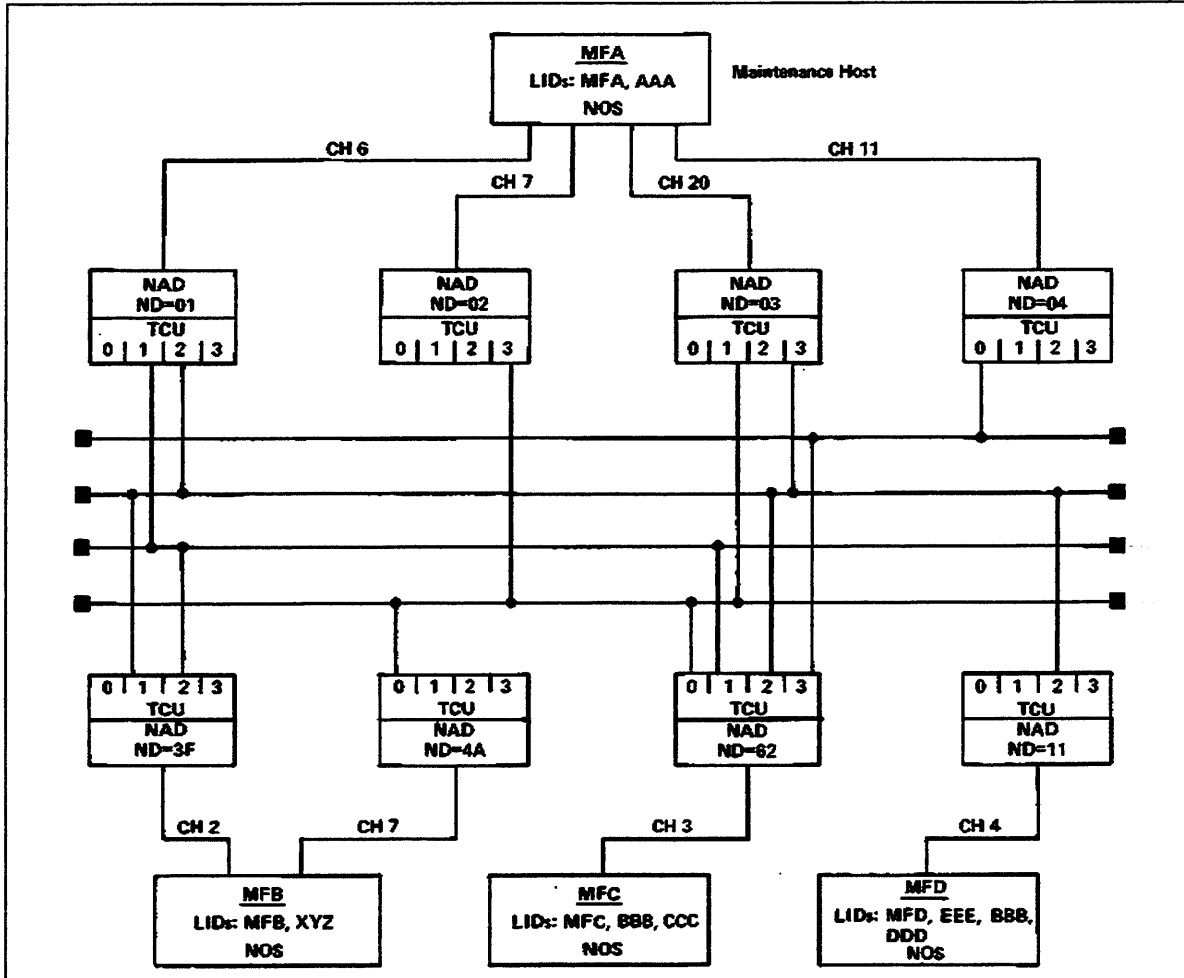


Figure 10-3. Sample LCN Network


```

LNDR    MAXNDRS=3
APPL    NAME=QTF ,MXCONS=4 ,ASTART=YES
APPL    NAME=QTFS ,MXCONS=1 ,MXCOPYS=8 ,ASTART=NO ,SVR=YES
APPL    NAME=PTF ,MXCONS=1 ,MXCOPYS=8 ,ASTART=NO
APPL    NAME=PTFS ,MXCOPYS=8 ,SVR=YES
APPL    NAME=USRAP ,MXCONS=6 ,ENABLED=NO
APPL    NAME=ITF ,MXCONS=2 ,MXCOPYS=1 ,ASTART=NO ,SVR=NO
NPID    PID=MFB ,ENABLED=YES ,MFTYPE=NOS2
PATH    ENABLED=YES ,LT=0110 ,RT=0110 ,RNAD=RN2 ,LNAD=LN1 ,AC=F0F0
PATH    ENABLED=YES ,LT=0001 ,RT=1000 ,RNAD=RN1 ,LNAD=LN2 ,AC=F0F0
PATH    ENABLED=NO ,LT=0001 ,RT=0100 ,RNAD=RN2 ,LNAD=LN3 ,AC=F0F0
PATH    ENABLED=YES ,LT=0100 ,RT=1000 ,RNAD=RN1 ,LNAD=LN3 ,AC=F0F0

NPID    PID=MFC ,MFTYPE=NOS2
PATH    ENABLED=NO ,LT=0110 ,RT=0110 ,RNAD=RN3 ,LNAD=LN1 ,AC=F0F0
PATH    ENABLED=NO ,LT=0001 ,RT=1000 ,RNAD=RN3 ,LNAD=LN2 ,AC=F0F0
PATH    LT=0101 ,RT=1010 ,RNAD=RN3 ,LNAD=LN3 ,AC=F0F0
PATH    LT=1000 ,RT=0001 ,RNAD=RN3 ,LNAD=LN4 ,AC=F0F0

NPID    PID=MFD ,MFTYPE=NOS2
PATH    LT=0010 ,RT=0010 ,RNAD=RN4 ,LNAD=LN1 ,AC=F0F0
PATH    LT=0001 ,RT=0010 ,RNAD=RN4 ,LNAD=LN3 ,AC=F0F0

LN1    LNAD    CH=6
LN2    LNAD    CH=7
LN3    LNAD    CH=20
LN4    LNAD    CH=11

RN1    RNAD    DD=0 ,ND=4A ,LOG=YES
RN2    RNAD    DD=0 ,ND=3F ,LOG=YES
RN3    RNAD    DD=0 ,ND=62 ,LOG=YES
RN4    RNAD    DD=0 ,ND=11 ,LOG=YES

```

Figure 10-4. Network Configuration Statements for the RHF on Mainframe MFA

APPL	NAME=QTF , ASTART=YES , MXCONS=4	
APPL	NAME=QTFS , MXCOPYS=4 , SVR=YES	
APPL	NAME=PTF , MXCOPYS=6	
APPL	NAME=PTFS , MXCOPYS=6 , SVR=YES	
APPL	NAME=ITF , MXCONS=2 , MXCOPYS=1 , ASTART=NO , SVR=NO	
NPID	PID=MFA , MFTYPE=NOS2	
PATH	RNAD=RN01 , LNAD=LN1 , LT=0010 , RT=0010 , AC=F0F0	
PATH	RNAD=RN03 , LNAD=LN1 , LT=0010 , RT=0001 , AC=F0F0	
NPID	PID=MFB , MFTYPE=NOS2	
PATH	RNAD=RN3F , LNAD=LN1 , LT=0010 , RT=0100 , AC=F0F0	
NPID	PID=MFC , MFTYPE=NOS2	
PATH	RNAD=RN62 , LNAD=LN1 , LT=0010 , RT=0010 , AC=F0F0	
LN1	LNAD	CH=4
RN01	RNAD	ND=01 , LOG=NO
RN03	RNAD	ND=03 , LOG=NO
RN3F	RNAD	ND=3F , LOG=NO
RN62	RNAD	ND=62 , LOG=NO

Figure 10-5. Network Configuration Statements for the RHF on Mainframe MFD

QTF Configuration Requirements

You must create the appropriate configuration files before you can perform queued file transfers using RHF or NAM.

RHF Configuration File Directives

For each copy of QTF that you want RHF to start automatically on a NOS host, you must include the following network configuration statement in the RCFGEN input file for that host:

```
APPL  NAME=QTF , MXCONS=4 , ASTART=YES
```

Refer to the APPL statement under RHF Configuration Files earlier in this section.

Each copy of QTF will establish at most four simultaneous connections. For most RHF configurations, a single copy of QTF should be sufficient.

To allow a NOS host to receive queued files, you must include the following network configuration statement in the RCFGEN input file for that host:

```
APPL  NAME=QTFS , SVR=YES , MXCOPYS=n
```

where n is the maximum number of QTFS servers you want to have active. Each copy of QTFS services one connection.

Refer to the appropriate reference manuals for other RHF implementations for information on how to configure non-NOS QTF and QTFS applications and connections.

NAM Configuration File Directives

In the NDLP input file, you must include QTFs OUTCALL statements for each PID to which you want to transfer queued files and one APPL statement for QTF:

```
QTF: APPL, PRIV, PRU, NETXFR, MXCOPY=n.
```

where n is the maximum number of QTF initiators you want to have active.

Parameter records in the NAMI job record file NAMSTRT specify which applications NAM automatically starts. Each default parameter record produced when QTF is installed includes one JOB(JOBQTF,QT) directive. To automatically start more than one copy of QTF, you can add additional job records to the NAMSTRT file and must add additional JOB directives to each parameter record desired. Refer to the NAM5 - Network Access Methods Version 1 subsection of the Special Product Installation Information section of the NOS Version 2 Installation Handbook for more information on NAMI job records.

Each copy of QTF will establish at most four simultaneous connections for file transfers. For many NAM configurations, a single copy of QTF should be sufficient. You might consider using more than one copy of QTF if you expect a large number of queued file transfers over slow communication lines.

To allow a NOS host to receive queued files, you must include in the NDLP input file for that host one or more QTFs INCALL statements and one APPL statement for QTFs:

```
QTFs: APPL, RS, PRIV, PRU, NETXFR, MXCOPY=n.
```

where n is the maximum number of QTFs servers you want to have active. Each copy of QTFs services one connection. Refer to the Network Access Method Network Definition Language Reference Manual for more information on APPL, INCALL, and OUTCALL statements.

The default NAMI parameter records include a JOB(JOBQTFs,QS) directive. The JOBQTFs job record creates or modifies permanent file ZZQTFs, which contains the job file submitted by NAM to start a new copy of QTFs when necessary.

QTF Procedure File

The QTF procedure file contains a set of default initialization commands for each of the two QTF variants, RHF and NAM. You can edit the commands in the procedure and capture the modified procedure in a new system file without rebuilding QTF, or you can include UPDATE modifications in the USER file for the RHP installation procedure. The QTF procedure is deck QTFPROC on the RHP1psrin program library file (refer to the NOS Version 2 Installation Handbook).

For the NAM variant, another option is to modify the JOBQTF record (on file NAMSTRRT), which is used to start QTF at NAM initiation.

Format of the call to the QTF procedure file is:

```
QTF,variant,I=infile.
```

Parameter	Description
variant	RHF or NAM, identifies which QTF variant to initiate. The default is RHF.
infile	File name from which initialization commands are to be processed. Default is to use the commands for the variant included in the QTF procedure file. QTF reads the initialization commands until end-of-record. Each command on the file is a separate line (equivalent to a K-display entry without the K. prefix). Specifying I=infile is equivalent to entering the command K.INCLUDE,FILE=infile.

Procedure File Example

Figure 10-6 shows some typical sets of initialization commands that you can include in the QTF procedure file. The example assumes you started with the default QTF procedure file.

```
.PROC,QTF*I, ...
:
:
.DATA,XXDEFI
.IF,$NAM$.EQ.$NOS,L1.
.*
.*  DEFAULT DIRECTIVES FOR NOS RHF VARIANT
.*
SCHED,MAXCONS=4.MAXIMUM OF FOUR CONNECTIONS.
CLASS,SC=A.          DEFINE THREE SELECTION CLASSES
CLASS,SC=B.          WITH IDENTICAL CRITERIA,
CLASS,SC=C.          ALLOWS UP TO 3 CONNECTIONS PER PID.
.ELSE,L1.
.*
.*  DEFAULT DIRECTIVES FOR NOS NAM VARIANT
.*
SCHED,MAXCONS=4.
CLASS,SC=A,FSI=1..4.
CLASS,SC=B,FSI=1..6.
CLASS,SC=C,FSI=1..6.
CLASS,SC=D,FSI=7,MAXIMUM=1. AT MOST ONE LARGE FILE.
DISABLE,SC=D,PID=M03. NEVER ALLOW LARGEST FILES TO M03.
.IF,(TIME.GT.0700).AND.(TIME.LT.1700),PRIME.
INCLUDE,F=PRIME.     MUST BE LAST DIRECTIVE
.ELSE,PRIME.
INCLUDE,F=OFF.      MUST BE LAST DIRECTIVE
.ENDIF,PRIME.
.ENDIF,L1.
.DATA,PRIME
.  PRIME-HOURS PARAMETERS
.  ENTER "K.INCLUDE,F=PRIME." TO CHANGE NAM PARAMETERS
DISABLE,SC=D.     NO LARGE FILES DURING PRIME TIME.
ENABLE,SC=B,PID=M03.  PERMIT MEDIUM FILES TO M03.
ENABLE,SC=C,PID=M03.
.DATA,OFF
.  OFF-HOURS PARAMETERS
.  ENTER "K.INCLUDE,F=OFF." TO CHANGE NAM PARAMETERS
ENABLE,SC=D.     ALLOW LARGE FILES EXCEPT TO M03.
DISABLE,SC=B,PID=M03.  DISALLOW MEDIUM FILES TO M03.
DISABLE,SC=C,PID=M03.
```

Figure 10-6. Procedure File Example

Mass Storage Extended Subsystem (MSE) 11

Introduction	11-1
MSE Utilities	11-2
Utility Description Terminology	11-3
Cartridge	11-3
SM Map	11-3
7990 Catalog	11-5
7990 Files	11-11
SSEXEC	11-12
Subfamily	11-12
Group	11-12
Volume	11-12
MSE Utilities	11-13
SSDEF	11-13
SSLABEL	11-14
SSLABEL Directives	11-15
OP=AS — Add SMs	11-16
OP=RS — Remove SMs	11-16
OP=AB — Add Cubicles	11-16
OP=RB — Remove Cubicles	11-16
OP=AM — Add Cartridges	11-17
OP=RM — Remove Cartridges	11-17
OP=RC — Restore Lost Cartridges	11-17
OP=FX — Fix Cartridge Labels	11-18
OP=IB — Control Cartridge Allocation	11-18
OP=FC — Free Cartridge Files	11-18
Parameters for SSLABEL Directives	11-19
SSLABEL Update Sequence	11-22
Restrictions to SSLABEL	11-23
SSLABEL Example	11-23
SSMOVE	11-24
SSMOVE Directives	11-27
Specific File (SF) Directives	11-27
Value Specifier Directives	11-28
Excluding Destage/Release Processing	11-32
Selection Algorithms	11-32
Destage and Release Algorithms	11-33
Decision Algorithm Hierarchy	11-34
Disk Space/Dump Tape Management	11-35
Restrictions to SSMOVE	11-35
Output Example	11-36
SSVAL	11-40
Release Processing	11-42
Problem Reporting	11-43
Error Detection and Classification	11-43
SM Map Analysis	11-43
7990 Catalog Analysis	11-44
PFC Analysis	11-45
Release Processing and Problem Fixing	11-46
Validation Report	11-47

Error Types 1 and 2	11-47
Error Types 3 and 4	11-47
Error Types 5, 6, and 7	11-47
Intersections	11-48
Validation Report Example	11-48
Typical SSVAl Runs	11-49
Restrictions to SSVAl	11-50
SSUSE	11-51
Basic Usage Report	11-52
Optional Report A	11-54
Optional Report B	11-56
Optional Report C	11-57
Optional Report D	11-58
SSDEBUG	11-59
SSDEBUG Directives	11-60
OP=RS — Read AUs	11-60
OP=RF — Read File	11-60
OP=RP — Release Frozen Chain Space	11-60
OP=RL — Remove Cartridge Entry From 7990 Catalog	11-60
OP=RC — Remove Cartridge Entry From SM Map	11-60
OP=CF — Change Flag in SM Map or 7990 Catalog	11-60
Parameters for the SSDEBUG Directives	11-61
Using SSDEBUG	11-64
Restrictions to SSDEBUG	11-65
y,z Relationship with SM Map Ordinal	11-65
Format of Disk Files Read From 7990 by SSDEBUG	11-66
7990 Stripe Header and Trailer	11-68
7990 Block Header and Trailer	11-69
SSBLD	11-70
SSBLD Statements	11-70
Node/Path Linkage Statements	11-70
SM/CU Linkage Statements	11-72
SSBLD Example	11-72
7990 Equipment Configuration Constraints	11-73
SSALTER	11-74
Console Input	11-74
MSE Operational Procedures	11-77
Initialization	11-77
Disk Space Management	11-78
7990 Space Management	11-79
Backup and Recovery	11-81
File Dumping	11-81
Full Dumps	11-81
Incremental Dumps	11-82
Dump Tape Management	11-82
SM Map/7990 Catalog Backup	11-83
File Reloading	11-84
Error Conditions and Corrective Actions	11-85
Clearing Permanent File Error Flags	11-85
Permanent File Recovery	11-85
Cartridge Restoration and Reuse	11-86
Removal of Faulty or Missing Cartridges	11-86

7990 Catalog/SM Map/PFC Problems	11-87
7990 Catalog/SM Map Mismatch	11-87
7990 Catalog Chain Problems	11-88
PFC/7990 Catalog Mismatch	11-88

Mass Storage Extended Subsystem (MSE) 11

Introduction

The Mass Storage Extended Subsystem (MSE) is the product consisting of the 7990 hardware, the channel interface, the diagnostics, and the operational software.

The 7990 hardware is a large-capacity online mass storage device, which is a cost-effective extension to the disk file storage system and an alternative to other types of magnetic tape storage. Storing files on the 7990 retains the security, data integrity, and online access capabilities provided by disk and reduces the operational and data integrity problems caused by storing, retrieving, and mounting tape volumes.

The 7990 hardware is composed of the following components:

Component	Description
Cartridge for the 7990	A plastic housing that encloses magnetic tape on which data is stored under program control. This cartridge is compatible with an IBM 3850 cartridge. Even though a cartridge for the 7990 is a type of magnetic tape, this manual reserves the terms tapes and magnetic tapes to reference both reel magnetic tapes and the cartridge tapes for ACS and CTS; tapes and magnetic tapes references never include cartridges for the 7990.
7991 Storage Module (SM)	The hardware unit that houses up to 312 usable data cartridges. The SM also consists of the cartridge accessor unit, which picks cartridges from and puts cartridges in their cubicles, and one or two data recording drives (DRD), which read data from and write data to the cartridges. The cubicles are assigned coordinate locations that are identified by the ordered pair (y,z), where the z axis is horizontal and the y axis is vertical.
7990 Control Unit (CU)	The hardware unit that is the controller for up to four 7991 SMs. The CU acts like a tape controller for up to four IBM channel interfaces and includes one or two data recording controllers (DRCs). The MSE software supports up to two IBM channel interfaces per CYBER mainframe.

To protect files from hardware and system failures, you can store copies of the same files on both cartridges for the 7990 and magnetic tapes. You cannot, however, use the 7990 hardware and optical disk to store copies of the same files.

MSE Utilities

The basic function of MSE is to store data on the 7990 and move it to disk upon request for access by an authorized user. Control of file movement between a disk and a 7990 is transparent to the applications programmer; however, there are operational changes and additions that site personnel should note. These include modifications to operational procedures in the areas of permanent file backup, permanent file recovery, and disk space management. In addition, new MSE utilities are introduced. Detailed information about the call and operation of each MSE utility, listed next, is included in this section.

Utility	Description
SSDEF	Creates the system files (SM maps and 7990 catalogs) necessary for MSE processing.
SSLABEL	Manages allocation of cubicles and assignment of cartridges in the SM.
SSMOVE	Controls the destaging of files (creating 7990 images) and the release of disk space.
SSVAL	Controls the release of 7990 space and analyzes the SM maps, 7990 catalogs, and PFC entries to identify and flag discrepancies within these three components.
SSUSE	Provides reports on the assignment and availability of cartridges and cubicles within an SM.
SSDEBUG	Corrects error conditions detected by SSVAL and recovers data from 7990 cartridges.
SSBLD	Builds the unit device table (UDT) for SSEXEC. The UDT describes the 7990 hardware configuration.
SSALTER	Provides dynamic modification of the 7990 hardware configuration.

Utility Description Terminology

Review the following terms; they are defined as used in the descriptions of the MSE utilities later in this section.

Cartridge

A cartridge is the 7990 storage component consisting of magnetic tape. The 7990 hardware always writes or reads 6652 8-bit byte records, called stripes, to or from the magnetic tape. Each cartridge has 27,087 stripes. The first 20 stripes contain a manufacturer's label. The second 20 stripes contain a 7990 cartridge label written by the SSLABEL utility. The remaining 27,047 stripes are organized into software-defined allocation units (AUs) of 14 stripes each for a total of 1931 AUs per cartridge.

An allocation unit is the smallest allocatable portion of a cartridge. One or more AUs may be allocated to contain data for a single file. The allocated AUs for each file are organized into volumes (definition follows), which are chained together in the 7990 catalog to identify the sequence of AUs that must be accessed in order to read a file. A head-of-chain (HOC) flag identifies the first volume in the chain, a link field identifies the next volume in the chain, and an end-of-chain (EOC) flag identifies the last volume in the chain.

The addition, removal, and reassignment of cartridges are managed by the MSE utilities. The cartridge labels, and also the 7990 catalogs and SM maps (definitions following), contain information concerning the location and content of the cartridges. The descriptions of the MSE utilities later in this section contain further information on cartridge management.

SM Map

An SM map is a direct access permanent file that contains information indicating how cubicles in an SM are assigned to a family and identifying the cartridges that reside in the SM. There is one SM map for each SM in the configuration. The permanent file name of the SM map is SMMAP_i, where *i* is the SM identifier (a letter from A to H); its user index is 377760₈; and its family is the default family on the mainframe on which SSEXEC executes (refer to the definition of SSEXEC later in this section).

An SM map contains an entry for each possible coordinate pair (*y,z*) that identifies a cubicle in the SM, from (0,0), which is bottom left, to (21,15), which is upper right. Certain coordinate pairs are reserved for customer engineering use, system use, and diagnostic purposes.

Each entry in an SM map has this format:

59	53	47	41	0
code	ord		familyname	
<i>l</i>	sub- family	id	unused	
cm		csn		

Field	Description																
code	Number from 0 to 6 indicating how the cubicle is assigned.																
	<table border="1"> <thead> <tr> <th style="text-align: left;">code</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Unassigned.</td> </tr> <tr> <td>1</td> <td>Reserved for customer engineer.</td> </tr> <tr> <td>2</td> <td>Reserved for system use.</td> </tr> <tr> <td>3</td> <td>Reserved for a different SM map.</td> </tr> <tr> <td>4</td> <td>Assigned to the cartridge scratch pool.¹</td> </tr> <tr> <td>5</td> <td>Assigned to a subfamily.</td> </tr> <tr> <td>6</td> <td>No cubicle exists at these coordinates.</td> </tr> </tbody> </table>	code	Description	0	Unassigned.	1	Reserved for customer engineer.	2	Reserved for system use.	3	Reserved for a different SM map.	4	Assigned to the cartridge scratch pool. ¹	5	Assigned to a subfamily.	6	No cubicle exists at these coordinates.
code	Description																
0	Unassigned.																
1	Reserved for customer engineer.																
2	Reserved for system use.																
3	Reserved for a different SM map.																
4	Assigned to the cartridge scratch pool. ¹																
5	Assigned to a subfamily.																
6	No cubicle exists at these coordinates.																
ord	Ordinal for this cubicle in the 7990 catalog of the subfamily to which the cubicle is assigned. This ordinal is referred to as the FCT ordinal (refer to the definition of 7990 catalog, next). This field is meaningful only if code=5.																
familyname	The 7-character name, in 6-bit display code, of the family to which the cubicle is assigned. This field is meaningful only if code=5.																
<i>l</i>	Linkage error flag that is set by the SSVAl utility when an SM map entry of a cubicle assigned to a family has no corresponding entry in the 7990 catalog.																
subfamily	Number from 0 to 7 identifying the subfamily to which the cubicle is assigned. This field is meaningful only if code=5.																
id	Letter from A to H identifying the SM.																
cm	Cartridge manufacturer code.																
	<table border="1"> <thead> <tr> <th style="text-align: left;">cm</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>A-</td> <td>Cartridge manufacturer is IBM.</td> </tr> <tr> <td>B-</td> <td>Reserved for future cartridge manufacturer.</td> </tr> </tbody> </table>	cm	Description	A-	Cartridge manufacturer is IBM.	B-	Reserved for future cartridge manufacturer.										
cm	Description																
A-	Cartridge manufacturer is IBM.																
B-	Reserved for future cartridge manufacturer.																
csn	The 8-character cartridge serial number of the cartridge assigned to the cubicle. If no cartridge is assigned, this field contains spaces.																

The zero entry in the SM map is the map header entry. In this entry the code field is 6, the leftmost 6 bits of the second word contain the SM identifier, and the remaining bits contain the permanent file name of the SM map.

¹. The pool is an area of the SM that stores scratch cartridges managed by the SSLABEL utility.

The SM map is updated whenever the SSLABEL, SSVAL, or SSDEBUG utility causes a change in cubicle or cartridge assignment. It is recommended that the SM map be backed up after every update to avoid problems such as:

- Mismatches between SM map and 7990 catalog entries
- Lost SM maps because of a disk failure or other problem
- Attempts to access cartridges that are no longer available
- Attempts to store cartridges in cubicles that are no longer empty


Thus, you should make a copy of the SM map on tape or another device or family in order to retain the latest version of the SM map. If a device containing SM maps is reloaded, the latest version of the SM maps should be recovered from the backup copy after the reload is completed. After recovering the SM maps, you should run the SSVAL utility to check that the entries in the SM maps and 7990 catalogs match. If there are inconsistencies, corrective action should be taken as described under Error Conditions and Corrective Actions later in this section.

7990 Catalog

A 7990 catalog is a disk-resident direct access permanent file that contains information describing which allocation units of each cartridge assigned to a particular subfamily are allocated to 7990 files and which AUs are available for allocation. There is one 7990 catalog for each subfamily of a family that can have 7990-resident files, and it resides on the master device for the subfamily. The permanent file name of the 7990 catalog file is SFMCAT_i and its user index is 37776_i₈, where *i* is the subfamily identifier (a number from 0 to 7). For example, file SFMCAT₃ and user index 37776₃₈ identify the 7990 catalog for subfamily 3.


A 7990 catalog is partitioned into subcatalogs, one subcatalog for each SM used by the subfamily. The maximum number of subcatalogs in a 7990 catalog is eight (the maximum number of SMs in a configuration). Each subcatalog consists of two parts, the file and cartridge table (FCT) and the allocation summary table (AST). The FCT has an entry for each cartridge assigned to the subfamily from the given SM. The maximum number of FCT entries in a subcatalog is 312 (the maximum number of cartridges in an SM). The AST contains information used by the allocation algorithm to select the cartridges on which a file will reside.

The first part of the 7990 catalog contains a header and at most eight subcatalog entries. The header identifies the family and subfamily of the 7990 catalog. Each subcatalog entry contains the length and location of its FCT and AST, the SM identifier, unallocated AUs available for small and large files, and the date and time of the last SSVAl run that resulted in the release of 7990 space assigned to orphan files. The format of the header is:



Field	Description
familyname	The 7-character name, in 6-bit display code, of the family for this 7990 catalog.
subfamily	Number from 0 to 7 identifying the subfamily for this 7990 catalog.
id	Number identifying this file as a 7990 catalog file. id is set to 1.

The format of the subcatalog entry is:



Field	Description
id	SM identifier for the subcatalog (a letter from A to H).
length	Number of FCT (and AST) entries in the subcatalog.
FCT loc	Location (beginning PRU number) of the FCT.
AST loc	Location (beginning PRU number) of the AST.
date time	Date and time of the last releasing of orphan files on the SM identified by id.
AUs small	Total number of free AUs available for small files on all cartridges in this SM.
AUs large	Total number of free AUs available for large files on all cartridges in this SM.
AUs group	Maximum number of free AUs available within one group in this SM.

Each cubicle from the given SM assigned to the subfamily has an entry in the FCT of the subcatalog. This entry contains the coordinates (y,z) of the assigned cubicle. If a cartridge has been assigned to the cubicle, the FCT entry also contains the csn of the cartridge, usage information, status flags, and information about each of the AUs of the assigned cartridge. The format of each FCT entry is:

59										55			53			51			47			41			35			32			29			17			11			2			0		
csn																		cm																											
y			z			first AU large			first AU small			flags			ord																														
ocl ₃			ocl ₂			ocl ₁			first stripe			stripe			pru			max AUs			cdp			oclu																					
unused																																													
unused															flawed AU																														
read stripe count									write stripe count																																				
soft read error									soft write error																																				
hard read error			demarked stripe			load count			load error																																				
cartridge link									cartridge link																																				
•																																													
•																																													
•																																													
cartridge link									cartridge link																																				

Field	Description
csn	The 8-character cartridge serial number, in 6-bit display code, of the cartridge assigned to this cubicle. If no cartridge is assigned, this field contains spaces.
cm	Cartridge manufacturer code.
cm	Description
A-	Cartridge manufacturer is IBM.
B-	Reserved for future cartridge manufacturer.
y	y coordinate of this cubicle.
z	z coordinate of this cubicle.
first AU large	First free AU available for large files in this cartridge.
first AU small	First free AU available for small files in this cartridge.

Field	Description														
flags	One of the following flags: <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>17</td> <td>Inhibit allocation flag indicating that space from the cartridge assigned to this cubicle is not to be allocated to a file. This flag is set or cleared by a directive to the SSLABEL utility.</td> </tr> <tr> <td>16</td> <td>Lost cartridge flag indicating that the cartridge assigned to this cubicle was not there the last time SSEXEC tried to pick it. This flag can be cleared by a directive to the SSLABEL utility.</td> </tr> <tr> <td>15</td> <td>Excessive write parity error flag set by SSEXEC when an error threshold has been passed. Space from this cartridge is not to be allocated. This flag can be cleared by a directive to the SSDEBUG utility.</td> </tr> <tr> <td>14</td> <td>Linkage error flag indicating that the SM map entry for the location (y,z) is inconsistent with the FCT entry. This flag is set by the SSVAL utility and cleared by the SSDEBUG utility.</td> </tr> <tr> <td>13</td> <td>Free cartridge flag. When this flag is set, SSEXEC inhibits further allocation. Files can be removed by running the SSVAL utility. This flag is set or cleared by a directive to the SSLABEL utility.</td> </tr> <tr> <td>12</td> <td>Reserved.</td> </tr> </tbody> </table>	Bit	Description	17	Inhibit allocation flag indicating that space from the cartridge assigned to this cubicle is not to be allocated to a file. This flag is set or cleared by a directive to the SSLABEL utility.	16	Lost cartridge flag indicating that the cartridge assigned to this cubicle was not there the last time SSEXEC tried to pick it. This flag can be cleared by a directive to the SSLABEL utility.	15	Excessive write parity error flag set by SSEXEC when an error threshold has been passed. Space from this cartridge is not to be allocated. This flag can be cleared by a directive to the SSDEBUG utility.	14	Linkage error flag indicating that the SM map entry for the location (y,z) is inconsistent with the FCT entry. This flag is set by the SSVAL utility and cleared by the SSDEBUG utility.	13	Free cartridge flag. When this flag is set, SSEXEC inhibits further allocation. Files can be removed by running the SSVAL utility. This flag is set or cleared by a directive to the SSLABEL utility.	12	Reserved.
Bit	Description														
17	Inhibit allocation flag indicating that space from the cartridge assigned to this cubicle is not to be allocated to a file. This flag is set or cleared by a directive to the SSLABEL utility.														
16	Lost cartridge flag indicating that the cartridge assigned to this cubicle was not there the last time SSEXEC tried to pick it. This flag can be cleared by a directive to the SSLABEL utility.														
15	Excessive write parity error flag set by SSEXEC when an error threshold has been passed. Space from this cartridge is not to be allocated. This flag can be cleared by a directive to the SSDEBUG utility.														
14	Linkage error flag indicating that the SM map entry for the location (y,z) is inconsistent with the FCT entry. This flag is set by the SSVAL utility and cleared by the SSDEBUG utility.														
13	Free cartridge flag. When this flag is set, SSEXEC inhibits further allocation. Files can be removed by running the SSVAL utility. This flag is set or cleared by a directive to the SSLABEL utility.														
12	Reserved.														
ord	Ordinal of the FCT entry in the subcatalog.														
ocl _i	Off-cartridge link specifying the ordinal of the FCT entry for the cubicle containing the next cartridge on which the file on the cartridge assigned to this cubicle resides. If a file does reside on multiple cartridges, the entire file must be contained within a 16-cartridge group. Also, since there are only three off-cartridge link fields, if a cartridge contains several files, only three can be contained on other cartridges.														

Field	Description
first stripe	First usable stripe following the cartridge label. This field is set to 40.
stripe	Number of stripes in each AU. This field is set to 14.
pru	Number of PRUs that can be written on each stripe. This field is set to 13.
max AUs	Maximum number of AUs in a cartridge. This field is set to 1931.
cdp	Cartridge division point between small and large files kept on the same cartridge. The cartridge division point is defined with the B=n parameter to an SSLABEL directive. The default is B=600, specifying that the first 600 AUs of a cartridge are available for small files.
oclu	Off-cartridge link usage field. When an ocl ₁ field is used, a corresponding bit is set in this field. Bit 0 is set for ocl ₁ , bit 1 is set for ocl ₂ , and bit 2 is set for ocl ₃ .
flawed AU	Number of flawed AUs on this cartridge.
read stripe count	Accumulated number of stripes read from this cartridge.
write stripe count	Accumulated number of stripes written to this cartridge.
soft read error	Accumulated number of soft read errors detected on this cartridge.
soft write error	Accumulated number of soft write errors detected on this cartridge.
hard read error	Accumulated number of hard read errors detected on this cartridge.
demarked stripe	Accumulated number of stripes demarked on this cartridge.
load count	Accumulated number of times this cartridge has been loaded for a read/write operation.
load error	Accumulated number of cartridge load errors detected using this cartridge.

Field	Description																								
cartridge link	There is a 30-bit cartridge link field for each AU in the cartridge. The bits are divided into the same fields and represent the same information for each AU.																								
	<table border="1"> <thead> <tr> <th>Bit(s)</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>29</td> <td>Free/busy flag indicating whether or not this AU is currently allocated to a file.</td> </tr> <tr> <td>28</td> <td>Unused.</td> </tr> <tr> <td>27</td> <td>AU conflict flag indicating an allocation conflict involving this AU. This flag is set by the SSVAL utility or by SSEXEC.</td> </tr> <tr> <td>26</td> <td>Frozen chain flag indicating a problem with this allocation chain. This flag is set by the SSVAL utility or by SSEXEC. This AU is not reused until this flag is cleared by a directive to the SSDEBUG utility.</td> </tr> <tr> <td>25</td> <td>Start of fragment flag indicating this AU is the beginning of a chain fragment.</td> </tr> <tr> <td>24</td> <td>Flawed AU flag indicating this AU is not to be allocated to a file.</td> </tr> <tr> <td>23</td> <td>Continuation AU flag indicating whether or not this AU is the first AU of a volume. If this flag is clear, the AU is the first AU of the volume; if set, the AU is not the first AU of the volume.</td> </tr> <tr> <td>22–21</td> <td>Off-cartridge link field indicating that the next volume of the file is on another cartridge. The next cartridge is identified by one of the ocl_i fields. The value of the off-cartridge link flag (1, 2, or 3) specifies which ocl_i field to use. The link field specifies the number of the first AU on the next volume of the file. This field is meaningful only if the chain control field indicates that this volume is the first or middle volume of the file.</td> </tr> <tr> <td>20–19</td> <td>Chain control field indicating whether this volume is the first (1), last (2), only (3), or middle (0) volume of the file.</td> </tr> <tr> <td>18–12</td> <td>File length field indicating the number of AUs in the file following the first accessed AU.</td> </tr> <tr> <td>11–0</td> <td>Link field indicating the first AU on the next volume of the file. If the continuation AU flag is clear, this field contains the ordinal of the first AU of the volume. If the continuation AU flag is set, this field contains the ordinal of the first AU within the volume.</td> </tr> </tbody> </table>	Bit(s)	Description	29	Free/busy flag indicating whether or not this AU is currently allocated to a file.	28	Unused.	27	AU conflict flag indicating an allocation conflict involving this AU. This flag is set by the SSVAL utility or by SSEXEC.	26	Frozen chain flag indicating a problem with this allocation chain. This flag is set by the SSVAL utility or by SSEXEC. This AU is not reused until this flag is cleared by a directive to the SSDEBUG utility.	25	Start of fragment flag indicating this AU is the beginning of a chain fragment.	24	Flawed AU flag indicating this AU is not to be allocated to a file.	23	Continuation AU flag indicating whether or not this AU is the first AU of a volume. If this flag is clear, the AU is the first AU of the volume; if set, the AU is not the first AU of the volume.	22–21	Off-cartridge link field indicating that the next volume of the file is on another cartridge. The next cartridge is identified by one of the ocl_i fields. The value of the off-cartridge link flag (1, 2, or 3) specifies which ocl_i field to use. The link field specifies the number of the first AU on the next volume of the file. This field is meaningful only if the chain control field indicates that this volume is the first or middle volume of the file.	20–19	Chain control field indicating whether this volume is the first (1), last (2), only (3), or middle (0) volume of the file.	18–12	File length field indicating the number of AUs in the file following the first accessed AU.	11–0	Link field indicating the first AU on the next volume of the file. If the continuation AU flag is clear, this field contains the ordinal of the first AU of the volume. If the continuation AU flag is set, this field contains the ordinal of the first AU within the volume.
Bit(s)	Description																								
29	Free/busy flag indicating whether or not this AU is currently allocated to a file.																								
28	Unused.																								
27	AU conflict flag indicating an allocation conflict involving this AU. This flag is set by the SSVAL utility or by SSEXEC.																								
26	Frozen chain flag indicating a problem with this allocation chain. This flag is set by the SSVAL utility or by SSEXEC. This AU is not reused until this flag is cleared by a directive to the SSDEBUG utility.																								
25	Start of fragment flag indicating this AU is the beginning of a chain fragment.																								
24	Flawed AU flag indicating this AU is not to be allocated to a file.																								
23	Continuation AU flag indicating whether or not this AU is the first AU of a volume. If this flag is clear, the AU is the first AU of the volume; if set, the AU is not the first AU of the volume.																								
22–21	Off-cartridge link field indicating that the next volume of the file is on another cartridge. The next cartridge is identified by one of the ocl_i fields. The value of the off-cartridge link flag (1, 2, or 3) specifies which ocl_i field to use. The link field specifies the number of the first AU on the next volume of the file. This field is meaningful only if the chain control field indicates that this volume is the first or middle volume of the file.																								
20–19	Chain control field indicating whether this volume is the first (1), last (2), only (3), or middle (0) volume of the file.																								
18–12	File length field indicating the number of AUs in the file following the first accessed AU.																								
11–0	Link field indicating the first AU on the next volume of the file. If the continuation AU flag is clear, this field contains the ordinal of the first AU of the volume. If the continuation AU flag is set, this field contains the ordinal of the first AU within the volume.																								

The 7990 catalog is updated whenever the SSLABEL, SSMOVE, SSVAL, or SSDEBUG utility causes a change in cartridge or cubicle assignment that affects the subfamily. Because the 7990 catalog for a subfamily resides on its master device, it will be backed up by PFDUMP whenever the master device for the subfamily is dumped. Consequently, when PFLOAD reloads all files, the 7990 catalogs are automatically recovered and all 7990-resident files as indicated in a recovered PFC will also have entries in the recovered 7990 catalog. Thus, no special operational procedures are needed to back up a 7990 catalog. It is possible, however, that the 7990 catalog will be inconsistent with the SM maps or cartridge labels. If such inconsistencies do exist, corrective action will have to be taken as described under Error Conditions and Corrective Actions later in this section.

7990 Files

In an MSE environment, permanent files can be categorized according to whether or not they reside on 7990. A disk file is a permanent file that resides on disk but not on 7990. A 7990 file is a permanent file that resides on 7990 and may or may not also reside on disk, depending on how the disk space is managed (refer to Disk Space Management later in this section). Depending on backup requirements (BR parameters), both disk and 7990 files can also have backup images on tape (refer to the NOS Version 2 Reference Set, Volume 3).

When a user defines a direct access file, initially it is a disk file. A disk file becomes a 7990 file when it is destaged to 7990; that is, an image of the file is created on 7990. Files are destaged through use of the SSMOVE utility, which is run periodically to manage disk space. When SSMOVE is run, files are destaged to 7990 and/or their disk space released depending on certain file characteristics (refer to SSMOVE later in this section). Thus, after an SSMOVE run a file can reside on disk, on 7990, or on both. If the file does have a 7990 image, the asa field in the PFC entry for the file indicates the location of the 7990 copy.

When a user attaches a 7990 file, it is staged to disk from 7990 (that is, a disk image is created) if the file is not on disk. If a direct access file is attached in write mode, the AFOBS flag is set in the PFC entry for that file. The current version of the file resides on disk only. If an indirect access file is replaced or appended to, the asa field is cleared in the PFC entry for that file. Setting the AFOBS flag or clearing the asa in the PFC ensures that the current version of the file will be copied to 7990 when that file is again selected for being released from disk.

If the file is purged, its disk space, but not its 7990 space, is immediately released. The SSVAL utility must be run to release 7990 space allocated to purged files. Thus, because a purged file has no PFC entry linking to its 7990 catalog entry, the 7990 image that still exists before SSVAL is run is called an orphan file. However, a user can never access an orphan file.

When a file is destaged to 7990, control information is written on each stripe to which the file data is written. This information is sufficient to identify the file.

SSEXEC

SSEXEC is the main processing program that is responsible for controlling MSE activities. The MSE utilities issue requests to SSEXEC to destage files from disk to 7990, purge unneeded 7990 files, label or relabel cartridges, update SM maps and 7990 catalogs, and so forth. In a multimainframe environment, two versions of this program exist: the mainframe to which the 7990 device is physically connected (the master mainframe) has a program called SSEXEC, and all other mainframes (the slave mainframes) have a program called SSSLV. Refer to section 13, Multimainframe Operations, for more information.

Subfamily

Each permanent file family consists of eight subfamilies, subfamily 0 through subfamily 7. The lower 3 bits of the user index identify the subfamily to which a user belongs. For example, a user whose index ends in 3 (or 011 in bit notation) belongs to subfamily 3. When the SSDEF utility is run to create MSE system files, the SM maps are created on the master device under user index 377760_8 (subfamily 0) and one 7990 catalog is created under each user index $37776i_8$ (subfamily i), $i=0,1,2,\dots,7$. When the SSLABEL utility is used to assign an SM, cubicle, or cartridge to a family, it is possible to specify assignment to particular subfamilies of the family.

Group

A group is a software-defined structure for allocating cartridges within a subfamily. A group of cartridges is a number of cartridges (up to 16) in which files belonging to a subfamily can reside. Any file can overflow from one cartridge in a group to any other cartridge in the same group. However, no file can overflow from a cartridge in one group to any cartridge in another group. Cartridges are placed in or removed from specific groups by using the SSLABEL utility. The SSVAl utility compares the 7990 catalogs (for each subfamily) with the same information obtained from the PFC entries.

Volume

A volume is a software-defined structure used to simplify the allocation and access of large 7990 files. Up to 128 consecutive AUs in a cartridge that have the same status (either allocated or not allocated) are considered to belong to the same volume. These consecutive AUs are composed of stripes that all have the same volume serial number (a physical identifier that is written to the tape on the cartridge). The MSE software fills cartridges (unallocated volumes) from the beginning to the end with no rewinding. This action minimizes repositioning of the tape in a cartridge for multivolume files.

When a file is staged back to disk, the control information on each stripe is verified. If a discrepancy is detected, an error message is issued, the file stage is aborted, an error flag is set in the PFC entry to indicate that the 7990 file could not be accessed, and the AU conflict flag is set in the 7990 catalog entry for the particular cartridge(s) and AU(s).

MSE Utilities

MSE utilities are described next. All of these utilities must be run from system origin jobs. In a multimainframe environment, the SSDEF, SSUSE, and SSBLD utilities can be run on any mainframe that has access to the family being processed. However, the remaining utilities must be run on the mainframe on which SSEXEC executes.

SSDEF

SSDEF creates the system files (SM maps and 7990 catalogs) that are necessary for MSE processing. If an SM is added to the 7990 hardware configuration, SSDEF is used to create the SM map for that SM. If a family is to be permitted to have 7990-resident files, SSDEF is used to create the eight 7990 catalogs for that family (one catalog for each subfamily).

NOTE

If the SM parameter is specified, SSDEF will create an SM map for the specified SM (refer to the definition of SM map earlier in this section). It is recommended that the SM map be copied on tape or on another device or family immediately after it is created.

The format of the SSDEF command is:

```
SSDEF , P1 , P2 .
```

P_i	Description
FM=familyname	Family for which 7990 catalogs are to be created, one catalog for each subfamily.
FM	Same as FM=system default family.
FM omitted	No 7990 catalogs are to be created. SM=id or SM must be specified.
SM=id	SM identifier of the SM for which an SM map is to be created; id is a letter from A to H.
SM	Same as SM=A.
SM omitted	No SM map is to be created. FM=familyname or FM must be specified.

Example:

```
SSDEF , SM=B .
```

SMMAPB, the SM map for SM B, is created, and its entries are as described in the previous definition of SM map earlier in this section. However, since SSDEF does not assign cubicles, the entries for cubicles available for use initially indicate that the cubicles are unassigned.

SSLABEL

SSLABEL manages cartridge assignment and cubicle allocation in an SM. The following functions are performed through use of directives to SSLABEL.

- Add an SM to a subfamily (AS directive).
- Remove an SM from a subfamily (RS directive).
- Add cubicles to a subfamily, the pool, or the reserved area (AB directive).
- Remove cubicles from a subfamily, the pool, or the reserved area (RB directive).
- Add cartridges to a subfamily or the pool (AM directive).
- Remove cartridges from a subfamily or the pool (RM directive).
- Restore an abnormally removed cartridge (RC directive).
- Repair a cartridge label or overwrite a family label (FX directive).
- Inhibit or allow further allocation of files to a cartridge (IB directive).
- Initiate freeing all files from an existing cartridge (FC directive).

Input to SSLABEL is via a directive file. SSLABEL reads the appropriate SM maps and 7990 catalogs to determine how to process each directive and then issues requests to SSEXEC to read and/or write cartridge labels and to update the SM maps and 7990 catalogs. SSLABEL generates a report detailing the action taken for each input directive. If the assignment information or cartridge label is not appropriate or conflicts with data in the SM map or 7990 catalog, the cartridge label information is included on this report and the cartridge is put into the exit tray. It may be possible to restore such a cartridge, as described under Cartridge Restoration and Reuse later in this section.

NOTE

SSLABEL updates the SM map for the specified SM (refer to the definition of SM map earlier in this section). It is recommended that the SM map be copied on tape or on another device or family immediately after each update.

The format of the SSLABEL command is:

```
SSLABEL, P1, P2.
```

P_i	Description
I=filename	File containing the directives to SSLABEL.
I	Same as I=INPUT.
I omitted	Same as I=INPUT.
L=filename	File on which listable output is to be written.
L	Same as L=OUTPUT.
L=0	No output file is to be generated.
L omitted	Same as L=OUTPUT.
Z	Directives are contained on the SSLABEL command. The I parameter is ignored.
Z omitted	Directives are contained on the file specified by the I parameter.

SSLABEL Directives

The directives to SSLABEL can be specified on a separate file (specified by the I parameter) or after the SSLABEL command (Z specified). If on the input file, each directive must be specified on a separate line using the OP=directive option. Parameters for a directive are on the same line, are separated by commas, and end with a period:

```
OP=directive, P1, P2, . . . , Pn.
```

Example 1:

```
SSLABEL, I=DIRFILE.
```

DIRFILE is the directive file; it contains the following directives:

```
OP=AM, N=4, PK=D.  
OP=RM, CN=44455566, FM, SB=1, PK=F.
```

Two directives to SSLABEL are specified. OP=AM adds four cartridges to the pool of SM A (SM parameter not specified). OP=RM removes the cartridge whose csn is 44455566 from subfamily 1 of the default family. (Refer to the following descriptions of directives and parameters.)

If the directives are contained on the command, they follow the command terminator. The first character following the terminator is the separator. Any character that does not appear in any of the directives can be used as the separator character. Each directive must be preceded by the separator and terminated by a period.

Example 2:

```
SSLABEL, Z. /OP=AM, N=4, PK=D. /OP=RM, CN=44455566, FM, SB=1, PK=F.
```

The slant is used as the separator. This command performs the same functions as those in the preceding example.

The following directives are available with SSLABEL. Some of these directives cause cartridges to be physically moved to and from cubicles and the input or exit tray (refer to OP=AM, OP=RM, OP=RC, and OP=FX in the following paragraphs). The remaining directives cause only logical operations to occur, updating the MSE system files. Descriptions of the parameters for these directives follow the directives.

OP=AS — Add SMs

OP=AS adds an SM to a subfamily. The 7990 catalog for the specified subfamily is updated to reflect that cartridges and permanent files for the subfamily can reside on the specified SM. This directive, however, does not manipulate cartridges or cubicles.

OP=RS — Remove SMs

OP=RS removes an SM from a subfamily. The 7990 catalog for the specified subfamily is updated to reflect that cartridges and permanent files for the subfamily cannot reside on the specified SM. Before OP=RS can be specified, all cubicles in the specified SM must have been removed previously from the family (refer to the OP=RB directive). This directive, however, does not manipulate cartridges or cubicles.

OP=AB — Add Cubicles

OP=AB adds an unassigned cubicle within an SM to a subfamily (PT=F), the pool (PT=P), or the reserved area of the SM (PT=R). More than one cubicle (N=n) can be added at a time. Specific cubicles (YI and ZI parameters) can be added, but they must be currently unassigned. For PT=R, YI and ZI must be used to add multiple cubicles; N=n is not valid. The SM map is updated to reflect the new assignment of cubicles. When PT=R is specified, the cubicle is reserved to a different SM map. Cubicles reserved for system use, storage module maintenance, or customer engineering use are reserved at compile time. Cubicles assigned to the pool are not assigned to a family or subfamily.

OP=RB — Remove Cubicles

OP=RB removes an assigned empty cubicle from a subfamily (PK=F), the pool (PK=P), or the reserved area of the SM (PK=R). More than one cubicle (N=n) can be removed at a time. SSLABEL reads the SM map and selects the first empty cubicle assigned to the subfamily, pool, or reserved area to be removed. Specific cubicles (YI and ZI parameters) can be removed, but they must be empty. The SM map is updated to reflect that the cubicles are unassigned.

OP=AM — Add Cartridges

OP=AM adds a cartridge to either a specific subfamily (PT=F) or the pool (PT=P). SSLABEL selects the first empty cubicle assigned to the subfamily or pool as the new location for the cartridge. More than one cartridge (N=n) can be added at a time or a specific cartridge (CN=csn) from the pool can be added to a subfamily. The SM map, 7990 catalog, and cartridge label are updated to reflect the new assignment of the cartridge. Cartridges added to a subfamily may be added to a specific group using the GR=n parameter. If cartridges are added to the pool, the GR=n parameter is not valid. The small and large boundaries for file size are specified by the B=n parameter with the OP=AM option. Default is B=600, implying that the first 600 AUs of the cartridge are available for small files. Cartridges added to the pool are not assigned to a family or subfamily. A customer engineer cartridge can be stored into location (0,0) or (0,15) by using the CC=loc parameter.

OP=RM — Remove Cartridges

OP=RM either removes an empty cartridge from a subfamily (PK=F) and moves the cartridge to the pool (PT=P) or it moves any cartridge from the pool (PK=P) and moves it to the exit tray (PT=D). Any cartridge currently assigned to a subfamily cannot be removed unless it is empty; that is, all AUs must be unallocated. To remove a cartridge that is not empty, first use the OP=FC directive to free up the cartridge.

If the cartridge specified by the CN=csn parameter is lost (does not reside in its assigned cubicle) and if the cartridge is assigned to a family, the LT parameter should be specified. This allows the appropriate entries in the SM map and 7990 catalog to be deleted even though the cartridge is not available to have its label updated. If LT is not specified, an error message is issued and SSLABEL aborts. SSLABEL will not update the SM map and 7990 catalog with the LT parameter unless the lost cartridge flag is set in the 7990 catalog. SSLABEL will set the lost cartridge flag if SSLABEL is run to remove a cartridge that is not present.

If the cartridge specified by the CN=csn parameter is lost and it is in the pool, then the LT parameter should not be used. Rather, SSDEBUG must be run to remove the SM map entry and SSLABEL must be run to reassign the cartridge to the pool.

A specified number (N=n) of cartridges may be removed from a group of a subfamily by using the GR=n parameter.

OP=RC — Restore Lost Cartridges

OP=RC restores to its proper cubicle a cartridge that was inadvertently removed from an SM. If restoration is successful, the lost flag in the 7990 catalog is cleared. If data recorded on the cartridge label does not agree with the information in the 7990 catalog and the SM map entry for the cubicle to which the cartridge is to be restored, the cartridge label information is reported and the cartridge is put into the exit tray for use in further processing of the cartridge.

NOTE

Not more than one lost cartridge can be restored at a time using the OP=RC directive. Also, since customer engineer cartridges do not exist in the SM map, OP=RC will not restore a cartridge to a customer engineer cartridge cubicle.

OP=FX — Fix Cartridge Labels

OP=FX writes a scratch label on a cartridge identified by the CN=csn parameter and adds the cartridge to the pool. This directive is intended for use when a cartridge label has been destroyed, but the cartridge itself is not physically damaged and can be reused. It can also be used when a cartridge with a family label is to be assigned to a different subfamily using the OP=AM directive, but it is not feasible to first remove the cartridge normally using the OP=RM directive. For example, if SSLABEL is run to add a cartridge to a subfamily, and a system failure occurs before the 7990 catalog and SM map are updated but after the cartridge is relabeled, then the cartridge label does not match the corresponding entries in the 7990 catalog and SM map. Hence, OP=RM cannot be used to remove the cartridge from the subfamily, but OP=FX can be used to rewrite the cartridge label and then OP=AM can be used to add the cartridge to a subfamily. However, if a family label is to be overwritten, the FM=familyname and SB=subfamily parameters must identify the family and subfamily to which the cartridge was assigned. The SM map and 7990 catalog are updated to reflect the new cartridge label.

OP=RM cannot be used with the LT parameter to remove pool cartridges. SSLABEL sets the SM map error flag because no 7990 catalog entries exist for pool cartridges. Use SSDEBUG to remove pool cartridges.

OP=IB — Control Cartridge Allocation

OP=IB sets or clears the inhibit allocation flag in the 7990 catalog entry for the specific cartridge (CN=csn). If the flag is set (ON), SSEXEC does not allocate new 7990 files to this cartridge. If the flag is cleared (OF), allocation of files to this cartridge is enabled.

OP=FC — Free Cartridge Files

OP=FC sets or clears the free cartridge flag in the 7990 catalog entry for the specified cartridge (CN=csn). If the flag is set (ON), SSEXEC will inhibit further allocation. Files can be removed by running the SSVAL utility. Refer to the free cartridge discussion under SSVAL later in this section. If the SM map entry is missing, SSDEBUG must be used to set the free cartridge flag. Refer to the discussion under SSDEBUG later in this section.

Parameters for SSLABEL Directives

The descriptions of the parameters to the SSLABEL directives follow. Not all parameters are valid with all directives, as indicated.

p_i	Description						
B=n	Number of AUs on this cartridge to be used to store small files. The remaining AUs (1931-n) are reserved for large files. The distinction between small and large files is controlled by an assembly parameter. B=n is valid only with OP=AM.						
B omitted	Same as B=600.						
CC=loc	Location into which the cartridge is to be stored.						
	<table border="1"> <thead> <tr> <th>loc</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Store cartridge into location (0,0).</td> </tr> <tr> <td>B</td> <td>Store cartridge into location (0,15).</td> </tr> </tbody> </table>	loc	Description	A	Store cartridge into location (0,0).	B	Store cartridge into location (0,15).
loc	Description						
A	Store cartridge into location (0,0).						
B	Store cartridge into location (0,15).						
	Any value of CC other than A or B is not accepted.						
CM=A-	Cartridge manufacturer code is A-, indicating IBM.						
CM	Same as CM=A-.						
CM omitted	Same as CM=A-.						
CN=csn	Cartridge serial number of the cartridge to be added, removed, or repaired; not valid if PK=pkloc is specified. If CN=csn is specified, n must be 1 if N=n is specified.						
CN	Cartridge serial number of the cartridge is not specified.						
CN omitted	Same as CN.						
FM=familyname	Family to which SSLABEL adds or from which it removes a cartridge or SM. With OP=FX, this parameter specifies the family to which the cartridge was assigned.						
FM	Same as FM=system default family.						
FM omitted	Same as FM=system default family.						
GR=n	Group to which SSLABEL adds or from which it removes a cartridge; $1 \leq n \leq 20$. With OP=AM, this parameter is ignored if PT=P is specified. GR=n is valid only with OP=AM or OP=RM.						
GR	Not permitted.						
GR omitted	Default groups are chosen sequentially. SSVAl runs most efficiently if SSLABEL chooses the default groups.						

P_i	Description
LT	SM map and 7990 catalog entries are to be updated, even though the cartridge is lost and its label cannot be updated; valid only with OP=RM.
LT omitted	If the cartridge is lost and OP=RM is specified, an error message is issued and SSLABEL aborts.
N=n	Number of cartridges or cubicles to be added, removed, or repaired; $1 \leq n \leq 100$; not valid if PT=R is specified. If CN=csn is specified, n must be 1.
N	Same as N=1.
N omitted	Same as N=1.
OF	The inhibit allocation flag or free cartridge flag in the 7990 catalog is to be cleared; valid only with OP=IB or OP=FC.
ON	The inhibit allocation flag or free cartridge flag in the 7990 catalog is to be set; valid only with OP=IB or OP=FC.
PK=pkloc	Location from which the cartridge or cubicle is to be picked; not valid if CN=csn is specified.
	pkloc Description
D	Cartridge is to be picked from the input tray. PK=D is valid only with OP=AM, OP=RC, or OP=FX.
F	Cartridge or cubicle is to be picked from the specified family (FM=familyname) and subfamily (SB=subfamily). PK=F is valid only with OP=RM or OP=RB.
P	Cartridge or cubicle is to be picked from the pool. PK=P is valid only with OP=AM, OP=RM, or OP=RB. PK=P is not valid if PT=P is specified.
R	Cubicle is to be picked from the reserved area of the SM. PK=R is valid only with OP=RB.
PK	Same as PK=P.
PK omitted	Same as PK=P.

P_i	Description										
PT=ptloc	Location into which the cartridge or cubicle is to be put.										
	<table border="1"> <thead> <tr> <th>ptloc</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Cartridge is to be put into the exit tray. PT=D is valid only with OP=RM.</td> </tr> <tr> <td>F</td> <td>Cartridge or cubicle is to be put into the specified family (FM=familyname) and subfamily (SB=subfamily). PT=F is valid only with OP=AM or OP=AB.</td> </tr> <tr> <td>P</td> <td>Cartridge or cubicle is to be put into the pool. PT=P is valid only with OP=AM, OP=RM, or OP=AB. PT=P is not valid if PK=P is specified.</td> </tr> <tr> <td>R</td> <td>Cubicle is to be put into the reserved area of the SM. PT=R is valid only with OP=AB.</td> </tr> </tbody> </table>	ptloc	Description	D	Cartridge is to be put into the exit tray. PT=D is valid only with OP=RM.	F	Cartridge or cubicle is to be put into the specified family (FM=familyname) and subfamily (SB=subfamily). PT=F is valid only with OP=AM or OP=AB.	P	Cartridge or cubicle is to be put into the pool. PT=P is valid only with OP=AM, OP=RM, or OP=AB. PT=P is not valid if PK=P is specified.	R	Cubicle is to be put into the reserved area of the SM. PT=R is valid only with OP=AB.
ptloc	Description										
D	Cartridge is to be put into the exit tray. PT=D is valid only with OP=RM.										
F	Cartridge or cubicle is to be put into the specified family (FM=familyname) and subfamily (SB=subfamily). PT=F is valid only with OP=AM or OP=AB.										
P	Cartridge or cubicle is to be put into the pool. PT=P is valid only with OP=AM, OP=RM, or OP=AB. PT=P is not valid if PK=P is specified.										
R	Cubicle is to be put into the reserved area of the SM. PT=R is valid only with OP=AB.										
PT	Same as PT=P.										
PT omitted	Same as PT=P.										
SB=subfamily	Subfamily to which SSLABEL adds or from which it removes a cartridge or SM; $0 \leq \text{sub} \leq 7$. With OP=FX, this parameter specifies the subfamily to which the cartridge was assigned.										
SB	Same as SB=0.										
SB omitted	Same as SB=0.										
SM=id	SM identifier of the SM to be used by SSLABEL; id is a letter from A to H.										
SM	Same as SM=A.										
SM omitted	Same as SM=A.										
YI=y ₁	Row of the SM to be added or removed; $0 \leq y_1 \leq 21$; valid only with OP=AB or OP=RB.										
ZI=z ₁	Column of the SM to be added or removed; $0 \leq z_1 \leq 15$; valid only with OP=AB or OP=RB.										
YI=y ₁ ,ZI=z ₁	y and z coordinates of the cubicle to be added or removed; $0 \leq y_1 \leq 21$, $0 \leq z_1 \leq 5$; valid only with OP=AB or OP=RB. The following coordinates (y,z) are reserved for customer engineering and/or diagnostic programming purposes: (0,0), (0,6), (0,15), (11,0), (11,1), (11,15), (15,0), (15,1), (21,6), and (21,15). The following coordinates are reserved for system use: (0,1), (0,14), (21,0), and (21,14).										

P_i	Description
YI=y ₁ ,ZI=z ₁ , YF=y ₂ ,ZF=z ₂	Rectangle of cubicles to be added or removed; cubicles with y coordinates between y ₁ and y ₂ and z coordinates between z ₁ and z ₂ are included; valid only with OP=AB or OP=RB. At most, 100 cubicles can be included in the rectangle. YF and ZF must both be specified, if either is specified. YF and ZF cannot be specified unless both YI and ZI are specified. If a reserved cubicle is included in the rectangle, the directive will be accepted even though the reserved cubicle cannot be added or removed. z ₁ =6 and z ₂ =6 are not valid.
YI and ZI omitted	With OP=AB, the next available cubicle closest to the top (for assignment to a family) or the bottom (for assignment to the pool) is to be selected. With OP=RB, the first empty assigned cubicle is to be selected.

SSLABEL Update Sequence

Each directive to SSLABEL updates SM maps, 7990 catalogs, and cartridge labels, whichever are appropriate, to reflect the changes in cartridge, cubicle, or SM assignment. Because the 7990 catalog is a disk-resident permanent file, it will be backed up on a dump tape whenever PFDUMP dumps the master device for its particular subfamily. Thus, it is not necessary for the analyst to back up the 7990 catalogs immediately after an SSLABEL run. However, the backup and recovery of SM maps do require special operational procedures, which should be performed immediately after an SSLABEL run (refer to the definition of SM map earlier in this section).

When SSLABEL is run to change the assignment of a cartridge, the update sequence consists of a series of steps to delete the old assignment information from the MSE system files, relabel the cartridge, and add the new assignment information to the MSE system files. If an interruption such as a system failure, SSLABEL abort, or SSEXEC abort prevents SSLABEL from completing the update sequence, the location of the affected cartridge and the status of the SM maps and 7990 catalogs depend on the point of interruption, as follows:

- If the cartridge label, 7990 catalog, and SM map do not all match, then the cartridge is put into the exit tray. The OP=RC directive cannot be used to restore the cartridge because of the inconsistency. However, OP=FX can be used to overwrite the cartridge label and add the cartridge to the pool, if the FM and SB parameters specify the family and the subfamily on the cartridge label.
- If the cartridge label, 7990 catalog, and SM map do match, then the cartridge may be returned to its original location, the new location, or the exit tray, depending on the exact point of interruption. If the cartridge is in the input tray, OP=RC can be used to restore the cartridge to the location indicated on the cartridge label.

Cartridge Restoration and Reuse, later in this section, describes the procedure for restoring cartridges found in the exit tray.

Restrictions to SSLABEL

The following restrictions apply to the SSLABEL utility.

- SSEXEC must be running when SSLABEL is run.
- Only one copy of SSLABEL can be run at a time.
- SSLABEL, SSVAL, and SSDEBUG cannot be run at the same time.

SSLABEL Example

Figure 11-1 shows the format of an SSLABEL report. The cartridge label information is included on the report because of a mismatch with the SM map. An error message is issued for the first directive on the command (in this case, the only directive).

```
SSLABEL REPORT FILE
SSLABEL, Z. /OP=RC, PK=D, SM=B.

    1  OP=RC, PK=D, SM=B.

    1  OP=RC, PK=D, SM=B.

    CM=A-
    CSN = 66157234
    FAMILY = SYSTST
    SUBFAMILY = 0
    SM = B
    Y = 2
    Z = 6
*** ERROR    8  DIRECTIVE 1
                UNEXPECTED SM, Y, Z, FAMILY OR SUBFAM.***
```

Figure 11-1. SSLABEL Report File

SSMOVE

SSMOVE manages disk and 7990 residence. That is, SSMOVE determines which files should be left on disk, which files should be released from disk and moved to 7990, and which files should be resident both on disk and on 7990.

The selection process includes algorithms that weigh certain file characteristics, as follows:

- Files are selected for destaging to 7990 based on file length, time since the last update, and the preferred residence specified by the user.
- Files are selected for release from disk based on time since the last access and the backup requirement specified by the user.

SSMOVE reads the PFC entries for a particular family and calculates release and destage values (refer to Selection Algorithms later in this section) for each file to determine its residence. If a file has both disk and 7990 images but is to reside only on 7990, SSMOVE releases the disk space for the file. For files that do not have a 7990 image, SSMOVE creates entries on the SSMOVE/SSEXEC communication file, MVOCOM, to identify the files to be destaged and to specify whether or not the file's disk space is to be released upon completion of the destage. SSEXEC then processes each destage and destage/release request on MVOCOM.

SSMOVE generates an output report that includes a list of input directives to SSMOVE, site-defined values used in the destage/release decision-making process, files selected for processing, and a summary of the number of files expected to reside on each device after destage/release processing.

The format of the SSMOVE command is:

SSMOVE, p₁, p₂, . . . , p_n.

p₁	Description
DN=device	Device number of the only disk from which files are to be destaged/released.
DN	SSMOVE destages and releases files from all devices in a specified family.
DN omitted	Same as DN.
FM=familyname	Family to be used by SSMOVE.
FM	Same as FM=system default family.
FM omitted	Same as FM=system default family.
I=filename	File containing the directives to SSMOVE.
I	Same as I=INPUT.
I=0	No input directives file exists. SSMOVE uses the default parameters.
I omitted	Same as I=INPUT.
L=filename	File on which listable output is to be written.
L	Same as L=OUTPUT.
L=0	No output file is to be generated.
L omitted	Same as L=OUTPUT.
LB=n	Large file boundary is used when sorting files for destaging. All files smaller than n PRUs are small files.
LB	Same as LB=365.
LB omitted	Same as LB=365.
LO=F	All files selected for staging, destaging, or releasing are to be included in the report file.
LO=P	Only files actually processed are listed in the report file. Refer to the PX parameter.
LO	Individual files are not to be listed in the report file.
LO omitted	Same as LO.

P1	Description														
NW	No wait. SSMOVE will not wait for completion of destage and release processing by SSEXEC.														
NW omitted	SSMOVE will wait for completion of destage and release processing by SSEXEC.														
PX=xxx	xxx is a character string that specifies which types of SSMOVE processing should not be done.														
	<table border="1"> <thead> <tr> <th>xxx</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Exclude file archiving or releasing disk space.</td> </tr> <tr> <td>B</td> <td>Exclude creating a backup copy by destaging a file from disk to the 7990.</td> </tr> <tr> <td>D</td> <td>Exclude direct access files.</td> </tr> <tr> <td>F</td> <td>Exclude freeing a file from the 7990 by clearing its asa value from the file's PFC entry.</td> </tr> <tr> <td>I</td> <td>Exclude indirect access files.</td> </tr> <tr> <td>S</td> <td>Exclude staging a file from the 7990 to disk.</td> </tr> </tbody> </table> <p>For example, the PX=ABFS parameter could be used to obtain a report of the results of an SSMOVE run using a given set of parameters without actually performing any of the selected actions.</p>	xxx	Description	A	Exclude file archiving or releasing disk space.	B	Exclude creating a backup copy by destaging a file from disk to the 7990.	D	Exclude direct access files.	F	Exclude freeing a file from the 7990 by clearing its asa value from the file's PFC entry.	I	Exclude indirect access files.	S	Exclude staging a file from the 7990 to disk.
xxx	Description														
A	Exclude file archiving or releasing disk space.														
B	Exclude creating a backup copy by destaging a file from disk to the 7990.														
D	Exclude direct access files.														
F	Exclude freeing a file from the 7990 by clearing its asa value from the file's PFC entry.														
I	Exclude indirect access files.														
S	Exclude staging a file from the 7990 to disk.														
PX	All selected SSMOVE processing is to be done.														
PX omitted	Same as PX.														
UI=userindex	Destage and release processing is restricted to files having user index userindex.														
UI	All user indexes are processed.														
UI=0	All user indexes are processed.														
UI omitted	All user indexes are processed.														

SSMOVE Directives

The directives to SSMOVE are specified on a separate file. Each directive must be specified on a separate line. Three types of input directives can be included in an SSMOVE directive file. They are: comment directives, specific file directives, and value specifier directives.

Comment directives have an asterisk (*) in the first column followed by any message or comment the user wishes to enter.

The specific file and value specifier directives are described in the following paragraphs:

Specific File (SF) Directives

SSMOVE uses SF directives to select a given file or set of files on which to perform the specified processing action. The format of the SF directive is:

```
SF, FN=filename, UI=userindex, PO=process.
```

Parameter	Description
FN=filename	The 1 to 7 characters that specify a permanent file or set of permanent files. Asterisks in the file name are wild card characters.
UI=userindex	User index of the file specified by the FN=filename parameter.
PO=process	Processing action to be performed on the file specified by the FN=filename parameter.
	process Description
	A Archive the specified file.
	B Create a backup copy by destaging the specified file from disk to the 7990.
	F Free the specified file from the 7990 by clearing its asa value in the file's PFC entry.
	S Stage the specified file from the 7990 to disk.

For example, the following SF directive selects all permanent files with a maximum of two characters in the file name beginning with Z from user index 002622₈ to be staged to disk.

```
SF, FN=Z*, UI=2622, PO=S.
```

The SF directive can be used without specifying the FN parameter to establish default values for the UI and/or PO parameters. These default values will be used for subsequent SF, FN=filename directives that are missing UI and/or PO parameters. Default values for UI and/or PO parameters can be redefined by using another SF directive without specifying the FN parameter.

Value Specifier Directives

SSMOVE uses the value specifier directives to redefine the values of the weight factors or thresholds (installation parameters) used in the algorithms that select files to be destaged or released. The site analyst uses these value specifiers to increase or decrease the importance of certain file characteristics used to determine which files are to be destaged and/or released. For example, specifying a large MN parameter prohibits SSMOVE from selecting small files for destage/release processing.

The general format of a value specifier directive statement is:

```
directive,qualifier(s),parameter(s).
```

Variable	Description														
directive	2-character value specifier directive keyword.														
	<table border="1"> <thead> <tr> <th>directive</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>BR</td> <td>Keyword for backup requirement directive.</td> </tr> <tr> <td>FR</td> <td>Keyword for file requirement directive.</td> </tr> <tr> <td>PR</td> <td>Keyword for preferred residence directive.</td> </tr> <tr> <td>SM</td> <td>Keyword for site management directive.</td> </tr> <tr> <td>WA</td> <td>Keyword for weight adder directive.</td> </tr> <tr> <td>WM</td> <td>Keyword for weight multiplier directive.</td> </tr> </tbody> </table>	directive	Description	BR	Keyword for backup requirement directive.	FR	Keyword for file requirement directive.	PR	Keyword for preferred residence directive.	SM	Keyword for site management directive.	WA	Keyword for weight adder directive.	WM	Keyword for weight multiplier directive.
directive	Description														
BR	Keyword for backup requirement directive.														
FR	Keyword for file requirement directive.														
PR	Keyword for preferred residence directive.														
SM	Keyword for site management directive.														
WA	Keyword for weight adder directive.														
WM	Keyword for weight multiplier directive.														
qualifier(s)	One or more qualifier keywords that specify the file type (direct/indirect access) and decision type (destage/release) to be used in the selection process. If all qualifier keywords are omitted, all four combinations of file types and decision types are used in the selection process.														
	<table border="1"> <thead> <tr> <th>qualifier(s)</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>IA</td> <td>Keyword for indirect access files.</td> </tr> <tr> <td>DA</td> <td>Keyword for direct access files.</td> </tr> <tr> <td>DS</td> <td>Keyword for destage decisions.</td> </tr> <tr> <td>RL</td> <td>Keyword for release decisions.</td> </tr> </tbody> </table>	qualifier(s)	Description	IA	Keyword for indirect access files.	DA	Keyword for direct access files.	DS	Keyword for destage decisions.	RL	Keyword for release decisions.				
qualifier(s)	Description														
IA	Keyword for indirect access files.														
DA	Keyword for direct access files.														
DS	Keyword for destage decisions.														
RL	Keyword for release decisions.														

Variable	Description
parameter(s)	Parameter keyword and a value (positive integer) for the specified directive keyword. Multiple parameter=value combinations can be entered in the same value specifier directive statement.
parameter(s)	Description
With BR	
Directive:	
N=wf	Weight factor to be used as the backup requirement value for decisions involving files with a BR=N attribute (no backup). ²
Y=wf	Weight factor to be used as the backup requirement value for decisions involving files with a BR=Y attribute (tape backup). ²
MD=wf	Weight factor to be used as the backup requirement value for decisions involving files with a BR=MD attribute (media dependent, tape, or 7990 backup). ²
With FR	
Directive:	
MN=min	Minimum file size in PRUs.
MX=max	Maximum file size in PRUs.
DD=dd	Number of days since the file was last accessed or updated.
TH=th	Destage or release threshold. If a file's calculated value is less than the specified threshold, the file is not a candidate for destage or release processing. TH=0 allows all files to be candidates for destage and release processing unless they are excluded by FR directive parameters or by other run-time parameter values.

². The file owner specifies the backup requirement attribute using the BR parameter on the DEFINE or CHANGE command (refer to the NOS Version 2 Reference Set, Volume 3).

Variable	Description
parameter(s)	<i>(Continued)</i>
parameter(s)	Description
With PR Directive:	
L=wf	Weight factor to be used as the preferred residence value for decisions involving files with a PR=L attribute (locked to disk preference). ³
D=wf	Weight factor to be used as the preferred residence value for decisions involving files with a PR=D attribute (disk preference). ³
M=wf	Weight factor to be used as the preferred residence value for decisions involving files with a PR=M attribute (7990 preference). ³
N=wf	Weight factor to be used as the preferred residence value for decisions involving files with a PR=N attribute (no preference). ³
With SM Directive:	
MG=mg	Master goal specifying the percent of master device disk space not to be exceeded.
SG=sg	Secondary goal specifying the percent of secondary device disk space not to be exceeded.
P1	Available for site use.
P2	Available for site use.

³. The file owner specifies the preferred residence attribute using the PR parameter on the DEFINE or CHANGE command (refer to the NOS Version 2 Reference Set, Volume 3).

Variable	Description
parameter(s)	<i>(Continued)</i>
parameter(s)	Description
With WA Directive:	
AG=wa	Weight adder to be added to the weighted file age.
LN=wa	Weight adder to be added to the weighted file length.
AC=wa	Weight adder to be added to the weighted file access count.
DV=sf	Scaling factor given to the divisor when calculating destage and release values.
With WM Directive:	
AG=wf	Weight factor given to the file age (days since the file was last accessed or updated).
LN=wf	Weight factor given to the file length.
AC=wf	Weight factor given to the file access count.

For example, the following value specifier directive specifies that the minimum size for an indirect access file to be destaged is three disk PRUs.

```
FR, IA, DS, MN=3.
```

The following value specifier directive requires that all files to be destaged or released must be less than 98765 PRUs long.

```
FR, MX=98765.
```

If a directive error is detected on the directive input file, SSMOVE ignores the incorrect directive and continues to run in report mode. A message is issued to the job dayfile and no processing is done.

Excluding Destage/Release Processing

SSMOVE uses FR directives to select candidates for destage and release processing. After selecting the candidates and making the final processing decisions, SSMOVE checks the PX parameter to determine if any processing should be excluded.

If FR,RL,MN=9999999 is specified, all files less than 9999999 PRUs will be given a negative release value. Any file less than 9999999 PRUs will not be released unless the file has been specifically selected for releasing with an SF directive.

If PX=A is specified, SSMOVE selects files for processing without considering that release processing is prohibited. When the final processing decisions have been made, any files selected for releasing will not be processed. That is, files selected for releasing only will not be processed; and files selected for both destaging and releasing will not be processed.

If PX=B is specified, destage processing is excluded just as release processing was in the previous description.

Destaging cannot be excluded with an FR,DS,MN=9999999 directive as releasing can. Even though all files less than 9999999 PRUs will be given a negative destage value and not selected for destaging, they will be made candidates for destaging in case any of them are selected for releasing. When the final processing decisions are made, any file that is a candidate for destaging and has been selected for releasing will be destaged and released. However, any destaging candidate not selected for releasing will not be destaged.

Selection Algorithms

SSMOVE determines which files to destage and/or release according to the following algorithms. Files that reside only on 7990 are not considered because they have been destaged and released previously. For all other files, SSMOVE checks the file length and excludes from further consideration any file whose length in PRUs is less than the minimum length threshold (refer back to the MN parameter) or greater than the maximum length threshold (refer back to the MX parameter).

If the current image of the file resides on both disk and 7990, SSMOVE uses the release algorithm to determine whether or not to release disk space; the destage algorithm is not used. If the current file resides on disk only, SSMOVE uses the destage algorithm to determine whether or not to destage the file. If the file is to be destaged to 7990, SSMOVE also determines via the release algorithm whether or not to release the file's disk space.

Destage and Release Algorithms

An analyst uses the FR directive to establish file requirements on files that are to be eligible for destage and/or release processing. These requirements include file length (minimum and maximum) and file age. A file's age is the number of days since the file was last updated (for destaging) or accessed (for releasing). Separate requirements can be established for the various combinations of file types (direct/indirect) and decision types (destage/release).

Files that meet the eligibility requirements are assigned separate destage and release values. SSMOVE calculates these values for each eligible file according to the following equation, which considers various file attributes from the file's PFC entry (length, age, access count, preferred residence, and backup requirements). These file attributes are weighted according to site-defined values (refer to Value Specifier Directives earlier in this section).

$$\text{value} = \frac{(a_1 + a_2 * \text{length}) * (b_1 + b_2 * \text{age}) * (\text{res}) * (\text{bkup})}{(c_1 + c_2 * \text{access}) * \text{sf}}$$

Variable	Description
a_1	Site-defined length weight adder specified by the WA directive LN parameter.
a_2	Site-defined length weight factor specified by the WM directive LN parameter.
length	Length of the file in PRUs.
b_1	Site-defined age weight adder specified by the WA directive AG parameter.
b_2	Site-defined age weight factor specified by the WM directive AG parameter.
age	Number of days since the file was last updated (for destage) or accessed (for release).
res	Site-defined preferred residence weight factor specified by the PR directive. The PR directive parameter depends on which preferred residence attribute the file owner specified.
bkup	Site-defined backup requirement weight factor specified by the BR directive. The BR directive parameter depends on which backup requirement attribute the file owner specified.
c_1	Site-defined access count weight adder specified by the WA directive AC parameter.
c_2	Site-defined access count weight factor specified by the WM directive AC parameter.
access	Number of times the file was accessed.
sf	Site-defined scaling factor specified by the WA directive DV parameter.

If a file's destage/release value exceeds a site-defined threshold (FR directive TH parameter), the file is selected for destage/release processing. Conversely, if the file's destage/release value does not exceed the threshold, the file is not considered for destage/release processing.

Files eligible for releasing are sorted according to their release values. The files on each disk with the greatest release values are selected to be released. When enough files have been selected and released, the site-defined disk space availability goals will be achieved.

Decision Algorithm Hierarchy

Several different functions of SSMOVE can lead to apparent contradictory processing decisions. The decisions are made using the following rules, which are listed in order of precedence.

1. SSMOVE does not consider any files for processing if they are excluded by the PX=I, PX=D, or UI parameters to SSMOVE or if the file has a PR=L attribute.
2. Files identified for specific processing using the SF directive are selected for the specified processing regardless of the file requirements controlled by the FR directive.
3. If DN=0, all files that meet the destage requirements are selected to be destaged. Otherwise, only files on the specified device are destaged.
4. No file with a BR=Y attribute will be selected to be released (including files specified for processing with an SF directive) if the file data has not been included in a full or incremental dump using PFDUMP.
5. If DN=0, all files that meet the release requirements are selected to be released, if needed, to meet the disk space availability goals. Otherwise, only files on the specified device are released.
6. If a file is selected for a specific processing action (such as destage) and the corresponding action is precluded by a PX parameter (in this case, PX=B), then the selected action is not performed. In all cases, however, the SSMOVE device status report is based on which files were selected for processing rather than whether the PX parameter actually allowed the selected processing to occur.

Disk Space/Dump Tape Management

As more disk-resident files are created and more 7990-resident files are staged to disk, it will be necessary to monitor the availability of disk space. It is recommended that SSMOVE be used as a periodic disk space management procedure to avoid frequent disk-full conditions (refer to Disk Space Management later in this section).

SSMOVE can also be used to reduce the amount of data written on dump tapes and thereby avoid maintaining large numbers of dump tapes. This is accomplished by destaging to 7990 and/or releasing disk space of files that need not be on the dump tape (refer to Dump Tape Management later in this section).

An analyst can specify what percent of disk space is to contain indirect access files and what percent is to contain direct access files. For direct access files, the analyst can further specify the percent of master device and/or secondary device disk space. These percent values are applied to all devices in the family on which SSMOVE is run or a specific device can be managed by using the DN parameter to SSMOVE. The specified percentages of disk space are considered to be goals in that other site-defined constraints and file-busy conditions may make it impossible to achieve these goals.

Restrictions to SSMOVE

The following restrictions apply to the SSMOVE utility.

- Only one copy of SSMOVE can be run at a time when it is in the process of destaging files. A second SSMOVE aborts if the first one has not completed.
- SSMOVE and PFDUMP should not be run at the same time.
- SSMOVE cannot be run when SSVAL is running.

Output Example

The SSMOVE report consists of the following five sections:

1. A list of directives (figure 11-2) from the input directive file plus any diagnostic messages.
2. A list of run-time parameter values (figure 11-3) used in the release and destage decision-making process. These values consist of the site-defined, assembly-time values as modified by the value specifier directives.
3. An optional list of the files (figure 11-4) selected for processing. This list shows the file name, file type, user index, length, date, access count, processing action, and destage and release values.
4. A device status report (figure 11-5) that includes a summary of the number of files and total file length expected to reside on each device as a result of release processing. Also, a subfamily report that includes a summary (for each subfamily) of the number of files and file length for files to be destaged, files that reside only on 7990, and files that reside on 7990 and disk.
5. A destage abandonment report (figure 11-6) that is not produced if the NW parameter is specified. This report provides a list of the number of files not destaged and an optional list of the abandoned files. It also provides an updated device status report reflecting only the files that were destaged.

```

yy/mm/dd. hh.mm.ss. PAGE 1
SSMOVE REPORT.
SSMOVE, I=BIGO, LO+F, FM=MTST.

1 FR, MN=0, DD=0.
2 SM, MG=0, SG=0.

```

Figure 11-2. SSMOVE Report: List of Directives

```

yy/mm/dd. hh.mm.ss. PAGE 2
SSMOVE REPORT.
RUN-TIME PARAMETER VALUES

* D E S T A G E *      * R E L E A S E *
DIRECT  INDIRECT      DIRECT  INDIRECT

FILE REQUIREMENTS

FR MN      0      0      0      0
MX 9999999 9999999 9999999 9999999
DD      0      0      0      0
TH      1      1      1      1

WEIGHT MULTIPLIER

WM AG      1      1      100     100
LN      1      1      1      1
AC      1      1      1      1

WEIGHT ADDER

WA AG      1      1      1      1
LN      1      1     1000     1000
AC      1      1      1      1
DV      1      1     1000     1000

PREFERRED RESIDENCE

PR L      0      0      0      0
D      1      1      1      1
M      1      1     100     100
N      1      1     10      10

BACKUP REQUIREMENT

BR N      1      1      1      1
Y      1      1      1      1
MD     5      5      1      1

SITE MANAGEMENT

SM MG      0      0      0      0
SG      0      0      0      0
P1      0      0      0      0
P2      0      0      0      0

```

Figure 11-3. SSMOVE Report: Run-time Parameter Values

SSMOVE REPORT.										
yy/mm/dd. hh.mm.ss. PAGE 4										
NAME	TYPE	UI	LENGTH	DATE	ACC-CT	ACTION(* = NOT DONE PER *PX* OPTION)	DES-VAL	REL-VAL		
DIR301	DIR.	30	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR305	DIR.	30	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR302	DIR.	30	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR303	DIR.	30	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR304	DIR.	30	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
IND3020	IND.	30	619	84.09.11	0	DESTAGE AND RELEASE.	3100	16190		
IND3019	IND.	30	618	84.09.11	0	DESTAGE AND RELEASE.	3095	16180		
IND3018	IND.	30	617	84.09.11	0	DESTAGE AND RELEASE.	3090	16170		
IND3017	IND.	30	616	84.09.11	0	DESTAGE AND RELEASE.	3085	16160		
IND3016	IND.	30	615	84.09.11	0	DESTAGE AND RELEASE.	3080	16150		
IND3015	IND.	30	614	84.09.11	0	DESTAGE AND RELEASE.	3075	16140		
IND3014	IND.	30	613	84.09.11	0	DESTAGE AND RELEASE.	3070	16130		
IND3013	IND.	30	612	84.09.11	0	DESTAGE AND RELEASE.	3065	16120		
IND3012	IND.	30	611	84.09.11	0	DESTAGE AND RELEASE.	3060	16110		
IND3011	IND.	30	610	84.09.11	0	DESTAGE AND RELEASE.	3055	16100		
IND3010	IND.	30	609	84.09.11	0	DESTAGE AND RELEASE.	3050	16090		
IND309	IND.	30	608	84.09.11	0	DESTAGE AND RELEASE.	3045	16080		
IND308	IND.	30	607	84.09.11	0	DESTAGE AND RELEASE.	3040	16070		
IND307	IND.	30	606	84.09.11	0	DESTAGE AND RELEASE.	3035	16060		
IND306	IND.	30	605	84.09.11	0	DESTAGE AND RELEASE.	3030	16050		
IND305	IND.	30	604	84.09.11	0	DESTAGE AND RELEASE.	3025	16040		
IND304	IND.	30	603	84.09.11	0	DESTAGE AND RELEASE.	3020	16030		
IND303	IND.	30	602	84.09.11	0	DESTAGE AND RELEASE.	3015	16020		
IND302	IND.	30	601	84.09.11	0	DESTAGE AND RELEASE.	3010	16010		
IND301	IND.	30	600	84.09.11	0	DESTAGE AND RELEASE.	3005	16000		
DIR313	DIR.	31	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR314	DIR.	31	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR312	DIR.	31	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR315	DIR.	31	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
DIR311	DIR.	31	1200	84.09.11	0	DESTAGE AND RELEASE.	6005	22000		
IND3120	IND.	31	619	84.09.11	0	DESTAGE AND RELEASE.	3100	16190		
IND3119	IND.	31	618	84.09.11	0	DESTAGE AND RELEASE.	3095	16180		
IND3118	IND.	31	617	84.09.11	0	DESTAGE AND RELEASE.	3090	16170		
IND3117	IND.	31	616	84.09.11	0	DESTAGE AND RELEASE.	3085	16160		
IND3116	IND.	31	615	84.09.11	0	DESTAGE AND RELEASE.	3080	16150		
IND3115	IND.	31	614	84.09.11	0	DESTAGE AND RELEASE.	3075	16140		
IND3114	IND.	31	613	84.09.11	0	DESTAGE AND RELEASE.	3070	16130		
IND3113	IND.	31	612	84.09.11	0	DESTAGE AND RELEASE.	3065	16120		
IND3112	IND.	31	611	84.09.11	0	DESTAGE AND RELEASE.	3060	16110		
IND3111	IND.	31	610	84.09.11	0	DESTAGE AND RELEASE.	3055	16100		

Figure 11-4. SSMOVE Report: Optional List of Files

yy/mm/dd. hh.mm.ss. PAGE 6

SSMOVE REPORT.

		(BEFORE) DEVICE STATUS		(AFTER) PERCENTS					
EO	BN	DT-N	TYPE	FILES /	PRU	FILES /	PRU	EXP. GOAL	FLAG.
2	7	DK-1	IND.	62	36952	0	0	0	0
2	7	DK-1	DTR.	23	25424	9	1008	1	0 **

** - DEVICE SPACE GOAL NOT MET

SUBFAMILY REPORT

SUB FAMILY	FILES TO DESTAGE			FILES ONLY ON 7990			FILES ON 7990					
	DIRECT	INDIRECT	PRU	DIRECT	INDIRECT	PRU	DIRECT	INDIRECT	PRU			
0	5	6000	20	12190	38	6363	20	12190	38	6363	20	12190
1	5	6000	20	12190	37	6352	20	12190	37	6352	20	12190
2	0	0	2	1201	37	10357	2	1201	37	10357	2	1201
3	0	0	0	0	42	12362	5	3010	42	12362	5	3010
4	0	0	0	0	37	6357	5	3010	37	6357	5	3010
5	0	0	0	0	37	6357	5	3010	37	6357	5	3010
6	0	0	1	3050	19	27915	69	56298	19	27915	69	56298
7	0	0	0	0	0	0	0	0	0	0	0	0

Figure 11-5. SSMOVE Report: Device Status Report

yy/mm/dd. hh.mm.ss. PAGE 7

SSMOVE REPORT.

DESTAGE ABANDONMENT REPORT

FILENAME	UI	CODE
CODE	FILES	REASON
1	0	NO SPACE
2	0	NO STORAGE MODULE AVAILABLE
3	0	NO CARTRIDGE OR GROUP AVAILABLE
4	0	FILE ALREADY DESTAGED
5	0	FILE BUSY / PFM PROBLEM
6	0	CATALOG ACCESS ERROR
7	0	OVERFLOW NOT LEGAL
8	0	GROUP FULL
9	0	DISK READ ERROR
90	0	CARTRIDGE LOST
11	0	CLOSED DESTAGE

		(BEFORE) DEVICE STATUS		(AFTER) PERCENTS					
EO	BN	DT-N	TYPE	FILES /	PRU	FILES /	PRU	EXP. GOAL	FLAG.
2	7	DK-1	IND.	62	36952	0	0	0	0
2	7	DK-1	DTR.	23	25424	9	1008	1	0 **

** - DEVICE SPACE GOAL NOT MET

SUBFAMILY REPORT

SUB FAMILY	FILES NOT DESTAGED			FILES ONLY ON 7990			FILES ON 7990				
	DIRECT	INDIRECT	PRU	DIRECT	INDIRECT	PRU	DIRECT	INDIRECT	PRU		
0	0	0	0	38	6363	20	12190	38	6363	20	12190
1	0	0	0	37	6352	20	12190	37	6352	20	12190
2	0	0	0	37	10357	2	1201	37	10357	2	1201
3	0	0	0	42	12362	5	3010	42	12362	5	3010
4	0	0	0	37	6357	5	3010	37	6357	5	3010
5	0	0	0	37	6357	5	3010	37	6357	5	3010
6	0	0	0	19	27915	69	56298	19	27915	69	56298
7	0	0	0	0	0	0	0	0	0	0	0

Figure 11-6. SSMOVE Report: Destage Abandonment

SSVAL

SSVAL either performs release processing or reports on problems with the current MSE system files. When release processing, SSVAL makes 7990 space available that is presently allocated to files that are no longer needed, or flags files to be staged to disk in order to free up cartridges. When problem reporting, SSVAL reports on irregularities or discrepancies found in the current 7990 catalogs and PFC entries for the specified family and, optionally, in certain SM maps. The function to be performed is determined by whether or not the RF parameter is specified, as described next.

NOTE

SSVAL updates the SM map for the specified SM. (Refer to the definition of SM map earlier in this section.) It is recommended that the SM map be copied on tape or on another device or family immediately after every update of the SM map.

The format of the SSVAL command is:

SSVAL, P₁, P₂, . . . , P_n.

P _i	Description
AM	The SM map for the SM specified by the SM=id parameter is to be analyzed in addition to the 7990 catalogs; not valid if RF=filename or RF is specified.
AM omitted	SM maps are not to be analyzed.
FM=familyname	Family to be analyzed; not valid if the RF parameter is specified.
FM	Same as FM=system default family; not valid if the RF parameter is specified.
FM omitted	Same as FM=system default family if the RF parameter is not specified. The family on the release data file is used if the RF parameter is specified.
FX=n	Error threshold. If the total error count is greater than n, neither release processing nor problem fixing is performed.
FX	Same as FX=0.
FX omitted	Same as FX=0.
L=filename	File on which listable output is to be written.
L	Same as L=OUTPUT.
L=0	No output file is to be generated.
L omitted	Same as L=OUTPUT.

P_i	Description
RF=filename	File that contains the release data file.
RF	Same as RF=ZZZZRDF.
RF omitted	Current versions of the 7990 catalogs are to be analyzed.
RL	Release processing is to be performed; valid only if the RF parameter is specified.
RL omitted	No release processing is to be performed.
SB=subfamily	Subfamily to be processed. Up to eight subfamilies can be selected by the numbers 0 through 7. For example, SB=723 selects subfamilies 2, 3, and 7.
SB	Same as SB=01234567.
SB omitted	Same as SB=01234567.
SM=id	SM identifier of the SM to be used. Up to 8 SMs can be selected by the letters A through H. For example, SM=ACG selects SM A, C, and G.
SM	Same as SM=ABCDEFGH.
SM omitted	Same as SM=ABCDEFGH.
ST=n	Scattered file criterion. Files are indicated as scattered if they are contained on at least n more cartridges than the minimum number needed to contain them.
ST	Same as ST=1. That is, files are scattered if they are contained on more than the minimum number of cartridges needed to contain them.
ST omitted	Same as ST=1.

Release Processing

If RF=filename or RF is specified, SSVAl determines which 7990 files are no longer needed and issues a request to SSEXEC to purge these files so their 7990 space can be reused. The procedure is for SSVAl to analyze copies of the 7990 catalogs and PFC entries for the specified family that are contained on the release data file (RDF) specified by the RF parameter. Those 7990 files described in a 7990 catalog but not having a PFC entry (that is, orphans) can be purged. During the analysis (refer to both 7990 Catalog Analysis and PFC Analysis later in this section), SSVAl keeps track of the error conditions it discovers, and if the error count is less than or equal to the threshold specified by the FX parameter, release processing is performed if RL is specified. That is, trouble-free orphans are purged and their 7990 space is made available for reuse. The current 7990 catalog is updated to reflect that these files no longer exist. A validation report is issued that lists the errors encountered, the number of trouble-free orphans, and the amount of released 7990 space. If RL is not specified, no release processing is performed but the validation report is issued, which lists the errors encountered, the number of trouble-free orphans, and the amount of releasable 7990 space.

The RDF used for this analysis is a file produced during a previous PFDUMP run and it contains versions of the 7990 catalogs and PFC entries that were current at the time of the dump. The site analyst chooses which RDF to use depending on how long after a file was purged he/she wants to wait before releasing 7990 space. For example, the site analyst might run SSVAl every week for release processing purposes and use the RDF from the previous week's full dump. There are some restrictions as to which RDFs can be used (refer to 7990 Space Management later in this section).

Unlike the purging of disk-only files, where the file space on disk is immediately made available for reassignment, the MSE software will not immediately free up a permanent file's 7990 space following a purge of the associated disk copy of the permanent file (assuming both a 7990 and a disk permanent file image). The reason for not freeing up the subject 7990 space is that a subsequent PFLOAD, following a disk failure, could reload the PFC entry for the associated permanent file. Therefore, 7990 space for purged permanent files that were resident on the 7990 must not be made available for reuse until the possibility of such a recovery action cannot occur.

The site analyst determines which dump tapes can and cannot be used for reloading files and PFC entries following a disk failure. When a PFDUMP is taken, two files are normally produced: the actual dump file used with PFLOAD to perform a recovery, and the RDF file to be used with SSVAl to make 7990 space available after purging files.

Using an RDF file allows the corresponding dump tape to be used (later dump tapes are also allowable), but earlier dump tapes are not allowed. Because the MSE software cannot control which dump tapes actually are used in reloading, it is the site analyst's responsibility to observe the above restriction.

The concurrent execution of PFDUMP with OP=S not specified and SSVAl with RL specified causes SSVAl to run very slowly. PFDUMP acquires and keeps the PF utility interlock for its duration, and this gives the appearance that SSVAl is hung, when it is actually just waiting. Therefore, operations personnel should not attempt to run both of these utilities with the above options at the same time.

If the RF parameter is not specified, SSVAl will set the free file flag for all files that totally or partially reside on all cartridges on which the free cartridge flag is set.

Problem Reporting

If RF=filename or RF is not specified, SSVAl reports on problems with the current 7990 catalogs and PFC entries for the specified family. If AM is specified, problems with SM maps are also included in the report. SSVAl examines the MSE system files and PFC entries and searches for problem chains and fragments, problem asa values, and SM map/7990 catalog mismatches. The procedures SSVAl uses to detect and classify inconsistencies and discrepancies are described next, under Error Detection and Classification. SSVAl keeps track of the error conditions, if any, and if the error count is less than or equal to the threshold specified by the FX parameter, problem fixing is performed. That is, SSVAl sets flags in the appropriate entries of the SM map, 7990 catalog, and/or PFC entries to prevent propagation of errors due to the inconsistencies or discrepancies found and to permit error recovery by the SSDEBUG utility. A count of the errors is recorded in the job dayfile. A validation report is issued, which lists the errors encountered, the number of trouble-free orphans, and the amount of releasable 7990 space.

NOTE

For release processing or problem fixing, SSVAl predicts what it expects to happen (if no problems occur) prior to doing the releasing or fixing. If SSVAl encounters problems, the reports may not reflect the actual status of the cartridges. In this case, run SSVAl in report mode following the release processing or problem fixing to verify the actual status of the 7990 catalogs and cartridges.

Error Detection and Classification

SSVAl detects and classifies errors according to the following procedures. During the SM map analysis, SSVAl detects and classifies errors in the SM map. During the 7990 catalog analysis, SSVAl detects errors with chains of AUs; during the PFC analysis, SSVAl classifies these errors. Whenever an error is encountered, the total error count is increased by 1. The action taken for each type of error is discussed under Release Processing and Problem Fixing later in this section.

SM Map Analysis

If the AM option is specified, SSVAl attempts to locate problems with SM map entries by comparing the 7990 catalogs and the SM map. For each coordinate pair (y,z) in a 7990 catalog entry, SSVAl locates the corresponding SM map entry. A type 1 error exists if the code field in the SM map entry is not 5 (assigned to a subfamily), or if the family, subfamily, or csn fields in the SM map entry do not match those in the corresponding 7990 catalog entry. SSVAl also scans the SM map for all other entries assigned to the subfamily and reports as a type 2 error any of these entries that does not have a corresponding 7990 catalog entry.

7990 Catalog Analysis

The 7990 catalog analysis locates problems with chains of AUs (refer to the definition of cartridge earlier in this section) and identifies on each HOC whether any of the following problems exist or whether any part of the 7990 files resides on a cartridge that is lost or has excessive write parity errors.

SSVAL scans the 7990 catalog for HOC entries that are allocated and follows each chain until it terminates. Normal termination occurs with an EOC entry. Abnormal termination occurs when no EOC is found, an AU links to an unallocated AU, or an AU links to an AU previously found in the chain being followed (looping chain).

During the chain scans, the following types of chains can be encountered; they are linkage problems that are identified on the HOC entry.

Chain	Description
Intersecting	More than one chain links to the same AU.
Scattered file	The number of cartridges used for the file exceeds the value specified by the ST parameter.

SSVAL also locates any AUs that are allocated but were not on any chain being followed. Such AUs are linked together to form partial chains without an HOC. These partial chains are called fragments and the first AU in a fragment is designated as the start of fragment. Each fragment chain is followed until it terminates. The abnormal termination conditions previously listed can also occur with fragments. Intersections can occur, but a fragment chain that intersects the start of another fragment chain is not an intersection; rather, one is the tail end of the other.

PFC Analysis

The PFC analysis is performed to classify the errors encountered on chains during the 7990 catalog analysis. For each PFC entry with `asa≠0` (the file has a 7990 image and the `asa` value identifies the first AU in the chain containing the file), SSVAl classifies the following errors. Error type 3 exists if the `asa` value is invalid. Error type 4 exists for any of the following reasons:

- The AU specified by the `asa` value is not allocated or is not an HOC entry.
- The chain does not terminate normally.
- The chain intersects with another chain or fragment.
- More than one PFC entry points to the chain.
- The chain includes a cartridge for which the lost or excessive parity error flag is set.

SSVAl also classifies the following error conditions:

- Error type 5 exists if an orphan chain terminates abnormally or intersects with other chains or fragments. Trouble-free orphans (chains without a PFC entry pointing to them and without linkage problems) are not classified as errors.
- Error type 6 exists if an orphan is a fragment.
- Error type 7 exists if a chain or fragment points to an unallocated AU.

SSVAl generates informational report messages if either of the following conditions is true.

- The system error flag is set in the PFC.
- The read error flag is set in the PFC.

Release Processing and Problem Fixing

If the total error count calculated during the analyses described previously exceeds the value specified by the FX parameter, then neither release processing nor action to flag or fix the detected error conditions is performed. The validation report, however, is produced. Otherwise, action taken depends on the parameter specified and the type of errors found, as follows.

If SSVAl was run for release processing purposes (RF=filename or RF specified), SSVAl issues a request to SSEXEC to release trouble-free orphans if RL was specified. If RL was not specified, no release processing is performed. If SSVAl was run for problem reporting purposes (RF omitted and FX=n, where n is greater than the number of validation errors), the following action is taken.

- For error type 1, the linkage error flag is set in the 7990 catalog entry.
- For error type 2, the linkage error flag is set in the SM map entry.
- For error type 3, there are three alternatives:
 - If the file also has a disk image, the asa field in the PFC entry is cleared. Thus, the good disk image will not be released and the file is accessible even if the 7990 image cannot be retrieved.
 - If the file does not have a disk image, no action is taken.
 - However, if the disk image can be reloaded from tape, it is recommended that the file be reloaded using PFLOAD with OP=Z so that the asa field will be cleared.
- For error type 4, the action taken is both that taken for error type 3 and that taken for error type 5, 6, or 7.
- For error type 5, 6, or 7, the frozen flag is set in the 7990 catalog entry for the initial AU on the problem chain or fragment. This enables the problem chain/fragment/AU to be made available to the SSDEBUG utility, but prevents these AUs from being overwritten until then. Thus, the SSDEBUG utility can be used to inspect or save data from the corresponding AUs or cartridges. If the AUs are not allocated, the inhibit flag is also set in the 7990 catalog for that cartridge so that no new files will overwrite any data that might be on the cartridge.

Validation Report

The validation report consists of a report heading and a series of report groups for each subfamily and SM being reported on. The heading identifies the subfamily, the SM, whether or not there are any problems, and the last purge date and time for the SM (the last time orphans on this SM were released).

There is one report group for each error detected, and the actual information recorded in a report group depends on the type of error, as described next. Each report group, however, contains the following items:

- Error type (a number from 1 to 7).
- Identification (refer to the particular error type, described next).
- Chain information (7990 catalog ordinal, AU number, A or U designation for allocated or unallocated, H or E designation for HOC or EOC).
- Error description.

After the last report group, the validation report lists the number of trouble-free orphans, the amount of released or releasable 7990 space, the total number of errors detected, and whether or not the MSE system files were updated.

Error Types 1 and 2

Error types 1 and 2 identify mismatches between the SM map and the 7990 catalog. In the validation report, the identification field lists the 7990 catalog ordinal, the y and z coordinates, and the csu of the cartridge in error. The chain field is blank because problem chains are not identified as either error type 1 or 2. The analyst should run the SSUSE utility to produce a detailed report of the appropriate SM map and 7990 catalog entries to determine the exact problem.

Error Types 3 and 4

Error types 3 and 4 identify problem chains and problem asa values. In the validation report, the identification field lists the permanent file name and user index of the affected file; the dump control date and time (from the PFC entry for the file) to identify the backup file, if any; and the letter N (no) or Y (yes) to indicate whether or not the file has a disk image. The chain field lists the 7990 catalog ordinal and volume number for all AUs in the affected chain. An A or U indicates whether each AU is allocated or unallocated, and an H or E identifies the HOC or EOC. An error description is printed for each error detected; one chain can have several errors.

Error Types 5, 6, and 7

Error types 5, 6, and 7 identify problem orphans, fragments, and problem unallocated AUs. In the validation report, the information reported is the same as for error types 3 and 4, except for the identification field. Instead of the permanent file identification, the word ORPHAN (error type 5), FRAGMENT (error type 6), or UNALLOCATED (error type 7) is printed. Error type 7 is an unallocated AU that is pointed to by a chain or fragment. Each such AU is also reported with the chain for the corresponding orphan or fragment.

Intersections

Intersections occur when more than one chain links to the same AU; they are classified as either error type 4 or 5. Thus, they are reported as explained previously. Intersections are also reported in a separate entry consisting of a heading and additional information identifying the 7990 catalog ordinals and AU numbers of the intersecting chains. Therefore, when intersections are reported, there are two entries for the affected subfamily and SM: one lists only the intersections and the other lists all the errors encountered.

Validation Report Example

Figure 11-7 shows the format of a validation report. SSVAL was run for problem reporting purposes (RF not specified), and no errors were detected.

```

SSVAL - VALIDATION REPORT                      SSVAL - VER *.*
      FAMILY = MTST

      SSVAL,FM=MTST,SB=0,SM=A.

              L = OUTPUT
              RF = 0
              AM = 0
              SM = A
              FM = MTST
              FX = 0
              RL = 0
              SB = 0
              ST = 0

SUBFAMILY = 0  SM = A    -- GOOD --    LPDT = ***** *****

RELEASABLE 7990 FILES    =    12

RELEASABLE 7990 AUS      =    800

TOTAL VALIDATION ERRORS =     0

CATALOGS NOT MODIFIED

**REPORT COMPLETE**

```

Figure 11-7. SSVAL Validation Report

Typical SSVAl Runs

As described previously, SSVAl is run either to make 7990 space available for reuse or to report on problems with the current MSE system files and/or PFC entries. The following examples show typical SSVAl runs that may be used periodically for these purposes.

Example 1:

The following SSVAl command causes 7990 space to be made available for reuse.

```
SSVAL, RF=DUMP1, RL.
```

DUMP1 is the release data file produced by a previous PFDUMP run from which SSVAl can identify all the 7990 files that were orphans at the time of the dump. If no error conditions are detected in the 7990 catalog and PFC entries contained on file DUMP1 (FX=0 by default), the orphans are purged and the 7990 space assigned to them is released. The last purge date and time field in the subcatalog (in the 7990 catalog) for each SM for each subfamily is updated to the time of the SSVAl run if any file from that SM is purged. This is to ensure that a subsequent SSVAl run does not release the same space a second time, as the 7990 space may be reallocated to a new file.

If errors were detected, they are listed on the validation report and release processing is not performed. This is indicated by the CATALOGS NOT MODIFIED message at the end of the validation report. The analyst can rerun SSVAl with the following command to ensure that release processing is performed (n is the number of errors detected during the previous SSVAl run).

```
SSVAL, RF=DUMP1, RL, FX=n.
```

After this SSVAl run has completed, an analysis of the current version of the SM map, 7990 catalog, and PFC entries should be made to determine whether or not the errors detected in the first SSVAl run exist in the current MSE system files and PFC entries. This is accomplished using the command:

```
SSVAL, FM=familyname, AM.
```

familyname is that on the release data file, DUMP1.

Example 2:

Upon completion of any device reload that includes recovery of a 7990 catalog and/or SM map, SSVAl should be run to determine whether any SM map/7990 catalog mismatches exist. The following call accomplishes this.

```
SSVAL, FM=familyname, SB=n1, n2, . . . , nh, AM.
```

familyname is that for which recovery was done; n₁, n₂, . . . , n_h are the affected subfamilies if just some of the devices of the family were reloaded.

SSVAL analyzes the SM map, 7990 catalog, and PFC entries for the subfamilies specified by the SB parameter and reports any discrepancies or inconsistencies. For example, if the recovery was not scheduled and SSLABEL was run to add or remove cartridges or cubicles from one of the affected subfamilies after the last incremental or full dump for these subfamilies, then the SM map will reflect the results of the SSLABEL run but the 7990 catalog will not. Such mismatches will be reported on the validation report.

If the following SSVAL command is then made, the errors detected during the previous run will be flagged, as described previously under Release Processing and Problem Fixing.

```
SSVAL,FM=familyname,SB=n1,n2,...,nn,AM,FX=n.
```

familyname and n₁,n₂,...,n_n are the same as in the previous SSVAL run; n is the number of errors detected during the previous run. After investigating the cause of these errors, the analyst can run the SSDEBUG utility to correct the error conditions (refer to SSDEBUG later in this section).

If a device reload includes recovery of all SM maps (default family, user index=377760_g), the latest SM maps should be recovered from backup copies. Then each family that has 7990-resident files should be analyzed using the following command to detect SM map/7990 catalog mismatches.

```
SSVAL,FM=familyname,AM.
```

It is recommended that the previous SSVAL run be made periodically to check whether any unexpected error conditions exist. As the site analyst becomes more familiar with MSE processing, these periodic SSVAL runs can be done less frequently.

Restrictions to SSVAL

The following restrictions apply to the SSVAL utility.

- SSVAL cannot be run when SSMOVE is in the process of destaging files.
- Only one copy of SSVAL can be run at a time.
- SSVAL, SSLABEL, and SSDEBUG cannot be run at the same time.

SSUSE

SSUSE reads data in the 7990 catalogs and SM maps and produces reports on the availability of space on 7990 cartridges and the allocation of cubicle space within an SM. The reports may not be completely up to date because the 7990 catalogs and SM maps can be updated while the reports are being generated.

The types of reports that SSUSE generates are:

Report	Description
Basic usage report	Usage report that lists general information about the use of each SM in a subfamily.
Optional report A	SM map report that lists the contents of an SM as described in the SM map.
Optional report B	Cartridge summary report that lists general status information for each cartridge entry in the 7990 catalog. The report identifies the available AUs and flags set for each cartridge in the 7990 catalog.
Optional report C	Detailed cartridge report that lists cartridge usage information for each cartridge entry in the 7990 catalog.
Optional report D	Detailed AU status report that lists AU status information for each entry in the 7990 catalog plus cartridge usage information.

The format of the SSUSE command is:

SSUSE, P₁, P₂, . . . , P_n.

P_i	Description
CM=A-	Cartridge manufacturer code is A-, indicating IBM.
CM	Same as CM=A-.
CM omitted	Same as CM=A-.
CN=csn	Cartridge serial number of a specific cartridge. The CN=csn parameter is valid only if the OP=D parameter is specified.
CN	Cartridge serial number is not specified.
CN omitted	Same as CN.
FM=familyname	Family to be reported on.
FM	Same as FM=system default family.
FM omitted	Same as FM=system default family.
L=filename	File on which listable output is to be written.
L	Same as L=OUTPUT.

Report	Description										
L=0	No output file is to be generated.										
L omitted	Same as L=OUTPUT.										
OP=op	Type of report to be generated. Multiple options can be specified (for example, OP=AB).										
	<table border="1"> <thead> <tr> <th>op</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Optional report A and basic usage report.</td> </tr> <tr> <td>B</td> <td>Optional report B and basic usage report.</td> </tr> <tr> <td>C</td> <td>Optional report C and basic usage report.</td> </tr> <tr> <td>D</td> <td>Optional report D and basic usage report.</td> </tr> </tbody> </table>	op	Description	A	Optional report A and basic usage report.	B	Optional report B and basic usage report.	C	Optional report C and basic usage report.	D	Optional report D and basic usage report.
op	Description										
A	Optional report A and basic usage report.										
B	Optional report B and basic usage report.										
C	Optional report C and basic usage report.										
D	Optional report D and basic usage report.										
OP	Basic usage report only is to be generated.										
OP omitted	Same as OP.										
SB=subfamily	Subfamily to be reported on. Up to eight subfamilies can be selected by the numbers 0 through 7. For example, SB=0273 selects subfamilies 0, 2, 3, and 7.										
SB	Same as SB=01234567.										
SB omitted	Same as SB=01234567.										
SM=id	SM identifier of the SM to be used. Up to eight SMs can be selected by the letters A through H. For example, SM=AGC selects SM A, C, and G.										
SM	Same as SM=A.										
SM omitted	Same as SM=A.										

Basic Usage Report

The basic usage report includes a title line, which identifies the family and SM being reported on. The statistics in this report contain data for each group followed by totals for SMs and subfamilies. The following items are listed.

- Number of cubicles reserved for each subfamily.
- Number of cartridges in each subfamily.
- Number of unassigned and unflagged AUs for small files.
- Number of unassigned and unflagged AUs for large files.
- Number of flagged cartridges.
- Number of flagged AUs.
- If any empty cubicles are assigned, they are listed first for each subfamily. This line shows a hyphen (-) as the group designator since groups are assigned only to cartridges.

Figure 11-8 shows the format of a basic usage report.

yy/mm/dd. hh.mm.ss. PAGE 1

SSUSE REPORT FILE
 SSUSE BASIC REPORT FAMILY = MTST

CUBE = CUBICLES CARTRIDGE FLAGS AU FLAGS
 CART = CARTRIDGES M = MISSING FA = FLAWED AND ALLOCATED
 I = INHIBIT ALLOCATION FU = FLAWED AND UNALLOCATED
 ** = SUBFAMILY TOTAL F = FREE CARTRIDGE SF = START OF FRAGMENT
 - = UNASSIGNED GROUP L = OFF CARTRIDGE LINKS AVAILABLE FC = FROZEN CHAIN
 P = EXCESSIVE PARITY ERRORS AC = AU CONFLICT
 E = MAP ERROR (AS DETECTED BY SSVAL)

SUB	SM	GR	CUBE	CART	---AVAILABLE---		NUMBER CARTRIDGES FLAGGED						--NUMBER AU FLAGGED--						
					AU (SMALL)	AU (LARGE)	M	I	F	L	P	E	FA	FU	SF	FC	AC		
0	A	-	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	A	1	2	2	1167	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
0	A	**	4	2	1167	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
1	A	-	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	A	1	2	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
1	A	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
2	A	-	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	A	1	2	2	1168	2606	0	0	0	2	0	0	0	0	0	0	0	0	0
2	A	**	5	2	1168	2606	0	0	0	2	0	0	0	0	0	0	0	0	0
3	A	-	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	A	1	2	2	1168	2586	0	0	0	2	0	0	0	0	0	0	0	0	0
3	A	**	5	2	1168	2586	0	0	0	2	0	0	0	0	0	0	0	0	0
4	A	-	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	A	1	2	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
4	A	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
5	A	-	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	A	1	2	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
5	A	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
6	A	-	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	A	1	3	3	5400	7	0	2	0	1	0	0	0	0	0	0	3	0	0
6	A	**	5	3	5400	7	0	2	0	1	0	0	0	0	0	0	3	0	0
7	A	-	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	A	1	2	2	1200	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
7	A	**	5	2	1200	2662	0	0	0	2	0	0	0	0	0	0	0	0	0

Figure 11-8. SSUSE Basic Report (Continued)

yy/mm/dd. hh.mm.ss. PAGE 2																			
SSUSE REPORT FILE																			
SSUSE BASIC REPORT SM AND SUBFAMILY TOTALS FAMILY = MTST																			
SUB	SM	GR	CUBE	CART	----AVAILABLE----		NUMBER CARTRIDGES FLAGGED					--NUMBER AU FLAGGED--							
					AU (SMALL)	AU (LARGE)	M	I	F	L	P	E	FA	FU	SF	FC	AC		
SM TOTALS FOR SPECIFIED SUBFAMILIES																			
**	A	**	39	17	13607	18509	0	2	0	15	0	0	0	0	0	0	0	3	0
SUBFAMILY TOTALS FOR SPECIFIED SM-S																			
0	**	**	4	2	1167	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
1	**	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
2	**	**	5	2	1168	2606	0	0	0	2	0	0	0	0	0	0	0	0	0
3	**	**	5	2	1168	2586	0	0	0	2	0	0	0	0	0	0	0	0	0
4	**	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
5	**	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0	0	0
6	**	**	5	3	5400	7	0	2	0	1	0	0	0	0	0	0	0	3	0
7	**	**	5	2	1200	2662	0	0	0	2	0	0	0	0	0	0	0	0	0

Figure 11-8. SSUSE Basic Report

Optional Report A

Optional report A includes a title line that identifies the SM being reported on. This report is made up of 15 pages, each page describing the SM map for specific y and z coordinates. Each page describes 22 y coordinates (rows) and 1 z coordinate (column). Column 6 (z=6) is not reported on because it contains the data recording drives. The information listed in optional report A includes the following items:

- y and z coordinates.
- Cartridge manufacturer code and CSN of the cartridge.
- Family and subfamily.
- Group and ordinal in group.
- Flags.

Figure 11-9 shows the format of optional report A.

		yy/mm/dd. hh.mm.ss. PAGE						3
SSUSE REPORT FILE								
SSUSE OPTIONAL REPORT A - STORAGE MODULE MAP FOR SM = A								
FLAGS: P = CARTRIDGE EXISTS IN POOL F = NOT ASSIGNED TO THIS FAMILY R = RESERVED FOR ALTERNATE SMMAP S = RESERVED FOR SYSTEM USE E = ERROR FLAG C = RESERVED FOR CUSTOMER ENGINEERING N = (Y,Z) DOES NOT EXIST IN SMMAP GPORD = ORDINAL IN GROUP								
Y	Z	CM	CSN	FAMILY	SUBFAMILY	GROUP	GPORD	FLAGS
0	0							C F
1	0							F
2	0							F
3	0							F
4	0							F
5	0							F
6	0							F
7	0							F
8	0							F
9	0							F
10	0							F
11	0							NF
12	0							NF
13	0							NF
14	0							NF
15	0							NF
16	0							F
17	0			MTST			0	
18	0			MTST			0	
19	0	A-51680502		MTST		1	1	
20	0	A-13727600		MTST		1	0	
21	0							S F

Figure 11-9. SSUSE Optional Report A

Optional Report B

Optional report B includes a title line identifying the subfamily, family, and SM being reported on. There is one line of output per cartridge within a subfamily and storage module. The report includes the following cartridge information about entries in the 7990 catalog.

- Group and ordinal in group.
- y and z coordinates.
- Cartridge manufacturer code and CSN of the cartridge.
- Free AUs for small and large files.
- Cartridge flags.
- Error conditions.

Figure 11-10 shows the format of optional report B.

```

yy/mm/dd. hh.mm.ss. PAGE 25
SSUSE REPORT FILE
SSUSE OPTIONAL REPORT B - CARTRIDGE SUMMARY REPORT SM = A SUBFAMILY = 6 FAMILY = MTST

NOTES:
FA = FLAWED AND ALLOCATED
FU = FLAWED AND UNALLOCATED
SF = START OF FRAGMENT
FC = FROZEN CHAIN
AC = AU CONFLICT
GPORD = ORDINAL IN GROUP

CARTRIDGE FLAGS:
M = MISSING
I = INHIBIT
F = FREE CARTRIDGE
L = LINK(FREE AU EXIST, NO OFF CARTRIDGE LINK)
P = EXCESSIVE WRITE ERRORS
E = MAP ERROR(DETECTED BY SSVAL)

-----ERROR CONDITIONS-----
-----NUMBER OF AU-----
GP GPORD Y Z CM CSN FREE AU SMALL LARGE CART FLAGS FA FU SF FC AC

```

GP	GPORD	Y	Z	CM	CSN	FREE AU SMALL	LARGE	CART FLAGS	FA	FU	SF	FC	AC
1	0	21	7	A-	13268794	1800	0	I L	0	0	0	3	0
1	1	20	7	A-	51329412	1800	3	I	0	0	0	0	0
1	2	19	7	A-	51328839	1800	4	L	0	0	0	0	0

Figure 11-10. SSUSE Optional Report B

Optional Report C

Optional report C is the detailed cartridge report. It lists the group and FCT ordinal for each 7990 catalog entry and the y and z coordinates and CSN for each cartridge. The report also includes the AU numbers of the first AUs for small and large files and the division point for each cartridge.

Figure 11-11 shows the format of optional report C.

yy/mm/dd. hh.mm.dd. PAGE 33										
SSUSE REPORT FILE										
SSUSE OPTIONAL REPORT C - DETAILED CARTRIDGE REPORT SM = A SUBFAMILY = 6 FAMILY = MTST										
FLAGS: I = INHIBIT ALLOCATION M = MISSING P = EXCESSIVE PARITY ERRORS E = MAP ERROR OCL = AVAILABLE LINK COUNT FCTORD = SFM CATALOG ORDINAL										
Y	Z	CM	CSN	GROUP	FCTORD	FLAGS	-----AU-----		CARTRIDGE DIVISION	OCL
							FIRST SMALL	FIRST LARGE		
21	7	A-13268794		1	16	I	1	0	1801	0
20	7	A-51329412		1	17	I	1	1916	1801	2
19	7	A-51328839		1	18		1	1832	1801	0

Figure 11-11. SSUSE Optional Report C

Optional Report D

Optional report D is the detailed AU status report. It lists the group and FCT ordinal of the cartridge, the CSN, the y and z coordinates of the cartridge, and the number of available off-cartridge links. The report also includes the octal numbers of the AUs with flawed AU, start of volume, and error flags noted. This report also specifically notes when a cartridge is empty (for use in determining the status of a cartridge that is to be made available).

Figure 11-12 shows the format of optional report D.

yy/mm/dd. hh.mm.ss. PAGE 6								
SSUSE REPORT FILE								
SSUSE OPTIONAL REPORT D - DETAILED AU STATUS REPORT SM = A SUBFAMILY = 1 FAMILY = MTST								
F = FLAWED AU (DEMARK FAILURE)								
V = START OF VOLUME								
E = ONE OF THE ERROR FLAGS SET (AU CONFLICT, FROZEN CHAIN, START OF FRAGMENT)								
FCTORD	Y	Z	CM	CSN	GROUP			
16	21	1	A-	51328833	1			
AU	XXX0	XXX1	XXX2	XXX3	XXX4	XXX5	XXX6	XXX7
100X	0041300731	0041270731	0041260731	0041250731	0041240731	0041230731	0041220731	0041210731
101X	0041200731	0041170731	0041160731	0041150731	0041140731	0041130731	0041120731	0041110731
102X	0041100731	0041070731	0041060731	0041050731	0041040731	0041030731	0041020731	0041010731
103X	0041000731	0040770731	0040760731	0040750731	0040740731	0040730731	0040720731	0040710731
104X	0040700731	0040670731	0040660731	0040650731	0040640731	0040630731	0040620731	0040610731
105X	0040600731	0040570731	0040560731	0040550731	0040540731	0040530731	0040520731	0040510731
106X	0040500731	0040470731	0040460731	0040450731	0040440731	0040430731	0040420731	0040410731
107X	0040400731	0040370731	0040360731	0040350731	0040340731	0040330731	0040320731	0040310731
110X	0040300731	0040270731	0040260731	0040250731	0040240731	0040230731	0040220731	0040210731
111X	0040200731	0040170731	0040160731	0040150731	0040140731	0040130731	0040120731	0040110731
112X	0040100731	0040070731	0040060731	0040050731	0040040731	0040030731	0040020731	0040010731
113X	0040000731	V0000621214	0040061131	00400601131	0040571131	0040561131	0040551131	0040541131
114X	0040531131	0040521131	0040511131	0040501131	0040471131	0040461131	0040451131	0040441131
115X	0040431131	0040421131	0040411131	0040401131	0040371131	0040361131	0040351131	0040341131
116X	0040331131	0040321131	0040311131	0040301131	0040271131	0040261131	0040251131	0040241131
117X	0040231131	0040221131	0040211131	0040201131	0040171131	0040161131	0040151131	0040141131
120X	0040131131	0040121131	0040111131	0040101131	0040071131	0040061131	0040051131	0040041131
121X	0040031131	0040021131	0040011131	0040001131	V0001771414	0041761214	0041751214	0041741214
122X	0041731214	0041721214	0041711214	0041701214	0041671214	0041661214	0041651214	0041641214
123X	0041631214	0041621214	0041611214	0041601214	0041571214	0041561214	0041551214	0041541214
124X	0041531214	0041521214	0041511214	0041501214	0041471214	0041461214	0041451214	0041441214
125X	0041431214	0041421214	0041411214	0041401214	0041371214	0041361214	0041351214	0041341214
126X	0041331214	0041321214	0041311214	0041301214	0041271214	0041261214	0041251214	0041241214
127X	0041231214	0041221214	0041211214	0041201214	0041171214	0041161214	0041151214	0041141214
130X	0041131214	0041121214	0041111214	0041101214	0041071214	0041061214	0041051214	0041041214
131X	0041031214	0041021214	0041011214	0041001214	0040771214	0040761214	0040751214	0040741214
132X	0040731214	0040721214	0040711214	0040701214	0040671214	0040661214	0040651214	0040641214
133X	0040631214	0040621214	0040611214	0040601214	0040571214	0040561214	0040551214	0040541214
134X	0040531214	0040521214	0040511214	0040501214	0040471214	0040461214	0040451214	0040441214
135X	0040431214	0040421214	0040411214	0040401214	0040371214	0040361214	0040351214	0040341214
136X	0040331214	0040321214	0040311214	0040301214	0040271214	0040261214	0040251214	0040241214
137X	0040231214	0040221214	0040211214	0040201214	0040171214	0040161214	0040151214	0040141214

Figure 11-12. SSUSE Optional Report D

SSDEBUG

SSDEBUG allows the analyst to update appropriate entries in the SM maps and/or 7990 catalogs and thereby resolve inconsistencies reported by the SSVAL utility. SSDEBUG can also copy data from selected 7990 files or cartridges to disk. This function of SSDEBUG is intended for use when errors such as unrecoverable read errors do not allow successful staging of a 7990 file.

NOTE

SSDEBUG updates the SM map for the specified SM. (Refer to the definition of SM map earlier in this section.) It is recommended that the SM map is copied on tape or on another device or family immediately after every update of the SM map.

Input to SSDEBUG is via a directive file. Refer to SSLABEL, earlier in this section, for a description of the directive file format.

The format of the SSDEBUG command is:

```
SSDEBUG, P1, P2, P3.
```

P_i	Description
I=filename	File on which directives are written.
I	Same as I=INPUT.
I omitted	Same as I=INPUT.
L=filename	File on which listable output is to be written.
L	Same as L=OUTPUT.
L=0	No output file is to be generated.
L omitted	Same as L=OUTPUT.
Z	Directives are contained on the SSDEBUG command. The I parameter is ignored.
Z omitted	Directives are contained on the file specified by the I parameter.

SSDEBUG Directives

The following directives are available with SSDEBUG. The descriptions of the parameters for these directives follow this listing of directives.

OP=RS — Read AUs

OP=RS reads selected AUs of a cartridge identified by its csn (CN=csn) or y and z coordinates (YI and ZI parameters). The SM parameter specifies the SM where the cartridge resides. The range of AUs to be read is specified by the SL and SU parameters. The AUs are written to the file specified by the PF parameter. All CYBER data plus 7990 and CRC data, including header and trailer information, is returned. The analyst can write an application program to reconstruct the permanent file by stripping off all CYBER or 7990 related control information. Refer to Format of Disk Files Read From 7990 by SSDEBUG later in this section for a description of the control information written to disk along with the disk file contents.

OP=RF — Read File

OP=RF reads the file whose alternate storage address is specified by the FO, SB, SM, and ST parameters. The file is written to the file specified by the PF parameter. All data, including header and trailer information, is returned (refer to Format of Disk Files Read From 7990 by SSDEBUG later in this section).

OP=RP — Release Frozen Chain Space

OP=RP clears flags in the 7990 catalog and releases 7990 space for the chain whose alternate storage address is specified by the FO, SB, SM, and ST parameters.

OP=RL — Remove Cartridge Entry From 7990 Catalog

OP=RL removes a 7990 catalog entry that is not linked properly to the SM map. The FO, SB, and SM parameters identify the 7990 catalog entry to be removed.

OP=RC — Remove Cartridge Entry From SM Map

OP=RC removes an SM map entry that does not have a corresponding FCT entry in the 7990 catalog. The SM, YI, and ZI parameters identify the SM map entry to be removed. The cartridge at the particular location (y,z) is put into the exit tray and the SM map entry is changed to unassigned. The SM map should be backed up immediately after the SSDEBUG run to ensure consistency with cartridge labels and 7990 catalogs (refer to the definition of SM map earlier in this section).

A scratch label is written on the cartridge prior to it being replaced in the exit tray, but the cartridge history information remains intact.

OP=CF — Change Flag in SM Map or 7990 Catalog

OP=CF sets or clears flags in the SM map or 7990 catalog. The FL and ON or OF parameters indicate the flag type and the action to be taken. The appropriate FCT entry or SM map entry is determined from the CN, FO, or YI and ZI parameters. For FCT flags (FL=AC, FC, or SF), the analyst can specify either a single AU using the ST parameter or a range of AUs using the SL and SU parameters.

Parameters for the SSDEBUG Directives

The descriptions of the parameters to the SSDEBUG directives follow. Not all parameters are valid with all directives, as indicated.

P_i	Description
CM=A-	Cartridge manufacturer code is A-; indicating IBM.
CM	Same as CM=A-.
CM omitted	Same as CM=A-.
CN=csn	Cartridge serial number of the cartridge to be used; not valid for OP=RS if YI=y or ZI=z is specified; and not valid for OP=CF if FO=ord, YI=y, or ZI=z is specified.
CN omitted	YI=y and ZI=z must be specified for OP=RS. FO=ord or YI=y and ZI=z must be specified for OP=CF.
FL=flag	Name of flag in SM map or 7990 catalog to be set or cleared. Valid only for OP=CF.
	flag Description
	AC AU conflict flag.
	EW Excessive write parity error flag.
	FC Frozen chain flag.
	FE Linkage error flag (in 7990 catalog FCT).
	IB Inhibit allocation flag.
	LC Lost cartridge flag.
	ME Linkage error flag (in SM map).
	SF Start of fragment flag.
FL omitted	FL=flag must be specified for OP=CF.
FM=familyname	Family to be processed.
FM	Same as FM=system default family.
FM omitted	Same as FM=system default family.

P_i	Description
FO=ord	7990 catalog ordinal indicating the file to be read or the chain whose space is to be released. Not valid for OP=CF if CN=csn, YI=y, or ZI=z is specified.
FO omitted	FO=ord must be specified for OP=RF, OP=RP, and OP=RL. CN=csn or YI=y and ZI=z must be specified for OP=CF.
OF	Flag specified by FL=flag is to be cleared. Valid only with OP=CF.
ON	Flag specified by FL=flag is to be set. Valid only with OP=CF.
PF=filename	File to which the 7990 image is to be copied. This file is defined under the user's current family and user index.
PF	Same as PF=MMMMBUG.
PF omitted	Same as PF=MMMMBUG.
SB=subfamily	Subfamily to be used; $0 \leq \text{subfamily} \leq 7$.
SB	Same as SB=0.
SB omitted	Same as SB=0.
SL=i	AU where OP=RS begins copying or OP=CF begins changing flags; $1 \leq i \leq 1931$; $i \leq j$ (refer to SU=j).
SL	Same as SL=1.
SL omitted	Same as SL=1.
SM=id	SM identifier of the SM to be used by SSDEBUG; id is a letter from A to H.
SM	Same as SM=A.
SM omitted	Same as SM=A.
ST=s	AU where OP=RF begins reading files or fragments or OP=RP begins releasing frozen chain space. If specified with OP=CF, ST=s specifies the AU number of an FCT flag to be changed and takes precedence over the range of AUs specified by the SL and SU parameters. AU numbers are meaningful with OP=CF only for flags FL=AC, FC, or SF.

P_i	Description
ST omitted	ST=s must be specified for OP=RF and OP=RP.
SU=j	AU where OP=RS ends copying or OP=CF ends changing flags; 1 ≤ j ≤ 1931; i ≤ j (refer to SL=i).
SU	Same as SU=1.
SU omitted	Same as SU=1.
YI=y	y coordinate of the cubicle where the cartridge to be read resides; 0 ≤ y ≤ 21. ZI=z must also be specified. CN=csn must not be specified with OP=RS. FO=ord or CN=csn must not be specified with OP=CF.
YI omitted	CN=csn must be specified for OP=RS. FO=ord or CN=csn must be specified for OP=CF. YI=y and ZI=z must be specified for OP=RC.
ZI=z	z coordinate of the cubicle where the cartridge to be read resides; 0 ≤ z ≤ 15 and z≠6. YI=y must also be specified. CN=csn must not be specified with OP=RS. FO=ord or CN=csn must not be specified with OP=CF.
ZI omitted	CN=csn must be specified for OP=RS. FO=ord or CN=csn must be specified for OP=CF. YI=y and ZI=z must be specified for OP=RC.

Using SSDEBUG

Suppose the following validation report (figure 11-13) was generated by running SSVAL for problem reporting purposes.

ERR	IDENTIFICATION	PROBLEMS	LPDT = yy/mm/dd. hh.mm.ss
		-CHAIN-	ERROR DESCRIPTION
TYP		FCT - AU - A/U - H/E	
5	ORPHAN	017-1737 AH-	ILL-FORMED CHAIN SCATTERED FILE (INFO)
		017-1738 A--	
		017-1739 A--	
		017-1740 A--	
		017-1741 A--	
		017-1742 A--	
		017-1743 A--	
		017-1930 A--	
		017-1931 A--	
		019-0000 U--	START OF FRAGMENT
7	UNALLOCATED	019-0000 U--	START OF FRAGMENT

Figure 11-13. SSDEBUG: Executing SSVAL in Problem-Fixing Mode

This partial validation report shows an error type 5 identified by ORPHAN with an error description of ILL-FORMED CHAIN (frozen chain). You can use SSDEBUG to set (OP=CF) and release (OP=RP) the frozen chain space or you can execute SSVAL in problem-fixing mode to set the frozen chain flag.

Enter the following SSDEBUG command to set the frozen chain flag for the range of AUs specified by the SL and SU parameters.

```
SSDEBUG, Z. /OP=CF, FL=FC, ON, FM=SMD, SB=4, SL=1737, SU=1931, FO=17, ST=1737.
```

When the SSDEBUG or SSVAL command is processed, enter the following command to release the frozen chain space.

```
SSDEBUG, Z. /OP=RP, SB=4, SM=A, ST=1737, FO=17, FM=SMD.
```

When this command is processed, you can verify that the frozen chain has been set by running SSVAL again to generate another validation report.

Restrictions to SSDEBUG

The following restrictions apply to the SSDEBUG utility.

- SSEXEC must be running when SSDEBUG is run.
- Only one copy of SSDEBUG can be run at one time.
- SSDEBUG, SSVAl, and SSLABEL cannot be run at the same time.

y,z Relationship with SM Map Ordinal

A one-to-one functional mapping exists between coordinate pairs and the ordinal entries within the SM map. To determine the coordinate pair (y,z) from the SM map ordinal (ord), use the following equations:

$$y = (352 - \text{ord}) / 16$$
$$z = 352 - \text{ord} - 16y$$

To determine the SM map ordinal (ord) from the coordinate pair (y,z), use the following equation:

$$\text{ord} = 352 - z - 16y$$

Format of Disk Files Read From 7990 by SSDEBUG

A local disk file created by SSDEBUG with the OP=RS or OP=RF option includes all CYBER and 7990 header and trailer information. Figure 11-14 shows the format of a disk file created when SSDEBUG reads a full stripe; figure 11-15 shows the format of a disk file created when SSDEBUG reads a partial stripe.

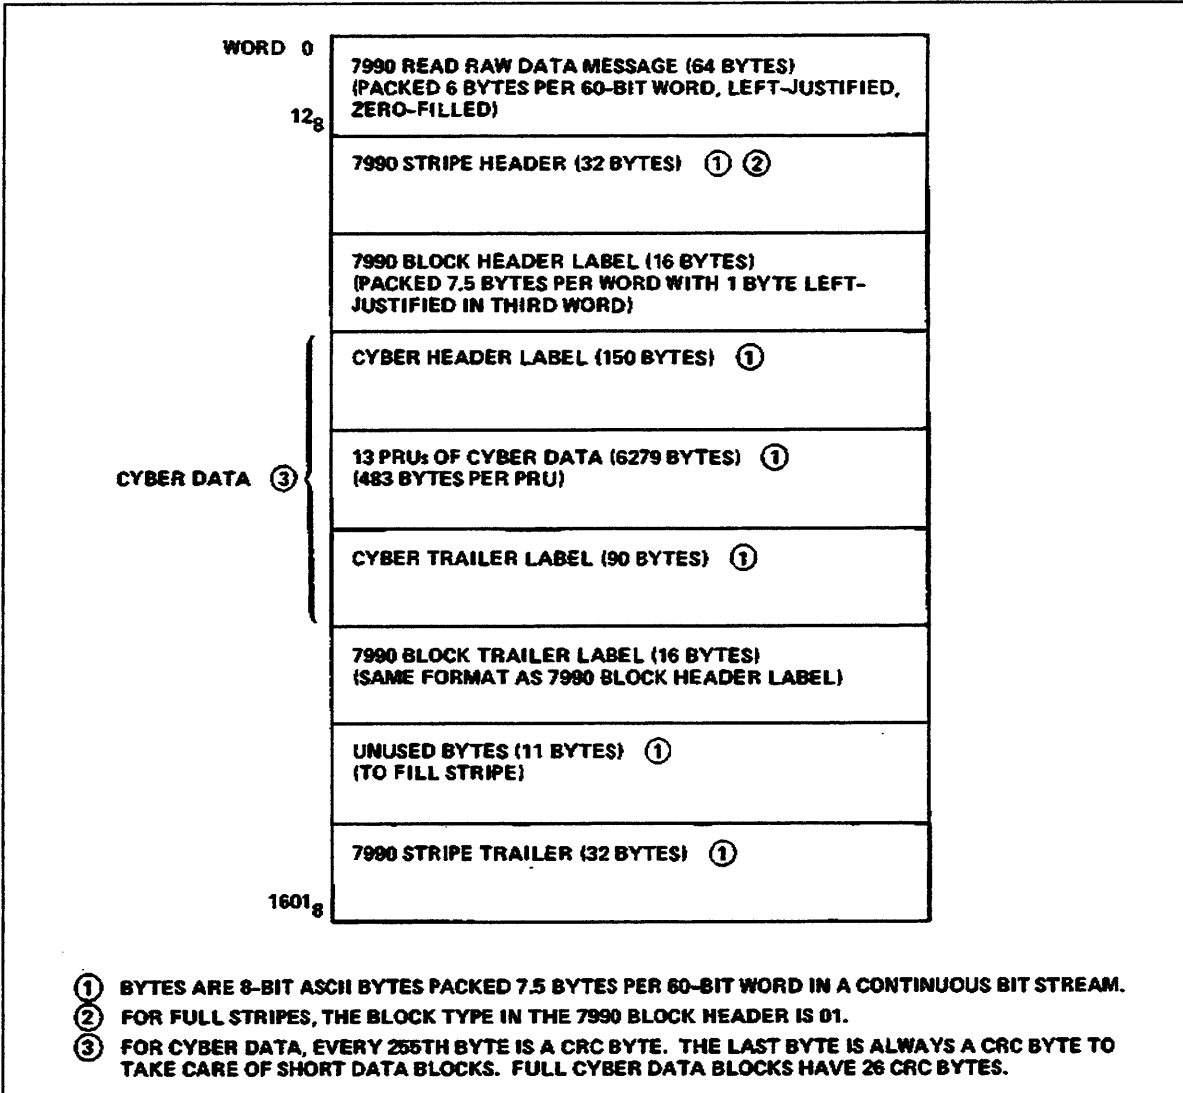


Figure 11-14. Full Stripe Disk File Format

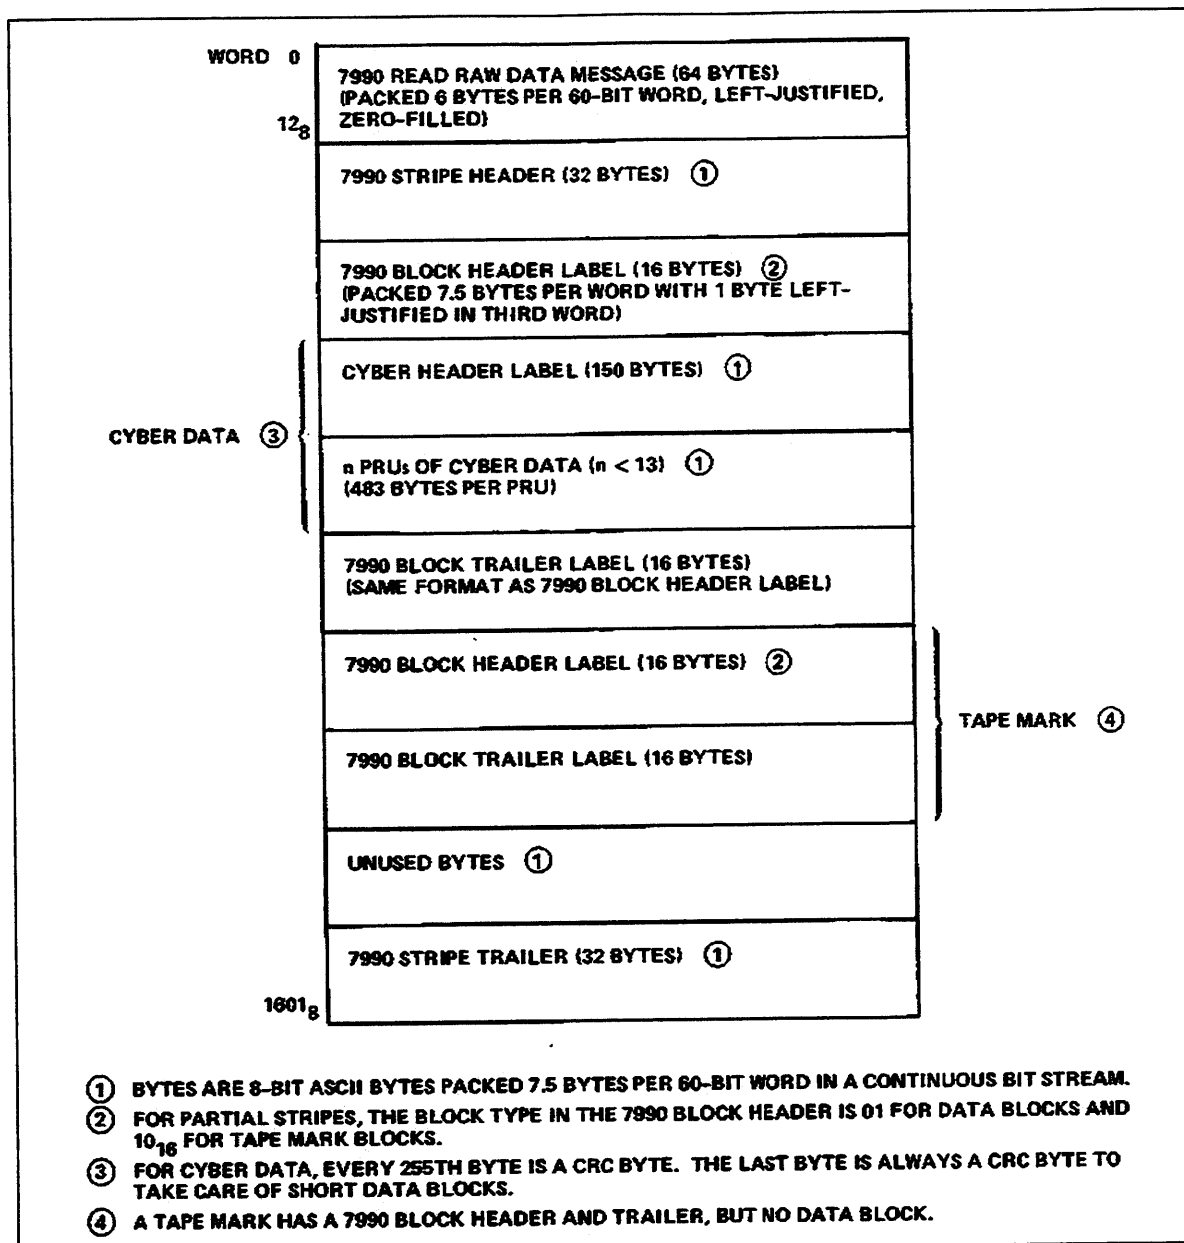


Figure 11-15. Partial Stripe Disk File Format

7990 Stripe Header and Trailer

The 7990 stripe header is a 32-byte block that occupies the first 32 bytes of the stripe; the 7990 stripe trailer occupies the last 32 bytes. The 7990 stripe header and trailer have the same format.

Bytes	Description
0-5	Volume serial number. A 6-character string specifying the volume serial number of the stripe.
6-7	Version number. Two characters specifying the version number associated with the volume serial number.
8-9	File number. Two characters specifying the number of the file within the volume.
A-B	Logical stripe number. Relative stripe number within the volume. The logical stripe number may not correlate with the physical stripe number due to demarked stripes.
C-D	Address of last block trailer. Byte number, within the stripe, of the first byte of the last block trailer.
E-1D	Reserved.
1E-1F	Checksum. Chosen such that the sum, modulo 65536, of all bytes in the header equals zero.

7990 Block Header and Trailer

Each 7990 block has a 16-byte header that precedes each block and a 16-byte trailer that follows each block, or partial block, of data on the stripe. The 7990 block header and trailer have the same format.

Byte(s)	Description												
0-1	Block type. The following block types are valid.												
	<table border="1"> <thead> <tr> <th>Block Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>01</td> <td>A complete block is contained in this stripe.</td> </tr> <tr> <td>02</td> <td>First block segment of a multistripe block.</td> </tr> <tr> <td>04</td> <td>Intermediate block segment of a multistripe block.</td> </tr> <tr> <td>08</td> <td>Last block segment of a multistripe block.</td> </tr> <tr> <td>10₁₆</td> <td>Tape mark.</td> </tr> </tbody> </table>	Block Type	Description	01	A complete block is contained in this stripe.	02	First block segment of a multistripe block.	04	Intermediate block segment of a multistripe block.	08	Last block segment of a multistripe block.	10 ₁₆	Tape mark.
Block Type	Description												
01	A complete block is contained in this stripe.												
02	First block segment of a multistripe block.												
04	Intermediate block segment of a multistripe block.												
08	Last block segment of a multistripe block.												
10 ₁₆	Tape mark.												
2-3	Block length. The number of bytes in a data block excluding the block header and trailer. The block length of a tape mark is zero.												
4-7	Block sequence number. The number of a given data block within a file.												
8-9	Block segment number. For complete blocks, the segment number is 0 (zero). For multistripe blocks, the segment number is the sequence number of the block segment within the data block.												
A-E	Reserved.												
F	Checksum. Chosen such that the sum, modulo 256, of all bytes in the header equals zero.												

SSBLD

SSBLD reads a specified direct access permanent file that contains SSBLD statements. These statements consist of 7990 component-oriented mnemonics that define the logical mapping from component to component within the 7990 configuration. SSBLD processes these statements and generates a direct access permanent file which SSEXEC uses as its unit device table (UDT).

The format of the SSBLD command is:

```
SSBLD,CF=confile,BF=bldfile.
```

Parameter	Description
CF=confile	Configuration file on which SSBLD statements are written under user index 377760 ₈ .
CF	Same as CF=SUDT.
CF omitted	Same as CF=SUDT.
BF=bldfile	Build file on which SSBLD generates the UDT for SSEXEC. The build file is under user index 377760 ₈ .
BF	Same as BF=BU DT.
BF omitted	Same as BF=BU DT.

SSBLD Statements

The 7990 hardware consists of a number of 7990 control units (CU) and a number of 7991 storage modules (SM). SSBLD uses two types of statements: those that link individually controllable nodes and paths within each CU, and those that link SMs to CUs. SSBLD statements must be processed in the following sequence or SSBLD aborts.

1. Statements that define the linkage between nodes and paths within each CU.
2. Statements that define the linkage between SMs and CUs.

Node/Path Linkage Statements

These SSBLD statements are order-dependent. There are seven types of statements for a given CU and they must appear together in the following sequence and format:

1. CUest=da,CH01=st,CH02=st,CH03=st,CH04=st.

est	EST ordinal of the CU (10 ₈ through 777 ₈).
da	Device address of CUest.
st	Channel status (ON or OFF).

The CYBER channel number is associated with the 7990 channel interface (CIF) number. The position (first, second, and so on) of the channel in the CU statement minus 1 defines the CIF number wired to that channel. For example, CH02 is wired to CIF1. Also, if a single digit channel is used to connect to a 7990, it must have a leading 0 in the CU statement. For example, channel 3 is entered as CH03.

2. CIFn=DTI0=st,DTI1=st,DTO0=st,DTO1=st.

n Device address of CIF (0 through 3).
 DTI Data transfer input.
 DTO Data transfer output.
 st Data transfer status (ON or OFF).

3. DTIn=DIF0=st,DIF1=st.

n Device address of DTI (0 or 1).
 DIF Device interface.
 st Device interface status (ON or OFF).

4. DTOn=DIF0=st,DIF1=st.

n Device address of DTO (0 or 1).
 st Device interface status (ON or OFF).

5. DIFn=DRC0=st,DRC1=st.

n Device address of DIF (0 or 1).
 DRC Data recording controller.
 st Data recording controller status (ON or OFF).

If two DRCs exist in a CU, their device addresses must be paired as 0 and 1 or 2 and 3.

6. DRCn=DRD0=st,...,DRD7=st. or DRCm=DRD8=st,...,DRD15=st.

n Device address of DRC (0 or 1).
 m Device address of DRC (2 or 3).
 DRD Data recording driver.
 st Data recording driver status (ON or OFF).

DRCs 0 and 1 can link only to DRDs 0 to 7, and DRCs 2 and 3 can link only to DRDs 8 to 15.

7. AIFn=AC0=st,AC1=st,AC2=st,AC3=st. or AIFn=AC4=st,AC5=st,AC6=st,AC7=st.

AIF Accessor interface.
 n Device address of AIF (0 or 1).
 AC Accessor.
 st Accessor status (ON or OFF).

An AIF can link to either ACs in the range 0 to 3 or 4 to 7, but not to ACs from both ranges.

SM/CU Linkage Statements

The SM statement defines the linkage between SMs and CUs; its format is:

```
SMid=st, ACn, CUest, CUest, DRDa=st, DRDa=st, DS=d, ST=s.
```

Parameter	Description
id	SM identifier (A through H).
st	Status (ON or OFF).
n	Device address of AC (0 through 3).
est	EST ordinal of the CU(s) connected to SMid.
a	Device address of the DRD(s) connected to ACn.
d	Number of DRDs for destaging per storage module (0 to 2). Default is 1.
s	Number of DRDs for staging per storage module (0 to 2). Default is 2.

SSBLD and SSEXEC use the common deck COMBUDT to define the amount of memory to reserve within SSEXEC for the unit device tables. Any attempt to change the number of SMs or CUs specified in the common deck will cause SSBLD to abort.

SSBLD Example

The first two SSBLD statements specify the total number of CUs and SMs in the following configuration. The first statement specifies CUs and the second statement specifies SMs. The rest of the statements define the 7990 configuration.

```
1
1
CU41=0, , CH23=ON, , .
CIF1=DTI0=ON, , DTO0=ON, .
DTI0=DIF0=ON, DIF1=ON.
DTO0=DIF0=ON, DIF1=ON.
DIF0=DRC0=ON, DRC1=ON.
DIF1=DRC0=ON, DRC1=ON.
DRC0=DRD0=ON, DRD1=ON, , , , , .
DRC1=DRD0=ON, DRD1=ON, , , , , .
AIF0=AC0=ON, , , .
SMA=ON, AC0, , CU41, DRD0=ON, DRD1=ON, DS=1, ST=2.
*COMMENT.
```

An asterisk (*) in the first column of an SSBLD statement specifies a comment statement. Comment statements are not sent to the dayfile.

7990 Equipment Configuration Constraints

Control paths and data paths within the 7990 configuration have certain constraints placed upon the allowable device addresses of the accessors and DRDs. The accessor numbers and corresponding DRD numbers are:

Accessor Number	DRD Numbers
0	0 or 1
1	2 or 3
2	4 or 5
3	6 or 7

Accessors and DRDs that are to be coupled to two controllers must have the same path numbers assigned from each controller. For example, if an accessor is plugged into accessor path 2 of CU1, then a DRD on the same storage module must be coupled to a DRC path of 4 or 5 from CU1. If a second controller is to be attached to this same storage module, then the other accessor port must also be plugged into accessor path 2 of CU2 and a DRD must be coupled to a DRC path of 4 or 5 from CU2.

SSALTER

SSALTER displays the current 7990 hardware configuration and allows an analyst to change the status of certain elements in the 7990 configuration. SSALTER works with SSEXEC to make the specified changes to the existing 7990 configuration.

The initial 7990 configuration file is built by the SSBLD utility and stored as a permanent file (default file name is BU DT) under user index 377760₈ on the default family. This configuration file is the unit device table read by SSEXEC during initialization.

SSALTER can be run at SSEXEC initialization time or at any time after SSEXEC has been initialized. SSALTER aborts if SSEXEC is not executing.

Console Input

SSALTER can be executed only from the system console. Input from the console requires the following sequence of entries and responses:

1. Call SSALTER by entering:

```
X.SSALTER.
```

The B,O display indicates the job sequence name (jsn) of SSALTER. To the right of the entry appears the message:

```
REQUEST *K* DISPLAY
```

2. Activate the K display for that control point by entering:

```
K,jsn.
```

In this entry, jsn is the job sequence name of SSALTER as specified on the B,O display.

3. Select the desired K display by entering one of the following:

```
K.SM=id.
```

This entry selects the SM K display (figure 11-16), which shows the 7990 configuration of the storage module and DRDs complete with status information (id is the SM identifier, A through H).

```
K.CUest.
```

This entry selects the CU K display (figure 11-17), which shows the status of various paths/nodes in the specified control unit (est is the EST ordinal of the CU).

```

-----
7990 SUBSYSTEM - SSALTER
-----
          PCU      SCU      SM/DRD
          ST       ST       ST

SM      A      42U              U

DRD     0       U              U

DRD     1       U              U

-----

VALID COMMANDS ARE -
K.SMI,CUNN=S.
K.SMI,DRDN,CUNN=S.
K.+
  I = SM IDENTIFIER
  N = DRD DEVICE ADDRESS
  NN = CU EST ORDINAL
  S = STATUS - U FOR ONLINE
                F FOR OFFLINE
                M FOR MAINTENANCE
  + = DISPLAY NEXT SM

```

Figure 11-16. SM K Display

When a DRD's status is changed to online, any cartridges in the DRD are placed in service cells if space is available; otherwise, they are ejected from the storage module. A clear cartridge is then loaded into the DRD, its cartridge label and CSN are checked, and it is unloaded. The same cartridge is reloaded, its label and CSN are checked again, and the cartridge is returned to its cubicle. The service cells are then checked for cartridges. If a cartridge is found, it is either returned to the cubicle matrix or ejected from the storage module, depending on the cartridge label's contents.

When a SM's status is changed to online, the accessor is first cleared of any cartridge it may contain. The cartridge is stored in a service cell if one is empty; otherwise, the cartridge is ejected from the storage module. The accessor then picks a clear cartridge from its cubicle, and immediately puts it back. If the SM's status is changed to online after SSEXEC has been initialized, the service cells are then checked for cartridges. Any cartridges found are ejected from the storage module.

```

-----
7990 SUBSYSTEM - SSALTER
-----
                CU
                ORD  ST
                CU42 ON

CIF1          DTI0  DTO0  DIF0  DRC0  DRC1  AIF0
-----  -----  -----  -----  -----  -----  -----
DTI0U          DIF0U DIF0U DRC0U DRD0U          SMA0U
DTI1U          DIF1U DIF1U DRC1U          DRD1U
DTO0U          DTI1  DTO1  DIF1
                -----  -----  -----
DTO1U          DIF0U DIF0U DRC0U          AIF1
                -----
                DIF1U DIF1U DRC1U          SMA0U

CH
CH  ST  CIF
24  ON  1
    
```

Figure 11-17. CU K Display

MSE Operational Procedures

Next is a summary of the operational procedures used in an MSE environment. These procedures include modifications to existing procedures and new MSE procedures, especially in the areas of permanent file backup and reloading and disk space management. Refer to the descriptions of the appropriate utilities for details about these operational procedures.

Initialization

MSE processing is initiated when the SSDEF utility is run to create the system files necessary to allow storage of files on 7990. SSDEF creates both an SM map for each SM that is to store file data on cartridges and eight 7990 catalogs (one for each subfamily) for each family that is to have 7990-resident files. Once these system files have been created, the SSLABEL utility can then be run to allocate portions of the SM to each subfamily and to initialize cartridges for use by files belonging to these subfamilies. Note that SSEXEC need not be executing when SSDEF is run, but SSEXEC must be executing when SSLABEL is run.

When permanent files are defined, they reside on disk until the SSMOVE utility is run. Depending on the parameters specified, SSMOVE can destage files to 7990 and release the disk space assigned to files so that it is available for general use.

Disk Space Management

The availability of disk space must be managed more frequently in an MSE environment than in a non-MSE environment. This is necessary because 7990-resident files must be staged to disk before a user can access them, and their disk space is not immediately released upon completion of the jobs that attached the files. Consequently, disks are apt to approach a disk-full condition unless the SSMOVE utility is used periodically to release disk space.

The SSMOVE utility destages selected files to 7990 and optionally releases their disk space. It is recommended that the site analyst use SSMOVE as part of a periodic disk space management procedure and also just prior to a full dump to reduce the amount of data dumped to tape. (Refer to Dump Tape Management later in this section.)

The E,M display can be used to monitor the tracks on a device or set of devices. When too few tracks (as defined by the installation) are available on the direct access devices for a family or subfamily, the SSMOVE utility should be run. SSMOVE reads the PFC entries for the specified family, selects certain files to reside on 7990, and decides which files should be released from disk.

Files can be destaged to 7990 without their disk space being released. This destage-only feature can be used prior to a routine backup dump to create 7990 images for files with the BR=MD attribute and thereby eliminate the need to include them on a backup dump tape. (Refer to Dump Tape Management later in this section.) Having a 7990 image also provides backup for the file in special cases where the permanent file is being updated at the same time PFDUMP is being run.

The destage-only feature can also be used to allow disk space to become available immediately when SSMOVE is run again. That is, SSMOVE can directly release disk space of files that have both disk and 7990 images but are to reside on 7990 only. However, for files that have no 7990 image, there is a delay before the disk space can be released because these files must first be destaged to 7990.

Permanent file PFGMSE1 contains sample procedures FREEDSK and DESTAGE to show how a site may use SSMOVE for disk space management. FREEDSK waits in the rollout queue for a low disk space threshold flag to be set. FREEDSK is then rolled in for execution and calls SSMOVE to release disk space. The DESTAGE procedure calls SSMOVE at set times to destage files to 7990 whose file size in PRUs falls within the range specified by the MN=min and MX=max parameters. Each site can change the SSMOVE parameters and directives in the sample procedures to satisfy their requirements.

NOTE

Before using the FREEDSK and DESTAGE procedures, you must call the SSCONTL procedure to build the programs necessary to monitor the low disk space threshold and the time.

7990 Space Management

The availability of 7990 space must be managed because 7990 space is not automatically released when a permanent file is purged. This allows you to perform a permanent file recovery to a time prior to the purging of the file. On a typical dump tape, the entry for a 7990-resident file includes the PFC and permit entries, but not the file data. A subsequent reload operation recovers only what was dumped; the PFC and permit information, but not the file data. Thus, 7990 space for a purged permanent file cannot be released until you no longer need to recover the file data, because the file data resides on 7990 only and typically not on any dump tape.

A 7990-resident file is stored only on the cartridges within the cubicles of SMs assigned to a subfamily. When destaging a file to 7990, SSEXEC selects the SM with the most space available. If no SM has enough free space to store the file being destaged, SSEXEC abandons that destage request, issues a message to the dayfile, and continues to destage other files that fit on the available 7990 space. To avoid this problem, use the SSVAl utility periodically to release 7990 space for purged files, thereby making it available for reuse.

When you want to reuse 7990 space currently allocated to files that have been purged, he/she runs SSVAl with the RF and RL parameters specified. The RF parameter specifies an RDF, which is an auxiliary file produced by the permanent file utility PFDUMP and which contains versions of the 7990 catalogs and PFC entries that were current at the time of the dump. SSVAl examines the RDF to identify unneeded 7990 files (those without PFC entries) and their 7990 space is released. Choose which RDF to use based on the date and time of the dump that produced it. That is, a particular RDF is used only if there is no need for a permanent file recovery to a time prior to the dump that produced the RDF.

Once an RDF is used by SSVAl to release 7990 space, you should discard it and any other RDFs for the same set of devices. These RDFs cannot be used as input to SSVAl in any future run. The MSE software does not permit an RDF to be specified by the RF parameter on an SSVAl call, unless it was produced by PFDUMP after the last SSVAl run that released 7990 space was completed. This prevents the inadvertent release of 7990 space that was previously released and reassigned to a file.

Example 1:

Assume that two PFDUMPs were run a week apart producing two release data files, RDF1 and RDF2. Assume, also, that according to both RDF1 and RDF2, file AFILE is an orphan (it has been purged but still resides on 7990). If the following call to SSVAl is made, the 7990 space assigned to AFILE is released and made available for reuse, if no errors were encountered.

```
SSVAL , RF=RDF1 , RL .
```

Now, suppose that a new file, BFILE, which did not exist when RDF1 and RDF2 were produced, is assigned to the 7990 space on which AFILE has resided. If SSVAl were to accept either RDF1 or RDF2 on a subsequent SSVAl call, the 7990 space allocated to file AFILE (which is now the space where BFILE resides) is released. However, any RDF created after the SSVAl run that purged AFILE would not have an entry for AFILE and, hence, SSVAl would not inadvertently release the 7990 space presently allocated to BFILE.

After SSVAl uses a release data file to release 7990 space, the corresponding backup tape from PFDUMP is the oldest tape that should be used for the first reel of a future incremental load. The first reel determines which files are to be reloaded according to the CIR on the tape (refer to PFLOAD in section 17, Permanent File Utilities). Use of an older tape as the first reel can allow the PFC entry for a purged file to be restored even though its 7990 space may have been reassigned. The MSE software does not prevent loading from the wrong dump tape, but does detect as an error an attempt to stage a file whose data was overwritten (see Example 2).

Example 2:

Assume that RDF1 and RDF2 are as in the previous example, and that file CFILE is an orphan according to RDF2, but not according to RDF1. The following SSVAl call releases 7990 space allocated to CFILE.

```
SSVAL , RF=RDF2 , RL .
```

Now suppose that the dump tape created by the PFDUMP run that produced RDF1 is used as the first reel of an incremental PFLOAD. Then the PFC and permit entries for file CFILE are reloaded, but, in fact, file CFILE does not exist on 7990. This will be detected as an error when a user attempts to attach CFILE.

Backup and Recovery

Permanent file backup and recovery methods in an MSE environment are similar to those in a non-MSE environment. It is assumed that most sites with disk-resident permanent files have a file backup procedure consisting of a full dump followed by a series of incremental dumps. The assumed reload procedure is to load from tapes starting with the most recent incremental dump tape and continuing through older incremental dump tapes until the most recent full dump tape is processed. These procedures can be modified as follows to accommodate 7990-resident files.

File Dumping

The normal backup dump procedure is to dump files based on the modification date field in the PFC entries for the files. New parameters to PFDUMP provide the capability to dump files that have been modified during a certain interval, which is useful in recycling old dump tapes (refer to Dump Tape Management later in this section). For files that reside only on 7990, you can either stage the file to disk and thereby include it in the dump or suppress staging and include only the file's PFC and permit information on the dump tape. If the device or family being dumped has 7990-resident files, the following procedures should be included with the normal dumping procedures.

Full Dumps

Full dumps should be performed as usual, but OP=S should be specified on the call to PFDUMP to suppress staging. Thus, all files on disk are dumped but only the PFC and permit entries for 7990 files without disk images are included in the dump. It is unnecessary and not feasible to dump the files that reside only on 7990. Specify OP=S and SD on the call to PFDUMP; use the RD=filename parameter at your own discretion.

Parameter	Description
OP=S	Selects no staging. If the SD option was inadvertently used on any previous nonincremental dump, an incremental dump should be taken prior to the full dump.
SD	Updates the dump date/time field on each master device dumped to permit files created or modified prior to this dump to be released when SSMOVE is run again. Refer to Destage and Release Algorithms earlier in this section.
RD=filename	Specifies that PFDUMP is to write a release data file to be used later by SSVAL. RD=filename is included or excluded on a call to PFDUMP depending on how often you want to purge unneeded 7990 files.

Incremental Dumps

Incremental dumps should be performed as usual. No staging should occur because any files that were created or modified since the last dump will still be on disk if a tape backup is required. The following parameters should be specified on the call to PFDUMP.

Parameter	Description
OP=M	Selects incremental dumps.
SD	Updates the date/time field on each master device dumped, as previously described for full dumps.
AD=yymmdd and AT=hhmmss	Selects the date just prior to the previous incremental or full dump of the device. Hence, all files not included in the last dump are dumped.

The following parameters should be used with caution.

Parameter	Description
OP=S	Selects no staging. If the SD option was inadvertently used on any recent partial dump, OP=S should not be specified so any file with a BR=Y attribute that had its disk space inadvertently released will be included in the dump.
RD=filename	Specifies that PFDUMP is to write a release data file to be used later by SSVAl, as previously described for full dumps.

Dump Tape Management

In a non-MSE environment, it is possible to make all previous dump tapes available for use after a full dump is completed. However, in an MSE environment, such dump tape recycling cannot occur because files that have resided only on 7990 for some time are probably not included on any recent dump tape. When a low percentage of files on a tape are active files, it is possible to redump any remaining files in order to recycle the tape by running PFDUMP and specifying OP=M, BD, and BT parameters, but not specifying the OP=S parameter.⁴ This performs a dump of all files that were last dumped before the date and time specified on the PFDUMP call. If necessary, files will be staged in from 7990; after the dump is completed, the disk space for these staged files is released. Thus, files that might not be on any dump tape made after the date specified by the BD parameter are dumped and tapes created before this date can be recycled.

4. It is possible to specify the AD and AT parameters also, thereby dumping files that were dumped during an interval after the date and time specified by AD and AT and before the date and time specified by BD and BT (refer to parameters for PFDUMP). This performs a dump of all files last dumped during the interval specified (including 7990-resident files) and allows your site to recycle all dump tapes created during this interval.

The number of dump tapes created can be reduced by using SSMOVE to reduce the amount of data written on these dump tapes. If SSMOVE is run prior to an incremental dump, files with the BR=MD attribute can be destaged to 7990 and excluded from the dump. That is, the incremental dump will exclude the files with a BR=MD attribute that are on both 7990 and disk. If SSMOVE is run prior to a full dump, disk space for files that do not have to be on disk can be released and these files are not written on the dump tape. The following procedure involving dumps and SSMOVE runs can be used for efficiency.

1. Run SSMOVE to destage files with the BR=MD attribute.
2. Take an incremental dump to back up files with the BR=Y attribute.
3. Run SSMOVE to release the disk images of files backed up in step 2.
4. Take a full dump to reduce the number of dump tapes involved if a reload is required.

SM Map/7990 Catalog Backup

An SM map is created for each SM in an MSE environment and is updated by the SSLABEL, SSVAl, and SSDEBUG utilities. (Refer to the definition of SM map earlier in this section.) It is recommended that the SM map be copied on tape or on another device or family immediately after every update of the SM map. PFDUMP can be used to do this by dumping user index 377760₈ with FM equal the default system family name, but only if SSEXEC is not running. When SSEXEC is executing, it has the SM map files attached in modify mode; hence, PFDUMP will not dump the SM maps.

If the device containing subfamily 0 of the default family is reloaded, the latest SM maps should be recovered from the backup copies. After recovering the SM maps, run the SSVAl utility with the AM and FM parameters specified to ensure that the SM map and 7990 catalog entries match. If they do not, corrective action should be taken as described in Error Conditions and Corrective Actions later in this section. Failure to retain or recover the latest SM maps can cause problems with future SSLABEL runs (refer to the definition of SM map earlier in this section).

A 7990 catalog is created for each subfamily of a family that can have 7990-resident files and is updated by the SSLABEL, SSMOVE, SSVAl, and SSDEBUG utilities. When PFDUMP copies permanent files to tape for a master device or all devices in a family, it ensures that corresponding 7990 catalogs are also dumped. (If SSEXEC has a 7990 catalog attached in modify mode, PFDUMP interfaces with SSEXEC so that the 7990 catalog is included on the dump tape. While PFDUMP is accessing the 7990 catalog, any file destage requests for that subfamily are delayed.) Consequently, if PFLOAD is used to reload all files, the 7990 catalogs are automatically recovered and entries for all 7990-resident files reflected in recovered PFC entries will exist in the recovered 7990 catalog.

If PFLOAD is run after the SSLABEL utility has updated the 7990 catalog and SM map but before the 7990 catalog has been backed up, the recovered 7990 catalog will not match exactly the corresponding SM map. In this case, recovery procedures (as described in 7990 Catalog/SM Map Mismatch later in this section) must be taken to bring the SM map, 7990 catalog, and cartridge label information into agreement.

File Reloading

PFLOAD runs should be made as usual. PFLOAD reads the CIR records from the most recent incremental dump tape to control loading of file data and PFC and permit information. For 7990-resident only files, PFLOAD reloads only the PFC and permit information (if PFDUMP dumped only this information). The following situations involve PFLOAD in an MSE environment.

- If one device needs to be reloaded, the DN or TD parameter specifies the device. The OP=Z and UD parameters should not be specified.
- If PFLOAD is used for other than a device recovery, both OP=Z and UD should be specified. The OP=Z parameter ensures that the PFC entry does not link to a 7990 file that may no longer exist, and UD ensures that the file is included in the next incremental or full dump.
- If a file has a nonzero asa value in its PFC entry (indicating that the file has a 7990 image), the SSCVER flag in the PFC entry will be set when the PFC and permit information are reloaded. If the file is successfully staged to disk when it is attached, this flag is cleared. Otherwise, the flag remains set. This can be used to indicate the possibility of an operational error, such as omitting OP=Z when a file or set of files is reloaded from an obsolete dump tape.
- If any unscheduled device reload that includes recovery of a 7990 catalog and/or SM map is made, SSVAl should be run with the AM and FM parameters specified. SSVAl examines the PFC entries for 7990-resident files, the 7990 catalog, and the SM map to detect any error conditions that exist.
- If dumping a file to tape from one user and then reloading the file to a different user, you must use the OP=Z parameter on the PFLOAD utility. Otherwise the PFC entry still points to the old SFMCAT, which causes an error indication for the file.

Error Conditions and Corrective Actions

It is possible for operational, hardware, or software malfunctions to cause a variety of error conditions to exist or appear to exist in the 7990 catalogs, SM maps, PFC entries, or cartridge labels. The SSVAl utility can analyze the 7990 catalogs, SM maps, and PFC entries to detect these errors, and the SSLABEL utility can identify problems with cartridge labels. The following paragraphs describe some specific error conditions and the actions that should be taken to correct them. Most of the errors are caused either by incorrect operational procedures that result in the recovery of the wrong version of an MSE system file, or by abnormal termination of SSEXEC or MSE utility processing, which can interrupt the updating of an MSE system file.

Clearing Permanent File Error Flags

The permanent file error flags that have been set by SSEXEC in the PFC can be cleared by the CE parameter of the CHANGE command. The job must be of system origin.

Permanent File Recovery

If a permanent file cannot be accessed or is included by SSVAl on the validation report, try the following procedures to recover a copy of the file.

1. Attach the file. If the file has a disk image or can be staged in from 7990, the file data will be available on disk and can be copied to another file. SSVAl should then be run as described in step 2.
2. If the file cannot be attached, reload the file from a dump tape using PFDUMP with OP=Z. To identify the PFDUMP run that saved a backup copy of the file, check the output from SSVAl or PFCAT, which includes date and time fields from the PFC entry for the file. Depending on the BR attribute (backup requirement) of the file, a backup copy on a dump tape may not exist.

If the file can be reloaded to disk, SSVAl should be run as follows:

```
SSVAL, FM=familyname, SB=subfamily, FX=n, AM.
```

familyname and subfamily identify the family and subfamily of the recovered file and n is sufficiently large that the fixing/flagging will be done. SSVAl sets the obsolete flag in the PFC entry because the file now has a disk image; thus the file is disassociated from the error condition. The error condition may still exist in the 7990 catalog or SM map, but you can analyze it and correct the error by using the SSDEBUG utility. Before using SSDEBUG to correct the error, you should take an incremental dump so the old PFC entry with an asa value pointing to a problem 7990-resident file will not be reloaded if a device reload occurs before the next scheduled backup dump.

3. If steps 1 and 2 fail, try to regenerate the file. Depending on the particular problem and status of the cartridges to which the file was destaged, the OP=RF or OP=RS directive to SSDEBUG might recover the file data.
4. If the file cannot be recovered, purge it so there is no PFC entry pointing to a problem 7990-resident file.

Cartridge Restoration and Reuse

If a cartridge is unexpectedly found anywhere outside its assigned cubicle, follow these steps to restore the cartridge so it can be used again:

1. Put the cartridge in the input tray of the SM to which it is assumed to belong and run SSLABEL with the OP=RC directive. If the cartridge label information agrees with the SM map entry for the cubicle with the coordinates (y,z) written on the cartridge label, the cartridge will be restored to that cubicle and the 7990 catalog entry for the cartridge is updated to clear the lost flag. Otherwise, the cartridge is put in the exit tray and the cartridge label information is included in the output from SSLABEL.
2. If the cartridge was not restored to its SM cubicle, examine the label information reported by SSLABEL. If the cartridge is assigned to a different SM, repeat step 1 using the correct SM. If the family name is blank, the cartridge has a scratch label and can be assigned normally with the OP=AM directive to SSLABEL.
3. If steps 1 and 2 fail, examine the SM map entry for the cubicle with the coordinates (y,z) written on the cartridge label to see how the cubicle is assigned. Examine the cartridge, if any, in this cubicle. Also, it might be useful to run SSVAL to ensure the SM map and 7990 catalog entries for the cartridge match or to identify affected permanent files, if any, if they do not match.
4. If step 3 does not produce any useful information and the cartridge cannot be restored, use the OP=FX directive to SSLABEL to rewrite the cartridge label and thereby make a scratch cartridge. It might be useful to first run SSDEBUG with the OP=RS directive to preserve the data on the cartridge.

Removal of Faulty or Missing Cartridges

Follow these steps to remove a cartridge that has unrecoverable errors, or to release the SM map entry for a cartridge that is missing.

For unallocated cartridges:

1. Attempt to remove the cartridge with the SSLABEL OP=RM directive.
2. If CARTRIDGE NOT FOUND is reported, repeat step 1 specifying the LT parameter.
3. If CARTRIDGE NOT EMPTY is reported, at least one AU on the cartridge is still allocated.

For allocated cartridges:

1. Use the SSLABEL OP=FC directive to inhibit further allocation of the cartridge.
2. Run SSVAL specifying the AM parameter. Examine the report for possible discrepancies between the SM map, 7990 catalog, and PFC entries related to this cartridge.
3. If no 7990 catalog entry corresponds to the SM map entry, use the OP=RC directive to SSDEBUG to remove the map entry.
4. If no discrepancy is reported by SSVAL, the affected files should be attached in write mode (if possible), reloaded using PFLOAD with the OP=Z parameter, or purged. Then SSVAL release processing should be done to release the 7990 space on the cartridge, and the cartridge can be removed with the SSLABEL OP=RM directive.

7990 Catalog/SM Map/PFC Problems

If error conditions exist in the 7990 catalogs or SM maps, the recovery procedure is to correct the appropriate entry rather than reload the entire file. Reloading other than the latest copy of the SM map is not feasible because it can contain entries for several families and these entries would not be up to date. Reloading an old copy of the 7990 catalog is not desirable because any permanent file that was destaged and released from disk since the last incremental dump will be lost. The following paragraphs describe corrective action for specific problems with 7990 catalogs, SM maps, or PFC entries.

7990 Catalog/SM Map Mismatch

The following examples of 7990 catalog/SM map mismatches can appear as error type 1 or 2 on the SSVAL output report, which identifies the mismatch and the permanent files or orphans affected. Before taking any corrective action, the analyst should ensure that operational errors did not result in loading the wrong version of the SM map or 7990 catalog.

- If a cartridge has a 7990 catalog entry indicating that all AUs on the cartridge are available for allocation but is not identified in any SM map entry (type 1 mismatch), use the OP=RL directive to SSDEBUG to remove the 7990 catalog entry for the cartridge.
- If a cartridge is identified in an SM map entry but has no 7990 catalog entry (type 2 mismatch), use the OP=RC directive to SSDEBUG to update the SM map entry to unassigned and empty status. This directive will write a scratch label for the cartridge, if any, and place the cartridge in the exit tray.
- If a cartridge has a 7990 catalog entry indicating that at least one AU on the cartridge is allocated, recover the disk space for the affected permanent file, if any, or purge the file. Run SSVAL to purge unneeded 7990 files and thereby make the cartridge unallocated. Use the OP=RL directive to SSDEBUG to remove the 7990 catalog entry for the cartridge.

NOTE

It sometimes takes two SSVAL/SSDEBUG runs to completely clear up type 1 or type 2 errors. Therefore, it is recommended that SSVAL be rerun with the AM parameter after an SSDEBUG run using OP=RC or OP=RL.

7990 Catalog Chain Problems

AUs that are allocated to a 7990-resident file are chained together in the 7990 catalog to identify the sequence of AUs that contain the file. (Refer to the definition of cartridge earlier in this section.) For most files, the AUs in the chain reside on the same cartridge. Hence, one disk access is sufficient to update information for the entire chain at one time. However, a 7990-resident file can reside on several cartridges, and several disk accesses may be necessary to update its chain. If SSEXEC is interrupted during such an update, the resultant chain will be incomplete. SSVAl identifies incomplete chains as error type 4, 5, or 6 on the validation report. The following are types of incomplete chains.

- If an incomplete chain is produced while a 7990-resident file is being purged, it is a chain fragment (one without an HOC entry).
- If an incomplete chain is produced while a file is being destaged to 7990, it is an ill-formed chain (one without an EOC entry).
- If an ill-formed chain is not corrected, the AU that should have been linked to it may be reused and become part of a good chain. Although the ill-formed chain may appear to be complete, the ill-formed chain and the good chain are intersecting.

The corrective action for any of these bad chain conditions is:

1. Run SSVAl with the FM parameter specified and use a sufficiently large FX=n value so that SSVAl updates the 7990 catalog entries affected by the problem chains.
2. Recover the permanent files, if any, that have 7990 images potentially affected by these error conditions. (The SSVAl report identifies these files.) Use the procedure described under Permanent File Recovery earlier in this section.
3. Take an incremental or full backup dump. This prevents PFC entries with inaccurate asa values or problem 7990 chains from being recovered if a device reload is necessary.
4. Run SSDEBUG with the OP=RP directive to make each AU on the problem chain available for allocation. Thus, the AUs can be included in a new chain without any problems. In the case of intersecting chains, all chains that have an AU in common with another chain should be released before any further destage attempts; otherwise, errors will occur when a user tries to attach a file that is contained on an intersecting chain.

PFC/7990 Catalog Mismatch

It is possible that an old version of a 7990 catalog can be reloaded because of incorrect operational procedures. If so, the asa value in the PFC entry for a file can be invalid according to the obsolete 7990 catalog or can point to a problem chain. If the correct version of the 7990 catalog is available on tape, it should be reloaded. If it is not available, the following action is appropriate. If the file also has a disk image, SSVAl should be run to clear the asa field in the PFC entry. If there is no disk image, the analyst should reload the file from tape, if possible, and then run SSVAl to clear the asa value.

Tape and Optical Disk Alternate Storage 12

Introduction	12-1
Terminology	12-1
Organization and Management of Destage Dumps	12-3
Tape Characteristics	12-3
Tape Data Format	12-3
Destaging Files to Tape	12-3
Tape Alternate Storage in a TMS Environment	12-5
Optical Disk Characteristics	12-6
Optical Disk Data Format	12-6
Destaging Files to Optical Disk	12-7
Releasing Disk Space	12-8
Restoring Destaged Files to Disk Residence	12-9
Modifying Destaged Files	12-10
Recycling Destage Dump Tapes	12-10
MAGNET Command	12-11
GENPFD Utility	12-11
Calling the GENPFD Utility	12-12
GENPFD Input Directives	12-13
File Selection Processing	12-16
Limitations	12-16

Tape and Optical Disk Alternate Storage 12

Introduction

This section describes the use of magnetic tape and optical disk as alternate storage media for NOS permanent files.

Terminology

The following terms apply to the use of tape and optical disk alternate storage.

Alternate Storage

The storage of permanent file data on external media other than mass storage. The storage media can be tape, MSE, or optical disk. When a file resides on alternate storage, the file's permanent file catalog (PFC) entry and permit data still reside on disk, but the disk space occupied by the file data can be released.

Alternate Storage Address (ASA)

The pointers in a file's PFC entry that point to the location of the file data on tape, MSE, or optical disk alternate storage.

Destage Dump Optical Disk Volumes

Optical disk volumes that contain files that were destaged from disk using both the REQOD command and the PFDUMP utility with the DT parameter specified.

Destage Dump Tapes

Tapes that contain files that were destaged from disk using the both the LABEL command and the PFDUMP utility with the DT parameter specified.

Destaging

The process of creating an alternate storage copy of a file on tape, MSE, or optical disk. Files are destaged to tape alternate storage when the site executes PFDUMP with the DT parameter specified and the file specified by the T parameter is assigned to a magnetic tape. Files are destaged to optical disk when the site executes PFDUMP with the DT parameter specified and the file specified by the T parameter is assigned to optical disk. Files are destaged to MSE when the site executes SSMOVE with the appropriate parameters.

Permanent File Catalog (PFC)

The system catalog containing the file name, owner identification, disk pointers, alternate storage pointers, and other attributes of a permanent file. Refer to Permanent File Catalog Entry in section 17, Permanent File Utilities, for the format of the PFC entry.

Primary VSN

Either the volume serial number (VSN) of a single-volume destage dump or the VSN of the first volume of a multivolume set of destage dumps. The VSN must be two alphanumeric characters followed by a four-digit decimal number; the number must be in the range 0000 through 4095.

Secondary VSN

Either the VSN of a single-volume destage dump or the VSN of the first volume of a multivolume set of destage dumps having a sequence number (the last four characters of the VSN) in the range 5000 through 9095. The existence of a secondary VSN assumes the existence of a primary VSN having the same two-character prefix and a sequence number of 5000 less. The primary and secondary VSNs, or multivolume sets, are assumed to contain identical data although the individual volumes of multivolume sets cannot be assumed to be identical.

Staging

The process of restoring file data to disk residence from a copy on alternate storage. Staging is initiated when a user executes an APPEND, ATTACH, GET, or OLD command for a file that is not currently disk resident.

Organization and Management of Destage Dumps

Destage dump tapes and destage dump optical disk volumes have different media characteristics and data formats. In addition, you use different procedures to manage destage dumps tapes and destage dump optical disk volumes. The destage dump tapes are described first.

Tape Characteristics

Destage dump tapes must be labeled 9-track tapes, ACS tapes, or CTS tapes with a tape format of either I or LI. Nine-track tapes should be written with a GE tape density. Although it is possible to have multivolume sets of tapes, create single volumes whenever possible to avoid the overhead of volume switching when restoring files to disk.

The recorded VSN of the first tape of a set of destage dump tapes must be in the following format:

ppnnnn

where pp is a prefix consisting of any two alphanumeric characters and nnnn is a decimal sequence number in the range from 0000 through 4095 or in the range from 5000 through 9095. If the sequence number is 5000 or greater, the secondary tape set is assumed to be a duplicate of a primary tape set that has the same prefix and a sequence number of 5000 less.

Tape Data Format

All destage dump tapes are written by PFDUMP and have exactly the same format as other tapes written by PFDUMP including file permits. However, no PFC-only files are present on destage dump tapes. You can also load, copy, or list files from a destage dump tape using the PFLOAD, PFCOPY, or PFATC utility, respectively.

Destaging Files to Tape

Destage dump tapes are created by executing PFDUMP with the DT (destage to tape) parameter specified. When the DT parameter is used, PFDUMP verifies the following:

- The dump file specified by the T parameter must be assigned to a labeled 9-track tape, ACS tape, or CTS tape that is recorded in either I or LI format. The VSN must be in the format ppnnnn, as described earlier in Tape Characteristics, and the sequence number must be in the range from 0000 to 4095.
- If a verify file is specified (VF parameter), it must be assigned to a labeled tape of the same type, density, and format as the tape specified by the T parameter. Its VSN must have the following relation to the VSN for the primary dump tape: the two-character prefix must be the same and the sequence number must be greater by 5000.

After the dump tape or tapes are successfully written, PFDUMP enters the primary dump tape VSN in the PFC entries of the dumped files. If a verify file is written, a flag is also set in the PFC entry indicating the existence of the secondary VSN. At this point, the disk space of the dumped files can be released, if necessary, as long as the file's backup requirements are satisfied.

Although general PFDUMP selection parameters can be used to directly select files for a destage dump, sites will normally use a separate file selection utility such as GENPFD to perform this function. (The GENPFD utility is described later in this section.) This procedure has the following advantages:

- Files can be prioritized by the time since they were last accessed or modified.
- Special user index or file name patterns can be special-cased according to site requirements.
- The GENPFD selection utility can select files that will reasonably fill a single tape volume without overflowing to additional volumes. The creation of single volume tapes is desirable because it avoids the overhead of volume switching when files are restored to disk.

When a file selection program such as GENPFD is used for destage processing, the following sequence of events should be followed:

1. Execute PFDUMP with the IP (inhibit processing) and DT (destage to tape) parameters specified and the summary file selected with the S parameter. Specifying PFDUMP parameters such as LS, US, BD, BT, etc., can be used to screen out desired classes of files. The TA=N and CA=N parameters can be used to exclude files that already have current alternate storage copies (from previous destage dumps).
2. Execute the file selection program (GENPFD) using the PFDUMP summary file as input. The file selection utility writes UI and PF utility directives for the selected files to the utility directives file.
3. Reexecute PFDUMP (with the same parameters as specified in step 1, but without specifying the IP parameter) using the utility directives file generated in step 2 as an input file.

When a destage dump is performed, files on user name SYSTEMX and on the MSE catalog user names SUBFAM0 through SUBFAM7 are unconditionally excluded even if they meet the selection criteria.

Example:

Destage enough files to fill one tape (alternate storage) with files from family TEST whose data has not been modified since 87/07/01 and which have not been destaged already. Create two copies of the dump tape on VSNs AA1234 and AA6234.

```
RETURN, SUMMARY.  
PFDUMP, FM=TEST, DT, OP=T, BD=870701, TA=N, IP, S.  
REWIND, SUMMARY.  
RETURN, UDIR.  
GENPFD, PO=D.  
REWIND, UDIR.  
LABEL, TAPE, VSN=AA1234, D=GE, W.  
LABEL, PFVER, VSN=AA6234, D=GE, W.  
PFDUMP, FM=TEST, DT, OP=T, BD=870701, TA=N, I=UDIR, VF.
```

In this destaging example, if GENPFD does not find enough files to fill a 2400-foot tape at GE density, the utility does not generate a UDIR file. This causes the second PFDUMP to abort without dumping any files.

Tape Alternate Storage in a TMS Environment

When a site is using Control Data's Tape Management System, special handling is required for permanent file destaging operations. For more information about TMS operations, see the NOS2 Tape Management System Site Operations Manual. The following information is provided as a brief background pertaining specifically to TAS.

Normally when tapes are written under TMS, they are automatically assigned as requested from a pool of scratch tapes maintained for that purpose. Because a TAS destage VSN must take the form *ppnnnn*, do not assign a destage dump tape from the TMS scratch pool because there is no way to ensure that an appropriate VSN is assigned, for either the primary or secondary VSN. Therefore, it is necessary to maintain (separately from the pool of scratch tapes) a pool of tapes with VSNs that conform to the TAS format.

To form a separate pool of TAS tapes, use the `SYSTEM=YES` directive when adding TAS tapes to the TMS database through TFSP. TMS does not allow VSNs with `SYSTEM` status to be part of the scratch pool.

Once a TAS pool of tapes is established, TMS offers two tools that work together as one to aid in automated management of TAS VSNs at a site: the `GETTASV` command and the `SETTASV` procedure. `GETTASV` selects TAS VSNs from the site's pool of TMS `SYSTEM` VSNs, then calls `SETTASV` to reserve and request the selected tapes.

The following example shows how `GETTASV` and `SETTASV` might be used in a destaging job to automate the process of assigning VSNs for destage dump tapes from the TMS database. In this example, cartridge tapes are used for destaging of files that belong to a particular family and have not been accessed since July 1, 1990.

<code>USER, usernam, passwd, family.</code>	User name, family are required.
<code>RESOURCE, CE=2.</code>	Ensure adequate resources.
<code>PFDUMP, IP, S, BD=900701, OP=A, FM=family, DT, TA=N.</code> <code>REWIND, SUMMARY.</code> <code>RETURN, UDIR.</code>	Select candidate files.
<code>GENPFD, PO=D.</code>	Select actual files to be destaged.
<code>IF, FILE(UDIR, .NOT.AS) .EXIT.</code>	Exit if no files were selected.
<code>TFSP, FM=family, SS=VSNS, OP=Z. /MREADVS=</code> <code>REWIND, VSNS.</code>	Select candidate VSNs.
<code>GETTASV, SS=VSNS, P=UR, N=2, D=CE, F=LI.</code>	Select TAS <code>SYSTEM</code> VSNs.
<code>*(GETTASV CALLS SETTASV AT THIS POINT)</code> <code>REWIND, UDIR.</code>	Reserve and request VSNs.
<code>PFDUMP, FM=family, I=UDIR, DT, VF.</code> <code>UNLOAD, TAPE, PFVER.</code>	Destage the files.

Optical Disk Characteristics

Although optical disks used as alternate storage can be multivolume, create single volumes whenever possible to avoid the overhead of volume switching when restoring files to disk.

The recorded VSN of the first volume of a set must be in the following format:

ppnnnn

where pp is a prefix consisting of any two alphanumeric characters and nnnn is a decimal sequence number in the range either from 0000 through 4095 or from 5000 through 9095. If the sequence number is 5000 or greater, the secondary VSN is considered to be a duplicate of a primary VSN that has the same prefix and a sequence number 5000 less.

The destage dump set must have the following characteristics specified on the REQOD command:

- The recorded file name must be PFARCHIVE (RFN=PFARCHIVE).
- The owner identification or user name must be SYSTEMX (FOI=SYSTEMX).
- The group identification or family name (FGI=name) must be the same as the family name you specify for the FM parameter on the PFDUMP command.

The REQOD command is described in NOS Optical Storage Subsystem QSE32835 Reference Manual.

Optical Disk Data Format

All destage dump optical disk volumes are written by PFDUMP and have exactly the same format as other dumps written by PFDUMP including file permits. However, no PFC-only files are present on destage dumps. You can also load, copy, or list files from a destage dump optical disk volume using the PFLOAD, PFCOPY, or PFATC utility.

Destaging Files to Optical Disk

The procedure for destaging files to optical disk is identical to that for destaging files to tape, except that:

- Rather than using the LABEL command to assign file TAPE to a magnetic tape, you use the REQOD command to assign file TAPE to an optical disk. The REQOD command and other aspects of optical disk usage are described in NOS Optical Storage Subsystem QSE32835 Reference Manual.
- You specify the U=MB parameter on the GENPFD command to calculate the destage file size in megabytes.

Refer to Destaging Files to Tape earlier in the section for more information on the destaging procedure.

Example:

Destage enough files to fill 102 megabytes of optical disk alternate storage with files from family TEST whose data has not been modified since 91/07/01 and which have not already been destaged.

```

RETURN, SUMMARY, UDIR.
PFDUMP, FM=TEST, DT, OP=T, BD=910701, CA=N, TA=N, IP, S.
REWIND, SUMMARY.
GENPFD, PO=D, U=MB.
REWIND, UDIR.
REQOD, TAPE, VSN=AB1234, RFN=PFARCHIVE, FOI=SYSTEMX, FGI=TEST.
PFDUMP, FM=TEST, DT, OP=T, BD=910701, CA=N, TA=N, I=UDIR.

```

As required to destage to optical disk, the FGI parameter of the REQOD command matches the FM parameter of the PFDUMP command, and both RFN=PFARCHIVE and FOI=SYSTEMX are specified for the REQOD command.

Releasing Disk Space

The disk space of a file that resides on alternate storage can be released if the file's backup requirements are met. The BR parameter on the DEFINE, SAVE, and CHANGE commands allows a user to specify a backup requirement of Y (yes), MD (media dependent), or N (none).

Parameter

Option	Description
BR=Y	<p>The disk space of a file with a backup requirement of BR=Y can be released if either of the following conditions is satisfied:</p> <ul style="list-style-type: none"> • At least two copies of the file data exist on alternate storage media. This could be one copy on either MSE or optical disk and one copy on tape or two copies (primary and secondary VSNs) on tape. • One copy of the file data exists on an alternate storage medium (MSE, tape, or optical disk) and the release inhibit date/time of either the file's master device or the file's residency device (as set by a PFDUMP with the SD parameter specified) is greater than the file's utility control date/time.
BR=MD	<p>The disk space of a file with a backup requirement of BR=MD can be released if the following condition is satisfied:</p> <ul style="list-style-type: none"> • One copy of the file data exists on MSE, tape, or optical disk alternate storage.

The PFREL utility is used to perform the releasing of disk space. As with PFDUMP, PFREL can be used with a file selection utility such as GENPFD in the following sequence of steps:

1. Execute PFREL with the IP (inhibit processing) parameter specified and the summary file selected with the S parameter. Files can be excluded by specifying selection parameters. In particular, OP=A can be used with the BD and BT parameters to exclude files that have been accessed recently.
2. Execute the file selection program using the PFREL summary file as input. The file selection utility writes UI and PF utility directives for the selected files to the utility directives file.
3. Reexecute PFREL (with the same parameters as specified in step 1, but without specifying the IP parameter) using the utility directives file generated in step 2 as an input file.

The PFREL utility can be used to release the disk space of files that reside on tape, MSE, or optical disk alternate storage. However, the MSE utility SSMOVE performs release processing only for MSE resident files. PFREL does not release the disk space of files cataloged on user name SYSTEMX or on the MSE catalog user names SUBFAM0 through SUBFAM7 even if the selection criteria are met.

Example 1:

The following example illustrates releasing disk space for destaged files.

Release the disk space for all files from family TEST that have not been accessed since 87/07/01 and for which backup requirements have been met.

```
PFREL,FM=TEST,OP=A,BD=870701.
```

Example 2:

Using the same file selection criteria as in example 1, release enough files so that 1000₈ tracks are available on device number 40 and 1200₈ tracks are available on device number 41.

```
RETURN, SUMMARY .
PFREL, FM=TEST, OP=A, BD=870701, S, IP.
REWIND, SUMMARY .
RETURN, UDIR .
GENPFD, PO=R, Z. /RL, 40, 1000 /RL, 41, 1200
REWIND, UDIR .
PFREL, FM=TEST, I=UDIR.
```

Restoring Destaged Files to Disk Residence

A file that does not reside on disk is restored to disk residence (staged) when a user accesses the file with an APPEND, ATTACH, GET, or OLD command. If the RT (real time) parameter is specified on an ATTACH, GET, or OLD command and staging is required, the job continues executing without waiting for the staging operation to complete. In this case, the caller must subsequently reissue the ATTACH, GET, or OLD command to access the file after the staging operation completes.

NOTE

To allow staging of files from 9-track tape and CTS tape alternate storage to disk, the site must enter the ENABLE,TAPE PF STAGING command (via DSD or IPRDECK).

To allow staging of files from ACS tapes, the site must enter the ENABLE,ACS TAPE PF STAGING command (via DSD or IPRDECK) and must have the ATF subsystem active.

In both cases, you must specify the SJ parameter on the MAGNET command in the MAG subsystem startup procedure.

Modifying Destaged Files

When you modify a file having a copy or copies on alternate storage media, the data on alternate storage is no longer current. Therefore, when a direct access file is attached in a writable mode, or an indirect access file is saved or replaced, the system removes the pointers to all alternate storage copies in the file's PFC entry. At this point, the file must be destaged again (to tape, MSE, or optical disk) if it is to reside on alternate storage.

Recycling Destage Dump Tapes

As files are modified and purged over a period of time, new destage dump tapes are created and the percentage of active files on older destage dump tapes decreases. Destage dump tape recycling consists of procedures to move the remaining active files' data on old destage dump tapes to new destage dump tapes so that the old tapes can be reused. Performing a destage PFDUMP (DT parameter) with staging enabled (OP=S not specified) of all files on a tape or set of tapes copies all file data to a new tape and changes all alternate storage references from the old tapes to the new tape.

Since only a subset of the files on tapes selected for recycling are typically active, a site normally wants to include other files to fill a new tape. The following example illustrates this procedure.

```
RETURN, SUMMARY.
PFDUMP, IP, S, DT, FM=TEST, TA=N, TA=AB0005, TA=AB0006.
REWIND, SUMMAY.
RETURN, UDIR.
GENPFD, OP=D, Z./SV, AB0005/SV, AB0006/DL, 2200/DT, 0
REWIND, UDIR.
LABEL, TAPE, VSN=AB0200, D=GE, W.
LABEL, PFVER, VSN=AB5200, D=GE, W.
PFDUMP, DT, FM=TEST, I=UDIR, VF.
```

The TA parameters on the first PFDUMP command select files residing on VSNs AB0005 and AB0006 (the tapes being recycled) and files not residing on tape alternate storage as candidates for destage. The SV directives cause GENPFD to select all files residing on VSNs AB0005 and AB0006 regardless of the modification date derived priority and priority group. GENPFD selects other files using the normal priority criteria to satisfy the destage limit specified by the DL directive. The DT (destage threshold) directive performs a recycle dump even if there are not enough files selected to fill the tape. The two LABEL commands cause two copies of the dump: one on VSN AB0200 and the other on AB5200.

To help determine which destage dump tape VSNs should be recycled, you may want to write a program that reads the PFCAT summary file and generates a list of how many files are still active on each destage dump tape VSN.

You can determine if any current files reside on a particular destage dump tape VSN by executing PFCAT with the TA=vsN parameter specified.

After you perform a recycle dump for a tape or a set of tapes, the tapes cannot be reused until a full or incremental PFDUMP has been performed on the affected devices. Without this PFDUMP, a permanent file reload could cause some of the reloaded PFC entries to refer to tapes that no longer exist. To prevent this problem, a tape recycling operation should always be immediately followed by a full or incremental PFDUMP.

MAGNET Command

The MAGNET command has two parameters that must be used in a tape alternate storage environment. The SJ parameter controls the maximum number of staging jobs that can be active concurrently to stage files to disk from tape alternate storage; and the SV parameter controls the maximum number of staging request VSNs that can be displayed on the DSD E,P display. Refer to the NOS Version 2 Operations Handbook for a description of the DSD E,P display.

The MAGNET command has the following format:

```
MAGNET , SJ=msj , SV=mvd .
```

Parameter Description

SJ=msj Specifies the maximum number of staging jobs that can be active at the same time ($0 \leq msj \leq 30_8$). The default is SJ=0; staging is not allowed.

NOTE

The SJ parameter must be specified or staging from tape alternate storage to disk will not be allowed.

SV=mvd Specifies the maximum number of requests for staging tape VSNs to display on the DSD E,P display ($1 \leq mvd \leq 77_8$). The default is SV=17₈. Requests for staging tape VSNs will appear first on the E,P display (before the normal tape requests) and can be distinguished by having a JSN of SYS and a user name of (STAGE).

GENPFD Utility

The GENPFD utility controls file selection processing. It reads a summary file created by one of the permanent file utilities and uses the information to generate file selection directives for subsequent permanent file utility operations.

Calling the GENPFD Utility

You call the GENPFD utility by executing the GENPFD command. The format of the command is:

```
GENPFD, I=input, S=summary, UD=udir, PO=opt, U=units, Z.directives
```

Parameter	Description						
I=input	The 1- to 7-character name of the file from which directives are read. The default file name is INPUT.						
S=summary	The 1- to 7-character name of the summary file previously created by a permanent file utility. The default file name is SUMMARY. This file is used as input to GENPFD.						
UD=udir	The 1- to 7-character name of the utility directives file to be created. This file is used for subsequent permanent file utility operations such as destaging files and releasing disk space. The default file name is UDIR.						
PO=opt	The processing option that specifies whether to select files for destage processing or releasing disk space. Acceptable values for opt are: <table border="1" data-bbox="451 808 1360 919"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Select files for destage processing.</td> </tr> <tr> <td>R</td> <td>Select files for releasing disk space.</td> </tr> </tbody> </table>	opt	Description	D	Select files for destage processing.	R	Select files for releasing disk space.
opt	Description						
D	Select files for destage processing.						
R	Select files for releasing disk space.						
U=units	Unit of measure for the destage limit and destage threshold directives (DL and DT). The default value is FT. Acceptable values for units are: <table border="1" data-bbox="451 1045 1360 1157"> <thead> <tr> <th>units</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>FT</td> <td>Measures in feet of tape.</td> </tr> <tr> <td>MB</td> <td>Measures in megabytes of storage.</td> </tr> </tbody> </table>	units	Description	FT	Measures in feet of tape.	MB	Measures in megabytes of storage.
units	Description						
FT	Measures in feet of tape.						
MB	Measures in megabytes of storage.						
Z	Input directives are contained on the GENPFD command following the terminator. If the Z parameter is specified, the I parameter is ignored. This eliminates the need to use a separate input file for the directives when only a few directives are needed.						
directives	When input directives appear on the GENPFD command, the first character following the command terminator is the separator character for all directives on the command. Any display code character that is not used in any of the directives, including a space, can be used as the separator character. Each directive must be preceded by a separator. The directives can extend to column 72 on the command. Continuation lines are not permitted.						

For example, (slant used for separator):

```
GENPFD, PO=D, Z. /D=PE/DL, 2200/SF, 3456
```


GENPFD Input Directives

The GENPFD input directives provide the capability to selectively choose files for processing. The GENPFD utility accepts the following input directives:

Directive	Description																		
D=density	Density of the destage dump tape that is written by the destaging operation. GENPFD uses the density in calculating the length of tape necessary for the selected files. The following table lists the acceptable values for density; the default is 6250.																		
	<table border="1"> <thead> <tr> <th>density</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AE</td> <td>38000 cpi (ACS tape)</td> </tr> <tr> <td>CE</td> <td>38000 cpi (CTS tape)</td> </tr> <tr> <td>HD</td> <td>800 cpi (9-track tape)</td> </tr> <tr> <td>PE</td> <td>1600 cpi (9-track tape)</td> </tr> <tr> <td>GE</td> <td>6250 cpi (9-track tape)</td> </tr> <tr> <td>1600</td> <td>1600 cpi (9-track tape)</td> </tr> <tr> <td>6250</td> <td>6250 cpi (9-track tape)</td> </tr> <tr> <td>38000</td> <td>38000 cpi (cartridge tape)</td> </tr> </tbody> </table>	density	Description	AE	38000 cpi (ACS tape)	CE	38000 cpi (CTS tape)	HD	800 cpi (9-track tape)	PE	1600 cpi (9-track tape)	GE	6250 cpi (9-track tape)	1600	1600 cpi (9-track tape)	6250	6250 cpi (9-track tape)	38000	38000 cpi (cartridge tape)
density	Description																		
AE	38000 cpi (ACS tape)																		
CE	38000 cpi (CTS tape)																		
HD	800 cpi (9-track tape)																		
PE	1600 cpi (9-track tape)																		
GE	6250 cpi (9-track tape)																		
1600	1600 cpi (9-track tape)																		
6250	6250 cpi (9-track tape)																		
38000	38000 cpi (cartridge tape)																		
DL,limit	<p>Destage size limit used with the PO=D option. The U=units option determines whether the limit is measured in feet of tape or megabytes.</p> <p>Specify a 1- to 10-digit number as the limit; the limit is assumed to be decimal unless a B radix is specified. This directive specifies the maximum combined size of all selected files. It is intended to prevent writing beyond the end of volume when selecting files for destaging with the PFDUMP utility. If selection of a file causes the limit to be exceeded, the file is skipped and the next lowest priority file is examined. The default limit depends on the U=units option. If U=MB, the default is 102 megabytes (an appropriate value for optical disk). If U=FT, the default varies with the D=density option. For 9-track tape densities, the default is 2200 feet, which allows for a significant number of tape errors on a 2400-foot volume of tape. For cartridge tape densities, the default is 500 feet, which allows for tape errors on a 530-foot cartridge.</p>																		

Directive	Description						
DT,threshold	<p>Destage size threshold used with the PO=D option. The U=units option determines whether the threshold is measured in feet of tape or megabytes.</p> <p>Specify a 1- to 10-digit number as the threshold; the threshold is assumed to be decimal unless a B radix is specified. This directive specifies the minimum combined size of all selected files. If the combined size of all selected files is less than the threshold, no utility directives file will be generated. This directive is used to ensure that enough files can be selected on a destage dump to reasonably fill the space available. The default value is 90 percent of the limit specified for DL.</p>						
F=format	<p>Format used for the destage dump tape. The default format is I. Acceptable values are:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Format</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>I</td> <td>Internal format.</td> </tr> <tr> <td>LI</td> <td>Long block internal format.</td> </tr> </tbody> </table>	Format	Description	I	Internal format.	LI	Long block internal format.
Format	Description						
I	Internal format.						
LI	Long block internal format.						
PG,p ₁ ,p ₂ ,p ₃	<p>Define preferred residence priority groups. The default is T,MN,D. The parameters p_n are strings of preferred residence mnemonics which define the priority groups used in file selection. The groups are defined in order of decreasing priority. The default value defines three groups: 1) a high priority group consisting of files with PR=T (tape alternate storage residence preferred); 2) a medium priority group consisting of files having PR=M (cartridge alternate storage residence on either MSE or optical disk preferred) and PR=N (no preferred residence); and 3) a low priority group consisting of files having PR=D (disk residence preferred). Files having PR=L (file locked to disk) are excluded from selection.</p> <p>It is possible to define other than three groups. For example, the following entry:</p> <pre style="margin-left: 40px;">PG, TMNDL</pre> <p>places all files in one group, effectively causing preferred residence to be ignored.</p> <p>The following entry defines four groups.</p> <pre style="margin-left: 40px;">PG, T, M, N, D</pre> <p>If a preferred residence type is not specified, files having that value are excluded from selection unless they are selected explicitly with an SF or SV directive. The following entry causes files with PR=D or PR=L to be excluded:</p> <pre style="margin-left: 40px;">PG, TMN</pre>						

Directive	Description
RL,dn,tracks RL,dt,tracks	<p>Define the track release limit for device number dn or device type dt. A device type is specified as a mass storage equipment mnemonic followed by a unit count. Files residing on a device are not released unless you enter an RL directive specifying either the device number or the device's type. A device number entry takes precedence over a device type entry if both apply. Files residing on the applicable device or devices are released until the number of available tracks reaches the specified limit. No release processing is performed for the device if the number of available tracks is equal to or greater than the specified limit. For example, the following device number entry:</p> <pre data-bbox="495 583 625 611">RL, 40, 300</pre> <p>specifies that files should be released on device 40 until there are 300 available tracks. If 300 tracks are already available, no files will be released.</p> <p>The following device type entry:</p> <pre data-bbox="495 793 641 821">RL, DQ2, 250</pre> <p>specifies a release limit of 250 tracks for each two-spindle 885 device for which a limit has not been specified by device number.</p>
SF,ui,f ₁ ,f ₂ ,...,f _n SF,ui	<p>Unconditionally select the specified files cataloged under user index ui, if found in the summary file. If no file names are specified, select all files cataloged under user index ui.¹</p>
SV,V ₁ ,V ₂ ,...,V _n	<p>Unconditionally select any files found in the summary file with tape alternate storage VSNs that match the specified VSNs. This directive is intended to be used when recycling destage dump tapes as part of destage processing. Specifying the VSNs of tapes to be recycled prevents files residing on those tapes from being excluded by the priority selection process.¹</p>
XF,ui,f ₁ ,f ₂ ,...,f _n XF,ui	<p>Exclude the selection of the specified files cataloged under user index ui. If no file names are specified, exclude all files cataloged under user index ui.</p>

1. The SF and SV directives cause GENPFD to treat the specified files as higher priority than others in the selection process. In some cases, GENPFD satisfies the specified destage or release limits without selecting all of the files.

File Selection Processing

The GENPFD utility sorts the candidate files into preferred residence priority groups as determined by the PG directive. Within each group, each file is assigned a priority based on the time elapsed since the last access (if PO=R) or last data modification (if PO=D) with older files receiving higher priorities. Files selected with SF directives or files residing on VSNs selected with SV directives are given a higher priority than other files to force selection.

For destage processing, files are selected beginning with the highest priority until the destage limit is reached. The permanent file utility parameters UI and PF are written to the utility directives file for each selected file unless the combined length of the selected files is less than the destage threshold (either the default or an explicit nonzero value specified in the DT directive).

For release processing, files residing on devices specified by RL directives are selected for release until the defined limit for each device is reached. The permanent file utility parameters UI and PF are written to the utility directives file for all files selected for release.

Limitations

GENPFD can process a maximum of 80 RL directives and a total of 1000 SF, SV, and XF directives. Each directive counts as one entry, whether or not it contains a file name or VSN. Each file name or VSN on a directive after the initial one counts as an additional entry.

Introduction	13-1
Linked SDM Operation	13-2
Deadstart	13-3
Shared Mass Storage	13-6
Configuring Shared Mass Storage Devices	13-6
Device Access Tables	13-6
Mass Storage Recovery Tables	13-7
Shared Buffered Devices	13-7
Device Initialization	13-8
Device Unload	13-9
Device Recovery	13-9
Device Checkpoint	13-10
MSE Processing	13-10
Independent SDM Operation	13-11
Deadstart	13-11
Shared Mass Storage	13-13
Configuring Shared Mass Storage Devices	13-14
Device Index Tables	13-15
Mass Storage Recovery Tables	13-15
Device Initialization	13-15
Device Unload	13-16
Device Recovery	13-16
Device Checkpoint	13-17
MSE Processing	13-17

Introduction

Multimainframe operation provides mechanisms by which more than one computer can share mass storage devices. Each mainframe in a multimainframe configuration has access to the shared mass storage devices and their preserved files. Preserved files are those files that are retained across a level 0 deadstart (nonsystem recovery). Types of preserved files are permanent files, queued files, and system dayfiles.

There are two types of multimainframe configurations:

- Linked shared device multimainframe (SDM).
- Independent shared device multimainframe.

The main difference between the two SDM configurations is their method of maintaining the critical device tables necessary to control shared mass storage and intermainframe communications. Linked SDM configurations maintain the critical device tables in extended memory, while independent SDM configurations maintain the tables on each shared mass storage device. Because the independent SDM configurations maintain the tables on the shared devices, it causes some performance degradation.

Each mainframe in a multimainframe configuration can operate in multimainframe mode or in standalone mode; however, two machines may not access the same device unless both are in multimainframe mode. A device is considered shared if it can be accessed by more than one of the mainframes; it need not be accessible to all mainframes in the complex.

The fact that a computer is operating as part of a multimainframe complex is not apparent to the user, except for some degrading of performance for the independent SDM. However, there are operational changes and additions that the operator should note. These include deadstart commands, displays, online commands necessary in the event of an interruption on one of the sharing mainframes, and the UNLOAD/MOUNT process for removable devices.

Recovery allows a machine to either join other machines operational in a multimainframe complex or to operate in a standalone mode. The standalone system is not allowed to use the same mass storage devices as other machines. In the event of a system interruption to one machine, it is possible to operate the remaining machines in a multimainframe environment.

Automatic detection of extended memory is not provided, because it is not possible to determine its absence and continue to run on all machine types. For example, a 6600 will hang if an attempt is made to execute an extended memory instruction without extended memory. Extended memory status is checked during deadstart when processing an extended memory entry in the EQPDECK.

Except where explicitly specified that controllers and/or equipment can be shared between mainframes, assume they cannot be shared. For example, NOS does not support sharing a two-channel tape controller between mainframes; nor does it support sharing mass storage controllers except as specified in this section.

Linked SDM Operation

In a linked SDM complex, up to seven computers can access mass storage devices. Any combination of 1 to 7 model 71, 72, 73, 74, 171, 172, 173, 174, 175, 176, 720, 730, 740, 750, 760, 865, 875, CYBER 180-class machines or 6000 Computer Systems mainframes can be used. Most mainframes require one CPU port connected to extended memory. The exception is CYBER 180-class machines which must be connected to extended memory through a low-speed port (LSP) or a distributive data path (DDP) instead of a CPU port. Shared mass storage devices supported are: 844 (single and double density), 885-11/12, extended memory, and DAS (SSDs, SABREs, and ELITEs).

Extended memory is the means for controlling shared mass storage and intermainframe communications. Each mainframe has a port into extended memory through which system activity is controlled. In order to control shared mass storage devices, several extended memory resident tables are required. The device access table (DAT) contains the logical description (family name/pack name and device number) of each mass storage device (shared or nonshared) that is accessible by any machine in the linked SDM complex. For each device in the DAT that is to be accessed by more than one machine, a corresponding MST (image of central memory resident MST) and TRT also reside in extended memory. In addition, a machine recovery table (MRT) exists in extended memory for each machine and device (that is, there are as many MRTs for each shared device as there are mainframes in the complex).

If shared buffered devices are configured, there also exists a buffer access table (BAT) and buffered device table (BDT) for each shared buffered device, as well as a buffer request table (BRT) for each mainframe.

If MSE is used, one mainframe in the linked SDM complex is physically connected to 7990 and is called the master mainframe; all other mainframes in the complex can access 7990 files and are called slave mainframes. SSEXEC runs on the master mainframe and is the main processing program that controls MSE activity. SSSLV runs on each slave mainframe and communicates with SSEXEC to retrieve files from 7990 in response to ATTACH requests by jobs running on the slave mainframes.

Deadstart

A linked SDM environment is defined at deadstart by CMRDECK and EQPDECK entries in each machine. Use the following entries to do this.

CMRDECK Entry	Description
MID=id.	2-character machine identification. This identification is used to associate a specific machine with its access to a shared device. It is required, however, even if no devices are shared during system operation. This id must be unique for each machine in the complex. If id is omitted, the default value is MID=AA.
EQPDECK Entry	Description
SHARE=est ₁ ,est ₂ ,...,est _n .	EST ordinals of the mass storage devices that will be shared with other machines in the linked SDM complex. The MST/TRT for each device resides in extended memory.
MAXSHARE=n,b.	Optional; defines the total number of devices to be shared by all the mainframes in the complex (n), and the total number of buffered devices to be shared by all the mainframes in the complex (b). This entry is required only if shared devices exist which are not configured on all mainframes in the complex. If this entry is not specified, the default values are the number of shared devices and shared buffered devices defined for this particular mainframe.
PRESET.	Initializes the flag register and the extended memory resident tables. You can specify this entry only at the system console. You must specify the PRESET entry on a level 0 deadstart on the first machine to be deadstarted in the linked SDM complex.

These entries are also described in section 3, Deadstart Decks.

Because in a linked SDM complex two or more machines can utilize the same mass storage devices, the device assignments and EQPDECKs of all machines are interrelated. Care must be taken to ensure proper EQPDECK settings for each machine to obtain the desired device configuration. The following items are important in obtaining this proper device configuration.

- Each machine must specify a unique machine identification (using the MID entry in the CMRDECK). This identification associates a particular machine with its files on a shared device. There are no external characteristics associated with this identification. However, if the machine identification specified on a machine deadstarting into a multimainframe complex is identical to a machine identification on a machine already operating, the deadstart process halts and an appropriate message is displayed.
- The assignment of shared mass storage devices should be made properly to ensure the recoverability of the device and the proper operation of the system (refer to Shared Mass Storage later in this section).

- If SHARE EQPDECK entries are present, an extended memory entry (DE or DP equipment type) must also be present in the EQPDECK. If none is found, an error message is given to the operator indicating that no link device has been defined.
- The first machine deadstarted (level 0 only) must specify the PRESET EQPDECK entry. This command causes tables to be preset in extended memory, and in so doing, assures that no other machine has arrived at the same point in the deadstart sequence and is attempting the same thing. These other machines should not have the PRESET entry specified during deadstart. All other machines that arrive at that point in the deadstart process display a message indicating that they are waiting for deadstart on the preset machine. A machine that does not preset extended memory has no means of detecting whether extended memory has been preset previously by another machine. Therefore, the operator must ensure that extended memory has been preset by a prior deadstart before deadstarting another machine that is to be in the multimainframe complex.
- If extended memory is initially placed in maintenance mode (half of extended memory used for online diagnostics), the first machine deadstarted in a complex must have the INITIALIZE and PRESET EQPDECK entries specified when deadstarting (level 0 only). The remaining machines must have the extended memory size correspondingly reduced in their EQPDECKs.

When a mainframe joins a linked SDM complex, it is associated with an identification that it utilizes during system operation but which is independent of the machine identification. This association is done during deadstart, when the machine investigates the multimainframe tables residing on the link device and places its machine identification in an empty slot of the seven that are available. Associated with each slot is a unique machine index and a unique machine mask, which the machine uses either to index itself into various MMF tables or to identify itself in these tables.

When performing a level 0 recovery deadstart on one of the machines, the MREC utility should be run on some or all of the remaining machines before the deadstart proceeds. The purpose of this utility is to clear interlocks held by the machine to be deadstarted that have not been cleared by CPUMTR. It can also recover mass storage space on a shared device that is currently not accessible because of a machine interruption (necessitating a level 0 deadstart). MREC may have to be run from more than one machine since it affects only shared devices (that is, devices specified on the SHARE EQPDECK entry) of the machine on which MREC is run. If the interrupted machine shares different devices with different machines, MREC must be run from enough machines to account for all devices shared with the interrupted machine. The interface to the MREC utility is described in section 8, K-Display Utilities.

If a level 3 recovery deadstart will be performed, extended memory must be intact as well as CMR. For a level 1 or 2 recovery in a linked SDM environment, extended memory must be intact.

Figure 13-1 illustrates the steps generally needed to deadstart a machine in a multimainframe complex.

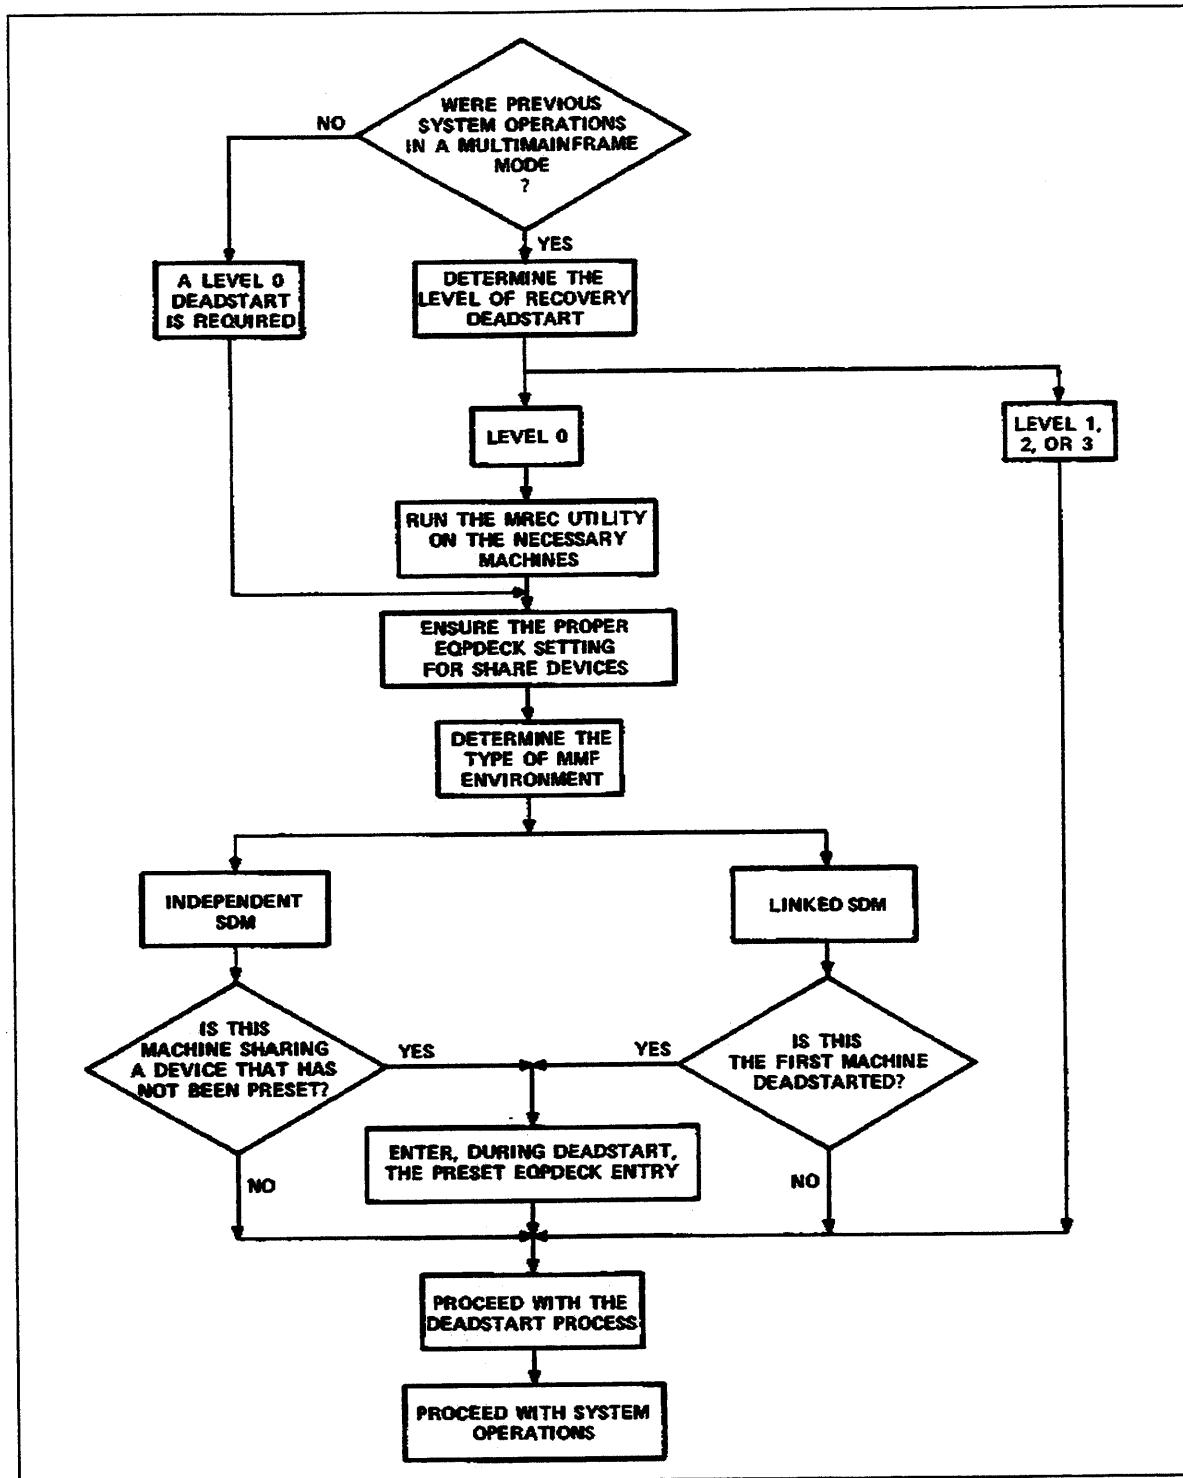


Figure 13-1. Deadstarting a Machine in a Multimainframe Complex

Shared Mass Storage

A major reason for operating a multimainframe complex is to be able to share mass storage devices between machines. Thus, the users of two systems may be able to utilize the same files if these files reside on a shared mass storage device. Having the ability to share files between systems also means that several additional procedures are required when operating a multimainframe complex. Additional consideration must also be given when configuring mass storage devices. This material, which is unique to multimainframe operation, is discussed in the following paragraphs.

Configuring Shared Mass Storage Devices

Some consideration should be given as to which devices should be made shared devices and which ones should not. Since a shared device contains preserved files that can be accessed by more than one machine, it must be physically connected to and logically defined (as shared) by each machine sharing the device. If one device of a family is defined as shared, then normally all devices in the family should be defined as shared for accessibility. If a removable device will be treated as shared, it must be defined as removable in each machine sharing it. If a shared device is not removable, additional use of the device (beyond preserved files) is defined by each machine, independently of any other machine definition (that is, each machine specifies whether the device is to contain such things as a copy of the system or temporary files).

There is no real advantage to having nonpreserved files on a shared device, unless spreading them across shared devices outweighs having them on a dedicated device. A key factor is drive and controller contention. Another factor is whether or not enough drives and accesses to those drives are available to warrant spreading temporary files across several devices from all systems. Also, additional overhead is incurred by the system for each shared device. This includes additional CMR space for each device (100₈ CM words) and additional time that is required to maintain the device-related tables (such as the MST or the TRT) on the link device. These factors must be weighed to determine how best to configure shared devices.

Device Access Tables

In a linked SDM complex, mainframes use DATs to determine device usage. The DATs identify the status of all shared and nonshared devices for the mainframes in the complex. One DAT is allocated in extended memory for each mass storage device.

In order to minimize configuration problems, shared removable equipment should be configured the same way on all machines in the complex. For example, if one system defines three shared units as three single-unit devices and another system defines the same units as one three-unit device, the first system can accommodate a two-unit device on these units, whereas the second system would consider it an error. Unless the configurations are the same on all machines, any devices mounted on those drives may not necessarily be recoverable on all machines.

NOTE

RESEX considers only the configuration of the machine on which it is executing in its overcommitment algorithm.

Mass Storage Recovery Tables

One problem that is created by having more than one machine sharing a mass storage device is that of recovering the mass storage space and interlocks of a machine, if it requires recovery processing. This problem is solved by defining a table that provides the information needed to recover the mass storage space of a machine and by having a machine recovery utility that performs the recovery. The table is called the mass storage recovery table (MRT). There is one MRT for each mainframe per device. It tells which tracks are interlocked and which tracks are first tracks of files local to a particular machine. The MRT is utilized by CPUMTR on another machine to clear track interlocks and by the MREC utility on another machine to recover the mass storage space of the interrupted machine. For specific information on the MRT, refer to the NOS Version 2 Systems Programmer's Instant.

Shared Buffered Devices

When performing I/O on a buffered device (for example, on a DAS device), the system uses a set of I/O buffers which are maintained in the UEM portion of central memory for the mainframe in question. Each of these buffers corresponds to a physical location on the disk, and the buffers are effectively used as a cache for the disk.

When a buffered device is configured as a shared device, the system must prevent the condition where two or more mainframes each have an I/O buffer which points to the same physical sector on the disk, but the contents of the two buffers are different. This problem is avoided by the use of the buffer access table (BAT). One BAT is allocated in extended memory for each shared buffered device.

The BAT is used to grant access permission for a given mainframe to access a given logical track in either read or write mode. While two or more mainframes can each have read access permission for the same track at the same time, only one mainframe can have write access permission for a track, and that write permission also locks out read access permission.

If one mainframe wants to obtain a BAT access permission that is not currently available (because another mainframe currently holds a conflicting BAT access permission), the first mainframe posts a request in the buffer request table (BRT). The second mainframe sees this request and releases the conflicting access as soon as possible. One BRT is allocated in extended memory for each mainframe.

In addition, due to the way error processing is performed on DAS devices, it is necessary to prevent more than one mainframe from performing error processing on the same shared DAS device at the same time. This interlocking is accomplished by the use of the buffered device table (BDT). One BDT is allocated in extended memory for each shared buffered device.

Device Initialization

To initialize a mass storage device, it is necessary first to prevent any new activity from starting up on the device, then to wait until all current activity has completed, and finally to interlock the device and proceed with initialization. To accomplish this on shared mass storage devices in a linked SDM environment, follow these steps:

1. Enter the DSD UNLOAD command. All machines sharing the device must unload it to prevent any new activity. (This command can be used for both removable and nonremovable shared devices.)
2. Enter the DSD INITIALIZE command on the machine from which the initialization is to take place. If you enter the INITIALIZE command from another machine, an error message is displayed. Refer to the NOS Version 2 Operations Handbook for error messages.
3. The machine from which the INITIALIZE command was entered monitors the status of the other machines that are sharing the device. Once they have unloaded the device and user activity has ceased, the device tables are updated and the device is checkpointed.
4. To activate the device on the other machines, enter the DSD MOUNT command from each machine. This command clears the UNLOAD status. If initialization is still in progress on another machine when you enter a MOUNT command, the MOUNT process is terminated with an error. Refer to the NOS Version 2 Operations Handbook for error messages.

Device Unload

In a linked SDM environment, unloading a device involves more than it does under a one-machine system. A device can be unloaded from a machine (referred to as a local unload), or it can be unloaded from the entire multimainframe complex (referred to as a global unload). A device can be physically removed from the complex only after a global unload has been accomplished. This is the general procedure to complete a global unload:

1. Enter the UNLOAD command from each machine. This is an indication to the machine that no new accesses should be initiated. You must enter this command from each machine sharing the device.
2. When all local unloads are set and user access has ceased, global unload status is set if the device is a removable device. This global unload status is displayed on all machines, indicating that there is no activity on the device from any machine and that the device may be physically unloaded.
3. You can then switch packs and enter the MOUNT command at the console to initiate recovery of the device. The MOUNT command clears the global unload status and the local unload status on the machine from which it was entered and indicates that this machine is now accessing the device. All other machines continue to ignore the device until you enter the MOUNT command on each machine. The MOUNT command does nothing if local unload status is not set on the machine.

Device Recovery

Deadstart and online recovery methods are similar in logic for recovering mass storage devices.

When a device is recovered in a linked SDM complex, the DAT in extended memory is interlocked while a check is made to see if an entry exists for this device. The presence of an entry indicates that another machine is also accessing the device. If an entry is found and the machine recovering the device has not been instructed to share it, an error is indicated and recovery halts with an appropriate message displayed. If the machine already accessing the device is not allowing it to be shared, the same error condition occurs. Therefore, if a device is being accessed, another machine can recover that device only if the recovering machine and the accessing machine use the device in a multimainframe mode.

NOTE

If two devices recovered on separate machines have the same family name/device number or pack name, there is no method of determining whether or not they are the same device if both are shared. If they are different devices, they are both destroyed when used.

Device Checkpoint

Local MST information for each machine that shares a mass storage device is maintained on the device. MST information for other machines may also be present on the device. The information for each machine is kept in one sector on the label track following the TRT sectors. Entries for up to 31 unique machine identifiers can exist.

During checkpoint, only the local MST information of the machine performing the checkpoint is updated. Since local MST information for many machines is kept on the device, updating of all these areas by one machine could cause a loss of information needed if the device were to be used with another system.

For shared devices in a linked SDM complex, duplication of checkpointing by more than one machine is prevented. If a machine attempts to checkpoint a shared device and determines that a second machine is performing the checkpoint (a checkpoint request bit is set in the local MST area of the second machine), no action is taken. Only one checkpoint bit is set at any given time for a device.

MSE Processing

Unlike most other mass storage devices, the 7990 is not shared by mainframes. Rather, 7990 is physically connected to and driven by one mainframe called the master mainframe. All other mainframes that share permanent files with the master mainframe are called slave mainframes. The disks to which 7990-resident files are staged must be shared devices and accessible by all mainframes.

SSEXEC runs on the master mainframe and is the main program that controls MSE processing activities. You can activate SSEXEC by entering the following DSD commands at the system console.

```
L.SUBSYST.
ENABLE,MASTER MSE.
L.END.
MSEffff.
```

SSSLV runs on each slave mainframe and is the program that supports file retrieval from a 7990 for jobs running on a slave mainframe. You can activate SSSLV by entering the following DSD commands at the system console.

```
L.SUBSYST.
DISABLE,MASTER MSE.
L.END.
MSEffff.
```

When a job running on a slave mainframe attaches a 7990-resident file that has no disk image, SSSLV is requested to stage the file to disk. However, SSSLV cannot access 7990 directly. SSSLV must communicate with SSEXEC, which stages the file to a shared disk and then notifies SSSLV when the file is available. SSSLV then causes the requesting job to be resumed so that the file ATTACH can complete.

The permanent files used for communication between SSEXEC and SSSLV are initialized and recovered automatically and require no backup or recovery processing. The name of the permanent file used by SSSLV running on mainframe *i* (*i*=1, 2, 3, 4) to send requests to SSEXEC is STBMNO*i*. The name of the permanent file used by SSEXEC to respond to all requests by all SSSLVs is MTBSPFN. These files belong to the family that contains extended memory and are saved under the user index 377760₈.

Independent SDM Operation

In an independent SDM complex, up to 16 computers can access shared mass storage devices. Any combination of from 1 to 16 model 71, 72, 73, 74, 171, 172, 173, 174, 175, 176, 720, 730, 740, 750, 760, 865, 875, CYBER 180-class machines, or 6000 Computer Systems mainframes can be used. Shared mass storage devices supported are 834, 836, 844 (single and double density) and 885-11/12 (single and double density).

Each shared device contains in its label track all of the tables necessary to control multimainframe access. These tables include the MST, TRT, device index table (DIT), and MRT for each mainframe sharing the device. Whenever a mainframe needs to access a shared device, it is necessary for the mainframe to ensure its copies of these tables are up-to-date.

One or more of the mainframes can have extended memory, but each mainframe's extended memory must be used as a nonshared device.

If MSE is used, one mainframe in the independent SDM complex is physically connected to 7990 and only that mainframe can access it.

Deadstart

An independent SDM environment is defined at deadstart by CMRDECK and EQPDECK entries in each machine. Use the following entries to do this.

CMRDECK Entry	Description
MID=id.	2-character machine identification. This identification is used to associate a specific machine with its access to a shared device. It is required, however, even if no devices are shared during system operation. This id must be unique for each machine in the complex. If id is omitted, the default value is MID=AA.
EQPDECK Entry	Description
ISHARE=est ₁ ,est ₂ ,...,est _n .	EST ordinals of the mass storage devices that will be independently shared with other machines in the multimainframe complex. The ISHARE entry must precede the PRESET entry. The MST/TRT resides on each shared device.
PRESET=est ₁ ,est ₂ ,...,est _n .	EST ordinal of the shared mass storage device to be preset. You can specify this entry only at the system console. Several machines at level 0 deadstart may need the PRESET entry. The PRESET entry is specified only on a level 0 deadstart by the first machine accessing the device and must follow the ISHARE entry. PRESET initializes the tables residing on the label track of the shared devices.

These entries are also described in section 3, Deadstart Decks.

Since in an independent SDM complex two or more machines can utilize the same mass storage devices, the device assignments and EQPDECKs of all machines are interrelated. Care must be taken to ensure proper EQPDECK settings for each machine to obtain the desired device configuration. The following items are important in obtaining this proper device configuration.

- Each machine must specify a unique machine identification (using the MID entry in the CMRDECK). This identification associates a particular machine with its files on a shared device. There are no external characteristics associated with this identification. However, if the machine identification specified on a machine deadstarting into a multimainframe complex is identical to a machine identification on a machine already operating, the deadstart process halts and an appropriate message is displayed.
- The assignment of shared mass storage devices should be made properly to ensure the recoverability of the device and the proper operation of the system (refer to Shared Mass Storage, next).
- Each machine must specify the ISHARE=est EQPDECK entry to inform the system that devices will be shared in the independent shared device mode.
- If you need to deadstart an interrupted mainframe in an independent SDM environment, you should put the nonfailing mainframes in STEP mode before deadstarting the interrupted mainframe.
- On a level 0 deadstart, any shared device that has not been shared prior to this deadstart must include the PRESET entry in its definition.
- On a level 0 deadstart, if a device was previously shared by the deadstarting machine, the MRT and DIT device resident tables must be zeroed out.
 - If the machines that share the device are not currently running, the device must be preset or initialized to zero out the tables.
 - If any machine that shares the device is running, the machine recovery utility (MREC) must be run from a running machine to remove the machine identification of any machines that are not currently running and to clear all interlocks held by these machines.
- On a level 0 deadstart, if a device was previously shared by a machine other than the deadstarting machine, the MRT and DIT device resident tables may or may not have to be zeroed out.
 - If another machine that is sharing the device is running, the device must not be preset.
 - If no machines that will be sharing the device are running and one or more of them will be deadstarted at a level 1, 2, or 3, the device must not be preset.
 - If all machines that will be sharing the device are not currently running and all will be deadstarted at level 0, the device can be preset or initialized.

When performing a level 0 recovery deadstart on one of the machines, the MREC utility should be run on some or all of the remaining machines before the deadstart proceeds. The purpose of this utility is to clear interlocks held by the machine to be deadstarted that have not been cleared by CPUMTR. It can also recover mass storage space on a shared device that is currently not accessible because of a machine interruption (necessitating a level 0 deadstart). MREC may have to be run from more than one machine since it affects only shared devices (that is, devices specified on the ISHARE EQPDECK entry) of the machine on which MREC is run. If the interrupted machine shares different devices with different machines, MREC must be run from enough machines to account for all devices shared with the interrupted machine. The interface to the MREC utility is described in section 8, K-Display Utilities.

In an independent SDM complex, as in a standalone system, critical tables that reside on the device must be intact for a level 1, 2, or 3 deadstart recovery.

Figure 13-1 illustrates the steps generally needed to deadstart a machine in a multiframe complex.

Shared Mass Storage

A major reason for operating a multiframe complex is to be able to share mass storage devices between machines. Thus, the users of two systems may be able to utilize the same files if these files reside on a shared mass storage device. Having the ability to share files between systems also means that several additional procedures are required when operating a multiframe complex. Additional consideration must also be given when configuring mass storage devices. This material, which is unique to multiframe operation, is discussed next.

Configuring Shared Mass Storage Devices

Some consideration should be given as to which devices should be made shared devices and which ones should not. Since a shared device contains preserved files that can be accessed by more than one machine, it must be physically connected to and logically defined (as shared) by each machine sharing the device. If one device of a family is defined as shared, then normally all devices in the family should be defined as shared for accessibility. If a removable device will be treated as shared, it must be defined as removable in each machine sharing it. If a shared device is not removable, additional use of the device (beyond preserved files) is defined by each machine, independently of any other machine definition (that is, each machine specifies whether the device is to contain such things as a copy of the system or temporary files).

There is no real advantage to having nonpreserved files on a shared device, unless spreading them across shared devices outweighs having them on a dedicated device. A key factor is drive and controller contention. Another factor is whether or not enough drives and accesses to those drives are available to warrant spreading temporary files across several devices from all systems. Also, additional overhead is incurred by the system for each shared device. This includes additional CMR space for each device (100₈ CM words) and additional time that is required to maintain the device-related tables (such as the MST or the TRT) on the link device. These factors must be weighed to determine how best to configure shared devices.

In an independent SDM complex, where device-related tables are maintained on the device itself, system performance degradation can occur if the shared devices are not properly configured. Since the device must be accessed every time its tables in CMR are not up to date, drive and controller contention is increased and system performance is affected. To reduce the impact on system performance you can:

- Put the system and temporary files on nonshared devices.
- Put only those fast-attach files that must be global on shared devices.
- Use multiple channels and multiple controllers to improve performance and provide additional paths to shared devices.

Device Index Tables

In an independent SDM complex, mainframes use DITs to determine device usage. The DITs contain the MIDs of the mainframes that can access the shared device. The shared device maintains a DIT for each mainframe that shares it. The DITs reside on each shared device within the MST on the device's label track.

Mass Storage Recovery Tables

One problem that is created by having more than one machine sharing a mass storage device is that of recovering the mass storage space and interlocks of a machine, if it requires recovery processing. This problem is solved by defining a table that provides the information needed to recover the mass storage space of a machine and by having a machine recovery utility that performs the recovery. The table is called the MRT. There is one MRT for each mainframe per device. It tells which tracks are interlocked and which tracks are first tracks of files local to a particular machine. The MRT is utilized by CPUMTR on another machine to clear track interlocks and by the MREC utility on another machine to recover the mass storage space of the interrupted machine. For specific information on the MRT, refer to the NOS Version 2 Systems Programmer's Instant.

Device Initialization

To initialize a mass storage device, it is necessary first to prevent any new activity from starting up on the device, then to wait until all current activity has completed, and finally to interlock the device and proceed with initialization. To accomplish this on shared mass storage devices in an independent SDM environment, follow these steps:

1. Enter the DSD UNLOAD command. All machines sharing the device must unload it to prevent any new activity. (This command can be used for both removable and nonremovable shared devices.)
2. Enter the DSD INITIALIZE command on the machine from which the initialization is to take place. Each entry of the INITIALIZE command toggles INITIALIZE on and off. If two machines in the complex enter the INITIALIZE command, the INITIALIZE command is toggled off and no error message is displayed.
3. The machine from which the INITIALIZE command was entered monitors the status of the other machines that are sharing the device. Once they have unloaded the device and user activity has ceased, the device tables are updated.
4. To activate the device on the other machines, enter the DSD MOUNT command from each machine. This command clears the UNLOAD status. If initialization is still in progress on another machine when you enter a MOUNT command, the machine waits for INITIALIZE to finish and then continues with the MOUNT process.

Device Unload

In an independent SDM environment, unloading a device involves more than it does under a one-machine system. A device can be unloaded from a machine (referred to as a local unload), or it can be unloaded from the entire multimainframe complex (referred to as a global unload). A device can be physically removed from the complex only after a global unload has been accomplished. This is the general procedure to complete a global unload:

1. Enter the UNLOAD command from each machine. This is an indication to the machine that no new accesses should be initiated. You must enter this command from each machine sharing the device.
2. When all local unloads are set and user access has ceased, global unload status is set if the device is a removable device. This global unload status is displayed on all machines, indicating that there is no activity on the device from any machine and that the device may be physically unloaded.
3. You can then switch packs and enter the MOUNT command at the console to initiate recovery of the device. The MOUNT command clears the global unload status and the local unload status on the machine from which it was entered and indicates that this machine is now accessing the device. All other machines continue to ignore the device until you enter the MOUNT command on each machine. The MOUNT command does nothing if local unload status is not set on the machine.
4. A pack that has not been previously mounted in this system must be preset with a MOUNT,est,P command.

Device Recovery

Deadstart and online recovery methods are similar in logic for recovering mass storage devices.

For a machine to recover a device in an independent SDM complex, the MID of the machine must be in the device's DIT. If the MID is not in the DIT, an error message is issued and recovery is impossible. When recovering a device in a standalone environment, the device is not checked if it is/was an independent shared device.

Device Checkpoint

Local MST information for each machine that shares a mass storage device is maintained on the device. MST information for other machines may also be present on the device. The information for each machine is kept in one sector on the label track following the TRT sectors. Entries for up to 31 unique machine identifiers can exist.

During checkpoint, only the local MST information of the machine performing the checkpoint is updated. Since local MST information for many machines is kept on the device, updating of all these areas by one machine could cause a loss of information needed if the device were to be used with another system.

In an independent SDM complex, the MST and TRT information is maintained on the device. If the MST and/or TRT information changes during a device access, the system updates the MST and/or TRT information making it unnecessary for the system to periodically checkpoint the device.

MSE Processing

Unlike most other mass storage devices, the 7990 is not shared by mainframes. Rather, 7990 is physically connected to and driven by one mainframe. The only machine that can access the 7990-resident files is the physically connected independent SDM. MSE processing is the same in an independent SDM environment as it is in a single mainframe environment. Refer to section 11, Mass Storage Extended Subsystem.

NAD Maintenance Utilities

14

Dump NAD Memory (DMPNAD)	14-1
Maintenance Host Facility (MHF)	14-2
Listing NAD Dumps	14-3

This section describes utilities that are used to maintain network access devices (NADs).

Dump NAD Memory (DMPNAD)

DMPNAD is a utility program that uses data obtained either directly from NAD memory or from a binary dump file (produced previously by MHF or DMPNAD) to generate a print file showing each NAD word in both hexadecimal and 7-bit ASCII code.

The format of the DMPNAD command is:

DMPNAD, P₁, P₂, . . . , P_n.

P_i is a parameter consisting of a keyword or a keyword equated to a value. All parameters are order independent.

P_i	Description
AC=aaaa	The 4-digit hexadecimal number specifying the remote NAD's access code (refer to the RHF K display in section 8, K-Display Utilities). The default is AC=0000.
B=binfile	The 1- to 7-character name of the file on which the binary output is to be written. The default is no binary output file.
B	Same as B=BINOUT.
CH=ch	The 1- or 2-digit octal number specifying the channel number of the local NAD (defined in the NAD's EST entry).
I=infile	The 1- to 7-character name of the file containing binary input data. DMPNAD dumps only one record from the file. The default is no binary input file.
I	Same as I=BININ.
L=outfile	The 1- to 7-character name of the file on which the list output is to be written. The default is L=OUTPUT.
LT=t ₀ t ₁ t ₂ t ₃	The 4-digit binary bit pattern specifying the TCUs that connect a local NAD to a remote NAD. t _n =1 enables TCU _n . For example, LT=1100 indicates that local TCUs 0 and 1 are enabled.
ND=nn	The 2-digit hexadecimal number specifying the remote NAD's address (refer to the RHF K display in section 8, K-Display Utilities).

The DMPNAD utility can be called from the console or a batch job. The calling job must be of system origin or you must be validated for system origin privileges. You normally include the channel number on the DMPNAD command to specify which local NAD is to be dumped. The dump can also be saved as a binary output file (B parameter) which can be listed later.

DMPNAD can dump either a local NAD (380-170 NAD connected to the host mainframe using a channel and configured in the EST) or a remote NAD (any 380 NAD accessible to a local NAD using a loosely-coupled network trunk). For a local NAD dump, the EST entry associated with the local NAD's channel number must be OFF or the controlware-not-loaded flag must be set. Since the local NAD dumping process stops all NAD activity, the NAD controlware must be reloaded using LOADBC (refer to section 1, Controlware Utilities) before the local NAD can be used by the operating system.

Remote NAD dumping is selected by specifying the remote NAD to be dumped (AC and ND parameters), the local NAD's channel number used to access the remote NAD (CH parameter), and the TCU used to connect the local NAD to the remote NAD (LT parameter). For remote NAD dumping, the local NAD must be loaded and running and its EST entry must be ON.

Remote NAD dumping operations can occur concurrently with RHF use of the local NAD. However, extreme care should be exercised when performing a remote NAD dump to ensure that the correct remote NAD is being dumped and that the remote NAD is not being used by the mainframe to which it is connected. After the remote NAD dumping operation, the NAD controlware must be reloaded into the remote NAD before it can be used by the remote mainframe.

When the I parameter is specified, DMPNAD will list a binary file that was saved earlier. Since no NAD access is required, no special privileges are required by the calling job.

Maintenance Host Facility (MHF)

MHF is a utility program that executes while RHF is active. MHF periodically copies NAD error logs to the binary maintenance log, loads microcode into local NADs during RHF initiation, and dumps and reloads local NADs that fail. MHF uses a copy of the network description table obtained directly from RHF to determine which NADs to log, dump, and/or load.

You control MHF operation by setting the AUTOLOAD, AUTODUMP, TRACE, and LOG parameters in the LNAV and RNAD directives for RCFGEN, which generates the RHF configuration file (refer to section 10, LID/RHF Configuration Files). You can control the remote NAD LOG parameter on RHF's network path table display (refer to section 8, K-Display Utilities). If you specify AUTODUMP=YES for local NADs, MHF saves NAD memory dumps as labelled binary records on permanent file NDFMid (id is the host mainframe identifier) under user name SYSTEMX. Each record label contains the record name NDMPLch (ch is the NAD's channel number), record number, and the date and time of the dump. MHF puts this dump information in the following sequence of job dayfile messages:

```
MHF, NAD DUMP RECORD NAME = NDMPLch
MHF, RECORD NO.0000    TIME = hh.mm.ss
MHF, FILE = NDFMid     DATE = yy/mm/dd
```

When loading a local NAD, MHF obtains the controlware from the system and uses a default set of controlware initialization parameters that it adjusts to suit the NAD's memory size.

Listing NAD Dumps

To list a NAD dump, you must be validated to obtain access to the NAD dump permanent file NDFMid (id is the host mainframe identifier).

You can either use ITEMIZE to locate a particular record on permanent file NDFMid or obtain the record number from MHF's dayfile, and then use DMPNAD to list the dump record. (ITEMIZE shows the record name, record number, date, and time corresponding to MHF's dayfile messages.)

The following example lists dump record 4 of permanent file NDFM64:

```
ATTACH, NDFM64 / UN=SYSTEMX.
```

```
ITEMIZE, NDFM64.
```

```
SKIPR, NDFM64, 3.
```

```
DMPNAD, I=NDFM64, L=LIST.
```

```
RETURN, NDFM64.
```

Check ITEMIZE output for the record number.

Skip to the fourth record.

List the fourth record on file LIST.

Network Organization	15-1
NAM Startup	15-5
Initiating NAM Without Operator Intervention	15-5
Altering the Memory File	15-6
NAM Shutdown	15-8
Network Control by HOP	15-9
Disabling and Enabling of Couplers	15-9
Status Message Formats for HOP	15-9
Unsolicited Status Reports for HOP	15-10
HOP Commands	15-11
NS Control Commands	15-11
Cancel Alternate NPU Load File Command	15-11
Change NPU Load File Command	15-12
NPU Load Status Command	15-12
Specific NPU	15-12
All NPUs	15-13
Status Response	15-13
NS Recent HISTORY Command	15-14
NVF Control Commands	15-15
Disable Host Element Command	15-15
Enable Application Command	15-16
Idle Host Element Command	15-16
NVF Recent HISTORY Command	15-18
Status Host Element Command	15-18
Particular Host Element	15-18
Network Applications	15-19
Active Network Applications	15-19
Terminal Devices	15-19
Users (Specified Application)	15-20
Network Control by NOP	15-21
Becoming a NOP	15-21
Status Message Formats for NOP	15-23
Unsolicited Status Reports for NOP	15-25
NOP Commands	15-26
ALERT Command	15-26
CONTROL Command	15-27
Disable Network Element Command	15-29
Single Element	15-29
Trunk	15-30
DISABLE Processing	15-30
Enable Network Element Command	15-32
Single Element	15-32
Trunk	15-33
Disabled Elements	15-33
Disabled Elements (All NPUs)	15-34
Disabled Terminals	15-34

SVCs	15-35
ENABLE Processing	15-35
NPU GO Command	15-36
NPU LOAD Command	15-37
Specified NPU	15-37
All NPUs	15-38
NPU Memory DUMP Command	15-39
Specified NPU	15-39
All NPUs	15-39
Send Message Command	15-40
Network Control by DOP	15-41
Becoming a DOP	15-41
Status Message Formats for DOP	15-43
Unsolicited Status Reports for DOP	15-45
DOP Commands	15-45
CS Information Command	15-46
NPU Diagnostic Test Command	15-47
TST Drop Option	15-47
TST Message Option	15-48
TST Request Option	15-48
Recent HISTORY Command	15-49
Report Unsolicited Status Command	15-50
Send Message Command	15-51
Status Network Element Command	15-52
Summary of Network Operation Commands	15-56
Control Commands for HOP	15-56
NAM K Display	15-56
NAM Startup	15-56
NAM Shutdown	15-56
NAM Mode Commands	15-57
NS Control Commands	15-58
NVF Control Commands	15-58
Control Commands for HOP and NOP	15-59
Control Commands for HOP, NOP, and DOP	15-59
Control Commands for COP	15-59

The commands described in this section are used when the network access method (NAM) is being initialized or is active. They initialize and control the network processing units (NPU's), device interfaces (DI's), trunks, lines, terminals, and application programs under network supervision.

NAM supports two types of networks: the 255x NPU network and the CDCNET network. Throughout this section, the term network refers to the 255x NPU network. However, whenever a statement is applicable to both network types, both types are named explicitly.

NAM controls the CDCNET network by sending commands to an application called Network Operator Utility (NETOU). This NAM application is also known as the network operations server. For further information on controlling the CDCNET network, refer to the CDCNET Network Operations Manual and the CDCNET Network Analysis Manual.

Network Organization

Initiating NAM (refer to the DSD NAM command in section 5, DSD Commands) causes NAM, and the communications supervisor (CS), network supervisor (NS), network validation facility (NVF), and terminal verification facility (TVF) to be loaded and initiated at control points. Refer to the NAM/CCP 3 Terminal Interfaces Reference Manual for more information on TVF. CS provides control and status of the NPU's and all of their communication lines and terminals. NS loads and dumps all of the NPU's in the communication network. NVF validates terminal users and applications as they attempt to access the network.

Multiple hosts may belong to the network. Hosts may assume one or more supervisory roles in the network. Each host will have a copy of NVF that will validate users connecting to and applications executing in that host. NVF provides these services for both the 255x NPU network and the CDCNET network. The supervisors CS and NS may execute in any or all of the CYBER hosts. Whenever an NPU needs to be loaded, it may request any copy of NS to perform the load (not always the same copy). Once an NPU is running, it may request any copy of CS to supervise it. Each NPU will have a preferred copy of CS, but it may request supervision from an alternate copy of CS if its preferred CS is not available. In any case, the network is always capable of functioning without any one or more of its hosts.

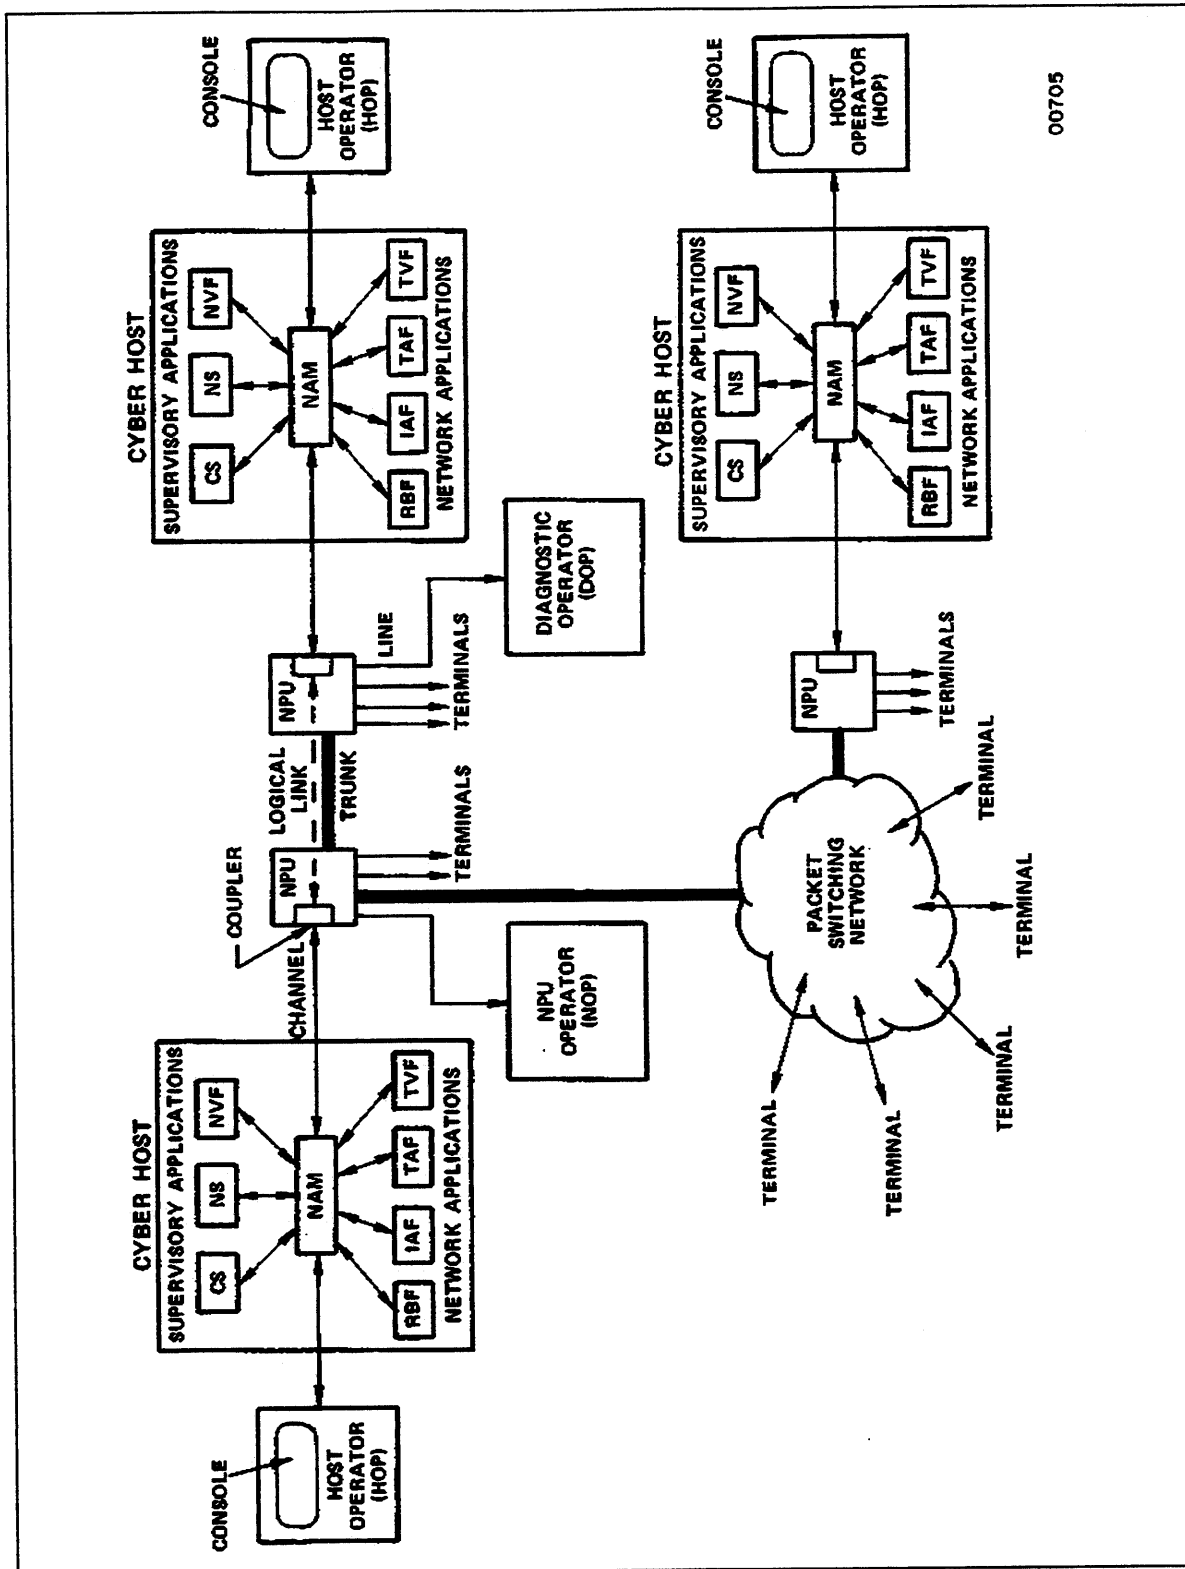
Whenever a CDCNET mainframe device interface (MDI) or mainframe terminal interface (MTI) must be loaded, it requests the services of a CDCNET job called INITMDI, which is responsible for loading and dumping the MDIs and MTIs. NAM automatically starts INITMDI upon demand.

Whenever a terminal device interface (TDI) or a network device interface (NDI) must be loaded, the MDI helping in the load process establishes a connection through NAM to the NAM application called the Network File Server (NETFS), which controls the loading and dumping of TDIs and NDIs.

The 255x NPU network consists of channels, couplers, lines, logical links, NPUs, terminals, and trunks (refer to figure 15-1) connected to the host computer.

- A channel is a data path on which a peripheral device controller can be accessed.
- A coupler is a hardware element that links an NPU to a channel.
- A line is the circuit that connects a terminal to an NPU.
- A logical link is a logical path connecting a coupler and an NPU or two couplers.
- An NPU is the communications controller that transmits data between terminals and a host computer or between two host computers.
- A public data network is an X.25 packet switching network (PSN). Terminals, NPUs, and foreign hosts can be connected to a PSN to exchange data.
- A terminal is a device by which a user supplies input messages to, and accepts output messages from, an application program.
- A trunk is the communication line connecting two NPUs.

The CDCNET network consists of channels, couplers, lines, logical links, device interfaces (MDIs/MTIs/NDIs/TDIs), terminals, and trunks. For definitions of CDCNET terminology, refer to the CDCNET Conceptual Overview Manual.



00705

Figure 15-1. Sample 255x NPU Network

Network host products (NHP) is the network software residing on the CYBER host. NHP recognizes four types of network operators: the host operator, the NPU operator, the diagnostic operator, and the CDCNET operator. A host operator (HOP) resides at the system console, initiates NAM, and can control 255x NPU network and CDCNET network elements being supervised by this host. An NPU operator (NOP) resides at a terminal and controls 255x NPU network elements (such as lines, logical links, NPUs, terminals, and trunks). A diagnostic operator (DOP) resides at a terminal and monitors and tests communication lines on the 255x NPU network.

A CDCNET operator (COP) resides at either the system console or a terminal and controls CDCNET network elements (such as lines, terminals, trunks, and DIs associated with CDCNET).

The HOP can always perform COP, NOP, and DOP functions regardless of the existence of COPs, NOPs, and DOPs. NOPs can perform all DOP functions regardless of the existence of DOPs. CS supports up to 10 concurrent NOP and DOP connections in addition to one HOP connection. NS and NVF support only the HOP connection.

The HOP can communicate with NAM and any of the supervisors (CS, NS, NETOU, or NVF) executing in its host. COPs can communicate only with NETOU. Refer to the CDCNET Network Operations Manual for additional information. NOPs and DOPs can communicate with any CS executing in any host to which the operator has access. However, a NOP or DOP can communicate with only one CS at a time. A NOP or DOP can monitor and control only those elements owned by NPUs being supervised by their CS.

Each NPU in the network is controlled by the HOP and, optionally, one NOP connected to its supervising CS. A HOP controls all NPUs supervised by CS. Each NOP connected to CS can control one or more NPUs.

NAM Startup

If the NPUs in the network were running version 3.5 or earlier Communications Control Program (CCP), you must master clear all local NPUs before initiating NAM. If the NPUs were running version 3.6 or 3.7 of CCP, then master clear NPUs only if the running version is not the desired version.

You initiate NAM by entering:

NAM.

or

NAMNOGO.

NAM¹ and NAMNOGO are names of procedure files that are called when the HOP initiates NAM.

The first time the network is initiated the NAMNOGO procedure must be used to create a memory file (refer to the NOS Version 2 Installation Handbook). The memory file contains the name of the master file, the last parameter record, and the network invocation number (NIN) used. The NIN is incremented by 1 every time NAM is started. The NIN is displayed in the NAM K display. If the network has been previously initiated, a memory file exists and the NAM procedure file may be used.

Initiating NAM Without Operator Intervention

If you initiate NAM by entering

NAM.

NAM² is brought up without further operator action. This results in starting the network using the previously specified master file and the parameter record RESTRT.

1. Any procedure file of the form NAMffff is allowed.
2. Any procedure file of the form NAMffff is allowed.

Altering the Memory File

If you initiate NAM by entering

```
NAMNOGO.
```

NAM displays the message:

```
ENTER DESIRED OPTION
```

You reply to this message with a CFO command in this format:

```
CFO,NAM.P1,P2,...Pn.
```

p_i is a parameter consisting of a keyword equated to a value. All parameters are order independent.

P_i	Description
MFN=filename	Permanent file name of the master file to be used. The master file contains all the information needed to initiate the network. Once this parameter is specified all subsequent network initiations use this master file until a new master file is specified. The master file resides in the system default family and must be permitted to SYSTEMX. If this parameter is not specified, NAMSTRT is assumed. Refer to NAM Procedure File in the NOS Version 2 Installation Handbook for further information about file NAMSTRT.
UN=username	User name for the master file. If this parameter is not specified, NETOPS is assumed.
PW=password	Password associated with the master file user name.

Pi	Description
RN=recordname ³	Name of the parameter record on the master file that is used for this network initiation. If this parameter is not specified, the previously specified parameter record is used. The parameter record specifies the characteristics of any particular network initiation. You can select one of the parameter record names listed in table 15-1 to obtain the given combination of features.
OIN=xx	Previous network invocation number. All network dump and trace files with a network invocation number less than or equal to the value specified by the OIN parameter are collected. If this parameter is not specified, the current network invocation number minus 1 is used.

You can enter all desired parameters with a single entry or enter several CFO commands. In this way you add (if the file is empty) or replace entries in the NAM memory file.

When you have finished adding or replacing data, enter the command:

CFO, NAM.GO.

After entering this command, you cannot add or replace additional parameters during this NAM initiation.

Table 15-1. Parameter Record Feature Descriptions

Feature Descriptions	INIT	RESTRT	RECOVR	MULTI	MINIT	MRECOV
This feature collects dumps/traces on tape before startup. ³	N ¹	Y ²	N	Y	N	N
This feature stops local NPUs at startup.	Y	N	N	N	Y	N
This feature stops local NPUs upon abnormal NHP termination.	Y	Y	N	N	N	N
This feature takes NPU dumps before initial load. ⁴	N	Y	N	Y	N	N

1. N indicates that the parameter record name causes the action not to occur.
2. Y indicates that the parameter record name causes the action to occur.
3. Dumps/traces are collected for NHP, 255x NPU, and CDCNET whenever present.
4. This feature has no effect on MDIs and MTIs.

3. Specifying a record name changes recordname only for this initiation.

NAM Shutdown

NAM is idled by assigning the K display to the application NVF, followed by entering the DISABLE,HOST or IDLE,HOST command.

The sequence of commands is:

```
K.AP=NVF.
```

```
K.DISABLE,HOST.
```

or

```
K.IDLE,HOST.
```

Use the IDLE command for a graceful shutdown, and the DISABLE command for an immediate shutdown.

Network Control by HOP

The commands described next are available only to the HOP. The HOP can perform all functions of the NPU operator, CDCNET operator, and diagnostic operator as well as those unique to the HOP. The commands unique to the HOP are processed by NS and NVF.

Disabling and Enabling of Couplers

The HOP enables either an NPU or MDI/MTI coupler using the DSD ON command and disables the coupler using the DSD OFF command. You can examine the NP equipment entry on the E,A display to determine the EST ordinal. When an active coupler is turned OFF, NAM stops communicating with that coupler. The coupler is not used again until it is turned ON while NAM is executing. When a coupler is turned ON, NAM automatically begins to communicate with the NPU or MDI/MTI connected across the coupler.

Status Message Formats for HOP

The HOP commands cause various status messages to be displayed at the system console. The normal response to a STATUS command consists of a single status line for each referenced element.

Following are the general formats for these status messages.

Message Type	Format
Application status line	APPL: name, status, conncount, mxcopies, jsn, restart.
Terminal status line	TERM: name, username, hh.mm.ss, applname/acn, ah.am.as, dt.
User status line	USER: name, termname, hh.mm.ss, applname/acn, ah.am.as, dt.

In the preceding status message general formats the following variables are used.

Variable	Description
acn	Connection number for the connected application.
ah.am.as	Time of the connection to the application (hours, minutes, and seconds).
applname	Name of the connected application.
conncount	Count of the active connections to this application.
dt	Device type.
hh.mm.ss	Time of the system or host login (hours, minutes, and seconds).
jsn	Job sequence name of the application.
mxcopies	Maximum number of copies of this application that can be running at one time.

Variable	Description										
name	Name of the specified element.										
restart	If the characters RS are present, they indicate that NVF will automatically start up a copy of this application to complete a connection request for it.										
status	Status of the application. The application program is initially set to either an enabled or disabled state. Enabled application programs become active when they start executing. The application program is in one of these states: <table border="1" data-bbox="396 470 1360 737"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>Application is active (you can establish a connection to the application).</td> </tr> <tr> <td>DI</td> <td>Application is disabled.</td> </tr> <tr> <td>DN</td> <td>Application is down (the application failed; NAM tables are being cleared up).</td> </tr> <tr> <td>EN</td> <td>Application is enabled.</td> </tr> </tbody> </table>	status	Description	AC	Application is active (you can establish a connection to the application).	DI	Application is disabled.	DN	Application is down (the application failed; NAM tables are being cleared up).	EN	Application is enabled.
status	Description										
AC	Application is active (you can establish a connection to the application).										
DI	Application is disabled.										
DN	Application is down (the application failed; NAM tables are being cleared up).										
EN	Application is enabled.										
termname	Name of the terminal at which the user resides.										
username	Login user name of the terminal user.										

Unsolicited Status Reports for HOP

Each time a significant event occurs, CS, NS, or NVF updates its recent history buffer (refer to Recent HISTORY Command later in this section) and logs the event in the NAM dayfile. CS, NS, or NVF then sends a report about the event, called the unsolicited status report, to the HOP. The HOP can get unsolicited status reports from CS, NS, or NVF by assigning the NAM K display to that particular application. Refer to section 8, K-Display Utilities, for information on how to assign the NAM K display. Unsolicited status reports are time-stamped whereas solicited status reports are not time-stamped.

If an unsolicited status report is important enough to require alerting the operator or system analyst, NS, CS, or NVF will display its application name on the alert line of the NAM K-display unless the NAM K-display is already assigned to that application. The operator or system analyst should then assign the NAM K-display to the application and follow the procedures indicated in the NOS Version 2 Operations Handbook for the appropriate alert message.

Refer to the CDCNET Network Operations Manual for NETOU's treatment of unsolicited information for the CDCNET operator.

HOP Commands

The HOP must assign the NAM K display to a supervisory application (CS, NS, or NVF) before entering any of the following commands. Refer to section 8, K-Display Utilities, for information on how to assign the NAM K display. CS, NS, and NVF support both the right and left K-displays. The right K-display provides help information.

All HOP commands consist of a command verb and, sometimes, one or more parameters. The commands must begin in the first character position after K. and contain no blanks. The terminator (.) is optional. Some command verbs and parameters can be shortened to a 2- or 3-character abbreviation. You can mix abbreviated and unabbreviated command verbs and parameters in all network commands. Parameters are order independent except where noted.

NS Control Commands

The HOP assigns the NAM K display to NS to enter the following commands. The NS control commands do not apply to the CDCNET network.

Cancel Alternate NPU Load File Command

After using a FILE command, the HOP can instruct NS to return to the use of the default NPU load file for loading all NPUs by specifying the NOFILE command.

The command format is:

NOFILE.

or

NO.

Change NPU Load File Command

The HOP can assign an alternate load file to NS for loading a specific NPU with the FILE command.

The command format is:

```
FILE,NPU=npuname,NLF=loadfile,UN=username,PW=password.
```

or

```
FI,NP=npuname,NL=loadfile,UN=username,PW=password.
```

Parameter Description

npuname	Name of the NPU.
loadfile	Name of the alternate load file loadfile.
username	Username of the user to which file loadfile belongs.
password	Password for the username to which file loadfile belongs.

The next load of the specified NPU will be from the alternate load file. You can return to the default NPU load file by using the NOFILE command.

NPU Load Status Command

The HOP can request the current load/dump status of a specific NPU or all NPUs using the STATUS command.

There are two forms of the STATUS command.

Specific NPU

The first form requests the current load/dump status of a specific NPU.

The command format is:

```
STATUS,NPU=npuname.
```

or

```
ST,NP=npuname.
```

All NPUs

The second form requests the current load/dump status of all NPUs.

The command format is:

STATUS, NPUS.

or

ST, NPS.

Status Response

The normal status response consists of four lines of information for each NPU, as shown below.

```

npuname      NN=nid  NBS=nid,nid,nid,nid,nid,nid      EC=errcnt
              LDS=ldcnt                LD=yy/mm/dd  LT=hh.mm.ss
dflag      dmpfile HALT ffff DD=yy/mm/dd DT=hh.mm.ss
message
    
```

Status	Description								
npuname	Name of the network processing unit.								
NN=nid	Node number of the NPU identified by npuname.								
NBS=nid,...,nid	Node numbers of up to six neighboring NPUs or couplers.								
EC=errcnt	Protocol error count. Number of dump/load service messages received from the NPU that violated the NPU dump/load protocol.								
LDS=ldcnt	Number of attempted loads of this NPU.								
LD=yy/mm/dd	Date in year, month, and day when the most recent load attempt of this NPU occurred.								
LT=hh.mm.ss	Time in hours, minutes, and seconds when the most recent load attempt of this NPU occurred.								
dflag	NPU memory dump indicator; one of the following: <table border="1" data-bbox="487 1354 1352 1512"> <thead> <tr> <th>dflag</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DUMP</td> <td>Dump NPU before a load.</td> </tr> <tr> <td>NODUMP</td> <td>No NPU dump before a load.</td> </tr> <tr> <td>blank</td> <td>NPU is not currently being loaded.</td> </tr> </tbody> </table>	dflag	Description	DUMP	Dump NPU before a load.	NODUMP	No NPU dump before a load.	blank	NPU is not currently being loaded.
dflag	Description								
DUMP	Dump NPU before a load.								
NODUMP	No NPU dump before a load.								
blank	NPU is not currently being loaded.								
dmpfile	Name of the most recent NPU dump file.								
ffff	Halt code (halt codes are described in the CCP 3 Diagnostic Handbook).								

Status	Description
DD=yy/mm/dd	Date in year, month, and day when the most recent dump of this NPU occurred.
DT=hh/mm/ss	Time in hours, minutes, and seconds when the most recent dump of this NPU occurred.
message	Message is the NS activity and its associated status (table 15-2).

Table 15-2 describes NS activity and its associated status.

Table 15-2. NS Activity

Activity	Status	Description
SAMLOAD	STARTED COMPLETED ABORTED - error	error is any of the messages listed in table 15-3.
DUMP/LOAD DUMP	REQUESTED STARTED COMPLETED ABORTED-error	error is any of the messages listed in table 15-3.
LOAD	modname COMPLETED ABORTED - error	modname is the name of the system module being loaded. error is any of the messages listed in table 15-3.

Table 15-3 lists the error messages associated with NS activity status.

Table 15-3. Load/Dump Status Messages

Error Message	Description
TIMEOUT	No response was received from the NPU during a preset time-out period.
RETRY TIMEOUT	Only abnormal responses were received from the NPU during a preset time-out period.
FILE ERRORS	Irrecoverable I/O error occurred on the network configuration file or network load file.
PREEMPTED	An incomplete SAM load or dump/load sequence was interrupted by a request from the same NPU to initiate a new SAM load.

NS Recent HISTORY Command

The HISTORY command requests a display of recent unsolicited status reports from NS (refer to Recent HISTORY Command later in this section for further information).

NVF Control Commands

The HOP assigns the NAM K display to the network validation facility (NVF) to enter the following commands.

Disable Host Element Command

The HOP uses the DISABLE command to force immediate termination of a particular application or all host network operations.

The command format is:

DISABLE, operation.

or

DI, operation.

Parameter Description

operation One of the following:

<u>operation</u>	<u>Keyword Abbreviation</u>	<u>Description</u>
APPL=applicationname	AP	applicationname is the name of the application to be disabled. If more than one copy of an application is allowed, then applicationname can be a 1- to 5-character root name specifying that all copies of the named application are affected.
HOST	HO	

Results of using the DISABLE,APPL command:

- The status of the named application is changed from enabled (EN), active (AC), or down (DN) to disabled (DI), thereby prohibiting the application from accessing the 255x NPU network or CDCNET network.
- If the application is active at the time the command is entered, NAM requests the application to immediately terminate access to the 255x NPU network or CDCNET network. If the application remains active for an undue period of time after entering this command, the HOP can resort to the DSD DROP or STOP command (refer to section 5, DSD Commands).

Results of using the DISABLE,HOST command:

- NAM requests that all active applications in the host immediately terminate their access to both the 255x NPU network and CDCNET network.
- When all applications (including CS, NS, and NETOU) terminate their access to the corresponding network, NAM and NVF also terminate.
- If NAM or NVF do not terminate, enter the STATUS,APPL command to determine which application has not terminated. You can then use the DSD DROP or STOP command to drop the application.

Enable Application Command

The HOP uses the ENABLE command to change the status of a network application from disabled (DI) to enabled (EN). The application status becomes active (AC) when the application begins to access the network.

The command format is:

```
ENABLE,APPL=applicationname.
```

or

```
EN,AP=applicationname.
```

Parameter	Description
applicationname	Name of the application to be enabled. If more than one copy of an application is allowed, then applicationname can be a 1- to 5-character root name specifying that all copies of the named application are affected.

Idle Host Element Command

The IDLE command terminates execution of an individual application or all host network operations.

The command format is:

```
IDLE,operation.
```

or

```
ID,operation.
```


Parameter Description

operation One of the following:

operation	Keyword Abbreviation	Description
APPL=applicationname	AP	applicationname is the name of the application to be idled. If more than one copy of an application is allowed, then applicationname can be a 1- to 5-character root name specifying that all copies of the named application are affected.
HOST	HO	

Results of using the IDLE,APPL command:

- The application is warned of an imminent shutdown.
- The application completes any operations in progress but does not allow any new connections.
- When all users are disconnected, the application terminates normally.

Results of using the IDLE,HOST command:

- Applications are warned of an imminent shutdown.
- New application requests to access either the 255x NPU network or CDCNET network are rejected.
- New application to application connections are rejected.
- New terminal connections to the host from either the 255x NPU network or the CDCNET network are rejected by NAM.
- CS warns all connected NOPs and DOPs that shutdown of the network is imminent. CS terminates when all NOPs and DOPs end their connections to CS. COPs must be warned separately using CDCNET network operator commands.
- NS ignores any new dump or load requests from NPUs. Any ongoing dump and load processes are completed normally before NS terminates.
- Requests for MDI/MTI loads and dumps continue to be honored and NAM continues to start INITMDI. However, requests for loading and dumping other DIs are rejected, since connection requests from MDIs to NETFS are refused.
- When all applications (including NS, CS, and NETOU) terminate their access to the network, NAM and NVF also terminate. If an application does not terminate, use the DISABLE command to request immediate termination of network activity in the host.

NVF Recent HISTORY Command

The HISTORY command requests a display of recent unsolicited status reports from NVF (refer to Recent HISTORY Command later in this section for further information).

Status Host Element Command

Host elements are the applications associated with the host, terminals currently connected to the host, and users currently logged in to the host. With the STATUS command a HOP can status a single application, terminal device, or user name; or all applications, terminal devices, or users connected to a given application.

There are five forms of the STATUS command. Refer to Status Message Formats for HOP, earlier in this section, for the format of the status line returned.

Particular Host Element

The first form requests the status of a particular host element.

The command format is:

STATUS,element.

or

ST,element.

Parameter Description

element One of the following. (If the following elementname is the name of an application and more than one copy of the application is allowed, then elementname can be a 1- to 5-character root name specifying that all copies of the named application are affected.)

<u>element</u>	<u>Keyword Abbreviation</u>	<u>Description</u>
APPL=elementname	AP	elementname is the name of the element whose status is desired.
TERM=elementname ⁴	TE	elementname is the name of the element whose status is desired.
UNAME=elementname	UN	elementname is the name of the element whose status is desired.

⁴ elementname may be the terminal name of a CDCNET terminal. Terminal names for CDCNET are assigned dynamically and are not associated with terminal names specified in the network configuration file.

Network Applications

The second form requests the status of all network applications in the host.

The command format is:

STATUS, APPLS.

or

ST, APS.

Active Network Applications

The third form requests the status of all active network applications in the host.

The command format is:

STATUS, APPLS, AC.

or

ST, APS, AC.

Terminal Devices

The fourth form requests the status of all terminal devices connected to the specified application or logged in under the specified user name.

The command format is:

STATUS, TERMS, termelement.

or

ST, TES, termelement.

Parameter Description

termelement One of the following:

termelement	Keyword Abbreviation	Description
APPL=elementname	AP	elementname is the name of the desired application or user.
UNAME=elementname	UN	elementname is the name of the desired application or user.

Users (Specified Application)

The fifth form requests the status of all users connected to the specified application.

The command format is:

STATUS, UNAMES, APPL=applicationname.

or

ST, UNS, AP=applicationname.

Parameter	Description
applicationname	Name of the desired application.

Network Control by NOP

The commands described later in this section are available to both HOPs and NOPs. The NOPs can perform all functions of the diagnostic operators as well as those unique to NOPs.

The controlling NOP (refer to CONTROL command later in this section) is responsible for controlling one or more NPUs and its associated couplers, lines, logical links, terminal devices, and trunks.

A controlling NOP can control any of the NPUs supervised by a CS and shares that responsibility with the HOP. The controlling NOP and HOP must, therefore, coordinate their activities.

CS does not control or report the status of CDCNET network elements; therefore, none of the following commands apply to CDCNET.

Becoming a NOP

To become a NOP, you must be authorized by site personnel to access the CS application and to become a controlling NOP.

After initiation of the network, you become a NOP as follows:

1. Log in as described in the NOS Version 2 Reference Set, Volume 2.
2. In response to the APPLICATION: prompt, enter:

CS

3. The system responds:

HELLO, YOU ARE NOW AN NPU OPERATOR

This message is followed by the prompt:

READY . .

You can now enter the NOP commands described later in this section.

You can cancel a command that results in a large volume of output by entering the interruption or termination sequence (refer to the NOS Version 2 Reference Set, Volume 3).

NOTE

Use caution when cancelling output because unsolicited status messages could also be discarded.

When you cancel output, the message

```
OUTPUT DISCARDED  
READY . .
```

or

```
OUTPUT DISCARDED . .
```

appears at the terminal at the point from which the data was discarded.

If the terminal is in page wait mode,⁵ you must press CR after every page of output before more output can be sent to the terminal.

To relinquish NOP status, enter one of the following.

- To log out and end the terminal connection to the host, enter:

```
GOODBYE
```

or

```
BYE
```

or

```
LOGOUT
```

- To initiate a new login dialog, enter:

```
HELLO
```

or

```
LOGIN
```

- To remain logged in and be prompted for another application, enter:

```
END
```

- To remain logged in and be switched automatically to the application named applname, enter:

```
END, applname
```

5. Use a hardcopy device for NOP terminals; this prevents the paging problem and provides a permanent log of network events.

Status Message Formats for NOP

Some of the NOP commands cause various status messages to be displayed on your terminal. The normal response to a STATUS command is one line of text returned when status is requested for a coupler, line, NPU, or terminal device. Two text lines can be returned when status is requested for logical links or trunks (if the CS to which you are connected is supervising NPUs at each end of the element, a line of text is returned for each end of the logical link or trunk). For multiple element status commands, a response is generated for all elements expressed or implied in the command. For example, a status request for all terminals on an NPU begins with a status report (text line) for the NPU. This is followed by a status report for the first active (AC) line on the NPU. A status report for each terminal device configured on the line follows. This is followed by a status report of the next active line and so on.

The response to the ENABLE or DISABLE command is a status report for each element affected by that command.

Following are the general formats for these status messages.

Message Type	Format
Coupler status line	CPLER: name, status, nodenum, npuname1/cplnum.
Line status line	LINE: name, status, linetype, npuname1/port.
Logical link status line	LLINK: name, status, RL=n, loglinktype, npuname1/nodenum, npuname2/nodenum.
NPU status line	NPU: name, status, nodenum, gostatus, dump.
Archetype switched virtual circuit (SVC) status line	SVC: atname, EN=ennum, DI=dinum, AC=acnum, stip.
Terminal device status line	TERM: name, status, dt/tc, /hostnid, linename/a1/a2, ldstat.
Trunk status line	TRUNK: name, status, linetype, npuname1/port, npuname2/port.

In the preceding status message general formats the following variables are used.

Variable	Description
acnum	Number of SVCs of the element that are in an active state (from 0 to 255).
atname	Name of the archetype SVC element.
a1	First level address (cluster address, configuration ordinal).
a2	Second level address (terminal address, stream number).
cplnum	Coupler number for the indicated NPU.
dinum	Number of SVCs of the element that are in a disabled state (from 0 to 255).
dt	Device type (CON, CR, LP, CP, PL).
dump	NPU dump flag (ON or OFF).
ennum	Number of SVCs of the element that are in an enabled state (from 0 to 255).
gostatus	Go status of the NPU (if waiting for GO, NEEDGO appears).
hostnid	Host (coupler) node number of the host to which the terminal is currently connected.
ldstat	Load status of the trunk (LDDI means the trunk is disabled for loading; LDEN means the trunk is enabled for loading).
linename	Name of the line to which the terminal device is connected.
linetype	Line type (refer to the NDL Reference Manual).
loglinktype	Logical link type (H-N for host to NPU, H-H for host to host).
name	Name of the element.
nodenum	Node number of a coupler or NPU.
npuname1	Name of an NPU that supports or is connected to the element (also the name of the NPU that reported the status).
npuname2	Name of the NPU at the other end of the logical link or trunk.
port	Port number of the line or trunk on the indicated NPU.

Variable	Description												
RL=n	Regulation level currently in effect.												
status	Network elements are initially set to either a disabled or enabled state. Initially enabled elements automatically become active during the course of normal network operation. Every network element (such as line, logical link, NPU, or trunk) is in one of the following states.												
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>Network element is active. An element becomes active when it is enabled and is handling network data traffic.</td> </tr> <tr> <td>DI</td> <td>Network element is disabled. The element cannot become active unless enabled by an operator command.</td> </tr> <tr> <td>DN</td> <td>Network element is down. A failure was detected that terminated data traffic for the element. If the failure condition ceases, the element returns to an enabled status.</td> </tr> <tr> <td>EN</td> <td>Network element is enabled. It was enabled by the operator and is in a physically operative condition.</td> </tr> <tr> <td>NC</td> <td>Network element is not configured. The NPU does not recognize the element.</td> </tr> </tbody> </table>	status	Description	AC	Network element is active. An element becomes active when it is enabled and is handling network data traffic.	DI	Network element is disabled. The element cannot become active unless enabled by an operator command.	DN	Network element is down. A failure was detected that terminated data traffic for the element. If the failure condition ceases, the element returns to an enabled status.	EN	Network element is enabled. It was enabled by the operator and is in a physically operative condition.	NC	Network element is not configured. The NPU does not recognize the element.
status	Description												
AC	Network element is active. An element becomes active when it is enabled and is handling network data traffic.												
DI	Network element is disabled. The element cannot become active unless enabled by an operator command.												
DN	Network element is down. A failure was detected that terminated data traffic for the element. If the failure condition ceases, the element returns to an enabled status.												
EN	Network element is enabled. It was enabled by the operator and is in a physically operative condition.												
NC	Network element is not configured. The NPU does not recognize the element.												
stip	SubTIP type (PAD or APPL).												
tc	Terminal class.												

Unsolicited Status Reports for NOP

Each time a significant event occurs, CS updates its recent history buffer (refer to Recent HISTORY Command later in this section) and logs the event in the NAM dayfile. CS then sends a report about the event, called the unsolicited status report, to all NOPs who have selected to receive unsolicited status reports from the NPU (refer to Report Unsolicited Status Command later in this section). Unsolicited status reports are time-stamped whereas solicited status reports are not time-stamped.

A single unsolicited status report occurs as a result of an element recovery or failure. An unsolicited status message is sent when an NPU joins the network. Batches of unsolicited status reports can occur as a result of an operator command to enable all elements of a given type (refer to Enable Network Element Command later in this section).

NOP Commands

All NOP commands consist of a command verb and, sometimes, one or more parameters. The commands must begin in the first character position and contain no blanks. The terminator (.) is optional. Some command verbs and parameters can be shortened to a 2- or 3-character abbreviation. You can mix abbreviated and unabbreviated command verbs and parameters in all network commands. Parameters are order independent except where noted. You can enter commands pertaining only to the NPUs supervised by the CS to which you are connected.

To enter the following commands you must control the NPU (refer to CONTROL Command later in this section).

ALERT Command

The ALERT command changes the threshold that CS uses for alerting the HOP and the controlling NOP of NPU saturation.

The command format is:

ALERT, CP=cpuse, BU=buff, element.

or

AL, CP=cpuse, BU=buff, element.

Parameter Description

cpuse	Percent of the NPU central processor currently in use (from 50 to 100 percent).
buff	Number of available NPU data buffers (from 0 to 500).
element	One of the following:

element	Keyword Abbreviation	Description
NPU=npuname	NP	npuname is the name of the NPU supervised by your CS.
NPUS	NPS	

If the CPU usage reported to CS by the NPU is greater than cpuse, the HOP and the controlling NOP are alerted. If the buffer availability reported to CS by the NPU is less than buff, the HOP and the controlling NOP are alerted. To turn off all messages, set CP=100 and BU=0.

CONTROL Command

After establishing communication with a CS, you can gain or lose control of one or more NPUs supervised by your CS with the CONTROL command.

The command format is:

CONTROL, option, status.

or

CO, option, status.

Parameter Description

option	Description
option	One of the following:
	Keyword
option	Abbreviation Description
NPU=npuname	NP npuname is the name of the NPU.
NPUS	NPS
AUTO	AU
status	ON or OFF.

The NPU=npuname option allows you to gain control of a specific NPU. The NPUS option allows you to gain control over all uncontrolled NPUs currently supervised by your CS. The AUTO option allows you to gain control over all uncontrolled NPUs currently supervised by your CS and all other uncontrolled NPUs that may later come under the supervision of your CS. The AUTO option can be cancelled by entering any valid CONTROL command with the OFF status specified.

Example:

Suppose there are four NPUs (1, 2, 3, and 4) available to be controlled in the network. The first NOP logs in to CS and controls NPU 1. The second NOP establishes a connection to CS and enters the CONTROL,AUTO,ON command. The second NOP becomes the controlling NOP and controls all of the uncontrolled NPUs (in this example 2, 3, and 4). When the first NOP relinquishes control of NPU 1, that NPU is then controlled by the second NOP. The second NOP also controls all of the new NPUs joining the network automatically.

If no other NOP exists for the NPUs specified, the response to a request for control consists of a list of the NPUs you now control followed by a READY.. prompt; otherwise, you receive a diagnostic message.

Since a HOP automatically controls all NPUs being supervised by a CS, the CONTROL commands are not used by the HOP. Whenever a NOP loses control of an NPU for any reason (including ending the connection to CS or the failure of some network element that was supporting the connection to CS), CS informs the HOP and all other connected NOPs by sending an unsolicited status report.

If any NOP has AUTO control status, the operators are informed of the identity of the controlling NOP. Whenever a CS gains supervision of a new NPU, the HOP, all NOPs, and all DOPs are informed that the NPU is being supervised.

Disable Network Element Command

The controlling NOP changes the status of a network element from enabled (EN), active (AC), or down (DN), to disabled (DI) by the **DISABLE** command.

There are three forms of the **DISABLE** command. Refer to **Status Message Formats for NOP**, earlier in this section, for the format of the status line returned.

Single Element

The first form disables a single element in a network.

The command format is:

```
DISABLE,element.
```

or

```
DI,element.
```

Parameter Description

element	One of the following:		
element	Keyword Abbreviation	Description	
LINE=elementname	LI	elementname is the name of the element to be disabled.	
LLINK=elementname	LL	elementname is the name of the element to be disabled.	
TERM=elementname	TE	elementname is the name of the element to be disabled.	
TRUNK=elementname	TR	elementname is the name of the element to be disabled.	

For a logical link or trunk, if the element is active when the command is entered, all data traffic through the element is terminated and the affected host application programs and terminal users are notified. For a line or terminal, if the element is active when the command is entered, all data traffic through the element is terminated and only the affected host application programs are notified.

Trunk

The second form disables a trunk for loading.

The command format is:

```
DISABLE, TRUNK=trunkname, LD.
```

or

```
DI, TRUNK=trunkname, LD.
```

Parameter Description

trunkname Name of trunk to be disabled for loading.

If the trunk is enabled, it will still be enabled after this command is issued, but it will be disabled for loading.

SVCs

The third form disables archetypical switched virtual circuits (SVCs) that are in an enabled state.

The command format is:

```
DISABLE, SVC=atname, NCIR=numcir.
```

or

```
DI, SV=atname, NC=numcir.
```

Parameter Description

atname Archetypical SVC name.

numcir Number of SVCs to change from an enabled state to a disabled state.

DISABLE Processing

This form of the DISABLE command allows you to decrease the number of enabled SVCs for a particular X.25 subTIP type. If the number of enabled SVCs is greater than or equal to numcir, then the number of disabled SVCs is increased by numcir and the number of enabled SVCs is decreased by numcir. If the number of enabled SVCs is less than numcir, then the number of disabled SVCs is increased by the number of available enabled SVCs and the number of enabled SVCs is set to zero. Disabling SVCs has no effect on active SVCs. If, for a particular SVC element, the command could not be executed, an unsolicited status report is sent to you giving the reason the command was not executed.

The NPU's affected by the command send an unsolicited status report indicating the new status for the element.

Trunks and logical links can have two owning NPUs (one at each end of the trunk or logical link). When two NPUs are involved and both are supervised by your CS, the `DISABLE,TRUNK` and `DISABLE,LLINK` commands are processed as follows:

- If one of the NPUs is local and one is remote, the command is forwarded to the local NPU.
- If both of the NPUs are local or both are remote, the command is forwarded to the NPU with the lower node number.
- If your CS supervises only one of the NPUs, the command is forwarded to that NPU.

In either case, the trunk or logical link enters the disabled (DI) state at the owning NPU that receives the `DISABLE` command and is detected as being down (DN) by the other NPU. Both of these changes are reported to CS, which informs all operators of the changes by sending unsolicited status reports.

If a single NPU owns both ends of the link, the command is forwarded to that NPU. This status change is also reported by sending unsolicited status reports.

Disabling a trunk can cause one or more logical links to go down. The owning NPUs report the failure of these links by sending unsolicited status reports.

NOTE

No attempt is made to protect you from disabling network elements that support your connection to CS. You may lose contact with CS as a result of such a command. When this happens, you may need to use another terminal or the host console to enable the disabled network element.

Enable Network Element Command

The controlling NOP changes the status of one or more network elements from disabled (DI) to enabled (EN), allowing the element(s) to become active for data traffic by the ENABLE command.

There are six forms of the ENABLE command. Refer to the Status Message Formats for NOP, earlier in this section, for the format of the status line returned.

Single Element

The first form enables a single element in the network.

The command format is:

ENABLE, element.

or

EN, element.

Parameter Description

Parameter	Description															
element	One of the following:															
	<table border="1"> <thead> <tr> <th><u>element</u></th> <th><u>Keyword Abbreviation</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>LINE=elementname</td> <td>LI</td> <td>elementname is the name of the element to be enabled.</td> </tr> <tr> <td>LLINK=elementname</td> <td>LL</td> <td>elementname is the name of the element to be enabled.</td> </tr> <tr> <td>TERM=elementname</td> <td>TE</td> <td>elementname is the name of the element to be enabled.</td> </tr> <tr> <td>TRUNK=elementname</td> <td>TR</td> <td>elementname is the name of the element to be enabled.</td> </tr> </tbody> </table>	<u>element</u>	<u>Keyword Abbreviation</u>	<u>Description</u>	LINE=elementname	LI	elementname is the name of the element to be enabled.	LLINK=elementname	LL	elementname is the name of the element to be enabled.	TERM=elementname	TE	elementname is the name of the element to be enabled.	TRUNK=elementname	TR	elementname is the name of the element to be enabled.
<u>element</u>	<u>Keyword Abbreviation</u>	<u>Description</u>														
LINE=elementname	LI	elementname is the name of the element to be enabled.														
LLINK=elementname	LL	elementname is the name of the element to be enabled.														
TERM=elementname	TE	elementname is the name of the element to be enabled.														
TRUNK=elementname	TR	elementname is the name of the element to be enabled.														

Trunk

The second form enables a trunk for loading. It will also enable the trunk if it is disabled.

The command format is:

```
ENABLE, TRUNK=trunkname, LD.
```

or

```
EN, TRUNK=trunkname, LD.
```

Parameter	Description
trunkname	Name of the trunk to be enabled for loading.

Disabled Elements

The third form enables all disabled elements of the specified type on an NPU.

The command format is:

```
ENABLE, element, NPU=npuname.
```

or

```
EN, element, NP=npuname.
```

Parameter	Description
element	One of the following:
element	Keyword Abbreviation
LINES	LIS
LLINKS	LLS
TERMS	TES
TRUNKS	TRS
npuname	Name of an NPU.

Disabled Elements (All NPUs)

The fourth form enables all disabled elements of the specified type on all NPUs.

The command format is:

ENABLE, element, NPUS.

or

EN, element, NPS.

Parameter Description

operation One of the following:

element	Keyword Abbreviation
----------------	-----------------------------

LINES	LIS
-------	-----

LLINKS	LLS
--------	-----

TERMS	TES
-------	-----

TRUNKS	TRS
--------	-----

Disabled Terminals

The fifth form enables all disabled terminals on a line.

The command format is:

ENABLE, TERMS, LINE=termline.

or

EN, TES, LI=termline.

Parameter Description

termline Name of the line to which terminals are connected.

SVCs

The sixth form enables SVCs that are in a disabled state.

The command format is:

```
ENABLE, SVC=atname, NCIR=numcir.
```

or

```
EN, SV=atname, NC=numcir.
```

Parameter Description

atname	Archetype SVC name.
numcir	Number of SVCs to change from a disabled state to an enabled state.

ENABLE Processing

The ENABLE,TRUNK and ENABLE,LLINK commands are forwarded to both of the owning NPUs (one at each end of the trunk or logical link) if both are supervised by your CS. Enabling a trunk can cause one or more logical links with a down status to become enabled or active. Status changes resulting from enabling a trunk or logical link are reported by the owning NPUs by sending unsolicited status reports.

Enabling a trunk can cause CS to become available to an NPU and, therefore, allow CS to gain control of an NPU.

The ENABLE command allows you to enable all lines or terminal devices on an NPU. When you enable all lines or terminal devices of a given type on an NPU or all terminals on a line, the NPU or line affected by the command sends an unsolicited status report telling you how many elements of the given type are currently disabled. The NPU then sends unsolicited status reports telling the new status of each element for which the command was executed. If, for a particular element, the command could not be executed, an unsolicited status report is sent to you giving the reason the command was not executed.

The ENABLE,SVC command allows you to increase the number of enabled SVCs for a particular X.25 subTIP type. If the number of disabled SVCs is greater than or equal to numcir, then the number of enabled SVCs is increased by numcir. If the number of disabled SVCs is less than numcir, then the number of enabled SVCs is increased by the number of disabled SVCs and the number of disabled SVCs is set to zero. Enabling SVCs has no effect on active SVCs. If, for a particular SVC element, the command could not be executed, an unsolicited status report is sent to you giving the reason the command was not executed.

NPU GO Command

After establishing a connection with CS, the newly supervised NPU may ask for a GO command from the controlling NOP.

The GO command format is:

GO,NPU=npuname.

or

GO,NP=npuname.

Parameter Description

npuname Network defined name of the NPU.

If there is a controlling NOP for the NPU when CS receives a GO request, the GO request is relayed to the controlling NOP and to the HOP. If there is no controlling NOP, the GO request is relayed to the HOP and to all currently connected NOPs. The HOP or controlling NOP (if there is one) can, as needed, request status information or enable/disable network elements. Thus, the controlling NOP can control the initial status or the network elements connected to or supported by a given NPU prior to that NPU beginning terminal operations.

The GO request message is repeated periodically until an operator takes control of the NPU and enters the GO command.

Upon entry of the GO command the NPU proceeds to activate lines and accept terminal connections.

NPU LOAD Command

The LOAD command causes an active (AC) NPU or all active NPUS to be stopped, reloaded, and reactivated.

There are two forms of the LOAD command.

Specified NPU

The first form causes the specified NPU to be stopped, reloaded, and reactivated.

The command format is:

```
LOAD, NPU=npuname, dumpoption.
```

or

```
LO, NP=npuname, dumpoption.
```

Parameter	Description
npuname	Name of the NPU.
dumpoption	Optional dump parameter; one of the following:
dumpoption	Keyword Abbreviation
DUMP	DU
NDMP	ND

If the DUMP option is specified, a memory dump of the NPU is taken prior to reloading. If the NDMP option is specified, the NPU is reloaded without taking a memory dump. If the dump status is not specified, taking a memory dump depends upon the setting of the dump flag. Refer to NPU Memory Dump Command, next, for more information. The DUMP or NDMP option is in effect only for the immediate load. It has no permanent effect on the NPU memory dump flag.

When an NPU executes a LOAD command, the CS connection to that NPU is lost during the dump/load process. The trunks and logical links supported by the NPU are detected as down (DN) by the NPUs at the other ends of the trunks and logical links. All of the appropriate operators receive corresponding unsolicited status reports.

All NPUs

The second form causes all active NPUs to be stopped, reloaded, and reactivated.

The command format is:

LOAD, NPUS, dumpoption.

or

LO, NPS, dumpoption.

Parameter Description

dumpoption Optional parameter; one of the following:

dumpoption	Keyword Abbreviation
DUMP	DU
NDMP	ND

If the DUMP option is specified, a memory dump of all active NPUs (supervised by this host) is taken prior to reloading. If the NDMP option is specified, the NPUs are reloaded without taking a memory dump. If the dump status is not specified, taking a memory dump depends upon the setting of the dump flag. Refer to NPU Memory Dump Command, next, for more information. The DUMP or NDMP option is in effect only for the immediate load. It has no permanent effect on the NPUs memory dump flag.

When all of the NPUs execute LOAD commands, the corresponding CS connections to the NPUs are lost during the dump/load process. The trunks and logical links supported by the NPUs are detected as down (DN) by the NPUs at the other ends of the trunks and the logical links. The appropriate operators receive corresponding unsolicited status reports.

NPU Memory DUMP Command

The DUMP command causes the memory dump flag to be reset for the specified NPU or NPUs. Initially, this reset condition is set independently for every NPU in the network and is specified on the NPU statement (refer to the NDL Reference Manual for more information on the NPU statements).

There are two forms of the DUMP command.

Specified NPU

The first form causes a memory dump of the specified NPU to be taken.

The command format is:

```
DUMP,NPU=npuname, status.
```

or

```
DU,NP=npuname, status.
```

Parameter	Description
npuname	Name of the NPU.
status	ON or OFF.

All NPUs

The second form causes a memory dump of all NPUs to be taken.

The command format is:

```
DUMP,NPUS, status.
```

or

```
DU,NPS, status.
```

Parameter	Description
status	ON or OFF.

If status is ON, a memory dump will be taken the next time the NPU is reloaded (unless overridden by the LOAD command). If status is OFF, a memory dump will not be taken the next time the NPU is reloaded (unless overridden by the LOAD command).

Send Message Command

A controlling NOP can send a message to the HOP or any terminal device user or group of terminal device users using the SEND command. The NOP can address the message to a single user or to all users on the specified line, logical link, NPU, or all NPUs under the supervision of the CS.

The HOP cannot use the HOST element to send a message.

The command format is:

```
SEND, element, MSG=message.
```

or

```
SE, element, MS=message.
```

Parameter Description

element	One of the following:		
<u>element</u>	<u>Keyword</u>	<u>Abbreviation</u>	<u>Description</u>
LINE=elementname	LI		elementname is the name of the element to which the message is addressed.
LLINK=elementname	LL		elementname is the name of the element to which the message is addressed.
NPU=elementname	NP		elementname is the name of the element to which the message is addressed.
TERM=elementname	TE		elementname is the name of the element to which the message is addressed.
HOST	HO		
NPUS	NPS		
message	Text message of no more than 50 characters. If more than 50 characters are entered, a diagnostic message is issued. The MSG=message parameter is order dependent and must appear as the last parameter of the command.		

CS, upon accepting a SEND command, does not issue the READY.. prompt until processing of the SEND command has completed, to prevent an NPU from being flooded with operator messages. CS sends a message to you as each NPU receives the SEND command. When processing is complete, CS sends the READY.. message.

If a terminal is specified as the network element, it must be an interactive terminal device. Otherwise, an error message is issued to the sender.

If the element referenced by a SEND command is not currently active or enabled, you receive an error message. This error message is followed by a status line for the element; the operator is thus informed why the message could not be delivered. The status text line is followed by the READY.. prompt.

Network Control by DOP

The commands described next are available to the HOPs, NOPs, and DOPs, but not COPs. All DOPs can obtain unsolicited status reports and use status and diagnostic commands.

Becoming a DOP

To become a DOP you must be authorized by site personnel to access CS.

After initiation of the network, you become a DOP as follows:

1. Log in as described in the NOS Version 2 Reference Set, Volume 2.
2. In response to the APPLICATION: prompt, enter:

```
CS
```

3. The system responds:

```
HELLO, YOU ARE NOW A DIAGNOSTIC OPERATOR
```

This message is followed by the prompt:

```
READY..
```

You can now enter the DOP commands described later in this section.

You can cancel a command that results in a large volume of output by entering the interruption or termination sequence (refer to the NOS Version 2 Reference Set, Volume 3).

NOTE

Use caution when cancelling output because unsolicited status messages could also be discarded.

When you cancel output, the message

```
OUTPUT DISCARDED
READY..
```

or

```
OUTPUT DISCARDED..
```

appears at the terminal at the point from which the data was discarded.

If the terminal is in page wait mode,⁶ you must press CR after every page of output before more output can be sent to the terminal.

To relinquish DOP status, enter one of the following:

- To logout and end the terminal connection to the host, enter:

GOODBYE

OR

BYE

OR

LOGOUT

- To initiate a new login dialog, enter:

HELLO

OR

LOGIN

- To remain logged in and be prompted for another application, enter:

END

- To remain logged in and be switched automatically to the application named applname, enter:

END, applname

6. Use a hardcopy device for DOP terminals; this prevents the paging problem and provides a permanent log of network events.

Status Message Formats for DOP

Some of the DOP commands cause various status messages to be displayed on your terminal. The normal response to a STATUS command is one line of text returned when status is requested for a coupler, line, NPU, or terminal device. Two text lines can be returned when status is requested for logical links or trunks (if the CS to which you are connected is supervising NPUs at each end of the element, a line of text is returned for each end of the logical link or trunk). For multiple element status commands, a response is generated for all elements expressed or implied in the command. For example, a status request for all terminals on an NPU begins with a status report (text line) for the NPU. This is followed by a status report for the first active (AC) line on the NPU. A status report for each terminal device configured on the line follows. This is followed by a status report of the next active line and so on.

The following are the general formats for these status messages.

Message Type	Format
Coupler status line	CPLER: name, status, nodenum, npuname1/cplnum.
Line status line	LINE: name, status, linetype, npuname1/port.
Logical link status line	LLINK: name, status, RL=n, loglinktype, npuname1/nodenum, npuname2/nodenum.
NPU status line	NPU: name, status, nodenum, gostatus, dump.
Archetype switched virtual circuit (SVC) status line	SVC: atname, EN=enum, DI=dinum, AC=acnum, stip.
Terminal device status line	TERM: name, status, dt/tc, /hostnid, linename/a1/a2.
Trunk status line	TRUNK: name, status, linetype, npuname1/port, npuname2/port.

In the preceding status message general formats the following variables are used.

Variable	Description
acnum	Number of SVCs of the element that are in an active state (from 0 to 255).
atname	Name of the archetype SVC element.
a1	First level address (cluster address, configuration ordinal).
a2	Second level address (terminal address, stream number).
cplnum	Coupler number for the indicated NPU.
dinum	Number of SVCs of the element that are in a disabled state (from 0 to 255).
dt	Device type (CON, CR, LP, CP, PL).
dump	NPU dump flag (ON or OFF).
ennum	Number of SVCs of the element that are in an enabled state (from 0 to 255).
gostatus	Go status of the NPU (if waiting for GO, NEEDGO appears).
hostnid	Host (coupler) node number of the host to which the terminal is currently connected.
linename	Name of the line to which the terminal device is connected.
linetype	Line type (refer to the NDL Reference Manual).
loglinktype	Logical link type (H-N for host to NPU, H-H for host to host).
name	Name of the element.
nodenum	Node number of a coupler or NPU.
npuname1	Name of an NPU that supports or is connected to the element (also the name of the NPU that reported the status).
npuname2	Name of the NPU at the other end of the logical link or trunk.
port	Port number of the line or trunk on the indicated NPU.
RL=n	Regulation level currently in effect.

Variable	Description												
status	Network elements are initially set to either a disabled or enabled state. Initially enabled elements automatically become active during the course of normal network operation. Every network element (such as line, logical link, NPU, or trunk) is in one of the following states.												
	<table border="1"> <thead> <tr> <th>status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>Network element is active. An element becomes active when it is enabled and is handling network data traffic.</td> </tr> <tr> <td>DI</td> <td>Network element is disabled. The element cannot become active unless enabled by an operator command.</td> </tr> <tr> <td>DN</td> <td>Network element is down. A failure was detected that terminated data traffic for the element. If the failure condition ceases, the element returns to an enabled status.</td> </tr> <tr> <td>EN</td> <td>Network element is enabled. It was enabled by the operator and is in a physically operative condition.</td> </tr> <tr> <td>NC</td> <td>Network element is not configured. The NPU does not recognize the element.</td> </tr> </tbody> </table>	status	Description	AC	Network element is active. An element becomes active when it is enabled and is handling network data traffic.	DI	Network element is disabled. The element cannot become active unless enabled by an operator command.	DN	Network element is down. A failure was detected that terminated data traffic for the element. If the failure condition ceases, the element returns to an enabled status.	EN	Network element is enabled. It was enabled by the operator and is in a physically operative condition.	NC	Network element is not configured. The NPU does not recognize the element.
status	Description												
AC	Network element is active. An element becomes active when it is enabled and is handling network data traffic.												
DI	Network element is disabled. The element cannot become active unless enabled by an operator command.												
DN	Network element is down. A failure was detected that terminated data traffic for the element. If the failure condition ceases, the element returns to an enabled status.												
EN	Network element is enabled. It was enabled by the operator and is in a physically operative condition.												
NC	Network element is not configured. The NPU does not recognize the element.												
stip	SubTIP type (PAD or APPL).												
tc	Terminal class.												

Unsolicited Status Reports for DOP

Each time a significant event occurs, CS updates its recent history buffer (refer to Recent HISTORY Command later in this section) and logs the event in the NAM dayfile. CS then sends a report about the event, called the unsolicited status report, to all DOPs who have selected to receive unsolicited status reports from the NPU (refer to Report Unsolicited Status Command later in this section). Unsolicited status reports are time-stamped whereas solicited status reports are not time-stamped.

DOP Commands

All DOP commands consist of a command verb and, sometimes, one or more parameters. The command must begin in the first character position and contain no blanks. The terminator (.) is optional. Some command verbs and parameters can be shortened to a 2- or 3-character abbreviation. You can mix abbreviated and unabbreviated command verbs and parameters in all network commands. Parameters are order independent except where noted. You can enter commands pertaining only to the NPUs supervised by the CS to which you are connected.

CS Information Command

You can access information about all operators and NPUs connected to and supervised by your CS with the INFO command.

The command format is:

INFO.

or

IN.

The response lines are an operator information line and an NPU information line for each operator (DOP and NOP) and for each NPU supervised by a CS, respectively.

The format of the operator information line is:

DOP: termname, username, acn.
NOP: termname, username, acn, auto.

The format of the NPU information line is:

NPU: npuname, status, nodenum, go-status, nopname.

In the preceding information line formats the following variables are used.

Variable	Description
acn	CS application connection number.
auto	Control status of the NOP (if the NOP is an auto control NOP, AUTO appears). Refer to the NOP CONTROL command described earlier in this section.
nodenum	Node number of the NPU.
nopname	Terminal name of the NOP that is controlling the NPU. If the NPU has no controlling NOP, the field is blank.
npuname	Name of the supervised NPU.
status	Status of the NPU. (AC if the NPU is active, MI if the NPU is in a mismatch condition).
termname	Name of the terminal at which the DOP or NOP resides.
username	User name of the DOP or NOP.
go-status	If the NPU is waiting for GO, NEEDGO is displayed.

NPU Diagnostic Test Command

The TST command allows any operator to request, interact with, or terminate the diagnostic program within any NPU supervised by a CS.

The command format is:

TST,NPU=npuname,option.

or

TS,NP=npuname,option.

Parameter	Description												
npuname	Name of the NPU containing the diagnostic program with which the operator wants to communicate.												
option	One of the following:												
	<table border="1"> <thead> <tr> <th>option</th> <th>Keyword Abbreviation</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DROP</td> <td>DR</td> <td></td> </tr> <tr> <td>MSG=message</td> <td>MS</td> <td>message is the command that controls the execution of the diagnostic program in the NPU.</td> </tr> <tr> <td>REQUEST</td> <td>RE</td> <td></td> </tr> </tbody> </table>	option	Keyword Abbreviation	Description	DROP	DR		MSG=message	MS	message is the command that controls the execution of the diagnostic program in the NPU.	REQUEST	RE	
option	Keyword Abbreviation	Description											
DROP	DR												
MSG=message	MS	message is the command that controls the execution of the diagnostic program in the NPU.											
REQUEST	RE												

CS keeps track of which NPUs are currently running diagnostic tests and which operators are using them. A single operator is permitted to perform diagnostic tests in several different NPUs simultaneously. All operators can interact with the diagnostics in a single NPU, but only one operator can run the diagnostics.

TST Drop Option

The DROP parameter cancels your diagnostic program reservation in the specified NPU. All diagnostics currently being performed by you in the specified NPU are terminated. If the diagnostic program is not executing, the TST command is ignored and an error message is returned.

TST Message Option

The MSG=message parameter allows you to interact with an NPU diagnostic program. It allows you to send a command message in the form of a text string to the diagnostic program in the specified NPU. A typical use of the message option is to initiate and terminate diagnostic tests. The option is accepted only if you previously entered and confirmed a REQUEST option; otherwise, an error message is received. Refer to the CCP 3 Diagnostic Handbook for more information.

The diagnostic program sends text messages through CS to you. The message is an unsolicited status report with this format:

```
FROM npuname/DIAGNOSTIC: TO termname:
    messagetext
```

Variable	Description
npuname	Name of the NPU.
termname	Name of your terminal.
messagetext	Message from the diagnostic program.

TST Request Option

The REQUEST parameter allows you to initiate interaction with an NPU diagnostic program. An NPU diagnostic program is assigned to you until one of these events occurs:

- You drop the diagnostic program (by sending a TST command with the DROP option).
- The connection between you and CS is terminated for any reason.
- Supervision of the NPU is lost by CS.

Recent HISTORY Command

The HISTORY command requests a display of the recent unsolicited status reports from a CS.

The command format is:

```
HISTORY, ALL.
```

or

```
HI, ALL.
```

If the ALL option is not used, only the last (most recent) page (based on your terminal's page size) of the unsolicited status reports is displayed.

Any DOP or NOP can enter the HISTORY command to obtain reports from CS.

A HOP can enter the HISTORY command to obtain reports from CS, NS, or NVF after assigning the NAM K display to the desired application (refer to section 8, K-Display Utilities, for more information on the NAM K display).

Each application maintains a separate record of the most recent significant NPU load, DOP, NOP, or HOP events. The occurrence of these various significant events prompts the sending of an unsolicited status report to all appropriate operators. All reports are temporarily available through the HISTORY command and are permanently recorded in the NAM dayfile.

A CDCNET operator can enter the DISPLAY_ALARM_HISTORY command to obtain reports from NETOU. Refer to the CDCNET Network Operations Manual for more information about this command.

Report Unsolicited Status Command

The REPORT command requests a CS to report all unsolicited status messages from an NPU.

The command format is:

REPORT, element, option

or

RE, element, option

Parameter	Description												
element	One of the following: <table border="1" data-bbox="422 630 1364 798"> <thead> <tr> <th>element</th> <th>Keyword</th> <th>Abbreviation</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>NPU=npuname</td> <td>NP</td> <td></td> <td>npuname is the name of the NPU.</td> </tr> <tr> <td>NPUS</td> <td>NPS</td> <td></td> <td></td> </tr> </tbody> </table>	element	Keyword	Abbreviation	Description	NPU=npuname	NP		npuname is the name of the NPU.	NPUS	NPS		
element	Keyword	Abbreviation	Description										
NPU=npuname	NP		npuname is the name of the NPU.										
NPUS	NPS												
option	ON or OFF. If option is not specified, OFF is assumed. If option is ON, CS reports all unsolicited status messages until you turn reporting off.												

Initially, HOPs do not receive reports for any NPUs supervised by a CS, but controlling NOPs receive reports for the NPUs they control. Other NOPs and DOPs do not receive any reports. However, any operator can use this command to request unsolicited status reports from any NPU supervised by a CS.

Send Message Command

A DOP can send a message to any of the other operators (HOP, NOPs, and other DOPs) connected to the same CS using the SEND command.

The command format is:

```
SEND, element, MSG=message
```

or

```
SE, element, MS=message
```

Parameter	Description									
element	One of the following: <table border="1" data-bbox="402 653 1346 871"> <thead> <tr> <th>element</th> <th>Keyword Abbreviation</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>TERM=termname</td> <td>TE</td> <td>termname is the name of the terminal to which the message is addressed.</td> </tr> <tr> <td>HOST</td> <td>HO</td> <td></td> </tr> </tbody> </table>	element	Keyword Abbreviation	Description	TERM=termname	TE	termname is the name of the terminal to which the message is addressed.	HOST	HO	
element	Keyword Abbreviation	Description								
TERM=termname	TE	termname is the name of the terminal to which the message is addressed.								
HOST	HO									
message	Text message of no more than 50 characters. If more than 50 characters are entered, an error message is issued. The MSG=message parameter is order dependent and must appear as the last parameter of the command.									

If a terminal device is specified as the network element, it must be an interactive terminal device currently connected to your CS. Otherwise, an error message is issued to the sender.

If you specify HOST, the message is sent to the HOP of the host to which you are connected.

If the element referenced by a SEND command is not currently connected to a CS, you receive a diagnostic message. The diagnostic message is followed by the READY.. prompt.

Status Network Element Command

You request a CS to display status information on the various 255x NPU network elements (such as couplers, NPUs, and trunks) with the STATUS command.

There are five forms of the STATUS command. Refer to the Status Message Formats for DOP, earlier in this section, for the format of the status line returned.

Particular Element

The first form requests the status of a particular element in the network.

The command format is:

STATUS, element.

or

ST, element.

Parameter	Description																								
element	One of the following:																								
	<table border="1"> <thead> <tr> <th>element</th> <th>Keyword Abbreviation</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>CPLER=elementname</td> <td>CP</td> <td>elementname is the name of the element whose status is desired.</td> </tr> <tr> <td>LINE=elementname</td> <td>LI</td> <td>elementname is the name of the element whose status is desired.</td> </tr> <tr> <td>LLINK=elementname</td> <td>LL</td> <td>elementname is the name of the element whose status is desired.</td> </tr> <tr> <td>NPU=elementname</td> <td>NP</td> <td>elementname is the name of the element whose status is desired.</td> </tr> <tr> <td>SVC=elementname</td> <td>SV</td> <td>elementname is the name of the element whose status is desired.</td> </tr> <tr> <td>TERM=elementname</td> <td>TE</td> <td>elementname is the name of the element whose status is desired.</td> </tr> <tr> <td>TRUNK=elementname</td> <td>TR</td> <td>elementname is the name of the element whose status is desired.</td> </tr> </tbody> </table>	element	Keyword Abbreviation	Description	CPLER=elementname	CP	elementname is the name of the element whose status is desired.	LINE=elementname	LI	elementname is the name of the element whose status is desired.	LLINK=elementname	LL	elementname is the name of the element whose status is desired.	NPU=elementname	NP	elementname is the name of the element whose status is desired.	SVC=elementname	SV	elementname is the name of the element whose status is desired.	TERM=elementname	TE	elementname is the name of the element whose status is desired.	TRUNK=elementname	TR	elementname is the name of the element whose status is desired.
element	Keyword Abbreviation	Description																							
CPLER=elementname	CP	elementname is the name of the element whose status is desired.																							
LINE=elementname	LI	elementname is the name of the element whose status is desired.																							
LLINK=elementname	LL	elementname is the name of the element whose status is desired.																							
NPU=elementname	NP	elementname is the name of the element whose status is desired.																							
SVC=elementname	SV	elementname is the name of the element whose status is desired.																							
TERM=elementname	TE	elementname is the name of the element whose status is desired.																							
TRUNK=elementname	TR	elementname is the name of the element whose status is desired.																							

All Elements

The second form requests the status of all of the elements of a particular type on one or all NPUs.

The command format is:

STATUS, linkelement npuelement, statuscode.

or

ST, linkelement, npuelement, statuscode.

Parameter	Description														
linkelement	One of the following: <table border="1" data-bbox="418 640 1344 955"> <thead> <tr> <th>linkelement</th> <th>Keyword Abbreviation</th> </tr> </thead> <tbody> <tr> <td>CPLERS</td> <td>CPS</td> </tr> <tr> <td>LINES</td> <td>LIS</td> </tr> <tr> <td>LLINKS</td> <td>LLS</td> </tr> <tr> <td>SVCS</td> <td>SVS</td> </tr> <tr> <td>TERMS</td> <td>TES</td> </tr> <tr> <td>TRUNKS</td> <td>TRS</td> </tr> </tbody> </table>	linkelement	Keyword Abbreviation	CPLERS	CPS	LINES	LIS	LLINKS	LLS	SVCS	SVS	TERMS	TES	TRUNKS	TRS
linkelement	Keyword Abbreviation														
CPLERS	CPS														
LINES	LIS														
LLINKS	LLS														
SVCS	SVS														
TERMS	TES														
TRUNKS	TRS														
npuelement	One of the following: <table border="1" data-bbox="418 1039 1344 1197"> <thead> <tr> <th>npuelement</th> <th>Keyword Abbreviation</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>NPU=npuname</td> <td>NP</td> <td>npuname is the name of the NPU.</td> </tr> <tr> <td>NPUS</td> <td></td> <td></td> </tr> </tbody> </table>	npuelement	Keyword Abbreviation	Description	NPU=npuname	NP	npuname is the name of the NPU.	NPUS							
npuelement	Keyword Abbreviation	Description													
NPU=npuname	NP	npuname is the name of the NPU.													
NPUS															
statuscode	Optional. If specified, only the elements with that status are reported. <table border="1" data-bbox="418 1281 1344 1476"> <thead> <tr> <th>statuscode</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>Active</td> </tr> <tr> <td>DN</td> <td>Down</td> </tr> <tr> <td>DI</td> <td>Disabled</td> </tr> <tr> <td>EN</td> <td>Enabled</td> </tr> </tbody> </table>	statuscode	Description	AC	Active	DN	Down	DI	Disabled	EN	Enabled				
statuscode	Description														
AC	Active														
DN	Down														
DI	Disabled														
EN	Enabled														

Terminal Devices

The third form requests the status of all terminal devices on a line or a logical link.

The command format is:

STATUS, TERMS, mode, statuscode.

or

ST, TES, mode, statuscode.

Parameter	Description										
mode	One of the following:										
	<table border="1"> <thead> <tr> <th><u>mode</u></th> <th><u>Keyword Abbreviation</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>LINE=elementname</td> <td>LI</td> <td>elementname is the name of the specified element.</td> </tr> <tr> <td>LLINK=elementname</td> <td>LL</td> <td>elementname is the name of the specified element.</td> </tr> </tbody> </table>	<u>mode</u>	<u>Keyword Abbreviation</u>	<u>Description</u>	LINE=elementname	LI	elementname is the name of the specified element.	LLINK=elementname	LL	elementname is the name of the specified element.	
<u>mode</u>	<u>Keyword Abbreviation</u>	<u>Description</u>									
LINE=elementname	LI	elementname is the name of the specified element.									
LLINK=elementname	LL	elementname is the name of the specified element.									
statuscode	Optional. If specified only the elements with that status are reported.										
	<table border="1"> <thead> <tr> <th><u>statuscode</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>Active</td> </tr> <tr> <td>DN</td> <td>Down</td> </tr> <tr> <td>DI</td> <td>Disabled</td> </tr> <tr> <td>EN</td> <td>Enabled</td> </tr> </tbody> </table>	<u>statuscode</u>	<u>Description</u>	AC	Active	DN	Down	DI	Disabled	EN	Enabled
<u>statuscode</u>	<u>Description</u>										
AC	Active										
DN	Down										
DI	Disabled										
EN	Enabled										

NPUs

The fourth form requests the status of all NPUs supervised by the CS.

The command format is:

STATUS, NPUS.

or

ST, NPUS.

SVCs

The fifth form requests the status of all archetype SVCs on a line.

The command format is:

```
STATUS, SVCS, LINE=line.
```

or

```
ST, SVS, LI=line.
```

Parameter	Description
line	Name of the specified line.

A CS does not keep track of the status of the various network elements. NPUs supply the information required to satisfy the STATUS commands by sending messages to CS. In the case of a multiple-element status request, the number of messages sent can become significant. You cannot enter another command until all the messages are delivered and a READY.. prompt is displayed. A controlling NOP, other NOPs, or the DOPs can cancel a STATUS command by entering the interruption or termination sequence (refer to the NOS Version 2 Reference Set, Volume 3 for details). A HOP can cancel a STATUS command by using the slant (/) key.

When status is requested for a coupler, line, NPU, or terminal the normal response to the STATUS command is one line of text. When status is requested for logical links or trunks, two text lines may be returned (if the CS to which you are connected is supervising NPUs at each end of the element, a line of text is returned for each end of the logical link or trunk). When status is requested for archetype SVCs on a line, a text line with the line status is returned and a text line for each active archetype SVC is returned.

Summary of Network Operation Commands

Following is a summary of the network operation commands.

Control Commands for HOP

The following commands are available only to the HOP.

NAM K Display

Assign the NAM K display by using one of the following (refer to section 8, K-Display Utilities, for additional information):

Command	Description
K.AP=CS.	Assigns the NAM K display to the communication supervisor (CS).
K.AP=NS.	Assigns the NAM K display to the network supervisor (NS).
K.AP=NVF.	Assigns the NAM K display to the network validation facility (NVF).
K.AP=NETOU.	Assigns the NAM K display to the network operator utility (NETOU).
K.*	Assigns the NAM K display back to NAM mode after it has been assigned to CS, NS, NVF, or NETOU.

NAM Startup

Initiate NAM by entering NAM.

NAM Shutdown

The sequence of commands used to shutdown NAM is:

K.AP=NVF.

K.DISABLE,HOST.

or

K.IDLE,HOST.

NAM Mode Commands

The HOP assigns the NAM K display to NAM mode before entering any of the following commands:

Command	Description
DB ⁷	Activates the inline debug code.
DE ⁷	Deactivates the inline debug code.
DU ⁷	Activates field length dump.
END	Clears the NAM K display.
FL	Changes NAM's maximum field length.
HELP	Displays information about the NAM mode commands on the right K-display screen.
IG	Causes NAM to ignore alert requests.
LB ⁷	Begins logging of the network traffic on the debug log file.
LE ⁷	Ends logging of the network traffic on the debug log file.
LR ⁷	Releases the debug log file.
RS ⁷	Dumps the statistics data to the statistics file, resets the counters to 0 (zero), and continues gathering statistics.
SEND	Sends a message to another HOP.
ST	Causes the NAM status display to appear on the left K-display screen.

7. These commands are supported by the network supervisory applications CS, NS, NVF, and NETOU. They are ignored by IAF, TAF, and many other network applications.

NS Control Commands

The HOP assigns the NAM K display to NS before entering any of the following commands:⁸

Command	Description
FILE	Assigns an alternate load file to NS for loading the specified NPU.
HISTORY	Requests a display of the recent unsolicited status reports.
NOFILE	Instructs NS to use the default NPU load file for loading all NPUs.
STATUS	Requests the current load/dump status of NPUs.

NVF Control Commands

The HOP assigns the NAM K display to NVF before entering any of the following commands:⁸

Command	Description
DISABLE	Disables an application or all host network operations for both the 255x NPU network and CDCNET network.
ENABLE	Enables an application.
IDLE	Terminates an application or all host network operations for both the 255x NPU network and CDCNET network.
HISTORY	Requests a display of the recent unsolicited status reports.
STATUS	Requests the status of network elements.

⁸. To use the abbreviated form of commands, you can type the first two letters of the command. For example, FI for FILE.

Control Commands for HOP and NOP

The HOP assigns the NAM K display to CS before entering any of the following commands. The NOP can enter these commands⁹ by logging into CS.

Command	Description
ALERT	Changes the threshold that CS uses for alerting the NOP and HOP of NPU saturation.
CONTROL	Allows you to gain or give up control of NPUs.
DISABLE	Changes the status of the specified elements to disabled.
DUMP	Causes the memory dump flag to be reset for the specified NPU.
ENABLE	Changes the status of the specified elements to enabled.
GO	Activates the lines and accepts terminal device connections on an NPU.
LOAD	Causes active NPUs to stop, reload, and reactivate.
SEND	Sends a message of no more than 50 characters to users.

Control Commands for HOP, NOP, and DOP

The HOP assigns the NAM K display to CS before entering any of the following commands. The DOP and NOP can enter these commands⁹ by logging into CS.

Command	Description
INFO	Accesses information about all operators and NPUs connected to CS.
HISTORY	Requests a display of the recent unsolicited status reports.
REPORT	Requests CS to report all unsolicited status messages from NPUs.
SEND	Sends a message of no more than 50 characters to the HOP or another NOP or DOP.
STATUS	Requests the status of network elements.
TST	Allows you to control interaction with an NPU diagnostic program in the specified NPU.

Control Commands for COP

Refer to the CDCNET Network Operations Manual for information about control commands intended for the CDCNET operator.

⁹ To use the abbreviated form of commands, you can type the first two letters of the command. For example, IN for INFO.

NAM Initialization (NAMI)	16-1
Network Dump Collector (COLLECT)	16-5
Load File Generator (LFG)	16-8
Input Directives	16-9
Network Load File	16-10
Summary Listing	16-10
NPU Dump Analyzer (NDA)	16-11
Input Directives	16-12
0 — Comments	16-12
1 — File 1 Register	16-12
2 — Macromemory Word	16-12
3 — Contiguous Block Data Structure	16-13
4 — Circular Buffer Data Structure	16-13
7 — Page Registers	16-14
8 — File 1 Registers	16-14
9 — Macromemory Words	16-14
A — LCB and TCB/CCB	16-15
B — Port Tables and MUXLCB	16-16
NDA Output File Content	16-27
NPU Dump Files	16-28
Dump Interpretation	16-28
PIP Dump Analyzer (LISTPPM)	16-28
AIP Trace With IAF	16-29
IAFTM	16-29
IAFTR	16-30
Network Description Files	16-31
IAF Stimulator	16-33
STIMULA	16-34
Session File Construction	16-34
Dynamic Login	16-39
Optional Think Time Specification	16-40
Task Definitions	16-41
Session File Processing	16-42
Multiple Sessions	16-46
Parameter Output Format	16-46
Postprocessing	16-48

NAM supports two types of networks: the 255x NPU network and the CDCNET network. Throughout this section, the term network refers to the 255x NPU network. However, whenever a statement is applicable to both network types, both types are named explicitly.

This section describes the following utilities and files involved in the operation of the networks and interactive subsystems of NOS.

- **NAM Initialization (NAMI).** This utility controls the network software startup process.
- **Network Dump Collector (COLLECT).** This utility produces tape files from the various NPU, CDCNET device interface, and host dump, trace, statistics, and list files that were generated by any abnormal termination of either the 255x NPU network or the CDCNET network.
- **Load File Generator (LFG).** This utility produces a direct access file used by the Network Access Method (NAM) to perform the downline load of network processing units (NPUs).
- **NPU Dump Analyzer (NDA).** This utility produces readable listings from dump files produced when NAM dumps NPUs.
- **PIP Dump Analyzer (LISTPPM).** This utility program converts binary records from the peripheral interface package (PIP) memory dump file into a report listed in byte format.
- **Application Interface Program (AIP) Trace.** This utility produces a trace file of the messages transferred between an application program and NAM. The information contained in this trace can be useful in tracking network problems and in debugging application programs.
- **Network description files.** Network description files are used by the Transaction Facility (TAF) to describe user access characteristics defined by terminal definition directives.
- **IAF stimulator.** The IAF stimulator is a collection of programs that artificially loads the system to analyze the effects of such a load on system performance and reliability.

Other utilities used with the CDCNET network are documented in the CDCNET Network Operations and Analysis manual.

NAM Initialization (NAMI)

The NAMI utility controls the network software startup process. The network software consists of the NAM Subsystem and several rollable jobs (NS, CS, NVF, TVF, COLLECT) and network applications (ATF, IAF, TAF, RBF, MCS, PTF, QTF, ITF). The NAMI utility starts NAM and the programs NS, CS, NVF, TVF, COLLECT, and other applications desired at the site such as RBF, QTF, and ITF by submitting jobs to the input queue. The NAMI utility can also start the NAM CDCNET applications NETOU, NETLS, and NETFS, although they are usually started by a connection request. The skeletons of these jobs are contained on a master startup file along with parameter records containing directives telling the NAMI utility what parameters to substitute in the job skeletons and what jobs to start. Refer to the NOS Version 2 Installation Handbook for descriptions of the master startup file, parameter records, and job skeleton records.

The advantage of having the NAMI utility start up the network software is that parameters that are common to more than one job or that change from one network startup to the next can be changed easily.

The console operator enters the procedure call NAM or NAMNOGO to call procedure file NAM or NAMNOGO respectively (refer to the NOS Version 2 Operations Handbook). These procedure files (saved under the SYSTEMX user name) contain the NAMI command, which calls the NAMI utility. The format of the NAMI command is:

NAMI, P₁, P₂, . . . , P_n.

P_i is a parameter consisting of a keyword or a keyword equated to a value. All parameters are order independent.

P _i	Description
GO	This parameter initiates the network startup process without operator intervention.
MFN=filename ¹	The permanent file name of the master file to be used. Once this parameter is specified, subsequent network startups use this master file until a new master file is specified. If this parameter is omitted, the previously selected master file is used. The master file must reside on the system default family.
OIN=nnn	The previous network invocation number (nnn) for which the dumps/traces will be collected during this run. If this parameter is omitted, the current network invocation number minus 1 is used.
PW=password ¹	The password assigned to the user name where the master file is stored. If this parameter is omitted, the previously selected password is used. This parameter is not required if UN=NETOPS is specified.

¹. This parameter must be specified for the initial network software startup. It can be specified either on a NAMI command or by an operator entry using the NAMNOGO procedure.

P_i	Description														
RN=recordname ²	The name of the parameter record on the master file that is used for this network initiation. Once this parameter is specified, subsequent network startups use this same parameter record until a new parameter record is specified. The following recordnames can be specified.														
	<table border="1"> <thead> <tr> <th>recordname</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>INIT</td> <td>Purge existing dumps/traces, stop but do not dump local NPUs at network startup, and stop local NPUs at abnormal host termination. Use INIT for the initial startup of single host networks.</td> </tr> <tr> <td>MINIT</td> <td>Purge existing dumps/traces, stop local NPUs at network startup, do not take NPU dumps before initial loading of NPUs, and do not stop local NPUs at abnormal host termination. Use MINIT for the initial startup of multihost networks.</td> </tr> <tr> <td>MRECOV</td> <td>Purge existing dumps/traces, do not take NPU dumps before initial loading of NPUs, and do not stop local NPUs at network startup or abnormal host termination. Use MRECOV for recovery startup of multihost networks after a failure.</td> </tr> <tr> <td>MULTI</td> <td>Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, and do not stop local NPUs at network startup or abnormal host termination. Use MULTI for the normal restart of multihost networks.</td> </tr> <tr> <td>RECOVR</td> <td>Purge existing dumps/traces, do not stop NPUs at network startup or termination, and do not dump NPUs before initial load. Use RECOVR for recovery startup of single host networks after a failure.</td> </tr> <tr> <td>RESTRT</td> <td>Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, stop local NPUs at abnormal host termination, and do not stop local NPUs before network startup. Use RESTRT for the normal restart of single host networks.</td> </tr> </tbody> </table>	recordname	Description	INIT	Purge existing dumps/traces, stop but do not dump local NPUs at network startup, and stop local NPUs at abnormal host termination. Use INIT for the initial startup of single host networks.	MINIT	Purge existing dumps/traces, stop local NPUs at network startup, do not take NPU dumps before initial loading of NPUs, and do not stop local NPUs at abnormal host termination. Use MINIT for the initial startup of multihost networks.	MRECOV	Purge existing dumps/traces, do not take NPU dumps before initial loading of NPUs, and do not stop local NPUs at network startup or abnormal host termination. Use MRECOV for recovery startup of multihost networks after a failure.	MULTI	Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, and do not stop local NPUs at network startup or abnormal host termination. Use MULTI for the normal restart of multihost networks.	RECOVR	Purge existing dumps/traces, do not stop NPUs at network startup or termination, and do not dump NPUs before initial load. Use RECOVR for recovery startup of single host networks after a failure.	RESTRT	Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, stop local NPUs at abnormal host termination, and do not stop local NPUs before network startup. Use RESTRT for the normal restart of single host networks.
recordname	Description														
INIT	Purge existing dumps/traces, stop but do not dump local NPUs at network startup, and stop local NPUs at abnormal host termination. Use INIT for the initial startup of single host networks.														
MINIT	Purge existing dumps/traces, stop local NPUs at network startup, do not take NPU dumps before initial loading of NPUs, and do not stop local NPUs at abnormal host termination. Use MINIT for the initial startup of multihost networks.														
MRECOV	Purge existing dumps/traces, do not take NPU dumps before initial loading of NPUs, and do not stop local NPUs at network startup or abnormal host termination. Use MRECOV for recovery startup of multihost networks after a failure.														
MULTI	Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, and do not stop local NPUs at network startup or abnormal host termination. Use MULTI for the normal restart of multihost networks.														
RECOVR	Purge existing dumps/traces, do not stop NPUs at network startup or termination, and do not dump NPUs before initial load. Use RECOVR for recovery startup of single host networks after a failure.														
RESTRT	Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, stop local NPUs at abnormal host termination, and do not stop local NPUs before network startup. Use RESTRT for the normal restart of single host networks.														

². This parameter must be specified for the initial network software startup. It can be specified either on a NAMI command or by an operator entry using the NAMNOGO procedure.

P_i	Description
RS=appldes	The 2-character application designator of the application to start up. (appldes can be more than two characters, but only the first two characters are used.) This parameter allows the specified application to be started when the network is already operational. It is not valid at network initiation. The application's startup job skeleton must reside on the parameter record/master file specified for this network initiation.

The following applications can be specified.

Application		
appldes	Name	Description
AT	ATF	Automated Tape Facility
CS	CS	Communication Supervisor
FS	NETFS	CDCNET Network File Server
LS	NETLS	CDCNET Network Log Server
NS	NS	Network Supervisor
OU	NETOU	CDCNET Network Operator Utility
PR	PSU	Printer Support Utility
PS	PTFS	Permanent File Transfer Facility Server
QS	QTFS	Queued File Transfer Facility Server
QT	QTF	Queued File Transfer Facility
RB	RBF	Remote Batch Facility
PI	FTPI	File Transfer Protocol Interpreter
TS	FTPS	File Transfer Protocol Server
TV	TVF	Terminal Verification Facility

UN=username ³	The user name where the master file is stored. If this parameter is omitted, the previously selected user name is used.
--------------------------	---

The first time the network startup process is executed, a permanent file is created using a network invocation number of 1 and parameters from the NAMI command. The initial call to NAMI must specify the filename, password, and username of the master file and parameter record to be used or the console operator must enter them using NAMNOGO.

On subsequent network startup operations, a permanent file exists and contains the names of the last used master file, parameter record, and network invocation number. Parameters from this permanent file will be used to start up the network unless overridden by parameters from the NAMI command or parameters entered by the console operator using NAMNOGO. The NAMI command parameters take precedence over the permanent file parameters, and operator entries take precedence over parameters from either the NAMI command or the permanent file.

3. This parameter must be specified for the initial network software startup. It can be specified either on a NAMI command or by an operator entry using the NAMNOGO procedure.

Network Dump Collector (COLLECT)

The COLLECT utility is a program that produces three local files: one of the various host dump, trace, statistics, and list files that result from any abnormal termination of the network, one of any NPU dump files, and one of any MDI/MTI/NDI/TDI dump files. The collected files may be copied to the three local files and then purged, or copied and not purged. The local files will then be copied to tapes. There is also an option to purge the collected files without copying them to the local files. The format of the COLLECT command is:

```
COLLECT, NIN=nnn, NOPURGE, NOSAVE.
```

Parameter Description

NIN=nnn	The 1- to 3-character network invocation number that indicates the upper limit of the invocation numbers to be collected. All files with an NIN value from 1 to nnn are copied to the local file. If NIN=nnn is omitted, the default value is 1.
----------------	--

NOTE

NPUs that stop when the host network software terminates are dumped during the subsequent network invocation. Therefore, the network invocation number of these dumps is 1 greater than the dumps of the host programs taken at the time of the termination.

NOPURGE	Specifies that no purging activity will take place after the collected files are copied to the local files. If NOPURGE is omitted, the collected files are purged after they are copied to the local files. NOPURGE should not be specified if NOSAVE is specified.
NOSAVE	Specifies that the collection process is not performed and the files indicated by NIN=nnn are purged. NOSAVE should not be specified if NOPURGE is used.

The COLLECT utility is automatically initiated by the network startup job (NAMI) upon each invocation of the network. (COLLECT may not be automatically initiated if the master file has been modified.) The utility receives an NIN value of the last network invocation number. The current network will be initiated with an incremented value of NIN. For example, on the fifth initiation of the network the COLLECT utility has a command call of

```
COLLECT, NIN=004.
```

and the network is initiated with an NIN of 005.

The COLLECT utility processes all dump, trace, statistics, and list files whose file names are in the form:

pptsnnn

Variable Description

pp	Description
pp	2-character product prefix.
pp	Description
AT	Automated Tape Facility
CS	Communication Supervisor
DI,...,DR,DS,...,D9	MDI/MTI/NDI/TDI Dumps
FS	Network File Server
IA	Interactive Facility
IN	INITMDI Program
IT	Interactive Transfer Facility
LS	Network Log Server
NI	Network Interface Package
NP	Network Processing Unit
NS	Network Supervisor
NV	Network Validation Facility
OS	Network Operator Server
PR	Printer Support Utility
PS	Permanent File Transfer Facility Server
QS	Queued File Transfer Facility Server
QT	Queued File Transfer Facility
RB	Remote Batch Facility
TI	File Transfer Protocol Interpreter
TS	File Transfer Protocol Server
TV	Terminal Verification Facility

Variable Description

t	Description
t	Type of file.
D	Dump file.
L	List file.
S	Statistics file.
T	Trace file.
s	Subtype 0, 1, 2, or 3 for the file.
nnn	Network invocation number in the range 000 through 999.

In the case of NPU dumps (pp=NP), ts is a unique hexadecimal number identifying the dump. In the case of DI dumps (pp = DI to DZ, or D1 to D9), ts is a unique character string identifying the dump. This character string is in the range from AA to 99. In the case of dayfile dumps, ts is DA. The files produced by the COLLECT utility each have one file with many records. Each dump, trace, statistics, or list file is copied to a local file as one or more records, where the first record copied is preceded by a 2-word record containing the original file name. This format enables an ITEMIZE or CATALOG of a COLLECT utility local file to show which collected files were copied to the local file.

The COLLECT utility assigns each file a file name of the form:

Form	Description
DUMPnnn	For host dump, trace, statistics, or list files. nnn is the network invocation number from the COLLECT command.
DUNPnnn	For NPU dump files. nnn is the network invocation number from the COLLECT command.
DMDInnn	For DI dump files. nnn is the network invocation number from the COLLECT command.

The local files are copied to magnetic tape (or other operator-assigned equipment) if the RESTRT or MULTI startup procedure is selected. Once the files are successfully copied, COLLECT will be executed again to purge the collected files.

The COLLECT utility purges all dump and trace files without collecting them on a local file if the INIT, MINIT, MRECOV, or RECOVR startup procedure is selected.

Load File Generator (LFG)

LFG is a utility program that reformats Communications Control Program (CCP) files for subsequent use by the network supervisor of NAM to load NPUs. LFG accepts one or more input files in single or multifile format and produces a single network load file (NLF) for use by the network supervisor.

LFG is invoked with this command:

```
LFG, P1, P2, . . . , Pn.
```

p_i is a parameter consisting of a keyword or a keyword equated to a value. All parameters are optional and order independent.

P_i	Description
BC=nn	This parameter specifies the maximum number of load service messages in a network load file record. It is used to optimize NPU loads. $1 \leq nn \leq 64$.
BC omitted	Same as BC=64.
I=infile	The 1- to 7-character name of the file on which the input directives are written.
I omitted	Same as I=INPUT.
L=outfile	The 1- to 7-character name of the file on which the summary listing is to be written.
L=0	Suppress the summary listing. If an error is detected an error message is written to file OUTPUT.
L omitted	Same as L=OUTPUT.
NLF=loadfile	The 1- to 7-character name of the network load file.
NLF omitted	Same as NLF=NLF.
Z	<p>An input directive is contained on the command after the terminator. The I parameter is ignored. This eliminates the need to use a separate input directive file when there are only a few input directives.</p> <p>When the input directive appears on the LFG command, the first character following the command terminator is the separator character. Any character that is not used in the input directive can be used as the separator character. The input directive must be preceded by a separator and terminated by a blank, period, or the end of the input line (column 72).</p> <p>For example, (slant used for separator):</p> <pre>LFG, Z . / *FILE, PFILE1, PFILE2 .</pre>
Z omitted	Input directives are not contained on the command. The system uses the I parameter.

Input Directives

Input directives have this format:

```
*FILE, filename1, filename2, ..., filenamen.
```

An input directive begins with an asterisk followed by the directive identifier **FILE** and a string of file names. The directive identifier and file names are separated by commas. If a specified file is empty, LFG ignores it. LFG scans a directive until it detects a terminator or the end of the directive line (column 72). A directive is terminated by a blank, period, or the end of the line. If an asterisk is not specified, LFG assumes ***FILE** and proceeds to read file names. LFG continues to read input directives until the end of the record.

Each file specified by filename contains a load partition or a program initiation control block (PICB) for LFG to reformat. These files are created by the CCP installation process (refer to the NOS Version 2 Installation Handbook). A load partition is any loadable software module that can be loaded into an NPU during one operation. For example, all of the software for CCP can be loaded into an NPU as one load partition. A PICB consists of a sequence of commands that controls the load or dump operations. For each step in a PICB that specifies a load operation there is a corresponding load partition.

The names of the input files can be specified on the LFG command (Z parameter) or in a separate file (I=infile parameter). In either case, these files must be local to the job that runs LFG. LFG does not rewind the input files.

If the input directives are written on a separate record in a batch job (I=INPUT), the job format is:

```
ujn.
USER, username, password, familyname.
LFG, I=INPUT.
--EOR--
*FILE, AFILE1, ..., AFILEn. }
*FILE, BFILE1, ..., BFILEn. }   file INPUT
*FILE, CFILE1, ..., CFILEn. }
--EOI--
```

Network Load File

The network load file produced by LFG is a random access file with the local file name specified by the NLF parameter of the LFG command. The default network load file name is NLF.

Summary Listing

The summary listing produced by LFG is a summary of the network load file contents. The summary is written on the file specified by the L parameter of the LFG command. The default summary file name is OUTPUT. Error messages produced by LFG are also written on this file. The summary listing format is:

No.	Partition Name	Input File	Records
1	PN1	AFILE1	3
2	PN2	AFILE2	7
3	PN3	AFILE3	2

Field	Description
NO.	Line number of the summary listing.
PARTITION NAME	Name of the load partition or PICB for LFG to reformat.
INPUT FILE	Name of the input file on which the load partition or PICB resides.
RECORDS	Number of records in the load partition or PICB.

NPU Dump Analyzer (NDA)

NDA is a utility program that produces a readable printout from the NPU dump file. The NPU dump file may contain one or more dump records, which are created by the network supervisor during the initiation or reloading of a failed NPU. NDA reports depict the state of an NPU at the time a dump is generated. NDA uses the following local scratch files during its processing: NEUFILE, ZZZZINP, and FDMP.

NDA is accessed by a command that allows selection of:

- Dumps from/to specific NPU macromemory addresses.
- Listing options.
- The input directive file.
- The output file.

The format of the command for the NPU Dump Analyzer is:

```
NDA, P1, P2, . . . , Pn.
```

Each parameter is a keyword or a keyword equated to a value. All parameters are optional and order independent. If a parameter is not specified, a default value is assumed. If the same keyword is specified more than once, the last (rightmost) occurrence of the keyword overrides all previous ones. Invalid or undefined parameters cause the run to terminate.

P_i	Description
BA=addr	The 1- to 6-digit hexadecimal address within the NPU macromemory at which the dump report begins. Default is BA=address of the actual beginning of the dump. This parameter is valid only if parameter LO=M, LO=E, or input directive 9 is specified.
EA=addr	The 1- to 6-digit hexadecimal address within the NPU macromemory at which the dump report ends. Default is EA=address of the actual end of the dump. This parameter is valid only if parameter LO=M, LO=E, or input directive 9 is specified.
I=infile	The 1- to 7-character name of the input directive file. If there is no input directive file, I=0 (zero) must be specified. Default is I=INPUT.
L=outfile	The 1- to 7-character name of the output file. Default is L=OUTPUT.

P_i	Description								
LO=opt	The 1- to 3-character listing option. Each character specifies a portion of the NPU dump to be listed; any unspecified portions are not listed. Default is LO=MR, if no input directive file is specified.								
	<table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>E</td> <td>All NPU words in the NPU dump record (all of macromemory, file 1 registers, and page registers). Redundant lines are not suppressed.</td> </tr> <tr> <td>M</td> <td>Macromemory only. Redundant lines are suppressed.</td> </tr> <tr> <td>R</td> <td>File 1 registers and page registers only.</td> </tr> </tbody> </table>	opt	Description	E	All NPU words in the NPU dump record (all of macromemory, file 1 registers, and page registers). Redundant lines are not suppressed.	M	Macromemory only. Redundant lines are suppressed.	R	File 1 registers and page registers only.
opt	Description								
E	All NPU words in the NPU dump record (all of macromemory, file 1 registers, and page registers). Redundant lines are not suppressed.								
M	Macromemory only. Redundant lines are suppressed.								
R	File 1 registers and page registers only.								
NDF=npudump	Local file name of the NPU dump file. The NPU dump file may contain multiple NPU dump records. Each record causes a separate NPU dump output to be generated. Default is NDF=NDF.								

Input Directives

Input directives must be contained in a single logical record on the input directive file. Each directive is in 80-column format. Column 1 is reserved for a 1-character keyword that indicates a change in the processing rules and defines the new rule to be used for the following directives. Each directive containing a keyword in column 1 causes NDA to begin a new listing page. Any comments contained in that directive are printed as the first line of the new page. Columns 71 through 80 are ignored on all directives. Comments are allowed in columns 11 through 70 on directives that contain a 1-character keyword.

Directive parameters are in fixed formats. The number of parameters and their respective columns within the directive varies with the 1-character keyword specified. Each parameter is a 5-digit hexadecimal value (leading zeros are required). Parameters are separated by either a comma or a blank. The remaining space between the last parameter and column 70 of a directive may contain optional comments; such comments appear on the line of the output file preceding the information produced by the directive.

0 — Comments

The 0 directive specifies that comments can be entered in columns 11 through 70 of this directive and succeeding directives until a different keyword is specified.

1 — File 1 Register

The 1 directive causes the contents of a specific file 1 register to be listed. Each file 1 register is specified by its 5-digit hexadecimal address in columns 2 through 6 of the next directive and succeeding directives until a different keyword is specified.

2 — Macromemory Word

The 2 directive causes the contents of a specific macromemory word to be listed. Each macromemory word is specified by its 5-digit hexadecimal address in columns 2 through 6 of the next directive and succeeding directives until a different keyword is specified.

3 — Contiguous Block Data Structure

The 3 directive causes identical data structures that are contiguously addressed in macromemory to be listed. Each directive contains five parameters in 5-digit hexadecimal format. Each parameter has a field name that is used in diagnostic messages.

Columns	Field Name	Description
2-6	PNTR	Macromemory location that contains the pointer to the first identical data structure.
8-12	SIZE	Size of the data structure in NPU words.
14-18	INDX	Index value of the first data structure to be listed. The first identical data structure has an index value of 1.
20-24	FIRST	Number assigned to the first data structure to be listed. The next data structure listing is assigned the next highest number.
26-30	LAST	Number assigned to the last data structure to be listed. The data structure listing terminates when it reaches this number.

4 — Circular Buffer Data Structure

The 4 directive causes circular buffers to be listed in chronological order. Two types of buffers can be listed; fixed-length data structures and variable-length data structures. Each directive contains seven parameters in 5-digit hexadecimal format. Each parameter has a field name that is used in diagnostic messages.

Columns	Field Name	Description
2-6	FWA	File 1 register location that contains the first word address of the circular buffer.
8-12	OLDEST	File 1 register location that contains the address of the oldest item in the circular buffer.
14-18	LWA	File 1 register location that contains the last word address of the circular buffer.
20-24	FLAG	File 1 register location used as a flag to indicate whether the buffer contains meaningful data. The buffer is not listed if this location contains zeros.
26-30	PTRN	Bit pattern that delimits a variable-length data structure.
32-36	MASK	Mask for isolating the appropriate bits to identify the delimiter.
38-42	SIZE	Data structure size. Nonzero indicates a fixed-length data structure. Zero indicates a variable-length data structure.

For variable-length data structures, the next occurrence of a data structure is determined by first masking each word to isolate the bit pattern and then comparing this pattern with the specified bit pattern until a match is found. Listing begins at the first match and ends at the second match or at the oldest item found. For fixed-length data structures, the mask and bit pattern are not used.

The circular buffer format starts by listing the oldest item in the buffer (specified in columns 8 through 12), then proceeds to the last word address (specified in columns 14 through 18), then goes to the first word address (specified in columns 2 through 6), and continues to the oldest item in the buffer.

7 — Page Registers

The 7 directive causes the contents of the page registers to be listed. The 7 directive does not have parameters.

8 — File 1 Registers

The 8 directive causes file 1 registers to be listed in lines of 16 consecutively addressed registers in hexadecimal format. The 8 directive does not have parameters. However, if the parameter LO=E is specified on the NDA command, the selected option will be in effect.

9 — Macromemory Words

The 9 directive causes NPU macromemory words to be listed in lines of 16 consecutively addressed words in hexadecimal format. Repetitive lines are suppressed and addressing continues with the next nonrepetitive line. This line is flagged with double asterisks to indicate the suppression. The 9 directive does not have parameters. However, if parameters BA=addr, EA=addr, or LO=E are specified on the NDA command, the selected options will be in effect.

A — LCB and TCB/CCB

The A directive causes line control blocks (LCB) to be listed with their associated terminal control blocks (TCB) or cluster control blocks (CCB). Each directive contains seven parameters in 5-digit hexadecimal format. Each parameter has a field name that is used in diagnostic messages.

Columns	Field Name	Description
2-6	LCBP	Macromemory location that contains the pointer to the first LCB.
8-12	LCBL	Size of the LCB in NPU words.
14-18	FTCB	Index value of the LCB word that contains the pointer to the first TCB/CCB. The first word in an LCB has an index value of 0 (zero).
20-24	TCBL	Size of the TCB/CCB in NPU words.
26-30	NTCB	Index value of the TCB/CCB word that contains the pointer to the next TCB/CCB. The first word in a TCB/CCB has an index value of 0 (zero).
32-36	NLCB	File 1 register location that contains the number of LCBs to be listed. If set to 0 (zero), only one LCB is listed.
38-42	LCBX	Index value of the first LCB to be listed. The first LCB has an index value of 0 (zero).

The LCBs and their associated TCB/CCBs are listed vertically in columns such that any line of the report represents an LCB and its associated TCB/CCBs. In this format, an LCB and 15 TCB/CCBs can be listed in one line. If more than 15 TCB/CCBs are associated with the LCB, the next page lists the next 15 TCB/CCBs with the LCB column blank to indicate the association of these TCB/CCBs with the previous LCB.

B — Port Tables and MUXLCB

The B directive causes port tables to be listed with their associated multiplex line control blocks (MUXLCBs). Only one MUXLCB can be associated with each port table. Each directive contains eight parameters in 5-digit hexadecimal format. Each parameter has a field name that is used in diagnostic messages.

Columns	Field Name	Description
2-6	PTTP	File 1 register location that contains the pointer to the first port table.
8-12	PTTL	Size of the port table in NPU words.
14-18	MUXP	Index value of the port table word that contains the pointer to the associated MUXLCB. The first word in a port table has an index value of 0 (zero).
20-24	MUXL	Size of the MUXLCB in NPU words.
26-30	PTRN	Bit pattern that identifies a valid MUXLCB.
32-36	MASK	Mask for isolating the appropriate bits to identify a valid MUXLCB.
38-42	MUXID	Index value of the port table word that contains the MUXLCB identifier. The first word in a port table has an index value of 0 (zero).
44-48	NPTT	File 1 register location that contains the number of port tables to be listed.

The validity of the pointer to a MUXLCB is determined by first masking the appropriate port table word to isolate the validating bit pattern and then determining if the specific bit pattern is present. If the MUXLCB pointer is invalid, MUXLCB is not listed.

The port tables and their associated MUXLCBs are listed vertically in columns such that any line of the report represents up to eight port tables, each with its associated MUXLCB.

Figure 16-1 is a sample of an NDA output file and its corresponding input directive file.

```

0 *** A COPY OF INPUT DIRECTIVES ***
  THIS IS A HEADING LINE FOR THE FIRST COMMENT PAGE.
  THIS IS A SAMPLE OF AN INPUT DIRECTIVE FILE SPECIFIED BY
  THE 1 PARAMETER ON THE NDA COMMAND.
  THE ZERO DIRECTIVE ALLOWS YOU TO PUT COMMENTS OR INTER-
  PRETATION INFORMATION ON A LISTING PAGE JUST BEFORE
  A PAGE OF THE FORMATTED DATA. FOR EXAMPLE, THE
  THE NAME ASSOCIATED WITH WORKLIST CONTROL BLOCK 4 CAN BE
  PRINTED.
1 FILE 1 REGISTER INFORMATION
  000F1 FIRST PORT TABLE
  000F3 MAXIMUM NUMBER OF LCBS
  00040 NEXT FRAME POINTER
2 MACROMEMORY INTERPRETATION
  00130 BYWLCB WORKLIST CONTROL BLOCK
  00151 BYCUMWLP CURRENT ACTIVE MONITOR WORKLIST CONTROL BLK
  00152 BYICB INTERNAL PROCESSING TCB
3 CONIGUOUS DATA STRUCTURE
  00161 00020 00001 00001 00003 NLI.A, CONSOLE, COUPLER, TCB
  00165 00002 00001 00001 00001 INITIALIZATION COMPLETE TABLE
  00150 00008 00008 00019 WORKLIST CONTROL BLOCK NUMBER
4 CIRCULAR INPUT BUFFER DUMP FROM CHECK TO CHECK PATTERN
  0003F 00040 00042 00030 00E00 0FF00 00000
  A PRINT ALL LCBS WITH THEIR TCBs -- SKIP LCBO
  A 0015F 0001F 00011 00030 00000 000F3 00001 LCB, TCB, AND SO FORTH
  A A NONZERO COLUMN 1 FORCES A NEW PAGE FOR THE CCB CHAIN
  B 0015F 0001F 00013 00010 00008 00000 00017 LCB, CCB, AND SO FORTH
  B PRINT PORT TABLE AND WORKCBS, TEST FOR BIT 12, PORT WORD 1
7 000F1 00008 00001 00010 01000 01000 00001 000F3
8 DUMP PAGE REGISTERS
9 DUMP FILE 1 REGISTER PER COMMAND LO= PARAMETER VALUE
  DUMP MACROMEMORY PER COMMAND BA=,EA=,LO= VALUES
  
```

Input directive file

Figure 16-1. Sample NDA Output File (Continued)

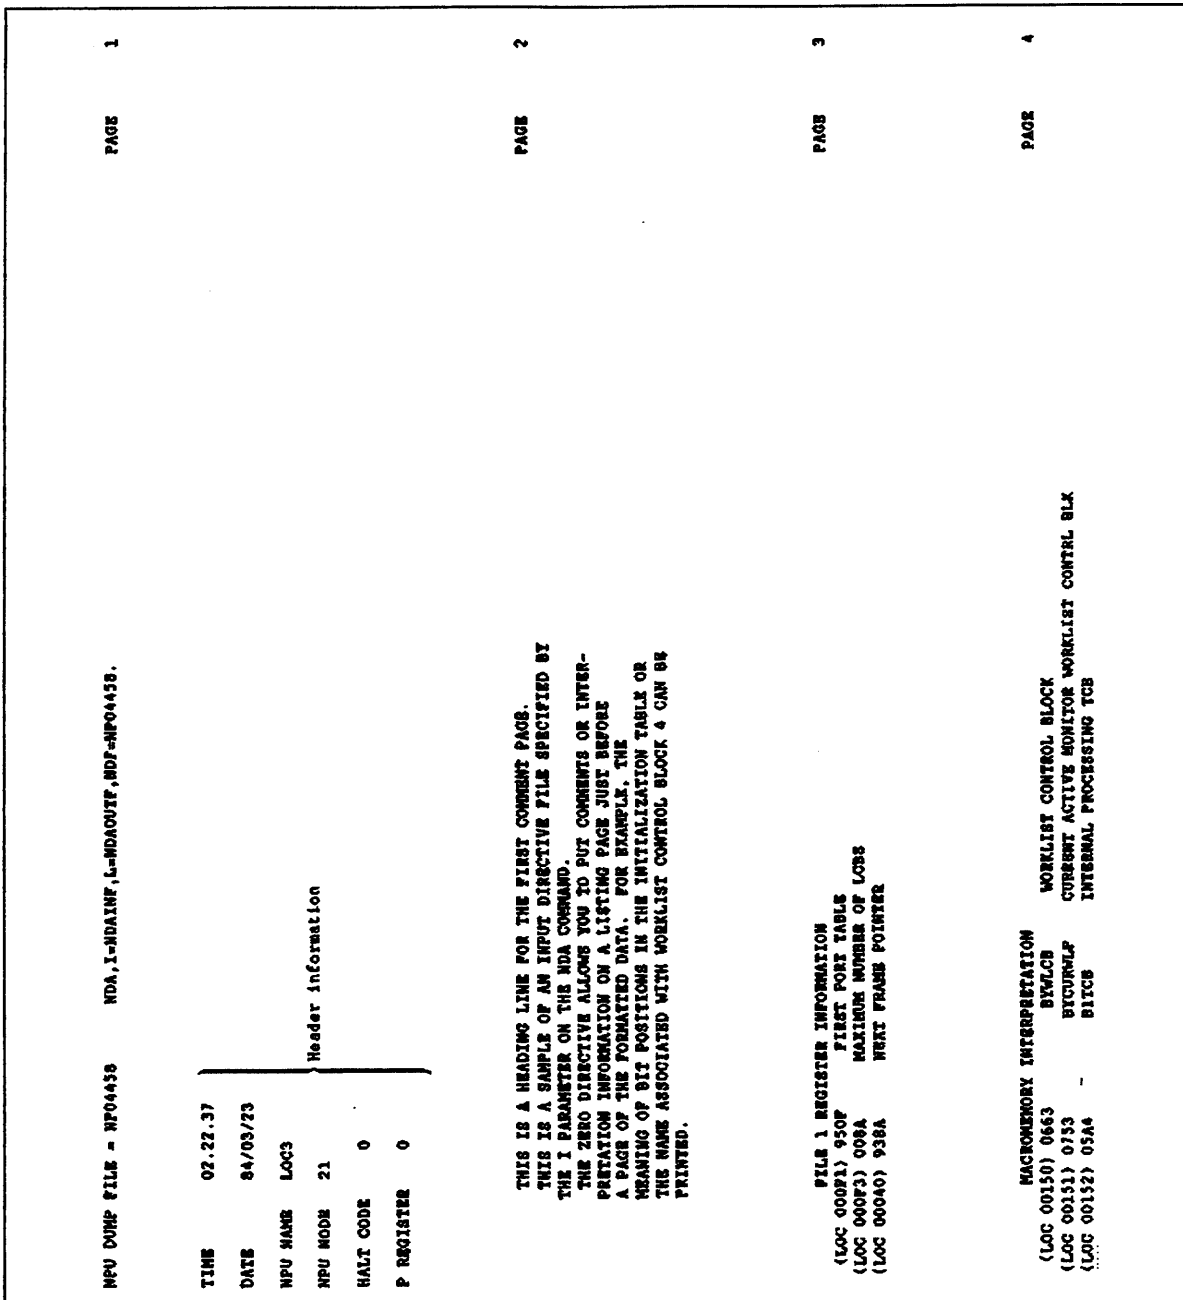


Figure 16-1. Sample NDA Output File (Continued)

CONTIGUOUS DATA STRUCTURE
MLIA, CONSOLE, COUPLER, TCS

OFFSET	1	2	3
0	0000	0000	995F
1	A95F	AAB8	0000
2	0000	0000	0000
3	0070	0070	0003
4	0000	0000	0007
5	0000	0000	0000
6	0000	0000	0000
7	0000	0000	000A
8	0000	0000	8012
9	0000	A000	0000
A	0000	0000	0000
B	0000	0000	0000
C	0000	0000	0000
D	0000	0000	0000
E	0000	0000	0000
F	0000	0000	0000
10	0000	0000	0000
11	0000	0000	0000
12	0000	0000	0000
13	0000	0000	002D
14	0000	0000	2E37
15	0000	0000	0000
16	0000	0000	0000
17	0000	0000	0000
18	0000	0000	0000
19	0000	0000	598A
1A	0000	0000	0800
1B	0000	0000	0000
1C	0000	0000	0000
1D	0000	0000	4951
1E	0000	0000	0000
1F	0000	0000	0000

INITIALIZATION COMPLETE TABLE

OFFSET	1
0	00FF
1	FFFF

WORKLIST CONTROL BLOCK NUMBER

OFFSET	8	9	A	B	C	D	E	F	10	11	12	13	14	15	16	17
0	0000	0000	0006	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
1	AC90	0000	EC89	0000	0000	0000	0000	0000	0000	CC30	0000	B770	AC80	0000	ED60	0000
2	AC90	0000	DFE6	0000	AD10	0000	0000	0000	0000	CC30	0000	B770	AC80	F3BC	ED60	0000
3	8003	0003	8003	4102	4001	4001	4001	8003	4102	8001	8005	8003	8001	0003	8003	4304
4	4B58	0000	525A	0000	0000	0000	0000	584D	424F	4250	4D4C	534D	5853	5248	4C49	4F4C
5	0000	0000	0000	0000	0000	0000	0000	06AB	0683	0683	068B	06C3	06CB	06D3	06DB	06E3
6	0000	0000	0000	0000	0000	0000	0000	0423	040C	082C	2827	0404	0411	0435	040F	0439
7	0000	0000	0000	0000	0000	0000	0000	3F9D	66F2	28D1	3F64	213D	8944	3885	80DA	366C

Figure 16-1. Sample NDA Output File (Continued)

OFFSET	CONTIGUOUS DATA STRUCTURE WORKLIST CONTROL BLOCK NUMBER	PAGE
0	18 0000	6
1	19 0000	
2	0000 0000	
3	0000 0000	
4	0000 0000	
5	0000 0000	
6	0000 0000	
7	0000 0000	

Figure 16-1. Sample NDA Output File (Continued)

LOC	938B 0803 CC1C 0803	CIRCULAR INPUT BUFFER DUMP FROM CHEK TO CHEX PATTERN	PAGE	7
6	0000	0000	0000	0002
7	0000	0000	0000	00A0
8	0000	0000	0000	3000
9	0009	0000	0000	0000
A	0004	0000	0000	0000
B	0003	0000	0000	0000
C	0001	0000	0000	0000
D	0000	0000	0000	0000
E	0000	0000	0000	0000
F	0000	0000	0000	0000
10	91A2	9177	0000	0000
11	0000	0000	0000	0000
12	0000	0000	0000	0000
13	0303	0000	0000	0000
14	0000	0000	0000	0000
15	0000	0000	0000	0000
16	0000	0000	0000	0000
17	0000	0000	0000	0000
18	0000	0000	0000	0000
19	0000	0000	0000	0000
1A	0000	0000	0000	0000
1B	0000	0000	0000	0000
1C	0000	0000	0000	0000
1D	0000	0000	0000	0000
1E	0000	0000	0000	0000
1F	0000	0000	0000	0000
20	0000	0000	0000	0000
21	0000	0000	0000	0000
22	0000	0000	0000	0000
23	0000	0000	0000	0000
24	0000	0000	0000	0000
25	0000	0000	0000	0000
26	0000	0000	0000	0000
27	0000	0000	0000	0000
28	0000	0000	0000	0000
29	0000	0000	0000	0000
2A	0000	0000	0000	0000
2B	0000	0000	0000	0000
2C	0000	0000	0000	0000
2D	0000	0000	0000	0000
2E	0000	0000	0000	0000
2F	0000	0000	0000	0000

Figure 16-1. Sample NDA Output File (Continued)

```

PRINT ALL LCBS WITH THEIR TCBS -- SKIP LCSD
LCB, TCN, AND SO FORTH
LCB009 TC8000 LCB00A TC8000 LCB00B TC8000 LCB00C TC8000 LCB00D TC8000 LCB00E TC8000 LCB00F TC8000 LCB00G TC8000 LCB00H TC8000 LCB00I TC8000 LCB00J TC8000
OFFSET 9800 0000 0800 981F 983E 985D 987C 989A 98B9 98D8 98F7 9915 9934 9953 9972 9991 99B0 99C9 99E8 99F7 9A15 9A34 9A53 9A72 9A91 9AB0 9AC9 9AE8 9AF7 9B15 9B34 9B53 9B72 9B91 9BB0 9BC9 9BE8 9BF7 9C15 9C34 9C53 9C72 9C91 9CB0 9CC9 9CE8 9CF7 9D15 9D34 9D53 9D72 9D91 9DB0 9DC9 9DE8 9DF7 9E15 9E34 9E53 9E72 9E91 9EB0 9EC9 9EE8 9EF7 9F15 9F34 9F53 9F72 9F91 9FB0 9FC9 9FE8 9FF7
LOC.. 0 1 2 3 4 5 6 7 8 9 A B C D E F 10 11 12 13 14 15 16 17 18 19 1A 1B 1C 1D 1E 1F 20 21 22 23 24 25 26 27 28 29

```

Figure 16-1. Sample NDA Output File (Continued)

A NUMBERED COLUMN 1 FORCES A NEW PAGE FOR THE CCB CHAIN
LOC, CCB, AND SO FORTH

LCB017 TCB000

PCB2

OFFSET	0	0000
LOC..	1	0000
	2	0000
	3	0000
	4	0000
	5	0000
	6	0000
	7	0000
	8	0000
	9	0000
	A	0000
	B	0000
	C	0000
	D	0000
	E	0000
	F	0000
	10	0000
	11	0000
	12	0000
	13	0000
	14	0000
	15	0000
	16	0000
	17	0000
	18	0000
	19	0000
	1A	0000
	1B	0000
	1C	0000
	1D	0000
	1E	0000

Figure 16-1. Sample NDA Output File (Continued)

PAGE 24

PRINT PORT TABLE AND WILDCARDS, TEST FOR BIT 17, PORT WORD 1

OFFSET LOC..	PORT000 MUX	PORT001 MUX	PORT002 MUX	PORT003 MUX	PORT004 MUX	PORT005 MUX	PORT006 MUX	PORT007 MUX
0	0000	9517	951F	9527	952F	9537	953F	9547
1	0000	A481	0280	A481	0280	0000	0000	28C0
2	0000	EGF0	0000	ACD0	0000	0000	0000	A860
3	0000	8807	0001	8807	0001	0000	0000	8803
4	0000	41FC	41FC	41FC	41FC	0000	0000	416C
5	0000	8AA7	0000	8AA7	0000	0000	0000	8827
6	0000	0000	0000	0000	0000	0000	0000	3200
7	0000	800F	0000	800F	0000	0000	0000	0000
8	0000	0000	0000	0000	0000	0000	0000	0000
9	0000	0000	0000	0000	0000	0000	0000	0000
A	0000	0000	0000	0000	0000	0000	0000	0000
B	0000	0000	0000	0000	0000	0000	0000	0000
C	0000	0000	0000	0000	0000	0000	0000	0000
D	0000	0000	0000	0000	0000	0000	0000	0000
E	0000	0000	0000	0000	0000	0000	0000	0000
F	0000	0000	0000	0000	0000	0000	0000	0000

OFFSET LOC..	PORT008 MUX	PORT009 MUX	PORT00A MUX	PORT00B MUX	PORT00C MUX	PORT00D MUX	PORT00E MUX	PORT00F MUX
0	9547	9557	955F	9567	956F	9577	957F	9587
1	28C0	0000	0280	A480	B290	2840	0000	B2A0
2	AE70	0000	0000	B290	A189	B2C0	0000	2840
3	416C	0000	0001	8807	DE44	B2C0	2000	B2A0
4	8827	0000	41FC	41FC	000A	0803	0000	0803
5	3200	0000	0000	8AA7	0000	4126	0000	4126
6	0000	0000	0000	0000	0000	0827	0000	0827
7	0000	0000	0000	800F	0000	3200	499F	3200
8	0000	0000	0000	0000	0000	0000	0000	499F
9	0000	0000	0000	0000	0000	0000	0000	0000
A	0000	0000	0000	0000	0000	0000	0000	0000
B	0000	0000	0000	0000	0000	0000	0000	0000
C	0000	0000	0000	0000	0000	0000	0000	0000
D	0000	0000	0000	0000	0000	0000	0000	0000
E	0000	0000	0023	0000	0000	0000	0000	0000
F	0000	0000	0000	0000	0000	0000	0000	0000

Figure 16-1. Sample NDA Output File (Continued)

DUMP PAGE REGISTERS										PAGE	38					
ADDRESS	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
000000	0000	0001	0002	0003	0008	0009	000A	000B	000C	000D	000E	000F	000C	000D	000E	000F
000010	0010	0011	0012	0013	0014	0015	0016	0017	0018	0019	001A	001B	001C	001D	001E	001F

DUMP FILE 1 REGISTER PER COMMAND LO- PARAMETER VALUE										PAGE	39					
ADDRESS	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
000000	0000	0000	0000	0000	0013	000F	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000010	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000020	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000030	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000040	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000050	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000060	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000070	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000080	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000090	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
0000A0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
0000B0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
0000C0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
0000D0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
0000E0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
0000F0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
000100	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000

Figure 16-1. Sample NDA Output File (Continued)

ADDRESS	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	T I T M	PAGE	40
000000	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400	5400			
000010	0011	0007	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000020	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000030	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000040	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000050	0001	0002	0003	0004	0005	0006	0007	0008	0009	000A	000B	000C	000D	000E	000F	0000			
000060	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000070	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000080	0001	0002	0003	0004	0005	0006	0007	0008	0009	000A	000B	000C	000D	000E	000F	0000			
000090	0001	0001	4101	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0000A0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0000B0	024C	024C	0240	0500	0000	010B	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0000C0	7FAF	7FAF	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0000D0	0140	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0000E0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0000F0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000100	0000	1400	17AC	0000	3C45	1400	17D3	0000	0000	1400	81F9	0000	3FAA	1400	821A	0000			
000110	0000	5400	828C	0E1D	2545	0000	6A46	0E1A	0000	5400	6A46	0E1B	0000	5400	828C	0E1C			
000120	8BCD	1400	823A	0000	0000	5400	828C	0E2A	0000	5400	828C	0E2B	0000	5400	828C	0E2C			
000130	0000	5400	828C	0E3D	0000	5400	828C	0E3A	0000	5400	828C	0E3B	0000	5400	828C	0E3C			
000140	1400	0000	1400	8B4D	1400	F02C	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000150	0663	0753	054A	0545	075B	0197	048E	078A	3E49	2222	88C5	2787	850F	9310	995F	99F9			
000160	0A9F	0A9F	0A87	0C1C	9177	08DE	ABC0	0036	1916	0614	3128	803F	0000	0000	019E	9251			
000170	020C	0A92	020C	0A92	0000	0000	0400	AFBF	0003	950F	4920	AA1	8AB0	8610	8E30	9E70			
000180	4510	4D30	5D70	7DF0	6D80	0000	03P0	0412	0434	0436	0478	049A	04BC	04DE	0500	0572			
000190	0544	0566	0000	9310	01FF	000A	0E15	0200	8A00	0000	0000	0000	0000	0000	0000	0011			
0001A0	0001	0005	000D	001D	007D	001D	000F	03FF	03FF	03FF	07FF	07FF	1FFF	3FFF	7FFF	7FFF			
0001B0	8F59	0000	CAB0	80F0	CAC9	GAL1	0000	043D	0000	0000	0000	0000	0000	0000	0000	0000			
0001C0	0000	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E	0B7E			
0001D0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0001E0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0001F0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000200	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF	3FFF			
000210	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000220	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000230	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000240	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000250	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000260	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000270	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000280	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000290	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0002A0	000F	000F	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0002B0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0002C0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0002D0	0400	0400	0800	0200	0100	0100	0200	C000	3000	3000	0000	0300	0300	0300	0300	0300			
0002E0	ABC0	FFFF	002F	FFFF	0607	F808	0605	0605	0409	095A	0400	0401	0402	060A	060E	060B			
0002F0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000300	0100	F0FF	0000	F0FF	0000	F8FF	0020	FFFF	8840	FFFF	0200	FFFF	0000	FFFF	0000	FFFF			
000310	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000320	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000330	000F	F000	F000	F000	F000	F000	F000	F000	F000	F000	F000	F000	F000	F000	F000	F000			
000340	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000350	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000360	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000370	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000380	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
000390	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0003A0	000F	000F	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0003B0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0003C0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0003D0	0400	0400	0800	0200	0100	0100	0200	C000	3000	3000	0000	0300	0300	0300	0300	0300			
0003E0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			
0003F0	0000	0000	0000	0															

NDA Output File Content

The NDA output file (see figure 16-1) begins with a copy of the input directive file, reformatted for printing. The input directive file is followed by the header information, and then the rest of the report is printed in the same order as the input directive file. A description of the header information follows.

```

NPU DUMP FILE = NPddnnn      NDA, command image
TIME             hh.mm.ss
DATE            yy/mm/dd
NPU NAME        npuname
NPU NODE        nodeid
HALT CODE       halt
P REGISTER      preg

```

Header	Description
NPddnnn	The name of the NPU dump record starting with NP, where dd is the hexadecimal dump number (00 through FF) and nnn is the network invocation number assigned at network startup.
command image	The NDA command parameters that were specified for this dump.
hh.mm.ss	Time of the dump in the form hour.minute.second.
yy/mm/dd	Date of the dump in the form year/month/day.
npuname	The 1- to 7-character name assigned to the NPU in the network configuration file.
nodeid	The 2-digit hexadecimal number assigned to the NPU in the network configuration file.
halt	The 1- to 4-digit hexadecimal number for the NPU halt code.
preg	The NPU's P-register contents at the time the NPU dump is generated in hexadecimal format.

NPU Dump Files

Each NPU dump file is written as a single-record direct access permanent file when an NPU dump is taken. Individual NPU dump files are named NPddnnn, where dd is a unique hexadecimal dump number consecutively assigned in the range 00 through FF and nnn is the network invocation number assigned at network startup.

Dump Interpretation

When a halt occurs, halt codes are usually sent to the network operator controlling the NPU. (Halt codes are described in the CCP 3 Diagnostic Handbook.) If a halt occurs after loading but before completion of initialization, or the system becomes trapped in a loop during initialization (before the CCP header prints), dump interpretation may be necessary to determine which halt has occurred, or in which subroutine of the initiation section the program is looping.

When interpreting the NPU dump listing to determine the cause of a halt or loop, first examine the contents of memory location 30_{16} . If the content of that location is nonzero, a halt has occurred; the halt code value is contained in that location.

If memory location 30_{16} equals zero after a halt during initialization, find the address table that begins at fixed memory address 150_{16} . The last nonzero entry in this table is the address of NPINTAB. (This is the table that is displayed at the end of a successful initialization. NPINTAB has a fixed address.) The NPISFL entry in the NPINTAB table contains the flags that mark the initialization subroutines that have completed running when the loop occurred.

PIP Dump Analyzer (LISTPPM)

The peripheral interface package (PIP) Dump Analyzer (LISTPPM) is a utility program that converts all available PIP dump binary records on the PIP memory dump file into a report to be listed in byte format. The PIP PP memory dump file is created as local file ZZZZZPP by NAM when the operator drops the network or NAM detects fatal network errors.

The format of the command for the PIP Dump Analyzer is:

```
LISTPPM, B=pipfile, L=outfile.
```

The parameters are order independent. Invalid or undefined parameters cause termination of the run.

Parameter	Description
B=pipfile	Local file name of the PIP PP memory dump file. Default is B=ZZZZZPP.
L=outfile	Local file name of the output file. Default is L=OUTPUT.

AIP Trace With IAF

The AIP trace utility produces a trace file of the messages transferred between IAF and NAM. This file, produced as a local file named ZZZZZDN at IAF's control point, contains information that can be useful in tracking network problems and in debugging application programs. However, the data in this file is in compressed format and hence the debug log file processor (DLFP) must be used to analyze the trace file. Refer to the NAM Version 1/CCP Version 3 Reference Manual for details about the trace output and the use of DLFP.

Select whether or not to use the AIP trace mode of operation when entering the DSD command, IAFfff, to initialize IAF. This command selects a procedure file that contains the IAFEX command. The T parameter on the IAFEX command determines whether or not AIP trace mode is to be in effect and also the number of messages to be contained on the trace file. The format of the IAFEX command is:

```
IAFEX, T=count.
```

Parameter	Description
T=count	5000 ≤ count ≤ 9999999. Message count specifying the number of messages that are to be logged on the trace file before that file is released to the system for processing. If T=* is specified, the trace file is processed only after IAF is terminated.
T	Same as T=16200.
T=0	AIP trace mode is not selected.
T omitted	Same as T=0.

Two procedure files are provided for use with the AIP trace, IAFTM and IAFTR. You can create other procedure files using different T parameters on the IAFEX command.

IAFTM

IAFTM is a procedure file that contains the IAFEX, T=* command, which causes the trace file ZZZZZDN to be processed only after IAF is terminated. The IAFTM procedure file also includes the commands necessary to process the ZZZZZDN trace file. Also, all messages logged on the trace file are written to the output file and then sent to a printer.

IAFTR

IAFTR is a procedure file that contains the IAFEX,T command, which causes the trace file ZZZZZDN to be processed after every 16,200 messages have been transferred between IAF and NAM. A local file called TRACIAF is required by the IAFEX,T command in order to process the trace file. TRACIAF contains a command record, which is copied to the first record of the ZZZZZDN trace file by IAF. Trace information is then written to file ZZZZZDN, starting at the next record. This makes file ZZZZZDN suitable for submission as a batch job. The format of the TRACIAF file should be:

```
ujn.
USER, username, password, familyname.
.
.
.
DLFP , . . . .
.
.
.
```

NOTE

When system activity is very high, the potential exists for these problems to occur if the AIP trace is used:

- If the system load remains at a very high level for a long period of time, a system hang due to PP saturation could occur.
- If mass storage requirements become great, the ZZZZZDN trace file could reach track limit. If so, IAF would stop and wait for resources to become available.

Both of the preceding problems are unlikely to occur in most situations. Also, the risk of reaching track limit can be minimized by selecting a message count on the IAFEX command that causes the trace file to be processed and released after the specified number of messages have been logged. This lowers the mass storage requirements. There are some problems that can only be solved by using the AIP trace. For these problems, do not hesitate to use the AIP trace since any risk involved is minimal.

Network Description Files

The network description files describe the physical and logical configuration of hardware and software elements that comprise the terminal network. They are important to the proper functioning of the TAF and NAM Subsystems of NOS as well as the QTF/PTF applications if the file transfer capability with NAM is a desired feature. Depending on which subsystems a site has, from one to three such files may be used on a single mainframe. Each file has the following unique function and/or construction.⁴

File Name	Description
LCFFILE and NCFFILE	Local configuration file and network configuration file that must be present if NAM is used. They are created by the network definition language processor (NDLP) from source text prepared according to specifications given in the Network Definition Language Reference Manual.
NCTFi ⁵	Network description file that must be used if the Transaction Facility (TAF) is used. The site analyst prepares this file according to the following specifications.

Under standard operating conditions, the network subsystems expect LCFFILE and NCFFILE to be permanent, direct access files stored under user name NETADMN. These files may be private with read permission given to user name NETOPS, semiprivate, or public. LCFFILE and NCFFILE are binary files created by NDLP; NCTFi is a text file composed of terminal definition directives. The text file is interpreted by code in common deck COMCRTN to produce tables used by TAF during system operation. TAF expects NCTFi to be a direct access public file stored under the SYSTEMX user index (377777₈).

A terminal definition directive can be up to 90 columns but only columns 1 through 72 are interpreted. Any line containing an asterisk in column 1 is a comment and is ignored.

⁴. For information about network description files associated with the CDCNET network, refer to the CDCNET Network Operations and Analysis manual.

⁵. i is the network file identifier as specified in the TAF configuration file (refer to the TAF Reference Manual).

The following format shows the terminal definition directive. No embedded blanks are allowed.

```
/terminalname,p1,p2,...,pn.
```

Parameter	Description
terminalname	The 1- to 7-character alphanumeric name that uniquely identifies the terminal; terminalname must be specified. When NCTFi is created, terminalname must agree with any automatic login user name.
p_i	Description
TT=*ID	Terminal type identified by terminal operator entry. A TAF terminal is defined with the previous parameter, as well as the following, which are unique to TAF.
DB=db	The 2-character data base name to be used by the terminal.
IS=stat	Initial on/off status of the terminal; stat can be either ON or OFF. Default is ON.
RS=n	Data base read security (from 0 to 7). Default is 0.
UA=n	24 bits defining the user argument area (from 0 to 77777777 ₈). Default is 0.
US=n	Data base update security (from 0 to 7). Default is 0.

To use the permanent file and queued file transfer capability with NAM, you must specify certain NDL statements when creating the local configuration file, as follows:

- Specify application statements with the PRU and NETXFR parameters for applications QTF, QTFS, PTF, and PTFS (QTFS and PTFS require the RS parameter).
- Specify the MXCOPY parameters for these applications depending on expected usage. Suggested values are MXCOPY=4 for QTF, and MXCOPY=14 for QTFS, PTF, and PTFS.
- Specify the INCALL and OUTCALL statements as needed.

Refer to the Network Definition Language Reference Manual for NDL statement formats.

IAF Stimulator

The IAF stimulator is a collection of central memory and peripheral processor programs that enters a hypothetical work load (called a session or script file) into the system to analyze the effects of such a load on response time and system reliability.

The IAF stimulator consists of a central memory program (STIMULA) and a peripheral processor program (1TS).

To use the stimulator:

1. Create the hypothetical load (refer to Session File Construction later in this section).
2. Process these jobs using stimulator commands at the system console (refer to Session File Processing later in this section).

Multiple IAF stimulator sessions are described under Multiple Sessions later in this section. Errors that can occur while the IAF stimulator software is in use are described in the NOS Version 2 Operations Handbook.

STIMULA

STIMULA is an internal stimulator that enters a hypothetical work load into the system through IAF. The stimulator software communicates directly with IAF using a channel. Therefore, 2550s are not used. Only interactive terminals can be analyzed with STIMULA; it does not check out other communication devices or paths. The format of the command is:

```
STIMULA, I=filename.
```

Parameter Description

filename	Local file to be used as the session file. If not specified, the initial K display (refer to figure 16-5) requests the session file name.
----------	---

Session File Construction

A session (or script) file is an indirect access permanent file containing the hypothetical job load that is to be analyzed by the IAF stimulator. A session file consists of a group of records; each record is a session and contains a hypothetical job that is composed of the following required entries made in the order shown:

1. Begin each internal stimulation session with the entry ANSWERBAC. (Other first lines may be used; however, ANSWERBAC is used in most cases.)
2. Enter each session entry on a separate line or card.
3. After ANSWERBAC, enter a family name entry, a user name entry, and a password entry. You may place these entries on the same line if they are separated by commas.
4. Specify the terminal commands and data.
5. End each session with the command BYE so that the login of the next session is performed correctly. The HELLO command must not be used unless it is followed by a blank line.

Figures 16-2, 16-3, and 16-4 are examples of session files and figure 16-5 shows the initial K display.


```

ANSWERBAC
FAMILY
USERNAM
PASSWRD
CATLIST      } Session record

DAYFILE
BYE
  *E-O-R*
ANSWERBAC
FAMILY
USERNAM
PASSWRD
CATLIST      } Session record

DAYFILE
BYE
  *E-O-R*
ANSWERBAC
FAMILY
USERNAM
PASSWRD
CATLIST      } Session record

DAYFILE
BYE
  *E-O-R*
$TASK0 50
OLD, FILE0
LIST
LIST          } Task record

  *E-O-R*
$TASK1 30
OLD, FILE1
LIST
LIST          } Task record

  *E-O-R*
$TASK2 20
OLD, FILE2
LIST
LIST          } Task record

  *E-O-F*

```

Figure 16-2. Sample Session File With Tasks for Use by STIMULA

```

ANSWERBAC
FAMILY
USERNAM
PASSWRD
CHARGE (7159,73K DFA)
> 1
FORTRAN
NEW, MATH
10 PROGRAM MATH
20 1 CONTINUE
30 READ 10,A
40 IF(A.EQ.0) STOP
50 B=A*A
60 C=SQRT(B)
70 IF(A.LT.0) D= -C**3
80 IF(A.GT.0) D=C**3
90 PRINT 20,A
100 PRINT 30,B
110 PRINT 40,D
120 GO TO 1
130 20 FORMAT (' A EQUALS ',F15.3)
140 30 FORMAT(' A SQUARED = ',F15.3)
150 40 FORMAT(' A CUBED = ',F15.3)
160 10 FORMAT(F15.3)
170 END
} Source line input
> >
RUN
2.0
0.0
} Program input
25 PRINT 50
155 50 FORMAT(' ENTER VALUE ')
LIST
RUN
1.0
2.0
3.0
4.0
0.0
160 10 FORMAT(F12.1)

```

Figure 16-3. Sample Session for Use by STIMULA (Continued)

```
LIST
RUN
10.0
11.1
13.3
0.0
160 10 FORMAT(F12.3)
LIST
RUN
15.555
40.0
50.0
0.0
LIST
RUN
10.0
20.0
30.0
40.0
0.0
ENQUIRE, T
DAYFILE
BYE
```

Figure 16-3. Sample Session for Use by STIMULA

```

ANSWERBAC
FAMILY, USERNAM, PASSWRD
BASIC
OLD
TESTB
LIS, 294
RUN
01000
REPLACE
BYE
*E-O-R*
:
additional
sessions
:
*E-O-F*

```

Figure 16-4. Sample Session File for Internal Stimulation

```

K.                STM
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                  MID=AA   NOS version

                  STIMULA

PARAMETER          DESCRIPTION

F = STPTEST       SESSION FILE NAME (1-7 CHARACTERS).
UI = 137          USER INDEX (0-377777B).
N = 1             NUMBER OF FILE ON SESSION FILE (0-377777B).

ENTER K.GO TO BEGIN PROCESSING OF SESSION FILE.

```

Figure 16-5. STIMULA K Display (Left Screen)

Dynamic Login

If each of the terminals to be analyzed is logged into the system with a unique user name and password, a unique session record for each terminal is stored in central memory at the IAF stimulator control point. Thus, a large amount of central memory could be unavailable to other components of the system. To alleviate this condition, dynamic login is possible, whereby one session record is generated with a unique user name and password that are dynamically translated into unique user names and passwords for each terminal at run time. The generated user names and passwords are then validated via system validation files. The format choices are:

```
username$$$
```

or

```
username$$$,password$$$
```

Parameter Description

username	User name; 1 to 4 alphabetic characters.
password	Password; 1 to 4 alphabetic characters.
\$	Special symbol 12/8/6 (6-bit display code 76 ₈), 1 (6-bit display code 34 ₈) punch; that is, 12/8/6 punch in one column and 1 punch in next column.

NOTE

The \$ denotes the special symbol previously described, not the character itself. The \$ character is translated into the IAF stimulator terminal numbers at run time.

Example:

P\$\$\$\$,P\$\$\$\$ is translated into P000,P000 for the first terminal, into P001,P001 for the second terminal, and so forth.

Optional Think Time Specification

Think time is the number of seconds that each line is delayed before it is sent through the interactive system. A random think time can be specified with a session file command. The format of the command is:

≥ number

Parameter Description

number Think time (octal): number ≤ 177₈.

To return the terminal to the initial think time (refer to figure 16-6), enter this command:

≥ ≥

K.	STM
hh.mm.ss. yy/mm/dd.	CDC NETWORK OPERATING SYSTEM
	MID=AA NOS version
	STIMULA
PARAMETER	DESCRIPTION
NT = 310B	NUMBER OF TERMINALS TO STIMULATE (1-310B).
LS = 15	LINE SPEED IN CHARACTERS/SECOND (1-1000).
IS = 15	INPUT SPEED IN CHARACTERS/SECOND (1-1000).
TT = 10	THINK TIME DELAY IN SECONDS (0-127).
TI = 8	RANDOM THINK TIME INCREMENT (A POWER OF 2 0-64). THINK TIME WILL VARY BETWEEN TT AND TT+TI.
AC = 4	NUMBER OF TERMINALS ACTIVATED EVERY AD SECS (1-64).
AD = 10	ACTIVATION DELAY IN SECONDS (1-127).
RC = 0	NUMBER OF TIMES TO REPEAT STIMULATION (0-31).
LD = 0	LOGOUT DELAY (0-4095).
RO = ON	OUTPUT RECOVERY (ON OR OFF).
LF = NO	LOPP ON SESSION FILE (YES OR NO).

Figure 16-6. STIMULA Parameters (Left Screen)

Task Definitions

A task is a predefined set of interactive commands designated by a task name and a task usage parameter. Tasks can be called throughout a session, and the task usage parameter is used to determine which task is the next most eligible task (that is, the task to be called next).

An ↑ character (6-bit display code 70₈) in the session calls the next most eligible task in the session file, which is determined by the following equation.

$$\frac{\text{Total initiations of task}_i}{(\text{Total initiations of all tasks})(\text{Task usage parameter for task}_i)}$$

For each task_{*i*} defined in the session file, the IAF stimulator calculates f_i , and the task with the smallest f value is the next most eligible task. Multiple ↑ characters are permitted on one line and are treated as if they occurred on separate lines. For example, ↑↑↑ calls the first three most eligible tasks, where eligibility is recalculated after each initiation of a task. If a task is called and the IAF stimulator cannot determine the next most eligible task in the time allotted, the first task defined in the session file is used.

Tasks for the IAF stimulator are defined in the session file as task records following the session records. Tasks are not permitted in the file before session records. The tasks should be defined in descending order according to the task usage parameter. This facilitates use of the first task defined as the default next most eligible task in situations where the IAF stimulator fails to calculate the next most eligible task on time.

The first line in a task record must be:

```
$task nnn
```

Parameter	Description
\$	Denotes the beginning of a new task (character position 1).
task	The 1- to 6-character task name.
nnn	Task usage percentage; $0 \leq nnn \leq 100$ (character positions 8, 9, 10). The task usage percentages of all tasks in the session file must total 100.

The last line in a task record must be a ↓ character (6-bit display code 71₈), which returns control to the calling session. An example of a session file with tasks defined is shown in figure 16-2.

Session File Processing

The IAF stimulator software is used to process the interactive session file following this procedure:

1. Ensure that an IAF stimulator entry is present in the EQPDECK and that the IAF stimulator has an ON status (refer to the section 3, Deadstart Decks). ON status must be set before IAF is active at a control point. If the status has been changed to OFF, idle IAF and then set ON status and reactivate IAF. The IAF stimulator entry must be on a channel that is not used by any other equipment that currently has a status of ON.
2. Use DSD to activate the IAF stimulator.

Enter STMffff, where ffff is 0 to 4 alphanumeric characters. This entry calls a procedure file STMffff, which must be a permanent indirect access file stored under the system user index (377777₈). The procedure file STM is called when you enter STM. The IAF stimulator is invoked by a STIMULA command within the STM procedure file. The IAF stimulator requests the K display.

3. Enter K,STM. Figure 16-5 appears on the left screen.
4. Enter the session file parameters individually as listed in the following format column or as a group on the same line in the format.

K.F=filename,UI=userindex,N=n,GO.

Format	Default	Description
K.F=filename	None	filename is the indirect permanent file name of the session file.
K.UI=userindex	0	userindex is the user index of file filename. The default family is assumed.
K.N=n	0	n is the number of files to be skipped on file filename before data is read from the session file; range is from 0 to 377777 ₈ .
K.GO	None	Indicates to the IAF stimulator that all the parameters have been entered or that none are to be entered.

Figure 16-6 appears on the left screen.

5. To display the mixed mode parameters, enter:

KK.

Figure 16-7 appears on the right screen. Then, enter:

K.MIXED.

Figures 16-8 (left screen) and 16-9 (right screen) appear.

K DISPLAY COMMANDS	
GO	BEGIN STIMULATION.
MIXED	DISPLAY PARAMETER MATRIX AND ENTER MIXED MODE PARAMETER ASSIGNMENT.

Figure 16-7. STIMULA K-Display Commands (Right Screen)


6. Enter the parameters (refer to figure 16-6) individually as listed in the following format column or as a group on the same line in the format.

K.NT=n, LS=s, IS=s, TT=n, TI=i, AC=m, AD=d, RC=n, LD=d, RO=o, GO.

Enter any mixed mode parameters (refer to figure 16-9) before the GO entry.

Format	Default	Description
K.NT=n	Current EST value	n is the current maximum number of terminals on which the Interactive Subsystem is initialized. It must be less than or equal to the number specified in the EST entry; the range is from 1 to 1000 ₈ .
K.LS=s	15	s is the line speed in characters per second at which you want to run the terminals; the range is from 1 to 1000. Any value above 60 characters per second causes a line speed of approximately 60 characters per second. If many terminals (600 ₈ , for example) are run, the maximum realizable line speed is determined by the overhead of 1TD and 1TS. The worst case can be determined by checking the 1TD cycle time dayfile message, which is issued when IAF is dropped.
K.IS=s	15	s is the typing speed in characters per second at which you want to run the terminals; the range is from 1 to 1000.
K.TT=n	10 seconds	n is the number of seconds (think time) that each line is to be delayed before it is sent to IAF; the range is from 0 to 127 seconds.
K.TI=i	8	i is the number of seconds by which the think time will vary; the range is from 0 to 64 and i must be a power of 2. If i is nonzero, the think time varies between n (the value of TT) and n+i.
K.AC=m	4	m is the number of terminals that are to be logged in every d seconds (the value of AD); the range is from 1 to 64. The AC and AD parameters provide a method to stagger the I/O commands and data to IAF.

Format	Default	Description
K.AD=d	10 seconds	d is the number of seconds of delay time until the login of another terminal is allowed; the range is from 0 to 127 seconds.
K.RC=n	0	n is the number of sessions that will be executed for each terminal; the range is from 0 to 31.
K.LD=d	0 seconds	d is the number of seconds of delay time after a logout is completed until the next login of the terminal is allowed; the range is from 0 to 4095 seconds.
K.RO=o	ON	When RO=ON, the output from IAF is recovered on file STIMOUT. The format of this file is:



Terminal Each byte is an ASCII character if less than 4000_g; otherwise, it is the running-second time.

When RO=OFF, the output is not recovered.

K.LF=o	NO	When LF=NO, looping occurs on the session assigned to the terminal. When LF=YES, the subsequent session is the next session in the session file.
K.MIXED	No mixed mode parameters entered	Specifies that mixed mode parameter assignments are required for specific input sessions. This command displays the stimulation parameter matrix (refer to figure 16-8) whose values can be changed as described in the right screen K display (refer to figure 16-9).
K.MX=s ₁ -s ₂ , ttys,ls,is, tt,ld,rc	Current value	Mixed mode input. s ₁ is the number of the first session to be changed, and s ₂ is the number of the last session to be changed. s ₂ is optional; the default is s ₁ . (Refer to figure 16-9 for other parameter descriptions.)

Format	Default	Description
K.CF=filename	None	Matrix change file. Indirect access file filename is accessed for mixed mode input directives, and must be saved under the same user index as the session file.
K.+	None	Page matrix display forward to the next screen.
K.-	None	Page matrix display backward to the previous screen.
K.GO	None	Indicates that all parameters have been entered. Stimulation is initiated.

Example:

```
K.MX=2-4,3,30,30,,,1.
```

Three terminals will be assigned to sessions 2, 3, and 4. The line speed and input speed for these terminals will be 30 characters per second. The think time and logout delay will not be changed from the current value. The terminals will have a repeat count of 1.

```

K.                               STM
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version

                      STIMULATION PARAMETER MATRIX

                      NUMBER OF TTYS ASSIGNED           200
                      MAXIMUM TTYS PERMITTED            200

CF = CHANGE                MATRIX CHANGE FILE (1-7 CHARACTERS).

SCRPT   TTYS   LS   IS   TT   LD   RC
  1      40    15   15   10    0    0
  2      40    15   15   10    0    0
  3      40    15   15   10    0    0
  4      40    15   15   10    0    0
  5      40    15   15   10    0    0

```

Figure 16-8. Parameter Matrix (Left Screen)

```

K DISPLAY COMMANDS
GO      BEGIN STIMULATION.
+      PAGE DISPLAY FORWARD.
-      PAGE DISPLAY BACKWARD.

MIXED MODE PARAMETERS

MX=SCRIPT, TTYS, LS, IS, TT, LD, RC.
  SCRIPT  SCRIPTS WHICH THE PARAMETERS WILL AFFECT.
          SCRIPT(N) OR SCRIPT(N)-SCRIPT(N+X).
  TTYS    NUMBER OF TTYS ASSIGNED TO SCRIPTS.
  LS      LINE SPEED ASSIGNED TO SCRIPTS.
  IS      INPUT SPEED ASSIGNED TO SCRIPTS.
  TT      THINK TIME ASSIGNED TO SCRIPTS.
  LD      LOGOUT DELAY ASSIGNED TO SCRIPTS.
  RC      REPEAT COUNT ASSIGNED TO SCRIPTS.
  NULL    PARAMETER WILL KEEP CURRENT VALUE.
CF=XXX   CHANGE MATRIX VIA PREDEFINED PARAMETERS.
  XXX = PERMANENT FILE CONTAINING CHANGE DIRECTIVES IN FORMAT
        MX=SCRIPT, TTYS, LS, IS, TT, LD, RC.

```

Figure 16-9. Mixed Mode Parameters (Right Screen)

Multiple Sessions

Each session on the session file is allocated to each terminal line. When each session has been allocated, the IAF stimulator goes to the first session and reallocates the session to the next group of terminals. This procedure is repeated until all the terminals are allocated to a session. If a repeat count is not designated (RC parameter), the line is shut off when the session is finished. When the session completes, the repeat count is checked.

The repeat count corresponds to the number of sessions to be executed on a given terminal. The current session is completed, and the LF parameter determines the next session to be executed. If LF=NO, the session currently assigned to the terminal is repeated. If LF=YES, the next session in the session file is executed. For example, if the repeat count is 4, the session currently assigned to the terminal is repeated four times (LF=NO), or the next four sessions on the session file (with repetition, if necessary) are executed (LF=YES).

Parameter Output Format

The parameter information (shown in figure 16-10) is placed on the output file for printing upon completion of the IAF stimulator run (refer to Session File Processing earlier in this section).

```

STIMULATOR PARAMETERS-
SESSION FILE NAME-           F=  SSSC1
USER INDEX-                   UI=  1
NUMBER OF THE FILE-           N=  1
NUMBER OF TERMINALS-          NT= 100B
LINE SPEED-                   LS=  30
INPUT SPEED-                  IS=  10
THINK TIME-                   TT=  3
RANDOM THINK TIME-            TI=  0
ACTIVATION COUNT-            AC=  1
ACTIVATION DELAY-            AD=  1
REPEAT COUNT-                 RC=  4
LOGOUT DELAY-                 LD=  0
OUTPUT RECOVERY-              RO=  ON
LOOP ON SESSION FILE-         LF=  NO
MATRIX CHANGE FILE-          CF=
      STIMULATION PARAMETER MATRIX

```

SCRPT	TTYS	LS	IS	TT	LD	RC
1	22	30	10	3	0	4
2	21	30	10	3	0	4
3	21	30	10	3	0	4

```

NUMBER OF TTYS ASSIGNED = 64
MAXIMUM TTYS PERMITTED  = 64

```

TASK	CALLS	COMPLETE	DESIRED %	ACTUAL %
TASK0	848	818	50.000	49.970
TASK1	509	492	30.000	29.994
TASK2	340	327	20.000	20.035
TOTALS	1697	1637	100.000	99.999

```

DEFAULT TASK CALLS = 15

```

Figure 16-10. Output Statistics

Postprocessing

The postprocessing of the IAF stimulator output is accomplished by reading the file STIMOUT and sorting the data according to terminal number. Use the DEMUX command to do this. As DEMUX sorts STIMOUT, it assigns an encounter number (starting at 1) to the output from each terminal session. DEMUX includes this number in the final line of each sorted terminal session output. The data for each terminal is separated by a page eject and an end-of-line indicator. The terminal data appears as it would on a terminal page. If the ASCII character on the STIMOUT file is 4000₈ or greater, the lower 10 bits are assumed to be time of a carriage return or the first character of output. If STIMOUT will be processed at a later time, you should save it.

Figure 16-11 shows an example of output from DEMUX.

```

ANSWERBAC
*****TTT.TTT
*****TTT.TTT
WELCOME TO THE NOS SOFTWARE SYSTEM.
COPYRIGHT CONTROL DATA 1978, 19XX.
yy/mm/dd.hh.mm.ss.
.
.
.
.
.
END OF THE LINE-----n

(Page Eject)
ANSWERBAC
.
.

```

Figure 16-11. Postprocessing Example

In figure 16-11, *****TTT.TTT is optional output, and the n in the previous listing is the decimal encounter number.

The format of the DEMUX command is:

DEMUX, P₁, P₂, . . . , P_n.

p_i is either a keyword or a keyword equated to a value.

P_i	Description
I=filename	Name of the file containing IAF stimulator data.
I omitted	Same as I=STIMOUT.
NR	Specifies that the input and output files are not to be rewound.
NT=n	Number of terminals to be processed (0 < n < 512). This tells DEMUX how much memory to allocate to process the input file. If no post radix is specified, octal base is assumed.
O=filename	Name of the file to receive translated session output.
O omitted	Same as O=OUTPUT.
SL=terminalnumber	Selective terminal number (within the range specified by the NT=n parameter) to be processed. terminalnumber is the terminal number minus 1 in the T display. If SL is omitted, all terminals within the range specified by the NT=n parameter are processed.
T	Specifies that the time of each carriage return and first output character is written with the processed data. Subtracting the two time values gives the response time. The time values can also be used to correlate data between each terminal's input and output.

Example 1:

Following is an example of a procedure file called when STM. is entered at the console. It executes the STIMULA command and performs postprocessing of the IAF stimulator output. The procedure file must be stored as a permanent indirect access file under the system user index (377777₈).

. PROC , STM .	Procedure file name. The IAF stimulator is executed.
STIMULA .	
DEMUX , NT=310 , O=NOS , T .	RESB is a user-created program to analyze DEMUX output.
GET , RESB / UN=username .	
RESB .	
ASSIGN , NE , TRASH .	The system dayfile is dumped to null equipment.
DFD , OP=I , L=TRASH .	
REQUEST (T) IDLE IAF	Processing stops until IAF is idled. Also, ASSIGN,STM,2 should be performed.
DFD , OP=I .	A partial dayfile dump is performed that includes IAF statistics.
ATTACH , ERRDATA / M=W .	Output of all failing jobs is copied to output.
COPYEI , ERRDATA , OUTPUT .	
EVICT , ERRDATA .	
RETURN , ERRDATA .	
EXIT .	
DIS .	

In the previous procedure file, output from failing jobs is copied from the file ERRDATA. In order for ERRDATA to contain this information, another procedure file is necessary to collect data on jobs that fail and place this data in ERRDATA. This procedure file (called ERRCHEK in the next example) should be called in the following manner after each IAF stimulator session.

```
GET , ERRCHEK / UN=username .
ERRCHEK .
```


Example 2:

Following is an example of this procedure file.

```
. PROC ,ERRCHEK*I .
IF ,EF.NE.0 ,EXIT .

DISPLAY ,EF .
ASSIGN ,MS ,OUTPUT .
ENQUIRE ,A .
DAYFILE .

ATTACH ,ERRDATA/UN=username ,M=W ,NA .
SKIPEI ,ERRDATA .
REWIND ,OUTPUT .
COPYEI ,OUTPUT ,ERRDATA .
RETURN ,OUTPUT ,ERRDATA .
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
COMMENT .FAIL*****
EXIT .
REQUEST ,FAIL .   ERRCHECK FAILED...

ENDIF ,EXIT .
```

If certain jobs normally cause compilation errors, additional procedure should be included to allow for these special cases.

Data is gathered on which program and which statement caused the error.

This data is placed on file ERRDATA, which is used later during postprocessing.

These comments warn users that failures have been encountered.

If serious problems have been encountered, DIS could be brought up at this point.

Care should be taken when comparing response times between different versions of the operating system with the IAF stimulator. Other factors may have a significant impact on response time. Be aware of these factors:

- The entire parameter set should be the same for each IAF stimulator run. The rate of activation of terminals, for example, can significantly impact response time.
- The hardware configurations must be identical. If an attempt is made to measure small differences, all devices that cannot be restored to their original state (normally temporary devices) should be deadstarted and fully initialized.
- If the possibility exists that permanent files have been destroyed, these files should be restored before each run. This can be accomplished easily by maintaining a copy of all required files under a user name not used during the stimulation. A procedure file can then be called to restore the files before each IAF stimulator run. This is important when stimulation loads are causing abnormal job failures. A REPLACE shortly after an unexpected abort may cause incorrect information to be written, for example.
- Job processing must be the same for each IAF stimulator run. When many terminals are to be analyzed, random job failures may be overlooked. These failures may, however, affect the entire results of the IAF stimulator run. To minimize this problem, it is necessary to check errors in all sessions quickly and easily. This can be done by using the procedures described earlier in this section.

By using the preceding techniques and procedures and observing the restrictions, it becomes easier to use the IAF stimulator to check performance and reliability and to obtain meaningful data from the system.

Introduction	17-1
Utility Description Terminology	17-2
Archive File	17-2
Alternate Storage	17-2
Destaging and Staging Files	17-3
User Index	17-4
Device Masks	17-5
Device Mask	17-5
Secondary Mask	17-7
Mass Storage Table	17-9
Catalog Track	17-9
Permanent File Catalog Entry	17-11
Track Reservation Table (TRT)	17-23
Family	17-26
Summary File	17-29
Calling the Utilities	17-36
Console Input	17-36
DIS Input	17-43
Batch Input	17-43
Terminal Input	17-44
Procedure Files	17-44
Parameters for the Permanent File Utilities	17-45
Description of Permanent File Utility Parameters	17-47
Preassigning the Archive File	17-61
Permanent File Utility Routines	17-62
Catalog Archive File (PFATC)	17-67
Catalog Permanent File Device (PFCAT)	17-67
Copy Archive File (PFCOPY)	17-72
Dump Permanent File (PFDUMP)	17-72
PFDUMP in a Secured System	17-73
Incremental Permanent File System Dump	17-73
Partial Permanent File System Dump	17-74
Full Permanent File Dump	17-74
PFDUMP File Interlocking	17-74
PFDUMP in an MSE, Tape, or Optical Disk Alternate Storage Environment	17-74
PFDUMP Protected Files	17-76
Load Permanent File (PFLOAD)	17-76
Catalog Image Record	17-78
PFLOAD Selection Process	17-79
Family Reconfiguration	17-81
PFLOAD in an MSE, Tape, or Optical Disk Alternate Storage Environment	17-82
PFDUMP/PFLOAD Example	17-82
Release Disk Space (PFREL)	17-89

Introduction

Five utility processors maintain the NOS permanent file system. They control the processes of dumping and loading permanent files, cataloging permanent files in the system and on backup storage (archive) files, copying archived files to a job as local files, and releasing the disk space of files that reside on tape or cartridge alternate storage.

The utility processors are overlays under the control of the permanent file supervisor (PFS). PFS processes the parameters in the utility command and loads the correct processing overlay. The overlays interact with the permanent file utility (PFU) routine, which manages the catalogs, permits, data allocation on a device, and the data transfer between the device and the overlay. Figure 17-1 is an overview of this procedure.

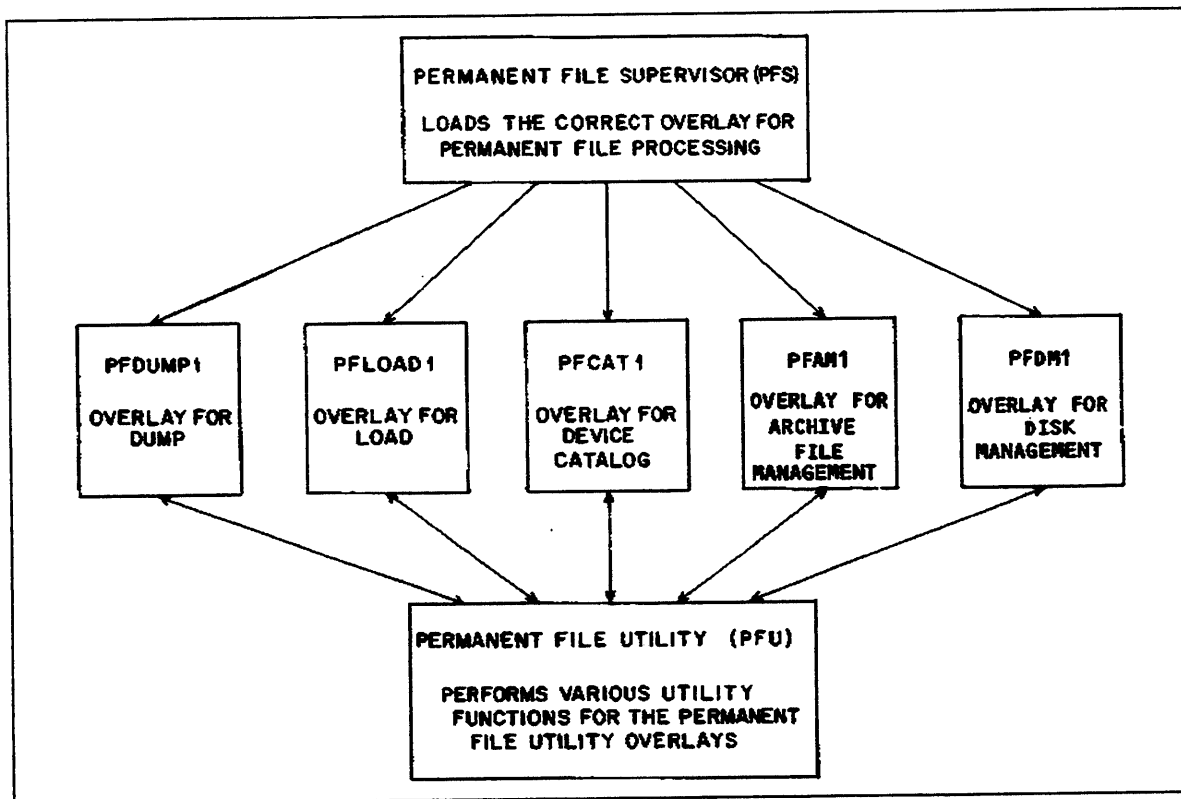


Figure 17-1. Functional Overview of the Permanent File Utility Overlays

The names and functions of the permanent file utilities follow. Detailed information about the call and operation of each utility is contained in this section.

Utility	Description
PFATC	Produces a cataloged directory of file information derived from an archive file.
PFCAT	Produces a cataloged directory of file information derived from catalog tracks on a permanent file device.
PFCOPY	Copies files from an archive file to a job as local files.
PFDUMP	Dumps files from a permanent file device to an archive file. Files created by this dump can be reloaded by the PFLOAD utility.
PFLOAD	Loads files from an archive file (created by PFDUMP) to a permanent file device.
PFREL	Releases disk space of files that reside on tape or cartridge alternate storage.

Utility Description Terminology

The descriptions of the permanent file utilities in this section include some terms you should be familiar with. These terms are defined next.

Archive File

All or some of the permanent files accumulated on mass storage can be dumped to a backup tape (or other type of backup medium) to protect the files from loss in case of a device malfunction or to free a device for temporary use during preventive maintenance. Each dump of permanent files is made to a file called an archive file; each permanent file dumped is called an archived file. Each archive file is a multirecord file in which each logical record is an archived file. If two or more dumps are made on one type of backup device, these archive files constitute a multifile archive file (refer to figure 17-2).

The archive file can be loaded back onto the permanent file system as a whole or in part and can also be used to generate reports. The individual archived files can be referenced for selective use. Refer to each utility for more information on how it uses the archive file.

Alternate Storage

The storage of permanent file data on external media other than mass storage (such as magnetic tape, MSE, or optical disk) is called alternate storage. When a file resides on alternate storage, the PFC entry and permits must reside on disk but the disk space occupied by the file data can be released.

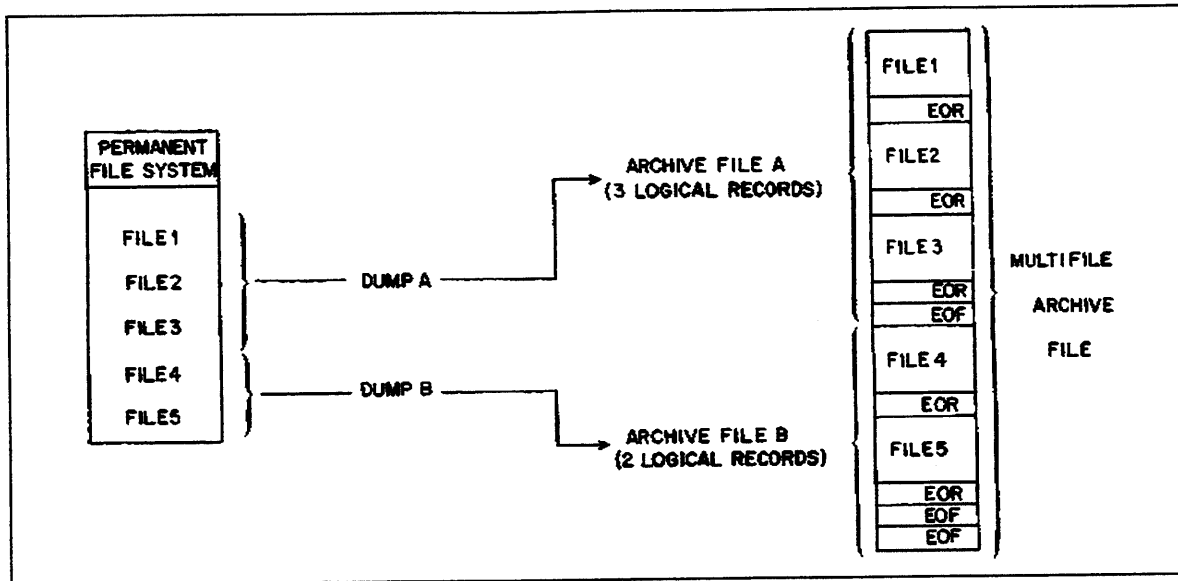


Figure 17-2. Example of Multifile Archive File Structure

Destaging and Staging Files

The process of copying file data to alternate storage from disk residence is called destaging the file. File destaging to either tape or optical storage is initiated by using the PFDUMP utility; destaging to MSE is initiated by using the SSMOVE utility.

The process of restoring file data to disk residence from a copy on alternate storage is called staging the file. File staging is initiated by entering an ATTACH, GET, OLD, or APPEND command for a file that is not currently disk resident.

User Index

A 17-bit user index is associated with each user name created on the user validation file. This index is entered through MODVAL (refer to User Validation in the Administration Handbook) with the UI parameter on the user name input directive (/username,UI=userindex) or MODVAL, by default, supplies the next available index.

Whenever you submit a job, the related user index is placed in the control point area along with the user name and other parameters that link hardware, files, and job. The permanent file manager (PFM) identifies the master device and catalog track for you by performing two masking operations that involve the user index and two sets of device parameters (device mask and number of catalog tracks) obtained from the mass storage table in central memory resident (CMR). One operation correlates the rightmost octal digit in the user index (bits 0 through 2) with the bit settings of the device mask for each device in the configuration to determine which device is your master device (refer to the device mask definition next). The other operation performs a logical AND between the remaining portion of the index and the number of catalog tracks on the master device to determine which track contains your catalog (refer to the catalog track definition later in this section).

The lower 3 bits (rightmost octal digit) of the user index are used to group users together into subfamilies. Each permanent file family (refer to the definition of family later in this section) consists of eight subfamilies, subfamily 0 through subfamily 7. Any user whose index ends in 0 belongs to subfamily 0, any user whose index ends in 1 belongs to subfamily 1, and so forth. The concept of subfamily is important in a Mass Storage Extended Subsystem (MSE) environment.

Device Masks

Two types of device masks exist for each device, the device mask (or primary mask) and the secondary mask.


Device Mask

The device mask is an 8-bit quantity used to identify the group of users that has the particular device as its master device; that is, it identifies the device that contains the group's file catalogs, all its indirect access files, and possibly some or all of the group's direct access files. (Assigning direct access files to a device is described next, under Secondary Mask.)

If bit i is set in the mask of the device, any user whose index ends in i has this device as a master device.

Example 1:

The device identified by the mask 261_8 is the master device for any user whose index ends in 0, 4, 5, or 7, because these bits are set in the mask.



When masks are assigned at device initialization time, the following rules must be observed.

- Within a family, the sum of all the device masks must be exactly 377_8 .
- Each bit position (0 through 7) must be set exactly once for the masks of the devices in one family.

Example 2:

For a family of three devices, 221_8 , 042_8 , and 114_8 are valid device masks because their sum is 377_8 and each bit is accounted for only once.

$$\begin{array}{r}
 1\ 0\ 0\ 1\ 0\ 0\ 0\ 1 = 221_8 \\
 0\ 0\ 1\ 0\ 0\ 0\ 1\ 0 = 042_8 \\
 0\ 1\ 0\ 0\ 1\ 1\ 0\ 0 = 114_8 \\
 \hline
 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1 = 377_8
 \end{array}$$

If the sum of the device masks is less than 377_8 , then at least 1 bit is not set in any of the device masks. Any user index ending in such a bit-position value does not reference a device. Therefore, if such a user tries to create a permanent file, the system issues an error message.

Example 3:

For a family of four devices, 142_8 , 020_8 , 010_8 , and 204_8 are not valid device masks because their sum is less than 377_8 . Any user whose index ends in 0 would have no master device and could not create a permanent file.

$$\begin{array}{r}
 0\ 1\ 1\ 0\ 0\ 0\ 1\ 0 = 142_8 \\
 0\ 0\ 0\ 1\ 0\ 0\ 0\ 0 = 020_8 \\
 0\ 0\ 0\ 0\ 1\ 0\ 0\ 0 = 010_8 \\
 1\ 0\ 0\ 0\ 0\ 1\ 0\ 0 = 204_8 \\
 \hline
 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0 = 376_8
 \end{array}$$

If the sum of the device masks is greater than 377_8 , then at least 1 bit is set for more than one device mask. Any user index ending in such a bit-position value references more than one master device.

Example 4:

For a family of four devices, 212_8 , 106_8 , 040_8 , and 021_8 are not valid device masks because their sum is greater than 377_8 . Bit 1 is set twice, which indicates that any user whose index ends in 1 would have two master devices, one with mask 212_8 and one with mask 106_8 .

$$\begin{array}{r}
 1\ 0\ 0\ 0\ 1\ 0\ 1\ 0 = 212_8 \\
 0\ 1\ 0\ 0\ 0\ 1\ 1\ 0 = 106_8 \\
 0\ 0\ 1\ 0\ 0\ 0\ 0\ 0 = 040_8 \\
 0\ 0\ 0\ 1\ 0\ 0\ 0\ 1 = 021_8 \\
 \hline
 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1 = 401_8 \\
 \\
 \\

 \end{array}$$

It is a necessary but not sufficient condition that all masks for one family total 377_8 . It is possible, but not allowable, that in one family the device masks total 377_8 but some bit is set in more than one device mask.

Example 5:

For a family of three devices, 261_8 , 115_8 , and 001_8 are not valid device masks because bit 0 is accounted for three times and bit 1 is not set at all. Any user whose index ends in 0 would have three master devices. Any user whose index ends in 1 would have no master device.

$$\begin{array}{r}
 1\ 0\ 1\ 1\ 0\ 0\ 0\ 1 = 261_8 \\
 0\ 1\ 0\ 0\ 1\ 1\ 0\ 1 = 115_8 \\
 0\ 0\ 0\ 0\ 0\ 0\ 0\ 1 = 001_8 \\
 \hline
 1\ 1\ 1\ 1\ 1\ 1\ 0\ 1 = 401_8 \\
 \\
 \\


 \end{array}$$

Secondary Mask

The secondary mask of a device is an 8-bit quantity used to identify groups of users who can place direct access files on the particular device. If bit i is set in the secondary mask of the device, any user whose index ends in i can place direct access files on this device.

Example 1:


A device whose secondary mask is 072_8 can contain direct access files of any user whose index ends in 1, 3, 4, or 5, because these bits are set in the mask.



Direct access files may or may not reside on a master device, depending on the user index and secondary mask. The appropriate bit must be set in the secondary mask of the device on which the user's direct access files are to reside. A direct access file is placed on a valid device with the most space available unless the user specifies a particular device.

Example 2:

The following situations occur for a device whose mask is 261_8 and secondary mask is 072_8 .



- This device is the master device for, and can contain direct access files of, any user whose index ends in 4 or 5 (bits 4 and 5 are set in both the device mask and the secondary mask).
- This device is the master device for, but cannot contain direct access files of, any user whose index ends in 0 or 7 (bits 0 and 7 are set in the device mask but not in the secondary mask).
- This device is not the master device for, but can contain direct access files of, any user whose index ends in 1 or 3 (bits 1 and 3 are set in the secondary mask but not in the device mask).
- This device is not the master device for, and cannot contain direct access files of, any user whose index ends in 2 or 6 (bits 2 and 6 are not set in either the device mask or the secondary mask).

Existing files can be defined as direct access permanent files, regardless of the secondary mask, for user indexes above 377700_8 (AUIMX). However, PFLOAD will enforce the secondary mask check when loading these files. Direct access files for MSE subfamily user indexes (377760 through 377767_8) are always loaded to their master devices regardless of previous residence or secondary mask values.

There are no restrictions on the number of devices in one family that can have the same bits set in the secondary mask or regarding the sum of the secondary masks for devices in the same family. The secondary mask must be a number between 0 and 377_8 .

Example 3:

For a family of three devices, 000 , 131_8 , and 326_8 are possible secondary masks. Any user whose index ends in 0, 3, 4, or 6 can place direct access files on the device whose secondary mask is 131_8 . Any user whose index ends in 1, 2, 4, 6, or 7 can place direct access files on the device whose secondary mask is 326_8 . No user can place direct access files on the device whose secondary mask is 000 . Any user whose index ends in 4 or 6 can place direct access files on two devices, and any user whose index ends in 5 cannot use direct access files.

$$\begin{array}{r}
 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0 = 000_8 \\
 0\ 1\ 0\ 1\ 1\ 0\ 0\ 1 = 131_8 \\
 1\ 1\ 0\ 1\ 0\ 1\ 1\ 0 = 326_8 \\
 \hline
 1\ 1\ 0\ 1\ 1\ 1\ 1\ 1 = 337_8 \\
 1\quad 1
 \end{array}$$

Mass Storage Table

The configuration of mass storage devices currently available to the system is defined by the CMR mass storage table (MST). Each logical device in this configuration has an entry in this table. Refer to the NOS Version 2 Systems Programmer's Instant for the MST format.

Catalog Track

Your catalog track is a track on your master device containing the catalog entries (definition follows) that define and specify the location of each permanent file you have created. Users are assigned by groups to catalog tracks according to their user index and number of catalog tracks on their master device.

The number of catalog tracks on a device is established when the device is initialized or by default. Default values are:

Type of Device	Default Number of Catalog Tracks
Extended memory (DE)	4
834 (DD)	40 ₈
836 (DG)	40 ₈
844 (DI, DJ, DK, DL)	40 ₈
885-11/12 (DM, DQ)	10 ₈
885-42 (DB)	10 ₈
887 (DF, DH)	10 ₈
895 (DC)	10 ₈
819 (DV, DW)	10 ₈
5832/5833/5838 (EA-EH)	10 ₈
9853 (DN)	10 ₈
DDP path to extended memory (DP)	4
Private device	1

The number of catalog tracks is always a power of 2; the maximum number is 200₈. Therefore, the possible numbers of catalog tracks on a device are:

1, 2, 4, 10₈, 20₈, 40₈, 100₈, 200₈

Reducing these numbers by 1 produces the following numbers, referred to as track masks.

0, 1, 3, 7, 17₈, 37₈, 77₈, 177₈

The track masks in binary form are:

```

          0
          1
         1 1
        1 1 1
       1 1 1 1
      1 1 1 1 1
     1 1 1 1 1 1
    1 1 1 1 1 1 1
  
```

These track masks, except the first one, have all bits set; when bits 3 through 9 of the user index for a particular user are ANDed with the track mask for his/her master device, the user's catalog track number on that device is produced.

Example:

For a family of two devices, the following device masks, number of catalog tracks, and track masks are valid.

Device	Device Mask		Number of Catalog Tracks	Corresponding Track Mask
1	221 ₈	or 10 010 001	40 ₈	40 ₈ -1=37 ₈ or 11 111
2	156 ₈	or 01 101 110	20 ₈	20 ₈ -1=17 ₈ or 1 111
	377 ₈	11 111 111		

A user whose index is 14224₈ is assigned device 1 as his master device because bit 4 (last digit in 14224₈) is set in the device mask for device 1. Therefore, device 1 contains this user's catalog track.

The binary form of 14224₈ is 001 100 010 010 100.

Therefore, the catalog track number for this user is 22₈ (on device 1) because the logical AND of bits 3 through 9 of the user index with the track mask for device 1 is:

```

          1 0 0 0 1 0 0 1 0
AND      1 1 1 1 1
-----
          1 0 0 1 0 or 228
  
```

The end of a catalog track is indicated by an end-of-information (EOI) on the device. All catalog tracks are linked in the track reservation table (TRT) and appear as one logical chain. When catalog entries exceed the logical track, a continuation track is reserved at the end of the catalog track chain in the track reservation table. The disk linkage bytes for the overflowed track point to the continuation track.

Permanent File Catalog Entry

Files in the permanent file system are referenced by dynamically updated permanent file catalog (PFC) entries on the catalog tracks of master devices. Whenever you create a permanent file, a catalog entry that specifies the characteristics of the file (access information about the file and its location) is entered on a catalog track on your master device. When you modify the file, the PFC entry is updated to reflect the modification.

The format of the PFC entry follows. The 4-character names to the right of each word are symbols used by the system to reference the PFC entry fields. These symbols are defined in common deck COMSPFM.

word	59	53	47	44	41	35	23	17	11	0	
0	filename						user index				FCFN, FCUI
1	file length					track			sector		FCLF, FCBT, FCBS
2	permit random index					creation date and time					FCRI, FCCD
3	access count					data modification date and time					FCAC, FCMD
4	cl	mode	ef	ec	dn	last access date and time					FCCT, FCAM, FCEF, FCEC, FCEO, FCAD
5	optical disk address					control modification date and time					FCOA, FCKD
6	pr	br	ss	ac	utility control date and time			expiration date		FCRS, FCBR, FCFS, FCAP, FCUD	
7	file password						expiration date				FCPW, FCXD
10	cafiags			cat		cartridge alt. storage address					FCAF, FCAT, FCAA
11	access categories					charge number					FCAL, FCFC
12	charge number										FCCN
13	project number										FCP1
14	project number										FCP2
15	tafiags		file seq. number			tape alt. storage VSN				FCFT, FCTS, FCTV	
16	user control word										FCUC
17	reserved for installation										FCIW

05389-1M

Field	Word	Bit(s)	Description
filename	0	59-18	Permanent file name.
userindex	0	17-0	User index of file creator.
file length	1	59-36	Length of the file in PRUs.
track	1	23-12	Beginning track of the file.
sector	1	11-0	Beginning sector of the file.
permit random index	2	59-36	Random disk address of first permit sector.
creation date and time	2	35-0	yymmddhhmmss in octal when this file was first entered on the permanent file system. The year (yy) is biased by 70.
access count	3	59-36	Total number of times this file has been accessed.

Field	Word	Bit(s)	Description
data modification date and time	3	35-0	yymmddhhmmss in octal when data in this file was last modified. The year (yy) is biased by 70. For direct access files this field is updated only when the file is attached in a modifiable mode.
ct	4	59-54	File category: 0 Private. 1 Semiprivate. 2 Public.
mode	4	53-48	Mode of access for semiprivate and public files: 0 Write, read, execute, append, modify, and/or purge. 1 Read and/or execute. 2 Append. 3 Execute. 4 Negate previous permission. 5 Modify. 6 Read and/or execute, allow modify. 7 Read and/or execute, allow append. 10 Update. 11 Read and/or execute, allow update.
ef	4	47-45	Error flag: 0 No error. 1 EOI changed by recovery.

Field	Word	Bit(s)	Description
ec	4	44-42	Error code: 0 No error. 1 Error in file data. 2 Error in permit entries for file. 3 Error in data and permit entries. 4 Error in file length. 5-7 Reserved.
dn	4	41-36	Device number (1 through 77 ₈) of the device on which the direct access file resides if other than the user's master device. If the file resides on the master device, dn is 0.
last access date and time	4	35-0	yymmddhhmmss in octal when this file was last accessed. The year is biased by 70.
optical disk address	5	59-36	Relative address of the file on optical disk. The complete address consists of this relative address, the version number of the archive file, and the optical disk VSN. The version number and VSN are specified in the cartridge alternate storage address field in word 10 ₈ .
control modification date and time	5	35-0	yymmddhhmmss in octal when this file's control information (catalog entry and permit record data) was last modified.

Field	Word	Bit(s)	Description
pr	6	59–57	Preferred residence ¹ for this file: <ol style="list-style-type: none"> 1 File locked to disk (PR=L specified by file owner with validation). 2 Disk residence preferred (PR=D specified by file owner). 3 Cartridge alternate storage (MSE or optical disk) residence preferred (PR=M specified by file owner). 4 No preferred residence (PR=N specified by file owner). 5 Tape alternate storage residence preferred (PR=T specified by file owner).
br	6	56–54	Backup requirement ² for this file: <ol style="list-style-type: none"> 1 No backup copy of file required. (BR=N specified by file owner). 2 Two backup copies required (BR=Y specified by file owner). 3 One backup copy required (BR=MD specified by file owner).
ss	6	53–48	Subsystem code for this file: <ol style="list-style-type: none"> 0 Null Subsystem. 1 BASIC Subsystem. 2 FORTRAN Subsystem. 3 FTNNTS Subsystem. 4 Execute Subsystem. 5 Batch Subsystem.

1. Refer to Permanent File Commands in the NOS Version 2 Reference Set, Volume 3 for details on the PR and BR parameters.

2. Refer to Permanent File Commands in the NOS Version 2 Reference Set, Volume 3 for details on the PR and BR parameters.

Field	Word	Bit(s)	Description
ac	6	47	Alternate CATLIST permission: ³ 0 Alternate users cannot obtain CATLIST information about the file (AC=N specified by file owner). 1 Alternate users can obtain CATLIST information about the file (AC=Y specified by file owner).
utility control date and time	6	35-0	yymmddhhmmss in octal set by PFM and PF utilities and used by PFDUMP in determining whether to dump the file when the OP=M option is specified. This field is updated when the data modification date and time field is updated or PFDUMP finds a file busy (attached in modifiable mode). This ensures that the file is selected on subsequent dumps even if the data modification date and time field does not change.
file password	7	59-18	Optional password.
expiration date	7	17-0	yymmdd in octal when this file's password expires. The year is biased by 70.

³. Refer to Permanent File Commands in the NOS Version 2 Reference Set, Volume 3 for details on the AC parameter.

Field	Word	Bit(s)	Description									
caflags	10	59–48	Cartridge alternate storage flags; refer to section 11, MSE, for details.									
			<table border="1"> <thead> <tr> <th>Flag</th> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AFLOK</td> <td>59</td> <td> <p>Interlock flag used during the process of destaging a file to MSE or optical disk:</p> <p>0 indicates that either the file is not in the destage process or the file was modified during the destage process. If you attempt to set an alternate storage address with the SETASA request, the request is rejected.</p> <p>1 indicates the file is in the destage process and is unmodified. After the destage process is complete, you can enter a SETASA request to set an alternate storage address in the PFC entry for the file. The SETASA request also clears this flag.</p> </td> </tr> <tr> <td>AFFRE</td> <td>54</td> <td> <p>Free up cartridge flag:</p> <p>0 means do not remove file from cartridge.</p> <p>1 means remove file from cartridge.</p> </td> </tr> </tbody> </table>	Flag	Bit	Description	AFLOK	59	<p>Interlock flag used during the process of destaging a file to MSE or optical disk:</p> <p>0 indicates that either the file is not in the destage process or the file was modified during the destage process. If you attempt to set an alternate storage address with the SETASA request, the request is rejected.</p> <p>1 indicates the file is in the destage process and is unmodified. After the destage process is complete, you can enter a SETASA request to set an alternate storage address in the PFC entry for the file. The SETASA request also clears this flag.</p>	AFFRE	54	<p>Free up cartridge flag:</p> <p>0 means do not remove file from cartridge.</p> <p>1 means remove file from cartridge.</p>
Flag	Bit	Description										
AFLOK	59	<p>Interlock flag used during the process of destaging a file to MSE or optical disk:</p> <p>0 indicates that either the file is not in the destage process or the file was modified during the destage process. If you attempt to set an alternate storage address with the SETASA request, the request is rejected.</p> <p>1 indicates the file is in the destage process and is unmodified. After the destage process is complete, you can enter a SETASA request to set an alternate storage address in the PFC entry for the file. The SETASA request also clears this flag.</p>										
AFFRE	54	<p>Free up cartridge flag:</p> <p>0 means do not remove file from cartridge.</p> <p>1 means remove file from cartridge.</p>										

Field	Word	Bit(s)	Description
<i>caflags (Continued)</i>			
		Flag	Bit
		Description	
		AFTMP	53
			Temporary error flag:
			0 indicates that no temporary error conditions have been detected.
			1 indicates that a temporary error condition prevented the MSE image of the file from being staged to disk. After the error condition is detected, this flag is cleared so that the file staging can be retried on the next request for the file.
		AFVER	52
			Verification flag:
			0 means that if the cartridge alternate storage address is not zero, a successful stage/ destage operation has occurred and system control errors should not be encountered.
			1 means that the PFC entry containing the cartridge alternate storage address was reloaded and a stage attempt has not yet been made. If a system control error is detected, the probable cause is that the cartridge alternate storage address is obsolete, probably because an obsolete dump tape was used during the reload.

Field	Word	Bit(s)	Description
<i>caflags (Continued)</i>			
		Flag	Bit
		Description	
		AFPDE	50
			Data error flag:
			0 means that no unrecoverable read errors have been detected.
			1 means that an unrecoverable read error prevented the MSE image of the file from being staged to disk.
		AFPSE	49
			System control error flag:
			0 means that no system error conditions have been detected.
			1 means that a system error condition prevented the MSE image of the file from being staged to disk.

Field	Word	Bit(s)	Description		
			Flag	Bit	Description
caflags (<i>Continued</i>)			na	48	Reserved.
cat	10	41–36	Cartridge alternate storage type:		
			0		File resides on disk only.
			1		Reserved.
			2		File has a copy on MSE.
			3		File has a copy on optical disk.
cartridge alt. storage address	10	35–0	Cartridge alternate storage address:		
			0		An MSE or optical disk copy of the file does not exist.
			Nonzero		Depends on the value of cat. If cat is 2, specifies the location of the beginning of the file on MSE. If cat is 3, specifies both the packed VSN of the optical disk containing the archive file (bits 23–0) and the version of the archive file (bits 35–24). The optical disk address field in word 5 contains the relative address of the file copy within the archive file.
al ⁴	11	38–36	File security access level:		
			0		Access level 0.
			1		Access level 1.
			:		
			7		Access level 7.

4. The 8 access levels and 32 access categories are referenced by names defined by or in micros in system common deck COMSMLS. The default names are LVL0 through LVL7 for access levels and CAT00 through CAT31 for access categories. Each site determines what the access levels and categories mean, and can change the corresponding names to clarify that relationship. For example, access level 3 may be defined as meaning secret, so the name for access level 3 can be changed from LVL3 to SECRET. If any of the COMSMLS micros are changed, all decks calling COMSMLS must be reassembled.

Field	Word	Bit(s)	Description
access categories	11	31–0	File security access categories: Each bit indicates whether the file has the corresponding access category as part of its category set. Bit 0 corresponds to access category 0; bit 1 to access category 1, and so forth. A file's category set can be any, all, or none of the 32 available categories.
charge number	12	59–0	Charge number to which this file is billed.
project number	13–14	59–0	Project number to which this file is billed.
taflags	15	59–48	Tape alternate storage flags; refer to section 12, Tape and Optical Disk Alternate Storage, for details.

Flag	Bit	Description
TFLOK	59	Interlock flag used during the destaging of a file to tape: 0 indicates that either the file is not in the destage process or the file was modified during the destage process. If you attempt to set an alternate storage address with the SETASA request, the request is rejected. 1 indicates the file is in the destage process and is unmodified. After the destage process is complete, you can enter a SETASA request to set an alternate storage address in the PFC entry for the file. The SETASA request also clears this flag.
TFACS	56	File on ACS cartridge tape.
TFCTS	55	File on either ACS or CTS cartridge tape.
TFLIF	54	File on tape recorded in LI format.
TFESV	53	Data error on secondary VSN flag.

Field	Word	Bit(s)	Description
<i>taflags (Continued)</i>			
		Flag	Bit
		Description	
		TFEPV	52
			Data error on primary VSN flag.
		TFSVN	51
			Secondary VSN not accessible flag.
		TFPVN	50
			Primary VSN not accessible flag.
		TFSVS	49
			Secondary VSN exists flag:
			0 indicates a secondary VSN does not exist.
			1 indicates a secondary VSN does exist.
		TFVER	48
			PFLOAD verification flag:
			0 means that if the tape alternate storage VSN pointer is not zero, a successful stage/destage operation has occurred and system control errors should not be encountered.
			1 means that if the PFC entry containing the tape alternate storage VSN pointer was reloaded and a stage attempt has not yet been made. If a system control error is detected, the probable cause is that the tape alternate storage VSN pointer is obsolete, probably because an obsolete dump tape was used during the reload.

Field	Word	Bit(s)	Description
file seq. number	15	41-24	Sequence number of the file on the dump tape produced by PFDUMP.
tape alt. storage VSN	15	23-0	<p>Tape alternate storage VSN pointer in the following format:</p> <p style="text-align: center;">12/cc, 12/nnnn</p> <p>where cc is any two alphanumeric characters and nnnn is a 12-bit binary integer.</p> <p>The tape's primary VSN is cc followed by nnnn converted to a decimal integer in the range of 0000 to 4095 with leading zeros. The secondary VSN, if it exists, is derived by adding 5000 to the primary VSN.</p>
user control word	16	59-0	User control information.

Track Reservation Table (TRT)

Every device in a permanent file system has a label track (usually track 0) that contains a label sector describing the device (family name, device mask, secondary mask, location of permit information, catalog information, and indirect files) and a number of physical record units (PRUs) containing the track reservation table (TRT). The TRT describes the physical layout of data on the device and is the key to allocating information on the device.

The TRT contains single-word entries that define track linkage and bit-setting controls for those tracks (refer to figure 17-3). Each word has four 12-bit linkage bytes and three sets of 4-bit control settings that match the 4 bytes (refer to figure 17-4). Linkage format is given in figure 17-5.

The numbering of the 12-bit linkage bytes or cells corresponds to the numbering of tracks on the device, with the first track starting at 4000. The entry in a cell references the next cell and its associated track. This next track either continues the information or starts a new sequence depending on the first-track bit setting.

Figures 17-3 and 17-4 show a sequential linkage of tracks (in figure 17-4, from track 4000 to track 4001 to 4002 to 4003). This numerical sequence is purely illustrative. The linkage could just as well have been:

4000 4002 4015 4012

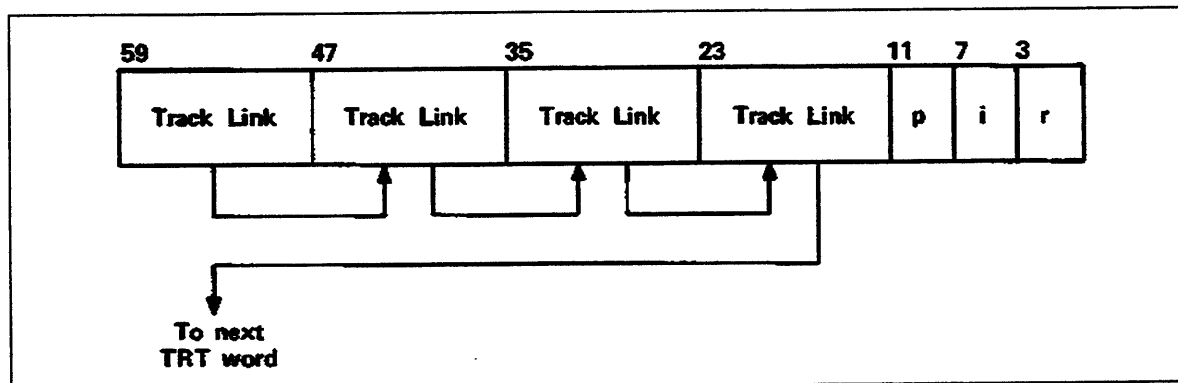


Figure 17-3. Track Reservation Table Word

Each track reservation table word contains the following entries:

Entry	Description
Track Link	Address of the next track that is a logical continuation of this file. (The track links are shown as sequential within a word, but this is not a requirement.)
p	Bit settings for identifying the first track of a preserved file (permanent file chain or queued file).
i	Bit settings for establishing interlock of a track.
r	Bit settings for track reservation.

The first group of control settings (bits 8 through 11) is used to identify those tracks that begin a sequence of file information. If any one of these bits is set, the associated track is the first track of a chain that may extend across a number of tracks. This chain can be a direct access file, an indirect access file data chain, a catalog chain, a permit chain, a system dayfile, or a queued file.

The second group of control settings (bits 4 through 7) is used to interlock tracks. If any one of these bits is set, the associated track cannot be accessed as long as this bit remains set. Whenever a file is accessed, the system automatically interlocks the catalog track containing the file. The interlocking capability is necessary because PFM may be processing several requests directed at one file simultaneously. Without interlock, these requests could overlap.

The third group of control settings (bits 0 through 3) is used to identify reserved tracks. A track is reserved either because it has data written on it or it is a flawed track. A flawed track is removed from availability by reservation at deadstart, by subsequent initialization, or by reformatting with the FORMAT utility for 844 devices.

Figure 17-4 shows matching of four reserve track bit settings with corresponding track link bytes. The same correspondence holds for the interlock bits (4 through 7) and the first-track bits (8 through 11).




Figure 17-4. Bit Settings for Track Link Bytes

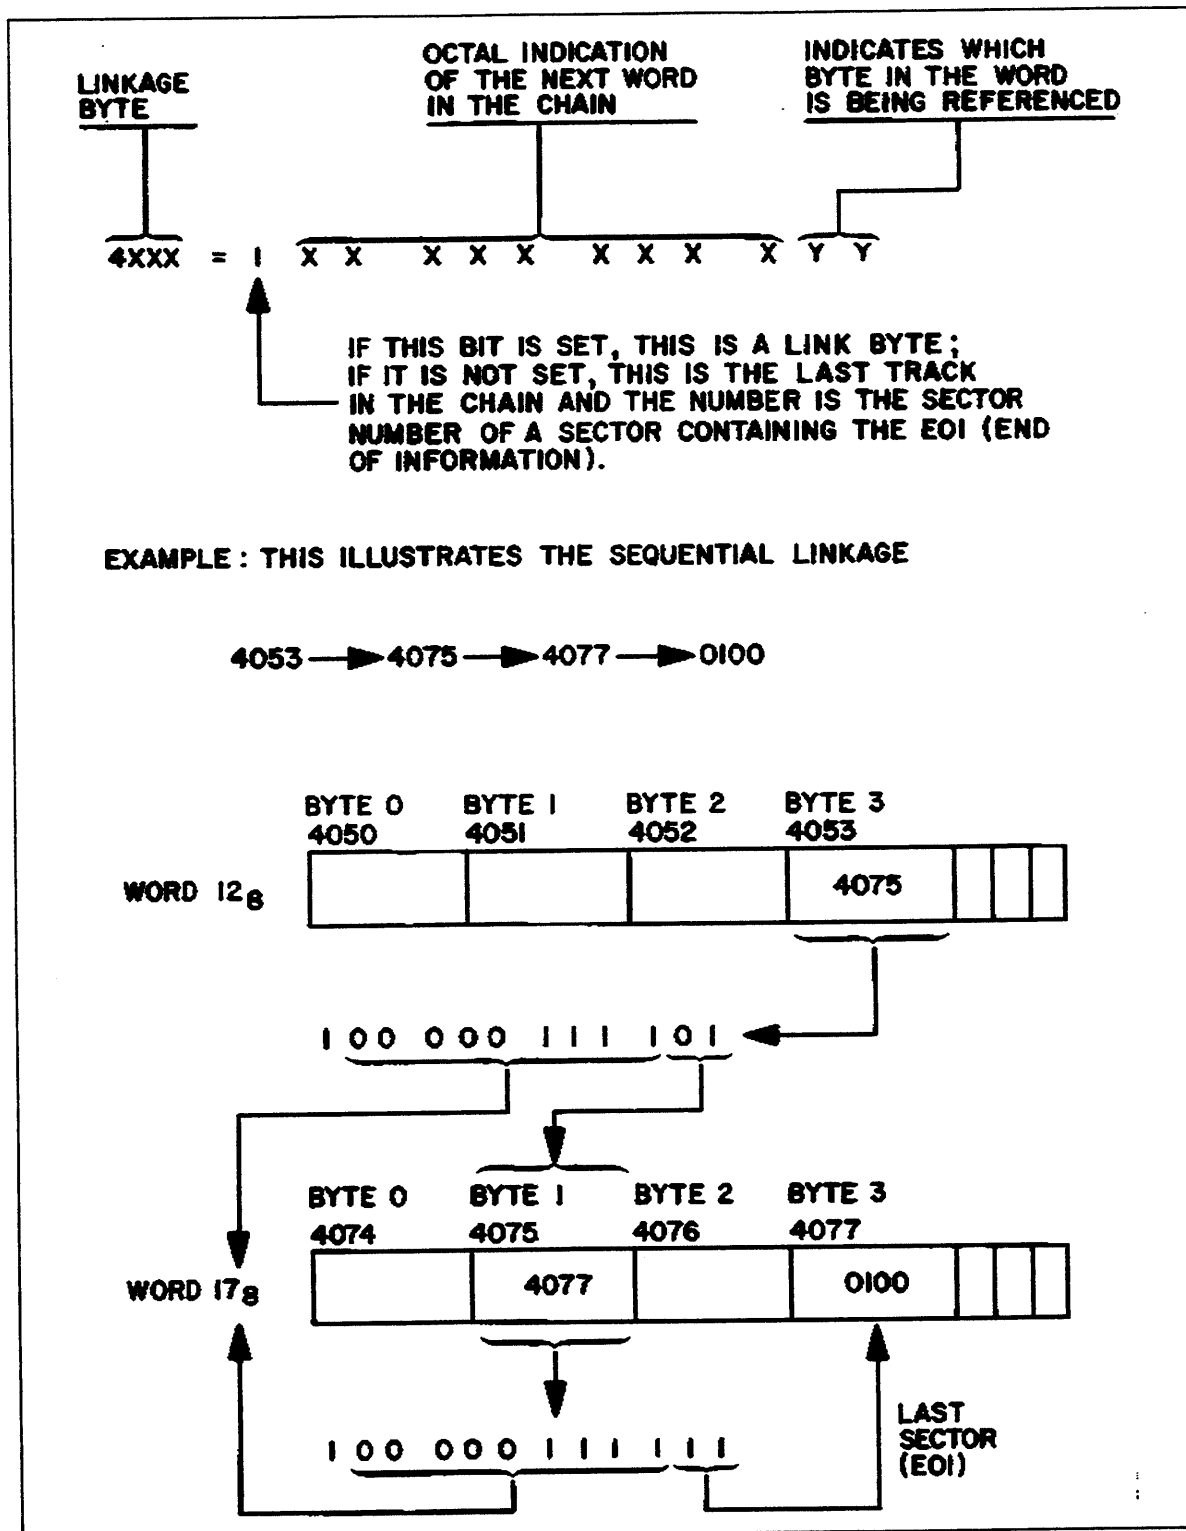


Figure 17-5. Linkage Format With Example

Family

A family is a collection of from 1 to 63 logical devices identified by a 1- to 7-character family name. Users and their files are grouped together according to the family available to them. Usually, the grouping is within the configuration on which these users normally run jobs. However, a family can be connected to any configuration and still relate to the same users and files.

An auxiliary device, identified by a 1- to 7-character pack name, is a single device that is not included in any family and can be accessed by validated users from any family. An auxiliary device provides users with an alternative to the master devices for storing and accessing permanent files. Use of auxiliary devices enables an analyst to provide special sets of permanent files for selected users or for designated periods. For example, an auxiliary device could be made available from 1200 to 1700 every day for any properly validated user.

A permanent file device is either a member of a family or an auxiliary device. Permanent files on a family device are accessed through user catalogs contained on a master device within the family. The user catalogs that reference permanent files on an auxiliary device are contained on that device; that is, an auxiliary device is a self-contained entity.

Families and auxiliary devices are defined at initialization time by the PF entry in the EQPDECK or by the DSD INITIALIZE command. Normally, a configuration has one family available. Additional families can be defined or introduced (on removable devices) in the same configuration. If more than one family is available in a configuration, you supply the family name at login or on the USER command. The default family is used if no family name is supplied.⁵ Any job can change its associated family name by using the USER command to specify the new family name. ⁶ A system origin job can use the FAMILY command to change the family name associated with the job. If the FAMILY command is included in any nonsystem origin job, the job aborts.

The FAMILY command is valid for unsecured systems only.

The format of the FAMILY command is:

```
FAMILY, familyname.
```

Parameter	Description
familyname	The 1- to 7-character name of a family of permanent file devices. If omitted, the default family name specified at deadstart is assumed.

5. The default family name is specified at deadstart time with the EQPDECK FAMILY entry and cannot be changed dynamically.

6. Refer to the IPRDECK entry in section 2, Deadstart Decks, and to the DSD command ENABLE, SECONDARY USER COMMANDS in section 5, DSD Commands; also refer to the USER command in the NOS Version 2 Reference Set, Volume 3.

If an alternate family of permanent file devices is introduced into the configuration without a user validation file, the job to create the user validation file could include a FAMILY command to identify the alternate family. If the familyname parameter is omitted, the default family name is assumed.

Figure 17-6 is an example of a typical set of configurations.

Example:

One configuration with six permanent file devices [equipment status table (EST) ordinals 6 through 13] is identified as system A. Three of the devices are grouped into a family with the name FAMA. They have device numbers 40, 41, and 42. The remaining three EST ordinals have been defined as removable. This means that the system will allow family and auxiliary devices to be introduced on this equipment during system operation.




Figure 17-6. Example of Families in Two Configurations

A second configuration with six permanent file devices is identified as system B. Three of the devices have been grouped into a family with the name FAMB. They also have device numbers 40, 41, and 42. The remaining three pieces of equipment have been defined as removable.

Two auxiliary devices are available to both systems. These have the pack names AUX1 and AUX2.

Users of FAMA would normally run jobs on system A. Users of FAMB would normally run jobs on system B.

The removable equipment on system A (EST ordinals 11, 12, and 13) could be used for two purposes:

- To mount auxiliary devices AUX1 and/or AUX2 as required by users of FAMA. If users of FAMB need access to AUX1 and/or AUX2, these auxiliary devices will have to be mounted on some combination of equipment defined by EST ordinals 11, 12, and 13 on system B.
- To mount FAMB devices when they are no longer accessible through system B. In this case, users of FAMB will have to transfer their access to system A. This transfer could involve moving disk packs from equipment defined by EST ordinals 6, 7, and 10 in system B to equipment defined by EST ordinals 11, 12, and 13 in system A, or system A could already have alternate channel connections to equipment defined by EST ordinals 6, 7, and 10 in system B but would define them as EST ordinals 11, 12, and 13 in its own system. If the users of FAMB access their files through system A, they submit their jobs to system A or dial into system A using a different telephone number. (They would dial the same number if communication gear is switched.) The ISF command must be used to make the validation file from FAMB available on system A.

Assuming that access to all of FAMB is transferred from system B to system A, the new array of system A will be as illustrated in figure 17-7. System A now has two families, its original default family FAMA and the newly attached family FAMB whose devices now have the EST ordinals 11, 12, and 13. Its device numbers (40, 41, and 42) remain the same. The device numbers happen to be the same as those used within FAMA but device numbers provide uniqueness only among devices within a family. Accordingly, total uniqueness of a device is provided by the combination of family name and device number.

Users of FAMA will still be able to access their files without specifying the family name since FAMA is still the default for system A. However, users of FAMB will have to specify to system A that they belong to FAMB either at login or on the USER command.

If another equipment (EST ordinal 14) were available on system A, either auxiliary pack AUX1 or AUX2 could be mounted on it and users from either family could access files on this pack.




Figure 17-7. Example of Transferring Family Access

Summary File

Any of the permanent file utilities can create a summary file containing utility processing information. The summary file is a machine-readable version of the information available on the utility output file. This summary file is used as input by the utility GENPFD; however, a site-written utility could use this information to generate selection directives for permanent file utility operations such as destaging files and releasing disk space, or to obtain permanent file statistics for billing, and so forth.

The summary file consists of a prefix table followed by a system identifier block, a device status block, an archive file identifier block, CIR entries block(s), and PFC entries block(s). The format of the prefix block allows you to use the CATALOG command (refer to the NOS Version 2 Reference Set, Volume 3) to display some information about the contents of the summary file and the environment in which the summary file was created.

Additional block types may be added in future releases of NOS; thus, any site-written utility should be coded to skip (and ignore) any block types other than the ones documented here.


The format of the summary file follows:

Prefix Table (all utilities)

	59	47	35	17	0
Word 0	7700B	0016B		0	
1	record name			0	
2	yy/mm/dd				
3	hh.mm.ss				
4	0				
5	0				
6	0				
7	0				
10	comment				
16					

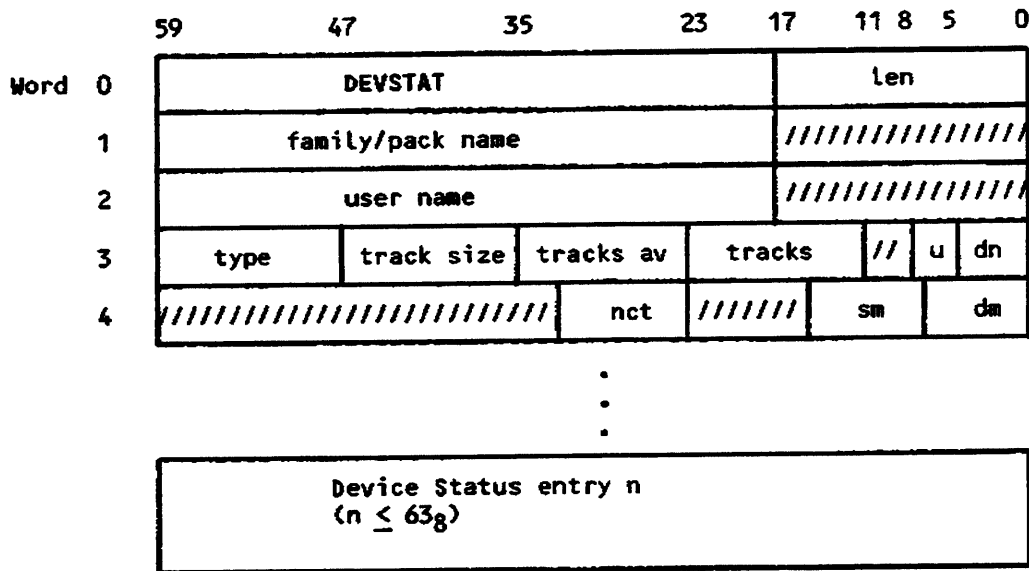
Field	Word	Bits	Description
record name	1	59-18	Name of the record recognized by common deck COMCSRT. The default value is the current date in the format yymmdd. You can use the SR parameter to specify the record name, if desired.
yy/mm/dd	2	59-0	Start date of the utility.
hh.mm.ss	3	59-0	Start time of the utility.
comment	10-16	59-0	Text indicating the utility name, the family or pack name (if applicable), and the system title from CMR words SYTL through SYTL+3.

System Identifier Block (all utilities)



Field	Word	Bits	Description
len	0	0-17	Length of this block excluding the control word (13 _g).
utility name	1	59-18	Name of the permanent file utility that created the file.
max block length	1	17-0	Maximum length of subsequent data blocks (1000 _g). This value can be used by a program interpreting the file to allocate memory.
yy/mm/dd	2	59-0	Start date of the utility.
hh.mm.ss	3	59-0	Start time of the utility.
system title	4-7	59-0	Contents of CMR words SYTL through SYTL+3.
system version	10-4	59-0	Contents of CMR words SVNL and SVNL+1.
packed date and time	12	35-0	Start date and time of the utility in packed format.
system level	13	17-0	Value of tag NOSLVL defined in COMSVR.

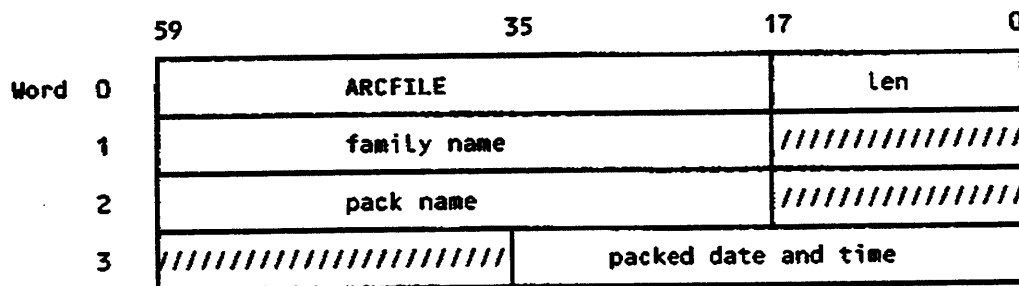
Device Status Block (PFCAT, PFDUMP, PFLOAD, and PFREL utilities only)



Field	Word	Bits	Description
len	0	11-0	Length of this block excluding the control word. Block length = (number of device status entries)*2+2.
family/pack name	1	59-18	Name of the family or auxiliary device to which the information pertains. For an auxiliary device, there will be only one device status entry and the device number will be zero.
user name	2	59-18	User name if this block is for a private auxiliary device; otherwise, this field is 0.
type	3	59-48	Equipment mnemonic from the device's EST entry.
track size	3	47-36	Number of sectors per track. This field allows computation of the space occupied by direct access files.
tracks av	3	35-24	Number of tracks available on this device. This field is used in disk space release processing.
tracks	3	23-12	Number of tracks on this device.

Field	Word	Bits	Description
u	3	8-6	Number of units on this device.
dn	3	0-5	Device number.
nct	3	32-24	Number of catalog tracks on this device.
sm	4	7-4	Secondary device mask.
dm	4	0-3	Device mask. The residency device number of a file that resides on a master device can be determined by matching the file's user index to a device mask.


Archive File Identifier Block (PFATC, PFCOPY, and PFLOAD utilities only)



Field	Word	Bits	Description
len	0	17-0	Length of this block excluding the control word (3).
family name	1	59-18	Name of the family to which the archive file belongs.
pack name	2	59-18	Name of the auxiliary device to which the archive file belongs.
packed date and time	3	35-0	Date and time from the archive file label in packed format.


If the permanent file utility does not find the archive file label or if an error is detected in the archive file label, a 3-word block is written with all fields zero.

Catalog Image Record Entries Block (PFATC, PFDUMP, and PFLOAD utilities only).



Field	Word	Bits	Description
len	0	17-0	Length of this block excluding the control word (2 to 1000 _g). Block length = (number of CIR entries)*2. If more than 400 _g CIR entries exist, more than one CIR block will be written to the summary file.
CIR entry	1-2	59-0	Catalog image record entry from the archive file. Refer to Catalog Image Record later in this section for the format of a CIR entry.

PFC Entries Block (all utilities)



Field	Word	Bits	Description
len	0	17-0	Length of this block excluding the control word (20 ₈ to 1000 ₈). Block length = (number of PFC entries)*21 ₈ .
PFC entry	1-20	59-0	PFC entry of the file processed by the permanent file utility. Refer to Permanent File Catalog Entry earlier in this section for the format of a PFC entry. For PFDUMP and PFATC, the first track pointer in the PFC entry indicates whether a file was disk resident before it was dumped. If PFDUMP staged the file to disk to dump data, the data is on the archive file, but the pointer is 0 (zero). The PFC only flag in the CATE block shows whether the data was dumped.
p	21	0	PFC only file flag (used only by PFDUMP and PFATC). This flag indicates that only the file's PFC entry and permits were written to the archive file.

Calling the Utilities

Permanent file utility operations can be initiated through console input or DIS input to a K display, batch input, terminal input, or procedure files. However, for batch or terminal input, the user must be validated for system origin privileges and the system must be in debug mode.

Console Input

It is not possible to preassign an archive file when using the console input sequence. Since a preassigned archive file is usually required, this type of call to the utilities should be avoided in most cases (refer to DIS Input later in this section). However, the console input sequence can be used when an archive file is not required, as with the PFCAT utility.

Input from the console requires the following sequence of entries and responses.

1. Call PFS by entering:

X.PFS.

The B,O display indicates the job sequence name (jsn) of PFS. To the right of the entry appears the message:

REQUEST *K* DISPLAY

2. Activate the K display for that control point by entering:

K,jsn.

jsn	Description
jsn	Job sequence name of PFS as specified on the B,O display.

The display shown in figure 17-8 appears on the left screen. Instructions on the display describe how to select the desired utility.

3. Select the desired utility by entering:

K.uo.

uo is one of the following:

uo	Description
AT	Catalog archive file.
CA	Catalog permanent files.
CP	Copy archived files to a job as local files.
DU	Dump permanent files.
LD	Load permanent files.
RL	Release disk space.

4. Activate the right K display by entering:

KK.

The display shown in figure 17-9 appears on the right and left screens. A description of the permanent file utility options is shown.

The parameter options available under the chosen utility appear on the left screen. Figure 17-10 shows the left screen, the first of 5 pages, after the DU utility (PFDUMP) has been called. Figures 17-11 through 17-14 show the following pages. In figure 17-10, the TCE appearing after the LIST OPTIONS VALID (in the description column) indicates that

T Files processed
 C Catalog image record (CIR) files
 E Errors

are the listings available for the PFDUMP utility. Different combinations of list options are available with the other utilities.

The CAMTIDPSYZ appearing after the OPTIONS VALID (in the description column) indicates that the following options are available for the dump utility. Different combinations of utility options are available with the other utilities.

Options Description	
C	Creation.
A	Last access.
M	Last modification.
T	Last data modification.
I	Indirect access files.
D	Direct access files.
P	Purge after dump.
S	Suppress staging of files.
Y	Process all files as if backup is required.
Z	Clear alternate storage information in PFC entries.

```

K. PFS.                jsn
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                    MID=AA   NOS version

                    PERMANENT FILE UTILITY OPTIONS

                    ENTER K.UO.

                    UO                DESCRIPTION

                    AT      PFATC (CATALOG TAPE)
                    CA      PFCAT (CATALOG PERMANENT FILES)
                    CP      PFCOPY (COPY TAPE)
                    DU      PFDUMP (DUMP PERMANENT FILES)
                    LD      PFLOAD (LOAD PERMANENT FILES)
                    RL      PFREL (RELEASE DISK SPACE)
    
```

Figure 17-8. Initial Permanent File Utilities K Display (Left Screen)

```

K. PFS.                jsn

                    OP OPTIONS                                LO OPTIONS

A - LAST ACCESS.      C - PERMANENT FILE
C - CREATION.          DEVICE CATALOG.
D - DIRECT ACCESS.    E - ERRORS.
E - EXTRACT CATALOG  S - SUMMARY.
I - INDIRECT ACCESS.  T - FILES PROCESSED.
L - LOAD LEVELING.
M - LAST MODIFICATION.
O - OMIT CATALOG IMAGE RECORD.
P - PURGE AFTER DUMP.
Q - ADD CATALOG AND PERMIT RECORDS.
R - REPLACE.
S - SUPPRESS FILE STAGING.
T - DATA MODIFICATION.
Y - PROCESS ALL FILES AS *BR=Y*.
Z - ZERO OUT ASA LINKAGE.

SPECIFY ONLY ONE OF (C, A, M, OR T).  OPTIONS ARE ENTERED AS -
SPECIFY ONLY ONE OF (I OR D).         LO = CET.    OR
                                        OP = CIP.
    
```

Figure 17-9. Permanent File Utilities K Display (Right Screen)

5. Select the desired parameter options by entering:

K.opt₁=val₁,opt₂=val₂,...,opt_n=val_n.

The opt_i=val_i are selected from the parameter option list on the left screen. The selected parameters replace the default values listed on the left screen (refer to Description of Permanent File Utility Parameters later in this section). A period is required at the end of each entry.

6. Initiate execution by entering:

K.GO.

The lower lines of the left display

N/A MEANS PARAMETER DOES NOT APPLY TO THIS UTILITY.

disappear and are replaced by:

DEVICE	MASK
DATE	TIME

Values appear with these identifiers as processing continues.

```

K. PFS.                jsn
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                        MID=AA   NOS version

PFDUMP                                PAGE 1 OF 5

      OPTION                                DESCRIPTION

FM = FAMILY                FAMILY NAME (1-7 CHARACTERS).
PN = 0                     PACK NAME (1-7 CHARACTERS).
DN = 1                     DEVICE NUMBER (0-77B).
TD = 2                     TRUE DEVICE NUMBER (0-77B).
UN = USER                 USER NAME (1-7 CHARACTERS).
UI = 12                   USER INDEX (0-377777B).
DD = -- N/A --           DESTINATION DEVICE NUMBER (0-77B).
DI = -- N/A --           DESTINATION USER INDEX (0-377777B).
LO = TCE                 LIST OPTIONS VALID - TCE
L = OUTPUT               OUTPUT FILE NAME (1-7 CHARACTERS).
OP = C                   OPTIONS VALID - CAMTIDPSYZ
EO  PROCESS              IF SPECIFIED, DO NOT PROCESS FILES WITH
                        MASS STORAGE ERRORS.

N/A MEANS PARAMETER DOES NOT APPLY TO THIS UTILITY.
  
```

Figure 17-10. PFDUMP K Display (Left Screen) (Page 1 of 5)

```

PFDUMP                                PAGE 2 OF 5

      OPTION                                DESCRIPTION

SD  NO SETTING           IF SPECIFIED, SET DISK SPACE RELEASE DATE.
UD  -- N/A --            IF SPECIFIED, UPDATE UTILITY CONTROL DATE.
DT  NO DESTAGE           IF SPECIFIED, DESTAGE TO TAPE DUMP.
IP  NO INHIBIT           IF SPECIFIED, INHIBIT PROCESSING.
AD = 85/01/14.          AFTER DATE (YYMMDD).
AT = 00.00.00.          AFTER TIME (HHMMSS).
BD = 0                  BEFORE DATE (YYMMDD).
BT = 0                  BEFORE TIME (HHMMSS).
LS = 0                  LOWER FILE SIZE LIMIT.
US = NO LIMIT           UPPER FILE SIZE LIMIT.  * = NO LIMIT.
COS = 0                 CATALOG ONLY SIZE LIMIT.  * = NO LIMIT.

N/A MEANS PARAMETER DOES NOT APPLY TO THIS UTILITY.
  
```

Figure 17-11. PFDUMP K Display (Left Screen) (Page 2 of 5)

PFDUMP		PAGE 3 OF 5
OPTION	DESCRIPTION	
T = TAPE	ARCHIVE FILE NAME (1-7 CHARACTERS).	
SF = 1	NUMBER OF FILES TO SKIP (0-99).	
N = -- N/A --	NUMBER OF FILES TO PROCESS (1-99).	
VF = VERIFY	VERIFY FILE NAME (1-7 CHARACTERS).	
MF = -- N/A --	MASTER FILE NAME (1-7 CHARACTERS).	
RD = RELEASE	RELEASE DATA FILE NAME (1-7 CHARACTERS)	
S = 0	SUMMARY FILE NAME.	
SR = 860901	SUMMARY FILE RECORD NAME.	
LA = LVL0	LOWER ACCESS LEVEL (1-7 CHARACTERS).	
UA = LVL7	UPPER ACCESS LEVEL (1-7 CHARACTERS).	
PVSN TAPE2	VSN OF PREVIOUS ARCHIVE TAPE (1-6 CHAR).	

N/A MEANS PARAMETER DOES NOT APPLY TO THIS UTILITY.

Figure 17-12. PFDUMP K Display (Left Screen) (Page 3 of 5)

PFDUMP		PAGE 4 OF 5		
USER INDEX AND FILE NAME SELECTIONS				
ENTER * TO MOVE FORWARD, / TO MOVE BACKWARD				
-23-	STUFF	FILE9	FILE23	FILE37
FILEA	-500-	FILE10	FILE24	FILE38
FILE2	ALL.	FILE11	FILE25	-12345-
FILE99	-2000-	FILE12	FILE26	ALL.
-24-	FILE1	FILE13	FILE27	-377777-
A	FILE2	FILE14	FILE28	PROFILC
AAA	FILE3	FILE15	FILE29	VALIDUS
ABCD	FILE4	FILE16	FILE30	
Z	FILE5	FILE17	FILE31	
-300-	FILE6	FILE18	FILE32	
ALL.	FILE7	FILE19	FILE33	
-467-	FILE8	FILE20	FILE34	
ALL.	FILESAB	FILE21	FILE35	
-470-	FILE8GH	FILE22	FILE36	

Figure 17-13. PFDUMP K Display (Left Screen) (Page 4 of 5)

PFDUMP	PAGE 5 OF 5	
MSE ALTERNATE STORAGE SELECTION		
RESIDENT		
TAPE ALTERNATE STORAGE SELECTIONS (21 ALLOWED)		
NON-RESIDENT	JQ1234	ZA0000
AA1000	PS6254	ZZ9095
AA6899	TB9005	100000
AB4095	UW2222	541288
AC5000	VX5432	
AC5001	VW7777	
AC5002	XX3131	

Figure 17-14. PFDUMP K Display (Left Screen) (Page 5 of 5)

The entries shown in figure 17-13 of the display are in columns sorted by user index and within each user index the entries are sorted by file name. If all files for a user index are selected, ALL. is displayed.

Figure 17-14 of the display shows alternate storage selections, if present. If the CA=R parameter is entered, RESIDENT is displayed as the MSE alternate storage selection. If the CA=N parameter is entered, NON-RESIDENT is displayed. VSNs entered with the TA=vsn parameter are displayed in columns sorted alphabetically. If the TA=N parameter is also entered, NON-RESIDENT is displayed before any VSNs. An entry of TA=R clears a NON-RESIDENT selection and any VSN selections and RESIDENT is displayed alone. When RESIDENT is selected, entry of VSNs is ignored since all VSNs are already implicitly selected.

DIS Input

The following DIS input sequence is recommended for calling PF utilities from the system console. Refer to section 4, DIS Operations.

1. Call DIS by entering:

X.DIS.

2. If the utility requires an archive file, it must be preassigned before calling the utility. This can be done with a LABEL or ASSIGN command or by using an existing procedure file.

3. Call the utility either by entering

PFS.

which leads to the same sequence of entries and responses previously described under Console Input, or by entering:

PFuo,opt₁=val₁,opt₂=val₂,...,opt_n=val_n.

Parameter	Description
uo	One of the following utility options.
uo	Description
ATC	Catalog archive file.
CAT	Catalog permanent files.
COPY	Copy archived files to a job as local files.
DUMP	Dump permanent files.
LOAD	Load permanent files.
REL	Release disk space.
opt _i =val _i	Desired parameter option and value for the selected utility (refer to Parameters for the Permanent File Utilities later in this section).

If an error in parameters is detected at initiation of a utility, the utility aborts unless the OP=K option is specified. If OP=K is specified, control returns to PFS and you may then enter the correct parameters using the K display.

Batch Input

To call a utility from a batch job, you must be validated for system origin privileges and the system must be in debug mode. The following sequence of commands is used to call a utility.

```
ujn.
USER,username,password,familyname.
PFuo,opt1=val1,opt2=val2,...,optn=valn.
```

uo and opt_i=val_i are the same as for DIS input.

Terminal Input

The format of terminal input for calling the permanent file utilities is substantially the same as that for batch input. You must be validated for system origin privileges and the system must be in debug mode. You enter the batch subsystem (or use the X command), call the desired utility, and enter the appropriate parameters with the command:

```
PFuo,opt1=val1,opt2=val2,...,optn=valn.
```

uo and opt_i=val_i are the same as for DIS input.

Procedure Files

Calling the PF utilities can be made simpler with site-defined procedure files. If these procedures are defined on file PROCFIL under user name SYSTEMX (user index 377777_g), they can be called from DSD. Since user name SYSTEMX is valid only from system origin jobs, the file PROCFIL should be write-permitted to another user name; this will allow easier maintenance of site-defined procedures.

For example, when called in the following manner, the procedure DUMPIT provides a full dump of device 1 of the system default family on archive tapes with VSNs of PFB11, PFB12, ..., PFB16.

```
X.BEGIN(DUMPIT,,DEVICE=1,SET=B)
```

The SETJOB command is used to set the UJN of the job so that the banner page indicates the dump operation performed. In this case it is FULL1B.

```
.PROC,DUMPIT,DEVICE=0,LISTOP=E,SET=A.
SETJOB,UJN=FULL_DEVICE_SET.

VSN,TAPE=PF_SET_DEVICE_1/PF_SET_DEVICE_2/PF_SET_DEVICE_3/
PF_SET_DEVICE_4/PF_SET_DEVICE_5/PF_SET_DEVICE_6.

LABEL,TAPE,D=GE,PO=W,W,FI=$FULL DUMP$.

PFDUMP,DN=DEVICE,LO=LISTOP.
```


Parameters for the Permanent File Utilities

Table 17-1 lists the parameters accepted by each permanent file utility. An X means that the parameter is accepted; a blank means that the parameter is not accepted. Refer to descriptions of parameters following the table for more information.

Table 17-1. Permanent File Utility Parameters

Parameter	PFATC	PFCAT	PFCOPY	PFDUMP	PFLOAD	PFREL
AD	X	X	X	X	X	X
AT	X	X	X	X	X	X
BD	X	X	X	X	X	X
BT	X	X	X	X	X	X
CA	X	X	X	X	X	X
CCA	X	X	X	X	X	X
COS				X		
CPF	X	X	X	X	X	X
CTA	X	X	X	X	X	X
CUI	X	X	X	X	X	X
DD					X	
DI					X	
DN		X		X	X	X
DT				X		
EO				X	X	
FM		X		X	X	X
I	X	X	X	X	X	X
IP				X		X
L	X	X	X	X	X	X
LA	X	X	X	X	X	X
LO	X	X	X	X	X	X

(Continued on next page)

Table 17-1. Permanent File Utility Parameters

Parameter	PFATC	PFCAT	PFCOPY	PFDUMP	PFLOAD	PFREL
<i>(Continued from previous page)</i>						
LS	X	X	X	X	X	X
MF			X			
N	X		X		X	
OP ¹	X	X	X	X	X	X
PF	X	X	X	X	X	X
PN		X		X	X	X
PVSN				X		
RD				X		
S	X	X	X	X	X	X
SD				X		
SF	X		X	X	X	
SR	X	X	X	X	X	X
T	X		X	X	X	
TA	X	X	X	X	X	X
TD		X		X	X	X
UA	X	X	X	X	X	X
UD					X	
UI	X	X	X	X	X	X
UN	X	X	X	X	X	X
US	X	X	X	X	X	X
VF				X		

1. The options valid with the OP parameter depend upon the utility being called, as explained in the OP=opt description.

Description of Permanent File Utility Parameters

Parameter	Description												
AD=date	Specifies the after date (AD) to be used when the OP=A, OP=C, OP=M, or OP=T option is selected. Specify the date using one of the following formats: <table border="1"> <thead> <tr> <th>date</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>YYMMDD</td> <td>Explicit date entered as six decimal digits. YY = year, MM = month, DD = day.</td> </tr> <tr> <td>BDn</td> <td>Date n days before current date. Specify n with one to five digit decimal number.</td> </tr> <tr> <td>Bn</td> <td>Date n days before the current date. Specify n with a one to six digit decimal number.</td> </tr> <tr> <td>O</td> <td>Clears a previous AD entry.</td> </tr> </tbody> </table> <p>When AT=HHMMSS is specified or neither the BD (before date) or BT (before time) parameters are specified, the default after date is the current date.</p>	date	Description	YYMMDD	Explicit date entered as six decimal digits. YY = year, MM = month, DD = day.	BDn	Date n days before current date. Specify n with one to five digit decimal number.	Bn	Date n days before the current date. Specify n with a one to six digit decimal number.	O	Clears a previous AD entry.		
date	Description												
YYMMDD	Explicit date entered as six decimal digits. YY = year, MM = month, DD = day.												
BDn	Date n days before current date. Specify n with one to five digit decimal number.												
Bn	Date n days before the current date. Specify n with a one to six digit decimal number.												
O	Clears a previous AD entry.												
AT=time	Specifies the after time (AT) or both the after date (AD) and after time (AT) to be used when the OP=A, OP=C, OP=M, or OP=T option is selected. If entered as an expression relative to the current date and time, both the after date and after time are determined and any previously specified AD=date entry is overridden. Specify the time and/or date using one of the following formats: <table border="1"> <thead> <tr> <th>time</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>HHMMSS</td> <td>Explicit time entered as six decimal digits. HH = hour, MM = minutes, SS = seconds.</td> </tr> <tr> <td>BDn</td> <td>Time and date n days before the current time and date. Specify n with one to five digit decimal number.</td> </tr> <tr> <td>BHn</td> <td>Time and date n hours before the current time and date.</td> </tr> <tr> <td>Bn</td> <td>Time and date n hours before the current time and date.</td> </tr> <tr> <td>BMn</td> <td>Time and date n minutes before the current time and date.</td> </tr> </tbody> </table>	time	Description	HHMMSS	Explicit time entered as six decimal digits. HH = hour, MM = minutes, SS = seconds.	BDn	Time and date n days before the current time and date. Specify n with one to five digit decimal number.	BHn	Time and date n hours before the current time and date.	Bn	Time and date n hours before the current time and date.	BMn	Time and date n minutes before the current time and date.
time	Description												
HHMMSS	Explicit time entered as six decimal digits. HH = hour, MM = minutes, SS = seconds.												
BDn	Time and date n days before the current time and date. Specify n with one to five digit decimal number.												
BHn	Time and date n hours before the current time and date.												
Bn	Time and date n hours before the current time and date.												
BMn	Time and date n minutes before the current time and date.												

Parameter	Description										
AT=time	(Continued)										
	<table border="1"> <thead> <tr> <th>time</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>BSn</td> <td>Time and date n seconds before the current time and date.</td> </tr> <tr> <td>O</td> <td>Clears a previous AT entry.</td> </tr> </tbody> </table> <p>The default is 000000 (midnight) when AD=date is specified and AT=time is not.</p>	time	Description	BSn	Time and date n seconds before the current time and date.	O	Clears a previous AT entry.				
time	Description										
BSn	Time and date n seconds before the current time and date.										
O	Clears a previous AT entry.										
BD=date	<p>Specifies the before date (BD) to be used when the OP=A, OP=C, OP=M, or OP=T option is selected. Specify the date using one of the following formats:</p> <table border="1"> <thead> <tr> <th>date</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>YYMMDD</td> <td>Explicit date entered as six decimal digits. YY = year, MM = month, DD = day.</td> </tr> <tr> <td>BDn</td> <td>Date n days before current date. Specify n with one to five digit decimal number.</td> </tr> <tr> <td>Bn</td> <td>Date n days before the current date. Specify n with a one to six digit decimal number.</td> </tr> <tr> <td>O</td> <td>Clears a previous BD entry.</td> </tr> </tbody> </table> <p>When BT=HHMMSS is specified, the default before date is the current date.</p>	date	Description	YYMMDD	Explicit date entered as six decimal digits. YY = year, MM = month, DD = day.	BDn	Date n days before current date. Specify n with one to five digit decimal number.	Bn	Date n days before the current date. Specify n with a one to six digit decimal number.	O	Clears a previous BD entry.
date	Description										
YYMMDD	Explicit date entered as six decimal digits. YY = year, MM = month, DD = day.										
BDn	Date n days before current date. Specify n with one to five digit decimal number.										
Bn	Date n days before the current date. Specify n with a one to six digit decimal number.										
O	Clears a previous BD entry.										
BT=time	<p>Specifies the before time (BT) or both the before date (BD) and before time (BT) to be used when the OP=A, OP=C, OP=M, or OP=T option is selected. If entered as an expression relative to the current date and time, both the after date and after time are determined and any previously specified BD=date entry is overridden. Specify the time and/or date using one of the following formats:</p> <table border="1"> <thead> <tr> <th>time</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>HHMMSS</td> <td>Explicit time entered as six decimal digits. HH = hour, MM = minutes, SS = seconds.</td> </tr> <tr> <td>BDn</td> <td>Time and date n days before the current time and date. Specify n with one to five digit decimal number.</td> </tr> <tr> <td>BHn</td> <td>Time and date n hours before the current time and date.</td> </tr> </tbody> </table>	time	Description	HHMMSS	Explicit time entered as six decimal digits. HH = hour, MM = minutes, SS = seconds.	BDn	Time and date n days before the current time and date. Specify n with one to five digit decimal number.	BHn	Time and date n hours before the current time and date.		
time	Description										
HHMMSS	Explicit time entered as six decimal digits. HH = hour, MM = minutes, SS = seconds.										
BDn	Time and date n days before the current time and date. Specify n with one to five digit decimal number.										
BHn	Time and date n hours before the current time and date.										

Parameter	Description										
BT=time	<i>(Continued)</i>										
	<table border="1"> <thead> <tr> <th>time</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Bn</td> <td>Time and date n hours before the current time and date.</td> </tr> <tr> <td>BMn</td> <td>Time and date n minutes before the current time and date.</td> </tr> <tr> <td>BSn</td> <td>Time and date n seconds before the current time and date.</td> </tr> <tr> <td>O</td> <td>Clears a previous BT entry.</td> </tr> </tbody> </table> <p>The default before time is 000000 (midnight) when BD=date is specified and BT=time is not.</p>	time	Description	Bn	Time and date n hours before the current time and date.	BMn	Time and date n minutes before the current time and date.	BSn	Time and date n seconds before the current time and date.	O	Clears a previous BT entry.
time	Description										
Bn	Time and date n hours before the current time and date.										
BMn	Time and date n minutes before the current time and date.										
BSn	Time and date n seconds before the current time and date.										
O	Clears a previous BT entry.										
CA=res	The CA parameter selects files according to residence on cartridge alternate storage.										
	<table border="1"> <thead> <tr> <th>res</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>R</td> <td>Select files that reside on cartridge alternate storage.</td> </tr> <tr> <td>N</td> <td>Select files that do not reside on cartridge alternate storage.</td> </tr> </tbody> </table> <p>When generating a summary file for input to a destage dump file selection utility such as GENPFD, CA=N should be specified if files residing on cartridge alternate storage are to be excluded from consideration.</p>	res	Description	R	Select files that reside on cartridge alternate storage.	N	Select files that do not reside on cartridge alternate storage.				
res	Description										
R	Select files that reside on cartridge alternate storage.										
N	Select files that do not reside on cartridge alternate storage.										

Parameter	Description						
CCA=res	The CCA parameter clears the effect of its corresponding CA parameter. This parameter is intended for K display use.						
	<table border="1"> <thead> <tr> <th>res</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>R</td> <td>Clear the selection of files that reside on cartridge alternate storage.</td> </tr> <tr> <td>N</td> <td>Clear the selection of files that do not reside on cartridge alternate storage.</td> </tr> </tbody> </table>	res	Description	R	Clear the selection of files that reside on cartridge alternate storage.	N	Clear the selection of files that do not reside on cartridge alternate storage.
res	Description						
R	Clear the selection of files that reside on cartridge alternate storage.						
N	Clear the selection of files that do not reside on cartridge alternate storage.						
COS=size	PFC only size threshold in sectors. Specifying COS=* sets an unlimited						
COS=*	PFC only size threshold (forces a data dump of all files). <p>This parameter determines whether PFDUMP will attempt to process a file as PFC only or force a dump of the file data. Processing files as PFC only reduces the time spent in dumping and loading operations and also reduces the size of the archive file. For files larger than or equal to the specified size, only the PFC entry and permits will be dumped if the file's backup requirement is met by a sufficient number of alternate storage copies. A file's backup requirement is met if the backup requirement is media dependent (BR=MD) and one alternate storage copy exists or if the backup requirement is yes (BR=Y) and more than one alternate storage copy exists. For all files smaller than the specified size, and for files larger than or equal to the specified size that are not backed up on alternate storage, the data will be dumped unless file staging is explicitly suppressed (OP=S) and the file is not disk resident. The COS parameter will not cause files with a backup requirement of none (BR=N) to be dumped.</p> <p>Specifying COS=* has the same effect as the OP=Y option for files with BR=MD except that OP=Y will not force a dump of the data of BR=Y files with multiple copies on alternate storage. Specifying COS=* is the only way to force a data dump of all files having multiple copies on alternate storage media.</p> <p>The COS parameter can be used on full and incremental dumps to dump the data of small disk resident files that are already backed up on alternate storage. This should reduce staging activity when permanent files are reloaded after a disk failure.</p>						
CPF=filename	This parameter clears the selection of the specified file. If that file was the only file selected for the user index, all files will be selected for the user index. To clear all file selections for the user index, use the CUI parameter.						

Parameter	Description								
CTA=res	The CTA parameter clears the effect of its corresponding TA parameter. This parameter is intended for K display use.								
	<table border="1"> <thead> <tr> <th>res</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>vsn</td> <td>Clear the selection of files that reside on the VSN specified by vsn.</td> </tr> <tr> <td>R</td> <td>Clear the selection of files that reside on all VSNs.</td> </tr> <tr> <td>N</td> <td>Clear the selection of files that do not reside on tape alternate storage.</td> </tr> </tbody> </table>	res	Description	vsn	Clear the selection of files that reside on the VSN specified by vsn.	R	Clear the selection of files that reside on all VSNs.	N	Clear the selection of files that do not reside on tape alternate storage.
res	Description								
vsn	Clear the selection of files that reside on the VSN specified by vsn.								
R	Clear the selection of files that reside on all VSNs.								
N	Clear the selection of files that do not reside on tape alternate storage.								
CUI=userindex	This parameter clears all file selections for the specified user index. This parameter is assumed octal unless the D radix or a nonoctal digit is specified.								
DD=dn	The 1- or 2-digit number that specifies the alternate device to which PFLOAD loads files when the device on which a file is to be loaded cannot be found, is not defined in the system, or cannot accept the file because of secondary mask restrictions. Default is 0 (no device is the alternate). This parameter is assumed octal unless the D radix or a nonoctal digit is specified.								
DI=userindex	The 1- to 6-digit number that specifies the destination user index to which PFLOAD loads all files being processed. Default is 0 (that is, files are loaded to the user index from which they were dumped). This parameter is assumed octal unless the D radix or a nonoctal digit is specified.								
DN=dn	The 1- or 2-digit octal number that specifies the device within the family to be cataloged, dumped, or loaded. For PFLOAD, refer to PFLOAD Selection Process under Load Permanent File (PFLOAD) later in this section for more information. For PFCAT or PFDUMP, the following distinctions are made. If the specified device is a master device, all files cataloged on it (whether or not they reside on the device) are processed if they also meet all other specified selection criteria. However, files that reside on the device that are cataloged elsewhere are not processed. If the specified device is not a master device, all files residing on it are processed if they also meet all other specified selection criteria. (For nonmaster devices the DN and TD parameters have the same effect.) Default is 0 for PFDUMP and PFLOAD (all devices are to be dumped or loaded). DN must be specified for PFCAT if an output file is selected (L and LO parameters). However, PFCAT does not require entry of the DN parameter if a summary file (S parameter) rather than an output file (L and LO parameters) is selected. In this case, all devices in the family will be processed.								

Parameter	Description
DT	Destage the selected files to tape or optical disk alternate storage, depending on whether the file specified by the T parameter is assigned to a magnetic tape or to an optical disk file. This parameter cannot be specified on a dump of an auxiliary pack. Refer to section 12, Tape and Optical Disk Alternate Storage, for examples.
EO	Specifies that PFDUMP or PFLOAD does not process files with mass storage errors. If EO is not specified, PFDUMP dumps files, and PFLOAD loads files, regardless of mass storage errors.
FM=familyname	The 1- to 7-character name of the family of permanent file devices to be cataloged, dumped, or loaded. This parameter option is not required if only one family of devices is active in the system. Default is the default system family name.
I=filename	The 1- to 7-character name of the file from which input directives are to be read. Default is INPUT. Any valid permanent file utility parameter or K display option can be specified in the directives file, except for the I parameter itself. The directives file is typically used to specify multiple user indices and file names for processing.
IP	Inhibit processing. This parameter specifies that PFDUMP or PFREL will generate an output file and/or summary file (as determined by the L, LO, and S parameters) listing the files that would have been processed by that utility without actually dumping files or releasing disk space. A selection utility (such as GENPFD) can use the summary file to generate a directives file specifying file names to be processed by PFDUMP or PFREL without processing inhibited. Default is that processing is not inhibited.
L=filename	The 1- to 7-character name of the file on which reports are to be written. Default is OUTPUT.
LA=level	The 1- to 7-character name that specifies the lower limit of the range of access levels to process. If this parameter is specified, the UA parameter must also be specified. If neither LA nor UA is specified, the default is that all access levels are selected. Refer to PFDUMP in a Secured System later in this section.
LO=opt	Character(s) specifying the type of information that the permanent file utility should include in its output report. Default is 0 (no options selected).
opt	Description
C	List all files in the catalog image record. This option is used only with PFATC, PFDUMP, and PFLOAD.
E	List errors.
S	List cumulative statistics. This option is used only with PFCAT.
T	List PFC data for all files processed.

Parameter	Description
LS=size	Lower size limit in sectors. When specified, only files of size greater than or equal to the lower size limit will be processed. For PFDUMP, the file size limit is checked before the COS parameter (if COS is specified).
MF=filename	The 1- to 7-character name of the master file to which PFCOPY copies all the files extracted from a designated archive file. Default is no name; that is, the archived files are copied as individual files retaining their permanent file names. This option allows you to extract a file from an archive file and change its name as a local file.
N=n	The 1- or 2-digit number that specifies the number of archive files on a multifile archive file to be processed. If n=0, one file is processed. Default is 1. This parameter is assumed decimal unless the B radix is specified.
OP=opt	The 1- to 7-character string specifying the utility options that control the processing of files. Many of the options require additional parameters to complete the definition of the OP selection. Default is 0 (no options selected).

opt	Description
------------	--------------------

K	Specifies that the K display should be brought up if an error is detected when processing input parameters and directives in PFS. If OP=K is not specified, the permanent file utility will abort if such an error is detected (which allows exit processing to be programmed within the calling procedure).
---	--

Only one of options A, C, M, or T can be used at a time. Each of these options requires one of the following to be specified.

- Specify BD/BT to select all files having a last access date/time (if OP=A), a creation date/time (if OP=C), a utility control date/time (if OP=M), or a data modification date/time (if OP=T) prior to this date/time.
- Specify AD/AT to select all files having a last access date/time (if OP=A), a creation date/time (if OP=C), a utility control date/time (if OP=M), or a data modification date/time (if OP=T) more recent than this date/time.
- Specify both BD/BT and AD/AT to select all files having a last access date/time (if OP=A), a creation date/time (if OP=C), a utility control date/time (if OP=M), or a data modification date/time (if OP=T) in the interval defined by these dates/times. The date/time specified by BD/BT must be more recent than the date/time specified by AD/AT to properly define an interval.

Parameter	Description																
OP=opt	<p><i>(Continued)</i></p> <p>See the descriptions of the AD, AT, BD, and BT parameters for more information.</p> <table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Make selection according to time of last access.</td> </tr> <tr> <td>C</td> <td>Make selection according to time of creation.</td> </tr> <tr> <td>M</td> <td>Make selection according to time of last modification.</td> </tr> <tr> <td>T</td> <td>Make selection according to time of last data modification.</td> </tr> </tbody> </table> <p>The following rules apply when OP=M is used. If the utility control date and time field meets the specified date/time criteria, the file is selected. If the date/time criteria are not met, the file can still be selected if the control modification date and time field meets the date/time criteria and is more recent than the utility control date and time field. Refer to Permanent File Catalog Entry earlier in this section for descriptions of the PFC fields.</p> <p>The T option is intended to allow recently modified files to be excluded when creating a summary file for input to a destage file selection utility such as GENPFD.</p> <p>Only one of the following two options can be used at a time. They can be used in conjunction with the A, C, M, or T option.</p> <table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Select direct access files only.</td> </tr> <tr> <td>I</td> <td>Select indirect access files only.</td> </tr> </tbody> </table>	opt	Description	A	Make selection according to time of last access.	C	Make selection according to time of creation.	M	Make selection according to time of last modification.	T	Make selection according to time of last data modification.	opt	Description	D	Select direct access files only.	I	Select indirect access files only.
opt	Description																
A	Make selection according to time of last access.																
C	Make selection according to time of creation.																
M	Make selection according to time of last modification.																
T	Make selection according to time of last data modification.																
opt	Description																
D	Select direct access files only.																
I	Select indirect access files only.																

Parameter	Description
OP=opt	(Continued) The following options are used only with PFDUMP. The P and S options cannot be specified at the same time.
opt	Description
P ⁷	Purge after dump. PFDUMP keeps a list of all the files it actually dumps on a scratch file. After the dump is complete and successful, PFDUMP then does a PURGE on each file that it actually dumped. If PFDUMP aborts or the operator drops PFDUMP before all of the files are dumped, none of the files are purged. Since the catalog track interlock is released between the dump and the purge, it is possible that a user could update a file during this interval; if this happens, the file is still purged. If this option is specified, backup requirement parameters are ignored when selecting files to process. Files with no backup requirement (BR=N) are processed, and the data for files residing on alternate storage media are dumped, regardless of file backup requirements. Files whose user indexes are greater than or equal to PGUI (default value is 300000 ₈) are not dumped or purged unless specific user indices are selected with the UI parameter. This is done to prevent files such as critical recovery files, applications, and system utilities from being purged from the system due to lack of use. If the UI parameter is used to select files cataloged on user indices greater than or equal to the PGUI threshold, the files are dumped and purged.
S ⁸	Suppress staging of files. If a file to be dumped resides on MSE, optical disk, or tape alternate storage and not on disk, it is not staged to disk and only its PFC and permit entries are copied to the archive file. If OP=S is omitted, the alternate storage file is copied to disk and included in the dump. This option should normally be specified for full dumps.
Y ⁷	This option causes all files to be processed the same as files with backup requirement BR=Y.
Z	Clear the cartridge and tape alternate storage information in the PFC entries written to the dump file.

7. Refer to Permanent File Catalog Entry earlier in this section for a description of the backup requirement field.

8. Refer to section 11, MSE, or to section 12, Tape and Optical Disk Alternate Storage, for more information.

Parameter	Description														
OP=opt	(Continued) The following six options are used only with PFLOAD.														
	<table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>R</td> <td>Load with the replace option. If R is specified, PFLOAD loads all selected files from the archive file even if files with the same name already exist in the permanent file system. If R is not specified, PFLOAD loads only those files selected from the archive file for which no corresponding files (files with identical file names and user indexes) already exist in the permanent file system.</td> </tr> <tr> <td>E</td> <td>Extract only the CIR. Refer to the PFLOAD description for information about the CIR. PFLOAD reads the CIR from the archive file, generates a random file and directory, and requests the next archive file without processing any of the files after the CIR on the first archive file.</td> </tr> <tr> <td>L</td> <td>Load each direct access file on the device with the most available space of those devices where the file is eligible to reside. This option overrides the DD parameter if it is also specified.</td> </tr> <tr> <td>O</td> <td>Do not read the CIR. PFLOAD does not read the CIR for the specified archive file but processes the records in this archive file. PFLOAD then terminates normally without requesting another archive file.</td> </tr> <tr> <td>V</td> <td>Request the next archive file by using the volume serial number (VSN) found in the PVSN field of the archive file. PFLOAD will automatically request a tape with this VSN when finished with the current archive file. Refer to the PVSN parameter description for more information.</td> </tr> <tr> <td>Z⁹</td> <td>Clear the alternate storage information in the PFC entry when the PFC entry is loaded, if the information points to an MSE, optical disk, or tape alternate storage image. Also suppresses the loading of PFC only files. Default is to leave the alternate storage information intact and to load PFC only files. OP=Z should normally be omitted if the device has to be recovered and should be specified if data for the file is being reloaded.</td> </tr> </tbody> </table>	opt	Description	R	Load with the replace option. If R is specified, PFLOAD loads all selected files from the archive file even if files with the same name already exist in the permanent file system. If R is not specified, PFLOAD loads only those files selected from the archive file for which no corresponding files (files with identical file names and user indexes) already exist in the permanent file system.	E	Extract only the CIR. Refer to the PFLOAD description for information about the CIR. PFLOAD reads the CIR from the archive file, generates a random file and directory, and requests the next archive file without processing any of the files after the CIR on the first archive file.	L	Load each direct access file on the device with the most available space of those devices where the file is eligible to reside. This option overrides the DD parameter if it is also specified.	O	Do not read the CIR. PFLOAD does not read the CIR for the specified archive file but processes the records in this archive file. PFLOAD then terminates normally without requesting another archive file.	V	Request the next archive file by using the volume serial number (VSN) found in the PVSN field of the archive file. PFLOAD will automatically request a tape with this VSN when finished with the current archive file. Refer to the PVSN parameter description for more information.	Z ⁹	Clear the alternate storage information in the PFC entry when the PFC entry is loaded, if the information points to an MSE, optical disk, or tape alternate storage image. Also suppresses the loading of PFC only files. Default is to leave the alternate storage information intact and to load PFC only files. OP=Z should normally be omitted if the device has to be recovered and should be specified if data for the file is being reloaded.
opt	Description														
R	Load with the replace option. If R is specified, PFLOAD loads all selected files from the archive file even if files with the same name already exist in the permanent file system. If R is not specified, PFLOAD loads only those files selected from the archive file for which no corresponding files (files with identical file names and user indexes) already exist in the permanent file system.														
E	Extract only the CIR. Refer to the PFLOAD description for information about the CIR. PFLOAD reads the CIR from the archive file, generates a random file and directory, and requests the next archive file without processing any of the files after the CIR on the first archive file.														
L	Load each direct access file on the device with the most available space of those devices where the file is eligible to reside. This option overrides the DD parameter if it is also specified.														
O	Do not read the CIR. PFLOAD does not read the CIR for the specified archive file but processes the records in this archive file. PFLOAD then terminates normally without requesting another archive file.														
V	Request the next archive file by using the volume serial number (VSN) found in the PVSN field of the archive file. PFLOAD will automatically request a tape with this VSN when finished with the current archive file. Refer to the PVSN parameter description for more information.														
Z ⁹	Clear the alternate storage information in the PFC entry when the PFC entry is loaded, if the information points to an MSE, optical disk, or tape alternate storage image. Also suppresses the loading of PFC only files. Default is to leave the alternate storage information intact and to load PFC only files. OP=Z should normally be omitted if the device has to be recovered and should be specified if data for the file is being reloaded.														

⁹ Refer to section 11, MSE, or to section 12, Tape and Optical Disk Alternate Storage, for more information.

Parameter	Description				
OP=opt	<p>(Continued)</p> <p>The following option is used only with PFCOPY.</p> <table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Q</td> <td>Select leading records. The archived file that is copied to the job as a local file includes two header records, one with the catalog entry for the file and the second with the permit information for the file.</td> </tr> </tbody> </table>	opt	Description	Q	Select leading records. The archived file that is copied to the job as a local file includes two header records, one with the catalog entry for the file and the second with the permit information for the file.
opt	Description				
Q	Select leading records. The archived file that is copied to the job as a local file includes two header records, one with the catalog entry for the file and the second with the permit information for the file.				
PF=filename	<p>The 1- to 7-character name of a permanent file to be processed. The selected file names appear on page 4 of the left K-display screen (refer to figure 17-13). To delete a permanent file name that has been entered, use the CPF parameter. When used from one of the utility commands, up to 10 file names can be entered on one line. The following parameter format is used for both the K-display and the utility commands.</p> <p style="text-align: center;">PF=filename1,PF=filename2,...,PF=filenamen</p> <p>Use of the PF parameter without a preceding UI or UN parameter is an error. Default is that permanent file names are not selection criteria.</p>				
PN=packname	The 1- to 7-character name of the auxiliary device to be cataloged, dumped, or loaded. The device must be mounted and available. Default is that pack name is not a selection criterion.				
PVSN=vsrn	The 1- to 6-character tape VSN to write into the PVSN field of the archive file. This value, stored in the PFDUMP label of the archive file, is usually the first tape VSN of your previous archive file. You can use this parameter to automate reloading of a device. Typically, you use PFLOAD to process the most recent set of archive file tapes, mount the next most recent set, and so on, completing the reload with your most recent set of full dump archive file tapes. Requesting of these sets may be automated by having each incremental set point to its previous set; that is, equating the PVSN value to the tape VSN on the first reel of the previous set. If you subsequently perform a PFLOAD, specifying OP=V will cause PFLOAD, upon completion of one set of archive file tapes, to automatically request the next set, using the PVSN value read from the PFDUMP label.				
RD=filename ¹⁰	The 1- to 7-character name of the release data file (RDF) to be created by PFDUMP, which identifies those MSE-resident files that are pointed to by PFC entries at the time of the dump. The MSE utility SSVAL uses this file to identify unneeded MSE-resident files whose space can be released. If RD alone is specified, the release data file created is named ZZZZRDF. Default is that PFDUMP does not create an RDF.				

10. Refer to section 11, MSE, for more information.

Parameter	Description								
S=filename	The 1- to 7-character name of a summary file generated by the permanent file utilities. Default is SUMMARY. The summary file contains the PFC entries for those files processed by the utility. Refer to Summary File earlier in this section for additional information and the format of the summary file. For the PFCAT utility, a summary file and an output file (L and LO parameters) cannot be specified simultaneously.								
SD ¹¹	Specifies that PFDUMP is to enter the date and time when the dump was initiated into the inhibit date/time field of the master device. This date/time entry is used by PFDUMP, PFRCL, and by the MSE utility SSMOVE to determine whether the disk image of a file can be released. If a file was last modified before this date/time, this PFDUMP is considered a backup copy for the purpose of satisfying the user-specified backup requirement (BR) for that file. Refer to PFDUMP in an MSE, Tape, or Optical Disk Alternate Storage Environment later in this section for more information.								
SF=n	The 1- or 2-digit number that specifies the number of archive files on a multifile archive file to be skipped before processing begins. Default is 0. This parameter is assumed decimal unless the B radix is specified.								
SR=recordname	The 1- to 7-character name of the summary file record generated by the permanent file utilities. Default is the current date in the form yymmdd. The specified record name will be written to the prefix table of the record generated. Refer to Summary File earlier in this section for the format of the prefix table.								
T=filename	The 1- to 7-character name of the file on which to store or read archive files. Usually, filename is a tape, but it can be a mass storage device. Default is TAPE.								
TA=res	The TA parameter selects files according to residence on the specified tape alternate storage VSN.								
	<table border="1"> <thead> <tr> <th>res</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>vsn</td> <td>Select tape alternate storage resident files that reside on the VSN specified by vsn.</td> </tr> <tr> <td>R</td> <td>Select tape alternate storage resident files that reside on all VSNs.</td> </tr> <tr> <td>N</td> <td>Select files that do not reside on tape alternate storage.</td> </tr> </tbody> </table> <p>When selecting files to be destaged to tape alternate storage, TA=N should be specified to prevent selection of files already destaged to other tapes.</p>	res	Description	vsn	Select tape alternate storage resident files that reside on the VSN specified by vsn.	R	Select tape alternate storage resident files that reside on all VSNs.	N	Select files that do not reside on tape alternate storage.
res	Description								
vsn	Select tape alternate storage resident files that reside on the VSN specified by vsn.								
R	Select tape alternate storage resident files that reside on all VSNs.								
N	Select files that do not reside on tape alternate storage.								

¹¹ Refer to section 11, MSE, for more information.

Parameter	Description
TD=dn	<p>The 1- or 2-digit octal number that specifies the device within the family to be dumped or loaded. This is the device number specified by the EQPDECK PF entry for the device. You can determine which device number to use within the family by referring to the DSD Family Status Display (E,F) described in the NOS Version 2 Operations Handbook.</p> <p>For PFLOAD refer to PFLOAD Selection Process under Load Permanent File (PFLOAD) later in this section for more information. For PFDUMP, the following distinctions are made. If the specified device is a master device, all files cataloged and/or residing on it are processed if they meet all other specified selection criteria. If the specified device is not a master device, all files residing on it are processed if they also meet all other specified selection criteria. (For nonmaster devices the DN and TD parameters have the same effect.) Default is 0 (all devices are to be dumped or loaded).</p>
UA=level	<p>The 1- to 7-character name that specifies the upper limit of the range of access levels to process. If this parameter is specified, the LA parameter must also be specified. If neither LA nor UA is specified, the default is that all access levels are selected. Refer to PFDUMP in a Secured System later in this section.</p>
UD ¹²	<p>Specifies that PFLOAD is to update the utility control date and time field in the PFC entry for the file being loaded. This ensures that the file will be considered for inclusion in the next incremental dump. UD should normally be specified when a particular user's files are being reloaded, but should normally be omitted when a full device is being reloaded.</p>
UI=userindex	<p>The 1- to 6-digit number that specifies the user index under which files to be processed are located. If UI is specified, DN need not be specified because the utility will locate the proper device. Default is that user index is not a selection criterion. This parameter is assumed octal unless the D radix or a nonoctal digit is specified.</p> <p>If no file selections are in effect for the affected user index at the time of entry, the UI parameter selects all files for the user index and also specifies the user index to be assumed for subsequent PF directives. If file selections are already in effect for the user index, the UI parameter simply sets the assumed user index for subsequent PF directives.</p>
UN=username	<p>The 1- to 7-character user name associated with the PN parameter. The user name must match information in the MST for the specified pack name. If UN is specified and PN is not, the utility will convert username to a user index. If, in addition, DN is not specified, the utility will locate the proper device in the family. Default is that user name is not a selection criterion.</p>

¹². Refer to section 11, MSE, for more information.

Parameter	Description
UN=username	<i>(Continued)</i> If no file selections are in effect for the affected user index at the time of entry, the UN parameter selects all files for the user index and also specifies the user index to be assumed for subsequent PF directives. If file selections are in effect for the user index, the UN parameter simply sets the assumed user index for subsequent PF directives.
US=size	Upper size limit in sectors. When specified, only files of size less than or equal to the upper size limit will be processed. For PFDUMP, the file size limit is checked before the COS parameter (if COS is specified).
VF=filename	The 1- to 7-character name that specifies that PFDUMP is to produce a verify file which is a duplicate of the archive file it creates. If VF alone is specified, the verify file is named PFVER. Default is that no verify file is written.

Example:

The following series of commands creates a two-file archive file (AA) and a matching verify file (BB). The VERIFY command compares the accuracy of the duplication. (Refer to the NOS Version 2 Reference Set, Volume 3 for a complete description of VERIFY.)

Command	Description
PFDUMP,T=AA,VF=BB.	The dump is written on archive file AA. A duplicate is written on the verify file BB. PFDUMP does not rewind after processing.
PFDUMP,T=AA,VF=BB.	A second dump is written after the first on both the archive (AA) and the verify (BB) files. Each dump produces a separate file on the archive and verify files.
VERIFY,AA,BB,N=0,A,R.	A binary comparison of AA and BB is performed. If words do not match, this command lists: <ul style="list-style-type: none"> Record number Word number within the record Words from both files that do not match N=0 specifies that the verify terminates when the first empty file is encountered. The A parameter specifies an abort if a mismatch is found. R rewinds both files before and after the verify.

Preassigning the Archive File

In all permanent file utilities except PFCAT and PFREL, the archive file must be preassigned. This can be done by using the LABEL or ASSIGN command. For example, file assignment can be made with the following commands.

```
LABEL, TAPE, VSN=PFDUMP, NT, D=1600.
```

or

```
ASSIGN, NT, TAPE.
```

If no archive file is present, PFDUMP will write on a local mass storage file. The other utilities will attempt to read an empty archive file and terminate. PFLOAD will automatically request the assignment of any incremental or full dump archive tapes needed after the first tape; the request will be for the same track type and density as the first archive tape.

Further automation may be added by indicating all tape VSNs during PFDUMP and/or PFLOAD, which automates the archive tape VSN requests within one archive file. This may be done through the VSN command (refer to Procedure Files earlier in this section). You can also use the PFDUMP parameter PVSAN and PFLOAD option OP=V, which automates the archive tape VSN requests between archive files. Refer to the PVSAN parameter description for more information.

NOTE

The archive file is not rewound or unloaded before or after processing by any of the utilities, except before a PFLOAD tape request.

Permanent File Utility Routines

The parameters available for each of the permanent file utility routines, discussed next, are listed in table 17-1.

The LO parameter allows each permanent file utility to produce a cataloged directory of file information. This catalog information is either derived from the catalog tracks on the permanent file device or from the archive tape. Although the header information may differ from utility to utility, the format of the catalog information remains the same. Following is the general format of the information listed for each file on the directory.

```

filename  type   cat           length  dn   cdate  ladate  dmdate  cmdate  udate
password  count  userindex mode  subsystem  time   time   time   time   time
exp date  level  pr br ac rs                cflags          cat asa  amsg
charge    project                tflags          tape vsn seqn
    
```

Figure 17-15 illustrates a typical page from a cataloged directory. Although this directory was produced using the command PFATC,LO=T., directories in similar format could be produced using the other utilities. Each of the following fields is shown and cross-referenced on the figure.

Header Information ¹³		Description
1	filename	Permanent file name.
2	type	Type of permanent file. This field can be either direct access (DIR) or indirect access (IND).
3	cat	File category. This field can be either PRIVATE, SPRIV, or PUBLIC. These categories are described in the NOS Version 2 Reference Set, Volume 3.
4	length	Length of the file in decimal PRUs.
5	dn	For direct access files, the device number of the mass storage device on which the file resides if other than the master device. If the file resides on the master device, this field is replaced by an *.
6	password	Password associated with the file.
7	count	Count that specifies the number of times the file has been accessed.
8	userindex	User index of the user on whose catalog this file resides.
9	mode	Permission mode. This field can be WRITE, MODIFY, UPDATE, APPEND, READ, READMD, READUP, READAP, EXECUTE, or NULL. These modes are explained in detail in the NOS Version 2 Reference Set, Volume 3.
10	subsystem	Subsystem under which the file was saved. Possible entries include FORT., FTNTS, BASIC, EXEC., or BATCH. If this field contains no entry, a subsystem is not associated with the file.

13. Numbers in figure 17-15 are circled.

Header Information ¹⁴	Description												
11	cdate time Date and time of file creation. The format is: yy/mm/dd. hh.mm.ss.												
12	ladate time Date and time of the last access to the file.												
13	dmdate time Date and time of the last data modification to the file.												
14	cmdate time Date and time of the last control information modification in the file's PFC entry or permit entries.												
15	ucdate time Utility control date and time set by PFM or a PF utility to ensure that the file will be included in the next incremental (OP=M) dump.												
16	exp date Expiration date associated with the file password.												
17	level Security access level assigned to the file.												
18	pr Preferred residence for the file. This field can be: <table border="1" data-bbox="568 787 1334 1071"> <thead> <tr> <th>Field</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>L</td> <td>Locked to disk.</td> </tr> <tr> <td>D</td> <td>Disk residence preferred.</td> </tr> <tr> <td>M</td> <td>Cartridge alternate storage (MSE or optical disk) residence preferred.</td> </tr> <tr> <td>N</td> <td>No preference.</td> </tr> <tr> <td>T</td> <td>Tape alternate storage residence preferred.</td> </tr> </tbody> </table>	Field	Description	L	Locked to disk.	D	Disk residence preferred.	M	Cartridge alternate storage (MSE or optical disk) residence preferred.	N	No preference.	T	Tape alternate storage residence preferred.
Field	Description												
L	Locked to disk.												
D	Disk residence preferred.												
M	Cartridge alternate storage (MSE or optical disk) residence preferred.												
N	No preference.												
T	Tape alternate storage residence preferred.												
19	br Backup requirement for the file. This field can be: <table border="1" data-bbox="568 1155 1334 1312"> <thead> <tr> <th>Field</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Y</td> <td>Two backup copies required.</td> </tr> <tr> <td>MD</td> <td>One backup copy required.</td> </tr> <tr> <td>N</td> <td>No backup copy required.</td> </tr> </tbody> </table> <p>Refer to PFDUMP in an MSE, Tape, or Optical Disk Alternate Storage Environment later in this section for more information.</p>	Field	Description	Y	Two backup copies required.	MD	One backup copy required.	N	No backup copy required.				
Field	Description												
Y	Two backup copies required.												
MD	One backup copy required.												
N	No backup copy required.												

14. Numbers in figure 17-15 are circled.

Header Information¹⁵		Description
20	ac	Alternate user CATLIST permission for the file specified by the file owner. This field can be Y (alternate user can obtain CATLIST information about the file) or N (alternate user cannot obtain CATLIST information about the file).
21	rs	Current residence for the file. Possible values are D, DC, DT, DCT, CT, C, and T. If a D is present, the file currently has a copy on disk. If a C is present, the file currently has a copy on cartridge alternate storage (MSE or optical disk). If a T is present, the file currently has a copy on tape alternate storage.
22	cflags	Cartridge alternate storage flags set for the file (refer to Permanent File Catalog Entry earlier in this section for details). If a flag is not set, a dash is listed. If an undefined flag is set, an asterisk is displayed. This field can be:
		Field Description
		D Data error flag.
		F Free up cartridge flag.
		L Locked to disk flag.
		O Obsolete flag.
		R Pseudorelease flag.
		S System control error flag.
		T Temporary error flag.
		V Verification flag.
23	cat	Cartridge alternate storage type. This field can be NONE (no cartridge alternate storage), MSE (file resides on MSE), or OD (file resides on optical disk).
24	asa	Cartridge alternate storage address of the file.
25	amsg	Appended message, *PFC ONLY, indicating when a PFC ONLY file (that is, an entry with PFC and permit entries, but no data) has been processed on the archive file.
26	charge	Charge number associated with the file.
27	project	Project number associated with the file.

¹⁵. Numbers in figure 17-15 are circled.

Header Information ¹⁶	Description																						
28	tflags Tape alternate storage flags set for the file (refer to Permanent File Catalog Entry earlier in this section for details). If a flag is not set, a dash is listed. If an undefined flag is set, an asterisk is displayed. This field can be: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th style="text-align: left;">Field</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>File on ACS tape.</td> </tr> <tr> <td>C</td> <td>File on either ACS or CTS tape.</td> </tr> <tr> <td>D</td> <td>Data error on primary VSN flag.</td> </tr> <tr> <td>E</td> <td>Data error on secondary VSN flag.</td> </tr> <tr> <td>F</td> <td>Tape recorded in LI format.</td> </tr> <tr> <td>L</td> <td>Destage in progress flag.</td> </tr> <tr> <td>N</td> <td>Secondary VSN not available flag.</td> </tr> <tr> <td>P</td> <td>Primary VSN not available flag.</td> </tr> <tr> <td>S</td> <td>Secondary VSN exists flag.</td> </tr> <tr> <td>V</td> <td>Verification flag.</td> </tr> </tbody> </table>	Field	Description	A	File on ACS tape.	C	File on either ACS or CTS tape.	D	Data error on primary VSN flag.	E	Data error on secondary VSN flag.	F	Tape recorded in LI format.	L	Destage in progress flag.	N	Secondary VSN not available flag.	P	Primary VSN not available flag.	S	Secondary VSN exists flag.	V	Verification flag.
Field	Description																						
A	File on ACS tape.																						
C	File on either ACS or CTS tape.																						
D	Data error on primary VSN flag.																						
E	Data error on secondary VSN flag.																						
F	Tape recorded in LI format.																						
L	Destage in progress flag.																						
N	Secondary VSN not available flag.																						
P	Primary VSN not available flag.																						
S	Secondary VSN exists flag.																						
V	Verification flag.																						
29	tape vsn VSN of the tape containing the tape alternate storage copy of the file.																						
30	seqn Sequence number on the tape containing the tape alternate storage copy of the file.																						

¹⁶. Numbers in figure 17-15 are circled.

PFATC. FILES ON ARCHIVE FILE. FAMILY = SYS964.										
ARCHIVE FILE CREATED yy/mm/dd. hh.mm.ss.					ARCHIVE FAMILY = TEST.					
FILE NAME	FILE TYPE	LENGTH	DN	CREATION	ACCESS	DATA MOD	CONTROL	UTILITY		
PASSWD	COUNT	INDEX	PERM.	SUBSYS	DATE/TIME	DATE/TIME	DATE/TIME	DATE/TIME	DATE/TIME	
EXPIRES	LEVEL	PR	BR	AC	RS	CARTRIDGE	FLAGS	TYPE	CARTRIDGE	
CHARGE NO.	PROJECT	NUMBER			TAPE FLAGS		TAPE VSN	SER.	NUMBER	
1	ARCPPD	DIR	PUBLIC	269	11	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ		15.22.31.	15.22.31.	15.22.31.	15.22.31.	15.22.31.
		LVLO	N Y	N DT				NONE		
								AA1234		0
2	SITRPT	IND.	PUBLIC		12	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
	PASSM					15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		1
3	LISTEXB	IND.	PUBLIC		8	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ	BATCH	15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		2
4	FSEPZ19	IND.	PRIVATE		3	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ		15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		3
5	DECK3	IND.	PUBLIC		2	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ		15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		4
6	SHORT	IND.	PUBLIC		11	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ		15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		5
7	PFHELPR	IND.	PUBLIC		6	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ		15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		6
8	PFACC	IND.	PUBLIC		36	87/04/13.	87/04/13.	87/04/13.	87/04/13.	87/04/13.
		0	3320	READ		15.22.29.	15.22.29.	15.22.29.	15.22.29.	15.22.29.
		LVLO	N Y	N DT				NONE		
								AA1234		7

Figure 17-15. Sample Directory Produced by PFATC,LO=T.

Catalog Archive File (PFATC)

PFATC produces a cataloged directory of file information derived from an archive file previously created by the PFDUMP utility. The format of the directory depends upon the parameters selected.

If LO=T (list all files processed) is specified, the directory produced is similar to that shown in figure 17-15. If LO=C (list all files in catalog image record) is specified and the archive file to be cataloged was produced by an incremental dump (refer to Incremental Permanent File System Dump later in this section), the directory produced is similar to that shown in figure 17-20.

Catalog Permanent File Device (PFCAT)

PFCAT produces a cataloged directory of file information derived from catalog tracks on a master device. The format of the directory depends upon the parameter options selected.

If LO=T (list all files processed) is specified, a directory of file information and a mass storage table report are produced. The directory is similar to that shown in figure 17-16. However, the files are listed according to user index, and totals are given after the files for each user index. The mass storage table report gives information about each mass storage device in the system. Figure 17-17 is an example of a mass storage table report. Messages issued with this report give information concerning the type and status of the device cataloged. The following status messages can be issued.

```
ACCOUNT INITIALIZE PENDING.
CATALOG TRACK OVERFLOW.
DAYFILE INITIALIZE PENDING.
ERRLOG INITIALIZE PENDING.
FORMAT PENDING.
I/O QUEUE INITIALIZE PENDING.
MAINLOG INITIALIZE PENDING.
PF INITIALIZE PENDING.
TOTAL INITIALIZE PENDING.
UNAVAILABLE FOR PF ACCESS.
8 WORD CATALOG ENTRIES.
```

The following types of devices can be listed.

```
ALTERNATE SYSTEM
REMOVABLE DEVICE
SYSTEM
```

If LO=S (list cumulative statistics for catalog) is specified, summary reports of file information for each user index and for the entire device are produced. Figure 17-18 lists a page of the information given for each user index. Figure 17-19 lists the general device information.

If TA=vsn (select tape alternate storage resident files that reside on vsn) is specified, PFCAT lists all current files that reside on a particular tape VSN. This can be used to determine if a specific tape may be recycled. Refer to section 12, Tape and Optical Disk Alternate Storage, for more information about recycling dump tapes.

CATALOG FILE										yy/mm/dd. hh.mm.ss.	
DIRECTORY OF PERMANENT FILE DEVICE 40 SYS964										PAGE 1	
CATALOG OF USER INDEX 1											
FILE NAME	FILE TYPE	LENGTH	DN	CREATION	ACCESS	DATA	MOD	CONTROL	UTILITY		
PASSWORD	COUNT	INDEX	PERM.	SUBSYS	DATE/TIME	DATE/TIME	DATE/TIME	DATE/TIME	DATE/TIME		
EXPIRES	LEVEL	PR	BR	AC	RS	CARTRIDGE FLAGS		TYPE	CARTRIDGE	ASA	
CHARGE NO.	PROJECT	NUMBER			TAPE FLAGS		TAPE	VSN	SEQ.	NUMBER	
1 TSTESTA	IND. PRIVATE			4	82/09/19.	86/04/17.	82/09/19.	87/05/11.	82/09/19.		
	150			WRITE	19.35.51.	10.45.07.	19.35.51.	13.43.30.	19.35.51.		
	LVLO	N	Y	N	D				NONE		
2 TESTB	IND. PRIVATE			5	82/09/19.	86/04/17.	82/09/19.	87/05/11.	82/09/19.		
	150			WRITE	19.35.51.	10.45.07.	19.35.51.	13.43.30.	19.35.51.		
	LVLO	N	Y	N	D				NONE		
3 TESTA	IND. PRIVATE			3	82/09/19.	86/04/17.	82/09/19.	87/05/11.	82/09/19.		
	150			WRITE	19.35.51.	10.45.07.	19.35.51.	13.43.30.	19.35.51.		
	LVLO	N	Y	N	D				NONE		
4 PROCB	IND. PRIVATE			1	75/09/16.	86/04/17.	75/09/16.	87/05/11.	75/09/16.		
	193			WRITE	01.16.10.	10.45.07.	01.16.10.	13.43.30.	01.16.10.		
	LVLO	N	Y	N	D				NONE		
5 PROCC	IND. PRIVATE			1	75/09/16.	86/04/17.	76/05/06.	87/05/11.	76/05/06.		
	195			WRITE	01.16.11.	10.45.07.	12.13.40.	13.43.30.	12.13.40.		
	LVLO	N	Y	N	D				NONE		
6 TESTC	IND. PRIVATE			1	75/04/22.	86/04/17.	76/06/02.	87/05/11.	76/06/02.		
	202			WRITE	23.40.27.	10.45.07.	19.25.46.	13.43.30.	19.25.46.		
	LVLO	N	Y	N	D				NONE		
7 IDA	IND. PRIVATE			6	75/09/17.	86/04/17.	76/06/02.	87/05/11.	76/06/02.		
	191			WRITE	01.43.02.	10.45.07.	19.25.59.	13.43.30.	19.25.59.		
	LVLO	N	Y	N	D				NONE		
8 TESTRJ	IND. PRIVATE			15	75/09/17.	86/04/17.	76/06/02.	87/05/11.	76/06/02.		
	188			WRITE	01.43.03.	10.45.07.	19.26.04.	13.43.30.	19.26.04.		
	LVLO	N	Y	N	D				NONE		

Figure 17-16. Sample Directory Produced by PFCAT,LO=T,DN=40,UI=4054

MASS STORAGE TABLE REPORT										yy/mm/dd. hh.mm.ss.	
FAMILY/PACK NAME SYST64											
DEVICE NUMBER	=		1	DEVICE TYPE	=		DJ				
DEVICE MASK	=		377	NUMBER UNITS	=		1				
SECONDARY MASK	=		377	USER NAME	=		-----				
D/A USER COUNT	=		0	EQUIPMENT STATUS	=		ACTIVE				
SYSTEM.											

Figure 17-17. Mass Storage Table Report

SUMMARY REPORT				yy/mm/dd. hh.mm.ss.
DIRECTORY OF PERMANENT FILE DEVICE			40 SYS964	PAGE 1
(* = DAF RESIDENT ON OTHER THAN MASTER DEVICE.)				
USER	INDEX	FILES	SECTORS	AVE.SEC.
	1000	4	21	5
	3000	30	1163	38
	6000	44	602	13
	60000	1	5	5
	2020	1	0	0
	4420	2	2	1
	60	7	75	10
	530	21	443	21
	160	2	240	120
	2160	14	25	1
	260	37	5024	135
	270	70	6352	90
	2750	251	1386	5
	377760	1	68	68
	2370	22	626	28
	377770	1	53	53
		508	16085	(GROUP TOTAL)
	1	34	153	4
	21	96	1256	13
	31	15	87	5
	51	5	5	1
	3511	1	21	21
	201	7	77	11
	22201	3	362	120
	3641	40	1788	44
	301	27	1020	37
	311	1	3	3
	3331	8	168	21
	43331	36	296	8
	2741	15	1399	93
		288	6635	(GROUP TOTAL)

Figure 17-18. Cumulative Statistics by User Index PFCAT,LO=S, DN=40.

SUMMARY REPORT			yy/mm/dd. hh.mm.ss.
DIRECTORY OF PERMANENT FILE DEVICE		40 SYS964	PAGE
INDIRECT ACCESS FILE HOLES			4
CATALOG TRACK	NUMBER HOLES	NUMBER SECTORS	
0	8	501	
2	52	849	
3	1	35	
4	6	68	
5	36	745	
6	1	1	
11	111	1787	
13	4	187	
17	1	18	
20	12	204	
25	12	202	
26	28	321	
32	142	2095	
34	39	1225	
36	603	7026	
TOTAL	1056	15264	

Figure 17-19. Cumulative Statistics for Entire Device PFCAT,LO=S,DN=40. (Continued)

SUMMARY REPORT				yy/mm/dd. hh.mm.ss.	
DIRECTORY OF PERMANENT FILE DEVICE			40 SYS964	PAGE	5
PERMANENT FILE STATISTICS SUMMARY					
MASTER DEVICE USAGE					
GROUP TOTALS					
GROUP	TOTAL FILES	TOTAL SECTORS	PERCENT OF TOTAL USAGE		
0	508	16085	24		
1	288	6635	10		
2	371	11546	17		
3	176	5506	8		
4	144	4317	6		
5	371	7277	10		
6	226	5276	7		
7	394	9513	14		
			TOTAL	IAF	DAF
TOTAL SECTORS			66155	25588	40567
TOTAL FILES			2478	2121	357
TOTAL USERS			113		
AVE. FILES/USER			21	18	3
AVE. SEC/FILE			26	12	113
DEVICE TYPE - DI					
PERCENT DEVICE USAGE			38	15	24
DIRECT ACCESS DEVICE USAGE					
DEVICE NUMBER	DEVICE TYPE	TOTAL FILES	TOTAL SECTORS	PERCENT USAGE	
40	DI	357	40567	24	

Figure 17-19. Cumulative Statistics for Entire Device PFCAT,LO=S,DN=40.

Copy Archive File (PFCOPY)

PFCOPY extracts files from an archive file and copies them to one or more files local to the job. The way the files are copied depends upon the parameter options selected.

If MF=filename is specified, all the files extracted from the archive file are copied to one master file (filename) local to the job and they do not retain their permanent file names.

In a secured system, the access level from the file's PFC entry is used to assign each file to be copied to an appropriate mass storage device. If no device can be found that allows the file to reside there, the PFCOPY utility skips the file and issues a diagnostic message.

Dump Permanent File (PFDUMP)

PFDUMP copies (dumps) permanent files to backup storage (an archive file). Dumps can be reloaded by the PFLOAD utility and can be accessed by the PFATC and PFCOPY utilities for cataloging and copying. PFDUMP issues messages to the dayfile indicating how many files were dumped and how many files were not dumped due to errors. The type of dump taken depends upon the purpose of the dump and is determined by the parameters selected. The defining characteristics of the three types of dumps follow.

Types	Characteristics
Incremental dump	OP=M is specified, and BD=yymmdd and BT=hhmmss are not specified.
Partial dump	Any dump that is not incremental and not full.
Full dump	No file selection options (date/time interval, file size, file names, etc.) except DN or TD are specified.

If PFDUMP encounters an unrecoverable parity error while writing the archive file, several options are available via the K display. These options are described on the K display. If the option to continue the dump on another reel of tape is selected (KGO), it may not be possible to recover in such a way as to ensure that the PFLOAD utility can successfully read the archive tape. At least the archived file being written at the time of the error cannot be reloaded correctly.

PFDUMP in a Secured System

In a secured system, PFDUMP determines the maximum range of access levels that can be dumped. If the LA and UA parameters have been used to select access level limits, these limits will be used. If no access level limits were selected, PFDUMP uses the device limits determined by taking the lowest lower access limit and the highest upper access limit of all the devices to be dumped. The range of possible access levels must be within the system access level limits or PFDUMP aborts the job and issues a diagnostic message. The range of possible access levels must also be within the equipment access level limits for the equipment (tape or mass storage) assigned to the archive file and verify file (if one is being written). If not within the equipment access level limits, PFDUMP aborts the job and issues a diagnostic message.¹⁷

Incremental Permanent File System Dump

An incremental dump copies those permanent files modified after a specified date; that is, the OP=M option is specified together with the date (AD) and time (AT) options but BD=yymmdd and BT=hhmmss are not specified. Other options specified on the PFDUMP call can restrict the dump even further, but caution should be used when doing so if the dumps are used as part of an incremental reload scheme.

For example, the command

```
PFDUMP,OP=MI,AD=790101,AT=100000.
```

causes only indirect access files (I parameter) modified after the specified date and time to be dumped. Normal procedure at a site is to follow up this dump with successive incremental dumps (usually with a periodic advancement of date and/or time). This produces a series of archive files containing successive updating of all files in the defined category. (These files can reside on one file: a multifile archive file.) The archive files can then be incrementally loaded (refer to the PFLOAD utility described later in this section) to return the most recently modified versions of the archived files to the permanent file system.

Each incremental dump writes a record (or records) at the beginning of the archive file it creates. The record contains catalog images (refer to Catalog Image Record later in this section) of all files active on the device(s) being dumped when the dump took place (files with no backup requirement (BR=N) are not included on the catalog image record unless OP=Y is specified). This enables a future incremental load of these files on a system or a device basis.

An example of incremental dumping is given under Load Permanent File (PFLOAD) later in this section.

17. The system access level limits are displayed as part of the left screen header. The device access level limits (for mass storage) and equipment access level limits (for magnetic tape) are displayed on the E.A display. Refer to the NOS Version 2 Operations Handbook for display formats.

Partial Permanent File System Dump

A partial dump copies permanent files according to any specified option(s),¹⁸ except those defining a full or incremental dump. For example, the command

```
PFDUMP, OP=D.
```

is a partial dump of all direct access permanent files.¹⁸ Similarly, a dump of all files created (OP=C) or accessed (OP=A) after a certain date is a partial dump.

A partial dump can dump files from either a certain device (DN or TD option) or all files in the system (DN and TD not specified) meeting the criteria of the specified options.

Full Permanent File Dump

A full dump copies all files in the system or those cataloged on a specified device;¹⁸ no file selection options are specified except possibly DN (for a master device) or TD.

PFDUMP File Interlocking

PFDUMP interlocks a catalog track during the time it takes to dump all files whose PFC entries reside on that catalog track. No permanent file commands are processed on files whose PFC entries reside there until PFDUMP begins processing the next catalog track but jobs with direct access files attached are unaffected. For this reason, full permanent file dumps should not be done during heavy system use. An incremental dump provides adequate backup with less interference to the running system. PFDUMP skips any direct access files attached by jobs in write, modify, or append mode. Files attached in a non-writable mode or in update mode are dumped.

PFDUMP in an MSE, Tape, or Optical Disk Alternate Storage Environment

Because a large percentage of the total permanent file data typically does not reside on disk in an MSE, tape, or optical disk alternate storage environment, performing full backup dumps of file data is impractical. Whether PFDUMP dumps the data of a selected file depends on the user-specified backup requirement, the PFDUMP parameters selected, and the number of backup copies that exist. If these criteria do not require a data dump, only the file's PFC entry and permits are written to the dump file.

The pertinent PFDUMP parameters are:

Parameter	Description
OP=Y	Treat all files as though the user specified a backup requirement of yes (BR=Y).
COS=size	PFC only size threshold. For all files with a file size greater than or equal to the specified size and for which the file's backup requirement is met by alternate storage copies, dump only the PFC and permits (do not dump the file data even if the data is available on disk).
OP=S	Suppress staging of file data from alternate storage if a data dump is selected and the file is not disk resident.

¹⁸. Files with a backup requirement of none (BR=N) are not dumped unless OP=Y is specified.

The decision hierarchy for PFDUMP operations is as follows:

- If OP=Y is specified, force the file's backup requirement to yes (BR=Y).
- If the file backup requirement is none (BR=N), do not dump the file.
- If the file size is less than the value specified by the COS parameter, select a data dump.
- If the file size is greater than or equal to the value specified by COS:
 - and the backup requirement is media dependent (BR=MD) and no alternate storage copy exists, select a data dump.
 - and the backup requirement is media dependent (BR=MD) and at least one alternate storage copy exists, select a PFC and permits-only dump.
 - and the backup requirement is yes (BR=Y) and less than two alternate storage copies exist, select a data dump.
 - and the backup requirement is yes (BR=Y) and two or more alternate storage copies exist, select a PFC and permits-only dump.
- If a PFC and permits-only dump is selected, write the file's PFC entry and permits to the dump file.
- If a data dump is selected:
 - and the file is disk resident, write the file's PFC entry, permits, and data to the dump file.
 - and the file is not disk resident and staging is not suppressed (OP=S omitted), stage the file data to disk from alternate storage and write the file's PFC entry, permits, and data to the dump file.
 - and the file is not disk resident and staging is suppressed (OP=S specified), write the file's PFC entry and permits to the dump file.

A typical set of PFDUMP parameters that a site might use for full and incremental backup dumps is:

OP=Y not specified.

COS set to some reasonably small value or zero (default is 0).

OP=S specified.

If only tape alternate storage is in use and secondary VSNs have been created on all tape destaging operations, the OP=S option does not have to be specified since the backup requirement for all files is met. Specifying a nonzero value for COS may reduce the amount of staging activity following a permanent file reloading operation. However, files larger than or equal to the COS value will have to be staged back to disk in such a situation since only the PFC and permits will be reloaded.

When a partial PFDUMP is performed to move files to a different user name, family, or system, it may be desirable to force a dump of the file data and clear the alternate storage information in the files' PFC entries depending on the situation. If files are being moved to a user name on the same family or to a different family on the same machine, there should be no need to dump the data of files residing only on tape alternate storage since the ability to stage files from tape does not depend on the family or user index under which a file is cataloged. If files are to be moved to a remote machine or if any of the files have an image on MSE or optical disk, the file data should be dumped and the alternate storage information should be cleared either by PFDUMP or on the subsequent PFLOAD. This can be accomplished by specifying the following parameters:

OP=S not specified.
 COS=* set to select unlimited PFC only size threshold.
 OP=Z specified.

Unlimited COS must be specified to select a data dump for all files having a backup requirement of yes (BR=Y). The OP=Z parameter could be specified on the PFLOAD instead of the PFDUMP, but it is recommended that it be done at dump time so that the operator does not have to remember to do it at load time.

PFDUMP Protected Files

When PFDUMP specifies purging of files (OP=P), files whose user indexes are greater than PGUI (default value is 300000_g) are not purged. This is done to prevent files such as critical recovery files, applications, and system utilities from being purged from the system due to lack of use.

Load Permanent File (PFLOAD)

PFLOAD loads archived files produced by the PFDUMP utility back into the permanent file system. The load can reestablish the permanent file system exactly as it was at the time of the dump, or can load only a desired subset of files on the archive file (as indicated by specified parameters). PFLOAD issues messages to the dayfile indicating how many files were loaded and how many files with errors were encountered.

If LO=T (list all files processed) is specified, the listing produced is similar to that shown in figure 17-15. If LO=C (list all files in the catalog image record) is specified, the listing produced is similar to that shown in figure 17-20.

The two types of loads available are incremental and nonincremental. An incremental load builds up (increments) an accumulation of the most recently modified versions of the files extracted from the archive files for loading. A series of archive files is read in the reverse order of creation. The CIR created by the most recent incremental dump is read and checked against the archived files on these files. If a file matches an entry on the CIR, that file is a candidate for loading. A nonincremental load does no CIR checking and uses only parameters specified on the PFLOAD call, if any, to select candidates for loading.


PFATC.	FILES ON CATALOG IMAGE RECORD.			FAMILY = SYS964.		
	ARCHIVE FILE	CREATED	yy/mm/dd.	hh.mm.ss.	ARCHIVE FAMILY = TEST.	
	FILE	USER	ACCESS	LAST ACCESS	DEVICE	
	NAME	INDEX	COUNT	DATE	TIME	NUMBER
1	ARCPD	3320	0	87/04/13.	15.22.31.	11
2	SITRPT	3320	0	87/04/13.	15.22.29.	
3	LISTEXB	3320	0	87/04/13.	15.22.29.	
4	FSEPZ19	3320	0	87/04/13.	15.22.29.	
5	DECK3	3320	0	87/04/13.	15.22.29.	
6	SHORT	3320	0	87/04/13.	15.22.29.	
7	PFHELPR	3320	0	87/04/13.	15.22.29.	
8	PFACC	3320	0	87/04/13.	15.22.29.	
9	UPROC	3320	0	87/04/13.	15.22.30.	
10	MAILGRP	3320	0	87/04/13.	15.22.30.	
11	ACPD1	3320	0	87/04/13.	15.22.30.	
12	COPYRF	3320	0	87/04/13.	15.22.30.	
13	STAGE	3320	0	87/04/13.	15.22.30.	
14	CPD6	3320	0	87/04/13.	15.22.30.	
15	WORK	3320	0	87/04/13.	15.22.30.	
16	MAG	3320	0	87/04/13.	15.22.30.	
17	OLDTIME	3320	0	87/04/13.	15.22.30.	
18	PRC	3320	0	87/04/13.	15.22.30.	
19	T1	3320	0	87/04/13.	15.22.31.	
20	FSEPROC	3320	0	87/04/13.	15.22.31.	
21	ARCGSY	3320	0	87/04/13.	15.22.31.	11
22	ARCPCS	3320	0	87/04/13.	15.22.33.	11

Figure 17-20. Catalog of Catalog Image Record

Catalog Image Record

Each incremental dump (OP=M with AD or AT specified) writes a CIR at the beginning of the archive file on which the permanent files are dumped. The other types of dumps (full and partial) do not produce a CIR. The CIR is composed of 2-word entries for every permanent file that exists on the device(s) being dumped at the time of the incremental dump (not just the files that are included in the dump).

The format of each entry is:



Field	Description
filename	Name of the permanent file.
userindex	User index under which filename was cataloged.
access count	Number of times filename was accessed.
dn	Device number of the device on which filename resides if it is a direct access file and resides on a device other than your master device. If it resides on the master device, dn=0. If the file is indirect access, dn is ignored.
date time	Date and time filename was last accessed.

When a file is loaded, this CIR information is placed in the permanent file catalog entry of the file being loaded. No other information can be changed over dumps and loads without updating the utility control date and time.

Figure 17-20 contains a partial listing of the CIR. This listing can be obtained after a selective dump if the LO=C option is specified on a PFDUMP, PFLOAD, or PFATC call. A listing similar to the one in figure 17-15 can be obtained after a selective dump if the LO=T option is specified on a PFDUMP, PFLOAD, or PFATC call. In both cases, the headings vary but the format of the information is the same.

PFLOAD Selection Process

Execution of the PFLOAD utility involves the following stages to determine which permanent files (from the archive file) should be loaded and on which devices they should reside.

1. If the DN, TD, or DD parameters are used, the device number specified refers to a device in the running system and must exist in the family being loaded or PFLOAD aborts and issues this message:

```
PFLOAD - DEVICE nn NOT FOUND.
```

2. If the specified parameters logically imply that no files can be selected (such as OP=I, DN=n where n is not a master device), PFLOAD aborts and issues this message:

```
PFLOAD - NO FILES SELECTED.
```

3. The user index mask for selecting files (based on the device mask of the device specified by the DN and DI parameters) must be such that it is possible for some files to be on the archive file. There must be some common bits between the device mask of the archive file (written in the header record by PFDUMP) and the device mask of the device specified by the DN parameter for PFLOAD. If no common bits exist, PFLOAD aborts and issues this message:

```
PFLOAD - SELECTED FILES NOT ON ARCHIVE FILE.
```

4. For nonincremental loads, this stage of checking is bypassed and the entire archive file to be loaded is passed on to the next stage.

The first step is to read the CIR from the most recent selective incremental dump file and place it on a random file. Then the archive files are read in the reverse order in which they were created (the most recently created file is read first).

Each file on an archive file is checked against the CIR on the random file for a match of file name and user index (these fields describe a file uniquely in a family). If a match is found, the file is a candidate for loading and moves to the next stage of checking. Also, the file name and user index are zeroed out in the corresponding CIR entry to prevent subsequent loading of older versions of the same file. If a match is not found, the file is skipped because it has already been accepted (it appeared on a previous archive file and a match with the CIR was found) or has been purged. This process passes on only the latest versions of files obtained from the archive files.

5. Only files that satisfy the criteria specified by the parameters on the PFLOAD command are selected (date/time, direct or indirect, user index, file name, and so forth). All other files are skipped.

6. The DN and TD parameters are used to select files from the archive file and to determine the device on which the files should be loaded. The device number specified refers to a device in the running system.

Indirect access files, if selected, are always loaded on the appropriate master device (based on user index). Direct access files, if selected, are loaded on the device on which they previously resided. This is the device number contained in the file's PFC entry on the archive file. If the device number is 0, the file previously resided on the master device and will be reloaded there.

It may not be possible to load a direct access file on the device on which it previously resided because the device number may no longer exist or its secondary mask may not allow it. (Refer to Device Masks earlier in this section for additional information.)

Either of these situations can occur as a result of a reconfiguration of the permanent file system. In either case, if a direct access file cannot be loaded, the DD parameter is checked. If an alternate device is specified and its secondary mask allows it, the file is loaded on the alternate device. Otherwise, the OP=L option is checked (see step 7). If the OP=L option is not specified, PFLOAD skips the file and issues this message:

```
PFLOAD - ALTERNATE DEVICE NOT FOUND, FN=filename, UI=userindex.
```

The DN parameter is used to select files from the archive file in the following manner. If the device specified by DN (in the running system) is a master device, all files (direct and indirect) that would be cataloged on the master device (files with user indexes corresponding to that device's device mask) are selected. If the device specified by DN is not a master device, only direct access files that previously resided on the device (files with a device number of DN in their PFC entry) are selected.

Similarly, the TD parameter is used to select files from the archive file in the following manner. If the device specified by TD is a master device, all files that were cataloged on it and all direct access files that previously resided on it are selected. (This corresponds to the way the TD parameter is processed by PFDUMP.) If the device specified by TD is not a master device, only direct access files that previously resided on it are selected.

If neither the DN nor TD parameters are specified, all files meeting other selection criteria are selected. Specifying both DN and TD has the same effect as specifying only TD.

In all cases, the device masks and master device status for the devices specified by the DN or TD parameters refer to the system in which PFLOAD is running, not the system in which the files were dumped. The user index and device number from the PFC on the archive file reflect the situation that existed at the time the files were dumped.

7. If the OP=L option is specified, the device selected in step 6 (or the failure to select a device) for the residence of a direct access file is ignored. PFLOAD determines the device, among those devices where the file may reside (based on the device mask), which has the most available space. The OP=L option is not applicable to indirect access files since they must reside on the appropriate master device.

8. When a device has been selected to receive a direct access file, PFLOAD determines if the file will fit on the device. If the file is too large, PFLOAD skips the file and issues this message:


PFLOAD - NO SPACE FOR FILE, FN=filename, UI=userindex.

9. In secured systems, PFLOAD verifies before loading a file that the device where the file is to reside (as selected in steps 7 and 8) is appropriate for the file's access level. If the device is not appropriate, PFLOAD skips the file and issues this message:

PFLOAD - NO DEVICE FOUND FOR FILE, FN=filename, UI=userindex.

If PFLOAD is loading direct access files to the device with the most available space (OP=L), PFLOAD selects the device with the most space that also allows the required access level.

10. PFLOAD determines the catalog track where the files to be loaded are to reside and builds an index file with a 1-word entry for each file on the track. The format of the entry is:



PFLOAD checks each candidate file against this entry to determine whether it is in the permanent file system. If it is not, the file is loaded. If it is, loading depends upon the OP=R option. If OP=R was specified on the PFLOAD call, the archive duplicate replaces the one in the permanent file system. If OP=R was not specified, archive duplicates are skipped.

Family Reconfiguration

When reconfiguring a permanent file family by adding or subtracting devices, or changing the device masks, use the PFLOAD OP=L parameter to greatly simplify the procedure. Follow this sequence to reconfigure a permanent file family:

1. PFDUMP, FM=familyname.
2. Initialize all devices in the family.
3. PFLOAD, OP=L, FM=familyname.

PFLOAD in an MSE, Tape, or Optical Disk Alternate Storage Environment

In an MSE, tape, or optical disk alternate storage environment, an incremental load as described previously is the typical reload procedure. PFLOAD uses the CIR from the first incremental dump file to control the loading of file data and PFC and permit entries. For files with alternate storage copies, the files' data may not have been copied to the dump tape (refer to PFDUMP in an MSE, Tape, or Optical Disk Alternate Storage Environment earlier in this section for more information on how PFDUMP decides whether to dump a file's data or to dump only the file's PFC entry and permits). PFLOAD reloads only what was dumped by PFDUMP. Thus, if the file data was included on the dump tape, it will be reloaded to disk. If the file data was omitted from the dump tape, only the PFC and permit information for the files will be reloaded. Refer to section 11, MSE, or to section 12, Tape and Optical Disk Alternate Storage, for further information on how the PFLOAD utility is used in an MSE or tape alternate storage environment.

PFDUMP/PFLOAD Example

In this example, permanent files created and modified on three devices are dumped nine times. The action runs from March 1, 1982 (82.3.1) to March 10, 1982 (82.3.10). To keep things simple, the time of day is not shown. Before the last dump, one of the devices becomes inoperable. The example then shows how this device can be loaded from the archive tapes. Following are the three devices.

Device	Description
40	A master device containing indirect access files only. Users with this master device have their direct access files written on 41 or 42.
41	A master device containing indirect and direct access files. Direct access files cataloged on this device may be written on 42 or this device.
42	A nonmaster device containing direct access files only. All files written on this device will be cataloged on 40 or 41.

File names consist of a letter and a number. The letter identifies which device contains the catalog entry for this file (A=device 40 and B=device 41). The number is the same as the device on which the file is written.

Each file name is followed by a date in parentheses. This indicates either the day on which the file was created or when it was last modified.

Figure 17-21 outlines the creation, modifications, and dumps that were made. There is one particular action per column. The top of the column identifies the action and the key parameters used by the utility. If an archive tape is produced, it is shown at the bottom of the column. All files in the system are listed in each column. Those files affected by the action are underscored.




Figure 17-21. PFDUMP/PFLOAD Example (Sheet 1 of 3)

DEVICE 40 MASTER INDIRECT	DUMP 3 made on 82.3.4 OP-M AD=82.3.2 DN=41	DUMP 4 made on 82.3.4 OP-M AD=82.3.2 DN=42	Modifications made on 82.3.5	DUMP 5 made on 82.3.6 OP-M AD=82.3.4 (all devices)	DUMP 6 made on 82.3.7 OP-M AD=82.3.4 DN=42
	A40(82.3.3)	A40(82.3.3)	A40(82.3.5)	A40(82.3.5)	A40(82.3.5)
	<u>B41(82.3.3)</u>	B41(82.3.3)	<u>B41(82.3.5)</u>	<u>B41(82.3.5)</u>	B41(82.3.5)
	A41(82.3.1)	A41(82.3.1)	A41(82.3.1)	A41(82.3.1)	A41(82.3.1)
	<u>B42(82.3.3)</u>	<u>B42(82.3.3)</u>	<u>B42(82.3.5)</u>	<u>B42(82.3.5)</u>	<u>B42(82.3.5)</u>
	A42(82.3.1)	A42(82.3.1)	A42(82.3.1)	A42(82.3.1)	A42(82.3.1)
DEVICE 41 MASTER DIRECT+INDIRECT	DUMP 3 tape CIR(82.3.4) B41(82.3.3) B42(82.3.3)	DUMP 4 tape CIR(82.3.4) B42(82.3.3)	DUMP 5 tape CIR(82.3.6) A40(82.3.5) B41(82.3.5) B42(82.3.5)	DUMP 6 tape CIR(82.3.7) B42(82.3.5)	
DEVICE 42 NONMASTER DIRECT	DUMP 3 tape CIR(82.3.4) B41(82.3.3) B42(82.3.3)	DUMP 4 tape CIR(82.3.4) B42(82.3.3)	DUMP 5 tape CIR(82.3.6) A40(82.3.5) B41(82.3.5) B42(82.3.5)	DUMP 6 tape CIR(82.3.7) B42(82.3.5)	

Figure 17-22. PFDUMP/PFLOAD Example (Sheet 2 of 3)

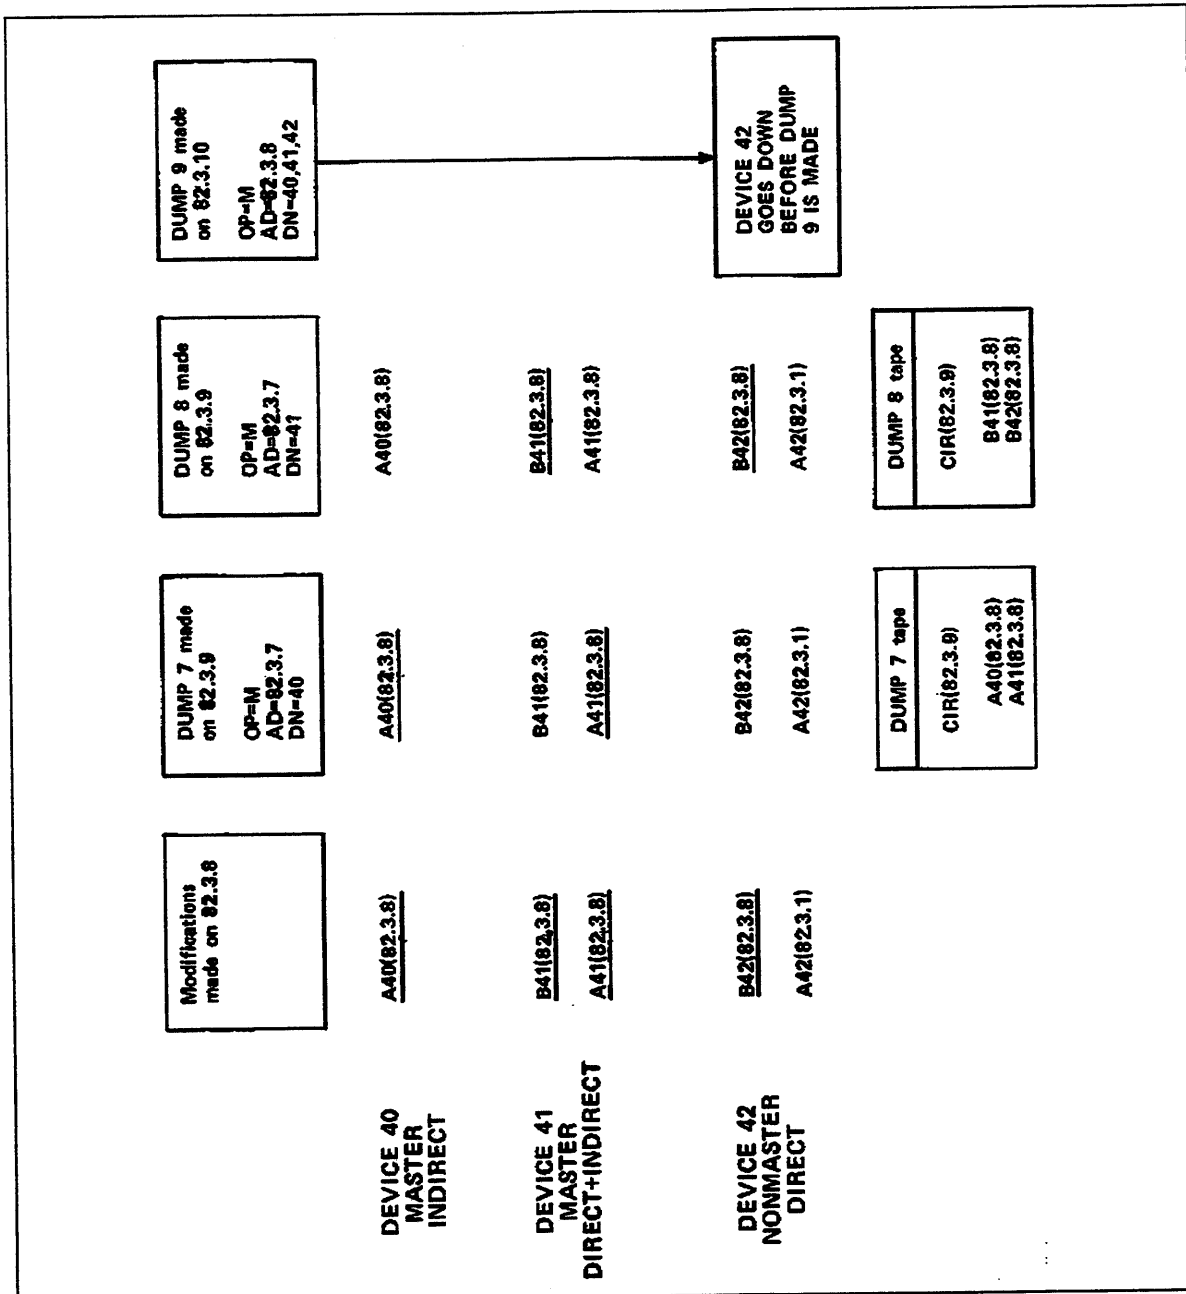


Figure 17-23. PFDUMP/PFLOAD Example (Sheet 3 of 3)

Actions during the 10 days in March in which the dumps and modifications were made run as follows:

Date	Action
82.3.1	Five permanent files are created on this date. File A40 (indirect) is written on device 40 and cataloged on device 40. File A42 (direct) is written on device 42 and cataloged on device 40. File A41 (direct) is written on device 41 and cataloged on device 40. File B41 (indirect) is written on device 41 and cataloged on device 41. File B42 (direct) is written on device 42 and cataloged on device 41.
82.3.2	DUMP 1. A full permanent file system dump. No CIR is written on the archive file.
82.3.3	Three files (A40, B41, and B42) are modified on this date. This is indicated by the modification date in parentheses following the file name.
82.3.4	DUMP 2. This is an incremental dump of device 40 that specifies permanent files modified after 82.3.3. The files for this device are the ones listed on its catalog track. The utility scans this catalog which contains entries for files A40, A42, and A41. File A40 is on device 40 and the other two are on the devices indicated by their names. Only file A40, which was modified on 82.3.3, qualifies for this dump. It is written on an archive tape after a CIR for this date (82.3.4). The CIR gives a current description of all five files in the permanent file system.
82.3.4	DUMP 3. This is an incremental dump of device 41 that specifies all files modified after 82.3.2. The catalog for this device contains entries for files B41 and B42. Both files were modified on 82.3.3 and qualify for this dump. They are written on an archive file after a CIR for this date.
82.3.4	DUMP 4. This is an incremental dump of device 42 that specifies all files modified after 82.3.2. The PFDUMP utility scans the catalogs on devices 40 and 41 to locate the files on 42. One file, B42 (82.3.3), falls in this category.
82.3.5	Three files are modified on this date.
82.3.6	DUMP 5. This is an incremental dump of each of the three devices. All files modified after 82.3.4 are specified by the AD option. Since no DN is specified, all devices are dumped in turn, producing a dump tape with three archive files.
82.3.7	DUMP 6. This incremental dump of device 42 specifies all files modified after 82.3.4. Since 42 is nonmaster, the dump routine scans the catalogs of the other two devices.
82.3.8	Four files are modified on this date.

Date	Action
82.3.9	DUMP 7. An incremental dump of device 40 is made with AD=82.3.7. Two files qualify.
82.3.9	DUMP 8. An incremental dump of device 41 is made with AD=82.3.7. The catalog on device 41 contains entries for files B41 and B42. File B41 is an indirect access file located on device 41, and file B42 is a direct access file located on device 42. Both files were modified after 82.3.7 and are dumped.
82.3.10	DUMP 9. An incremental dump of devices 40, 41, and 42 is specified, but device 42 becomes inoperable before the dump is made.

To reestablish permanent files on device 42 with the most recent modifications available, the device is first initialized and then incrementally loaded, beginning with the most recent incremental dump tape. Archive dump 8 is the most recent incremental dump tape and accordingly contains the most recent CIR. This reel is assigned first and then the CIR read onto a random file. The dump tapes are read in reverse order with each file on each tape being checked against the CIR. Dump tapes 2 and 3 are omitted since they do not involve device 42.

For each catalog track on device 42 that is to receive file entries, PFLOAD builds and maintains an index file. Before each file can be loaded, PFLOAD checks the index file to see if the file has already been loaded. If it has, the archive duplicate is skipped; if not, the archive file is loaded and another entry goes into the index file.

After initialization of device 42, the loading runs as follows:

Archive File	Action						
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">DUMP8 tape</td></tr> <tr><td style="padding: 2px;">CIR(82.3.9)</td></tr> <tr><td style="padding: 2px;">B41(82.3.8)</td></tr> <tr><td style="padding: 2px;">B42(82.3.8)</td></tr> </table>	DUMP8 tape	CIR(82.3.9)	B41(82.3.8)	B42(82.3.8)	<p>The analyst assigns DUMP8 tape and enters the command:</p> <pre style="margin-left: 40px;">PFLOAD, T=DUMP8, DN=42.</pre> <p>The utility reads the CIR from this tape onto a random file. It then looks for files for device 42 on this tape. B42 qualifies, and its entry in the CIR is removed. The file B42 is loaded on device 42. The catalog of this file on device 41 is created. After loading this file, the utility makes an entry in the index file as follows:</p> <table border="1" style="margin-left: 40px; width: 60%; border-collapse: collapse;"> <tr> <td style="padding: 5px;">B42</td> <td style="padding: 5px; text-align: right;">userindex</td> </tr> </table>	B42	userindex
DUMP8 tape							
CIR(82.3.9)							
B41(82.3.8)							
B42(82.3.8)							
B42	userindex						
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">DUMP7 tape</td></tr> <tr><td style="padding: 2px;">CIR(82.3.9)</td></tr> <tr><td style="padding: 2px;">A40(82.3.8)</td></tr> <tr><td style="padding: 2px;">A41(82.3.8)</td></tr> </table>	DUMP7 tape	CIR(82.3.9)	A40(82.3.8)	A41(82.3.8)	<p>The utility then requests the next reel.</p> <p>The analyst assigns this tape, and the utility reads the archive file and checks it against the CIR loaded from DUMP8 above. No match is found. The next reel is requested.</p>		
DUMP7 tape							
CIR(82.3.9)							
A40(82.3.8)							
A41(82.3.8)							
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">DUMP6 tape</td></tr> <tr><td style="padding: 2px;">CIR(82.3.7)</td></tr> <tr><td style="padding: 2px;">B42(82.3.5)</td></tr> </table>	DUMP6 tape	CIR(82.3.7)	B42(82.3.5)	<p>This tape is assigned and then compared with the CIR. B42 is a file from device 42. However, its entry in the CIR was removed when it was read from DUMP8.</p>			
DUMP6 tape							
CIR(82.3.7)							
B42(82.3.5)							
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">DUMP5 tape</td></tr> <tr><td style="padding: 2px;">CIR(82.3.6)</td></tr> <tr><td style="padding: 2px;">A40(82.3.5)</td></tr> <tr><td style="padding: 2px;">B41(82.3.5)</td></tr> <tr><td style="padding: 2px;">B42(82.3.5)</td></tr> </table>	DUMP5 tape	CIR(82.3.6)	A40(82.3.5)	B41(82.3.5)	B42(82.3.5)	<p>This tape is assigned and then compared with the CIR. B42 is from device 42 but it has already been removed from the CIR and the utility skips this version.</p>	
DUMP5 tape							
CIR(82.3.6)							
A40(82.3.5)							
B41(82.3.5)							
B42(82.3.5)							
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">DUMP4 tape</td></tr> <tr><td style="padding: 2px;">CIR(82.3.4)</td></tr> <tr><td style="padding: 2px;">B42(82.3.3)</td></tr> </table>	DUMP4 tape	CIR(82.3.4)	B42(82.3.3)	<p>This tape is assigned and compared with the CIR originally read from DUMP8 tape. It has been removed from the CIR and this file is skipped.</p>			
DUMP4 tape							
CIR(82.3.4)							
B42(82.3.3)							

Archive File**Action**

DUMP1 tape
A40(82.3.1)
A42(82.3.1)
A41(82.3.1)
B41(82.3.1)
B42(82.3.1)

This tape, produced by a full dump, contains all the files that were on the permanent file system when the dump was made (82.3.1). These files are checked against CIR (82.3.9). A match is found with A42. This file is loaded.

Device 42 now has the same files with the same update status that they had on 82.3.9 when the last incremental dump was made.

Release Disk Space (PFREL)

The PFREL utility releases the disk space of files that reside on alternate storage if the files' backup requirements are met.¹⁹ PFREL can be called directly or used with a file selection utility such as GENPFD in the following sequence.

First, execute PFREL with inhibit processing (IP parameter) and the summary file (S parameter) selected. Files can be excluded by specifying selection parameters. In particular, OP=A can be used with BD and BT to exclude files that have been recently accessed. For this example, enter the following command.

```
PFREL, IP, S, OP=A, BD=yymmdd, BT=hhmmss, FM=familyname.
```

Second, execute the file selection utility GENPFD using the summary file generated in step 1 as input. A utility directives file will be generated containing the UI and PF parameters for the selected files. Refer to section 12, Tape and Optical Disk Alternate Storage, for a description of the GENPFD utility.

Third, reexecute PFREL without specifying the IP parameter and using the utility directives file generated in step 2 as input. For this example, enter the following command.

```
PFREL, I=dirfile.
```

The PFREL utility will not release the disk space of files cataloged on user name SYSTEMX or on the MSE user names SUBFAM0 through SUBFAM7 even if the selection criteria are met.

¹⁹. Refer to PFDUMP in an MSE, Tape, or Optical Disk Alternate Storage Environment earlier in this section for a discussion of backup requirements.

Queued File Utilities	18-1
Calling the Queued File Utilities	18-2
Console Input	18-2
Batch Input	18-11
Terminal Input	18-11
Parameters for the Queued File Utilities	18-11
Description of Queued File Utility Parameters	18-14
File Preassignment	18-26
Active and Inactive Queued Files	18-26
Queued File Utility Routines	18-27
DFLIST	18-27
DFTERM	18-27
LDLIST	18-30
QALTER	18-30
QDUMP	18-36
QFTLIST	18-40
QLIST	18-40
QLOAD	18-41
QMOVE	18-44
QREC	18-49
Dayfile Dumping Utilities	18-55

Several utility programs control queued input and output files and system, account, error log, and binary maintenance log dayfiles. The utility programs are divided into the following two groups, according to the type of selective processing they do.

- The queued file utilities are 10 utilities under the control of the queued file supervisor program QFSP. These utilities select for processing queued files or dayfiles that share certain characteristics you specify such as device residence, file destination, and job sequence name.
- The dayfile dumping utilities are five independent utilities that dump all or selected parts of the active system dayfile, account dayfile, error log, or binary maintenance log. These utilities each process the contents of a single file according to criteria you specify.

Queued File Utilities

The names and functions of the queued file utilities follow. The rest of this section includes detailed information about the call and operation of each utility.

Utility	Description
DFLIST	Lists dayfiles that have been made permanent files by the DFTERM utility.
DFTERM	Terminates an active or inactive dayfile and retains it as a direct access permanent file.
LDLIST	Lists queued files present on a QDUMP dump tape.
QALTER	Alters routing information associated with active queued files; purges active queued files.
QDUMP	Dumps queued files to a tape or disk.
QFTLIST	Lists detailed information about active queued files.
QLIST	Lists inactive queued files.
QLOAD	Loads files dumped by QDUMP as queued files.
QMOVE	Moves queued files from one disk device to another.
QREC	Deactivates or activates selected queued files; purges inactive queued files.

Calling the Queued File Utilities

All queued file utility operations, except DFLIST, LDLIST, and QLIST, can be initiated through console input to a K display. All queued file utilities except DFLIST, DFTERM, QALTER, and QFTLIST can be initiated through entries made at the system console (under DIS or DSD control, including procedure files), card input (batch), or terminal entries. DFTERM, DFLIST, QALTER, and QFTLIST operation can be initiated only through entries made at the system console. However, in all cases, for batch and terminal input, you must be validated for system origin privileges, the system must be in debug mode, and you must observe certain parameter order dependencies. Since LDLIST, QDUMP, and QLOAD require preassignment of the dump file, only DIS, batch, or terminal input is appropriate for these utilities. DFLIST, DFTERM, LDLIST, QDUMP, QFTLIST, and QLIST can access files on an inaccessible device. QMOVE can read files from an inaccessible device, but the job may get rolled out if QMOVE attempts to move files to another inaccessible device. If this happens, drop the job and attempt a QDUMP.

Console Input

Input from a console through a K display requires the following sequence of entries and responses. These procedures do not apply to DFLIST, QLIST, or LDLIST, because K displays are not available for these utilities.

1. Call QFSP by entering:

```
X.QFSP.
```

The B,O display indicates the job sequence name of QFSP. To the right of the entry this message appears:

```
REQUEST *K* DISPLAY
```

2. Activate the K display for that control point by entering:

```
K, jsn.
```

Parameter	Description
jsn	Job sequence name of QFSP as specified on the B,O display.

The QFSP K display shown in figure 18-1 appears on the left screen. Instructions on the display describe how to select the desired utility.

```
K. QFSP.          jsn
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                    MID=AA   NOS version

                    QUEUED FILE SUPERVISOR

ENTER K. UTILITY.

UTILITY            DESCRIPTION

QDUMP              QUEUED FILE DUMP PROCESSOR.
QLOAD              QUEUED FILE LOAD PROCESSOR.
QMOVE              QUEUED FILE MOVE PROCESSOR.
QREC               PROCESS QUEUED FILES.
DFTERM            TERMINATE DAYFILES.
QFTLIST           LIST ACTIVE QUEUED FILES.
QALTER            ALTER ACTIVE QUEUED FILES.
```

Figure 18-1. QFSP K Display (Left Screen)

3. Select the desired utility by entering:

K.utility

Parameter	Description
utility	One of the utilities shown in the K display, figure 18-1.

The parameters available under the chosen utility appear on the left screen as the initial K display for the utility. Figures 18-6, 18-14 through 18-16, 18-19 through 18-21, 18-23 through 18-25, and 18-28 through 18-30 exhibit initial K displays for the various utilities. Different parameters are available with the different utilities. (Refer to Description of Queued File Utilities later in this section.)

It is also possible to call each utility directly by substituting the following procedure for steps 1, 2, and 3.

a. Call the desired utility by entering:

X.utility

Parameter	Description
utility	One of the utilities shown in figure 18-1.

The B,O display indicates the job sequence name of the chosen utility. This message also appears:

REQUEST *K* DISPLAY

b. Activate the initial K display for the chosen utility by entering:

K,jsn.

Parameter	Description
jsn.	Job sequence name as specified on the B,O display.

c. Commands and directives are available to aid in the use of the utilities and direct processing. A right screen display lists and defines the commands and directives. Activate this K display by entering:

KK.

Figures 18-2 through 18-5 show the resultant QFSP K display on the right screen.

```

K. utility          jsn

K DISPLAY COMMANDS                                PAGE 1 OF 4

GO      PROCEED WITH PROCESSING.
OUT     RELEASE FILE SPECIFIED BY *L* PARAMETER TO PRINTER.
RESET   RESET PARAMETERS TO DEFAULT VALUES.
STOP    TERMINATE INPUT, END RUN.

+       PAGE LEFT SCREEN FORWARD.
-       PAGE LEFT SCREEN BACKWARD.
(       PAGE RIGHT SCREEN FORWARD.
)       PAGE RIGHT SCREEN BACKWARD.

```

Figure 18-2. QFSP K Display (Right Screen) (Page 1 of 4)

```

K DISPLAY COMMANDS                                PAGE 2 OF 4

CLEAR   SET BOTH SCREENS TO FIRST PAGE OF K DISPLAY.
        (QREC, DFTERM, QFTLIST, QALTER).
ERROR   TOGGLE STATUS OF ERRORED FILE PROCESSING.
        (QDUMP, QLOAD, QMOVE).
FAMILY  IF ENTERED, ENTRIES WILL BE MADE IN THE FOT FOR
        THE INACTIVE FAMILIES AND THE FILES WILL BE QUEUED.
        (QREC, QLOAD, QMOVE).

LID     IF ENTERED, ENTRIES WILL BE MADE IN THE LID TABLE
        FOR THE INACTIVE LIDS AND THE FILES WILL BE QUEUED.
        (QREC, QLOAD, QMOVE).
LIST    DISPLAY JSN-S OF ACTIVE/INACTIVE QUEUED FILES ON
        RIGHT SCREEN. (QREC, DFTERM, QALTER, QFTLIST).
REWIND  IF ENTERED, DUMP/LOAD FILE WILL BE REWOUND WHEN
        *GO* COMMAND IS ENTERED. (QDUMP AND QLOAD).

```

Figure 18-3. QFSP K Display (Right Screen) (Page 2 of 4)

DIRECTIVE FORMATS		PAGE 3 OF 4
DA=Yymmdd	PROCESS THIS DATE ONLY.	
DA=Yymmdd-Yymmdd	PROCESS ALL DATES IN SPECIFIED RANGE.	
DA=ALL	PROCESS ALL DATES.	
UI=XXXXXX	PROCESS THIS UNDER INDEX ONLY.	
UI=XXXXXX-YYYYYY	PROCESS ALL INDICES IN SPECIFIED RANGE.	
TUI=XXXXXX	PROCESS THIS USER INDEX ONLY.	
TUI=XXXXXX-YYYYYY	PROCESS ALL INDICES IN SPECIFIED RANGE.	
FS=XXXXXX	PROCESS THIS FILE SIZE ONLY.	
FS=XXXXXX-YYYYYY	PROCESS ALL SIZES IN SPECIFIED RANGE.	
	77777B INCLUDES ALL LARGER SIZES.	
ID=XX	PROCESS THIS ID ONLY.	
ID=XX-YY	PROCESS ALL ID-S IN SPECIFIED RANGE.	
FC=F1	PROCESS THIS FORMS CODE ONLY.	
FC=F1/F2/F3	PROCESS ALL FORMS CODES IN LIST.	
FC=F1-F2	PROCESS ALL FORMS CODES IN THE RANGE.	

Figure 18-4. QFSP K Display (Right Screen) (Page 3 of 4)

DIRECTIVE FORMATS		PAGE 4 OF 4	
DS=DT	FILE DESTINATION=DISPOSITION TYPE.		
DS	DESCRIPTION		
BC	BATCH.		
RB	REMOTE BATCH.		
DT	DESCRIPTION	DT	DESCRIPTION
IN	INPUT.	PU	PUNCH.
PL	PLOT.	PR	ANY PRINTER.
P2	512 PRINTER.	LR	580-12 PRINTER.
LS	580-16 PRINTER.	LT	580-20 PRINTER.
LX	5870 NIP.	LY	5970 NIP.
WT	WAIT QUEUE.	SF	SPECIAL QUEUE TYPES.
ALL	ALL QUEUES.	PRINT	ALL PRINT QUEUES.
NONE	NO QUEUES.		

Figure 18-5. QFSP K Display (Right Screen) (Page 4 of 4)

Following is a list of the queued file utility commands included on the first two pages of the QFSP K display (refer to figures 18-2 and 18-3) and a description of each.

Command	Description
CLEAR	Used in conjunction with the LIST command; returns the right screen K display to the list of processing commands initially displayed (refer to figure 18-2) and returns the left screen K display to the first page.
ERROR	<p>If the QDUMP utility is active, enabling this command causes files on which unrecoverable read errors were encountered to be dumped with the errors listed on the output file. Disabling this command causes these files to be ignored. The default under QDUMP is enabled.</p> <p>If the QLOAD utility is active, enabling this command causes such files to be loaded with the errors listed on the output file. Disabling the command causes such files to be ignored. The default under QLOAD is disabled.</p> <p>If the QMOVE utility is active, enabling this command causes such files to be moved with errors listed on the output file. Disabling this command causes such files to be ignored. The default under QMOVE is disabled.</p> <p>Entry of this command toggles setting of this command.</p>
FAMILY	Specifies whether or not files are queued when creation or destination families are not active [not in family ordinal table (FOT)]. Default is not to queue these files. When the FAMILY command is specified, entries are made in the FOT for inactive families and the files are queued. This command applies only to the QLOAD, QMOVE, and QREC utilities.
GO	Directs the active utility to proceed with the processing of the entered parameters. When the processing is complete, the left screen K display is reset to the default values.
LID	Specifies whether or not files are queued when the creation logical identifier (LID) is not in the LID table. Default is not to queue these files. When the LID command is specified, entries are made in the LID table and the files are queued. This command applies only to the QLOAD, QMOVE, and QREC utilities.

Command	Description
LIST	<p>If the DFTERM utility is active, this command displays a list of all permanent dayfiles (refer to figure 18-7). The list of files displayed is also written on the output file specified by the DFTERM L parameter.</p> <p>If the QALTER or QFTLIST utility is active, this command displays a list of active queued files on the right screen K display (refer to figure 18-11). The list of files is also written on the output file specified by the QALTER or QFTLIST L parameter.</p> <p>If the QREC utility is active, this command displays a list of inactive queued files on the right screen K display (refer to figure 18-31). The list of files displayed is also written on the output file specified by the QREC L parameter.</p>
LIST=qft	Displays detailed information about the active queued file at QFT ordinal qft on the right screen K display (refer to figure 18-12) and writes the information to the output file specified by the QALTER or QFTLIST L parameter. This form of LIST is accepted only by QALTER and QFTLIST.
OUT	Sends the output file specified by the L option of each utility to the print queue for immediate printing. Any information that has been or is currently displayed on the right screen K display via the LIST command is also included with the disposed output. In addition, if the LIST command is currently active, the list displayed on the right screen is disposed, and the initial right screen K display shown in figure 18-2 is returned automatically.
RESET	Resets all options displayed on the left screen K display to their default values.
REWIND	Rewinds the load or dump file specified by the FN parameter before processing. This command applies only to the QLOAD and QDUMP utilities.
STOP	Terminates the active utility and ends the K-display interaction.
+	Pages the left screen forward.
-	Pages the left screen backward.
(Pages the right screen forward.
)	Pages the right screen backward.

The third and fourth pages of the QFSP K display (refer to figures 18-4 and 18-5) show directive formats. These directives are described in detail under Parameters for the Queued File Utilities later in this section.

4. Select the desired parameters by entering:

```
K.opt1=val1,opt2=val2,...,optn=valn
```

The opt_i=val_i are selected from the parameter list on the initial K display for the chosen utility.

The option parameter entries are issued to the job dayfile, which is included in the output file specified by the L parameter for each utility (refer to Parameters for the Queued File Utilities later in this section).

If the FM parameter is specified for QDUMP, QMOVE, or QREC (indicating the processing of queued files for only the specified family of devices), the initial left screen display (refer to figures 18-14 through 18-16, 18-23 through 18-25, and 18-28 through 18-30) is replaced by the secondary left screen display (refer to figures 18-17, 18-26, and 18-32). This display closely resembles the initial left screen for the utility but restricts processing to the family of devices specified. Each device in the specified family containing inactive queued files is indicated. In addition, any parameter entered for the utility is reflected in this display. Parameters not entered remain at default values.

If the FM parameter is not specified, the initial K display remains on the left screen. However, the display is modified to reflect the parameters entered. Parameters not entered remain at the default values.

The displays are modified in this manner to provide an opportunity to double-check the entry. If it is necessary to modify the entry, reenter those parameters that are to be changed. Entry of incorrect data causes an error message to appear at the bottom of the left screen K display. In this case, the portion of the entry up to the point where the error was encountered is processed; the remainder of the entry (from left to right) is ignored.

5. Initiate execution by entering:

```
K.GO.
```

When processing is complete, the left screen K display is reset to default values, the message

```
utility COMPLETE (utility is the name of the utility processed)
```

appears at the bottom of the screen, and the right screen returns to the original display.

6. Terminate use of the chosen utility and end K-display interaction by entering:

```
K.STOP.
```

An output file is generated indicating the queued files or dayfiles processed and any charges that were made to these files. Examples of output files generated by each utility are shown in figures 18-7, 18-18, 18-22, 18-27, and 18-33.

It is also possible to call each utility and specify appropriate parameters without using K displays; this is the only way to call DFLIST. You do this using a single keyboard entry in the following format.

```
X.utility(opt1=val1,opt2=val2,...,optn=valn,I=filename,PO=N)
```

Parameter	Description
utility	Any queued file or dayfile utility.
opt _i =val _i	Parameters (refer to Parameters for the Queued File Utilities later in this section).
I=filename	Name of alternate input file (optional). This file may contain K-display utility commands (GO, STOP, RESET, and so on) and parameters.
PO=N	Specifies that no K-display input is allowed. The use of the PO=N parameter and proper parameter selection allows the utility to run without additional operator intervention. This parameter is not needed and should not be entered for QLIST, LDLIST, and DFLIST.

The specified utility processes parameters on the command before the alternate input file, if any. If the utility is QLIST, LDLIST, or DFLIST or if PO=N is specified, the requested queued file/dayfile processing proceeds after all command and alternate input file parameters have been processed (unless a STOP command is encountered on the alternate input file). If PO=N is omitted from the command and the utility is not one previously mentioned, the utility requests the K display after processing the command and alternate input file (unless a STOP command is encountered on the alternate input file). The K display then reflects the status of all parameters after processing of the command and alternate input file is complete. If the utility encounters an error while processing the command or alternate input file, a request for the K display appears on the DSD B,O display (except for QLIST, LDLIST, and DFLIST). You may then activate the K display, enter the correct parameters, and continue processing. However, if errors are encountered and the utility was invoked by a nonsystem origin job, a message is issued to the dayfile, and utility processing is terminated.

Batch Input

To call the utilities from a batch job, you must be validated for system origin privileges, and the system must be in debug mode. Use the following sequence of commands to call a utility.

```
ujn.
USER, username, password, familyname.
utility, opt1=val1, opt2=val2, . . . , optn=valn, I=filename, PO=N.
```

Parameter	Description
utility	Any of the following utilities: QREC, QDUMP, QLOAD, QMOVE, QLIST, or LDLIST.
opt _i =val _i	Parameters (refer to Parameters for the Queued File Utilities later in this section).
I=filename	Name of alternate input file (optional). This file may contain K-display utility commands (GO, STOP, RESET, and so on) and parameters.
PO=N	Specifies that no K-display input is allowed. The use of the PO=N parameter and proper parameter selection allows the utility to run without additional operator intervention. This parameter is not needed and should not be entered for QLIST or LDLIST.

The DFLIST, DFTERM, QALTER, and QFTLIST utilities cannot be called from a batch job. They must be initiated from the system console.

Terminal Input

The format of terminal input for calling the queued file utilities is substantially the same as that for batch input. You must be validated for system origin privileges, and the system must be in debug mode. At login, you enter the batch subsystem, call the desired utility, and enter the appropriate parameters with the command:

```
utility, opt1=val1, opt2=val2, . . . , optn=valn, I=filename, PO=N
```

The utility, opt=val, I=filename, and PO=N parameters are the same as described for batch input.

Parameters for the Queued File Utilities

Tables 18-1 and 18-2 specify the parameters accepted by each queued file utility. (DFLIST is not included because it does not accept any parameters.) An X indicates that a parameter is accepted; a blank indicates that a parameter is not accepted. Refer to the descriptions of parameters following the table for more information.

Table 18-1. Queued File Utility Parameters (DFTERM through QFTLIST)

Parameter	DFTERM	LDLIST	QALTER	QDUMP	QFTLIST
BC		X	X	X	X
DA		X		X	
DD		X			
DF		X	X		X
DN	X		X	X	X
FC		X	X	X	X
FM	X		X	X	X
FN		X		X	
FS		X		X	
FT	X				
FU		X		X	
ID		X	X	X	X
JSN			X	X	X
L	X	X	X	X	X
LA		X	X	X	X
LD		X	X	X	X
LO			X		X
MI		X		X	
NAL			X		
NDC			X		
NDF			X		
NEC			X		
NF		X		X	
NFC			X		
NID			X		
NLD			X		
NM	X				
NPR			X		
NRC			X		
NUN			X		
OP	X		X		
PO	X		X	X	X
RB		X	X	X	X
SC		X			
TF				X	
TP				X	
TUI		X		X	
UA		X	X	X	X
UI		X		X	
UJN			X		X
UN			X		X

Table 18-2. Queued File Utility Parameters (QLIST through QREC)

Parameter	QLIST	QLOAD	QMOVE	QREC
BC	X	X	X	X
DA	X	X	X	X
DD		X	X	
DF		X	X	
DN	X		X	X
FC	X	X	X	X
FM	X		X	X
FN		X		
FS	X	X	X	X
FT				
FU	X	X	X	X
ID	X	X	X	X
JSN	X	X	X	X
L	X	X	X	X
LA	X	X	X	X
LD	X	X	X	X
LO				
MI	X	X	X	X
NAL				
NDC				
NDF				
NEC				
NF		X		
NFC				
NID				
NLD				
NM				
NPR				
NRC				
NUN				
OP		X	X	X
PO		X	X	X
RB	X	X	X	X
SC		X		
TF		X	X	X
TP			X	
TUI	X	X	X	X
UA	X	X	X	X
UI	X	X	X	X
UJN				
UN				

Description of Queued File Utility Parameters

Parameter	Description
BC=dt	Queued files destined for local batch devices can be selected for processing by disposition type (dt).
dt	Description
ALL	Files of all disposition types.
IN	Input files.
LR	580-12 printer files.
LS	580-16 printer files.
LT	580-20 printer files.
LX	5870 printer files.
LY	5970 printer files.
NONE	No files.
PH,PU	Punch files.
PL	Plot files.
PR	Files that can print on any printer.
PRINT	All printer files.
P2	512 printer files.
SF	Installation-defined special files; cannot be used with QALTER or QFTLIST.
WT	Queued files with a wait disposition.

The BC entry can be specified more than once, with each successive entry reversing the previously established condition. By default, all disposition types are selected; so any entry actually clears the automatic selection. For example, BC=PU means local batch-destined punch files are not to be processed. However, a second BC=PU entry reenables processing of local batch-destined punch files.

If BC=ALL or BC=NONE is specified, processing of all disposition types is either enabled or disabled, respectively. If all disposition types are currently enabled, entering BC=ALL has no effect. Likewise, entering BC=NONE has no effect if all disposition types are disabled.

Parameter	Description
DA=yymmdd	Processing date in the form year month day. If one date is specified, only queued files created on that day are processed. If two dates, separated by a hyphen, are specified (for example, 820130–820213), all queued files created within the specified range (including the end dates) are processed. If DA=ALL is specified, all queued files are processed. If this parameter is omitted with utilities other than QLIST, queued files created five days prior to the current date are processed. If this parameter is omitted with QLIST, all inactive queued files, regardless of creation date, are processed.
DD=dn	Destination device to which files are to be loaded or moved. The DF parameter must be specified before the DD parameter. If DF is specified, DD must also be specified. Since QMOVE requires DF, DD must also be specified.
DF=familyname	Family of devices to which files are to be loaded or moved. With QMOVE, this parameter must be specified. With QALTER or QFTLIST, DF must be used with the UN parameter. The utility will process files belonging to the remote terminal identified by these parameters.
DN=dn	The 1- or 2-digit device number (0 through 77 ₈). With utilities other than DFTERM, this parameter specifies the device to be processed, dumped, or loaded. The FM parameter must be specified before the DN parameter. Default is all devices. With DFTERM, this parameter specifies the device on which the inactive dayfile resides or on which the new dayfile will reside if the active dayfile is terminated. Default is the device on which the current dayfile resides. If the FM parameter is specified, the default for DF is zero.

Parameter	Description								
FC=fc	Forms code for output files considered for processing. This parameter does not deselect input files.								
	<table border="1"> <thead> <tr> <th>fc</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ALL</td> <td>All files are considered for processing regardless of forms codes.</td> </tr> <tr> <td>fc₁/fc₂/fc₃</td> <td>fc_i is 2 alphanumeric characters or **, indicating null forms codes. Up to three forms codes or ** can be specified; files with the specified forms code(s) are considered for processing.</td> </tr> <tr> <td>fc₁-fc₂</td> <td>fc_i is 2 alphanumeric characters or **. The 6-bit display code value of fc₁ must be less than or equal to the 6-bit display code value of fc₂. fc₂ can be ** only if fc₁ is also **. Files in the range fc₁ through fc₂ are considered for processing; if fc₁ equals **, files with no forms code are (also) considered for processing.</td> </tr> </tbody> </table>	fc	Description	ALL	All files are considered for processing regardless of forms codes.	fc ₁ /fc ₂ /fc ₃	fc _i is 2 alphanumeric characters or **, indicating null forms codes. Up to three forms codes or ** can be specified; files with the specified forms code(s) are considered for processing.	fc ₁ -fc ₂	fc _i is 2 alphanumeric characters or **. The 6-bit display code value of fc ₁ must be less than or equal to the 6-bit display code value of fc ₂ . fc ₂ can be ** only if fc ₁ is also **. Files in the range fc ₁ through fc ₂ are considered for processing; if fc ₁ equals **, files with no forms code are (also) considered for processing.
fc	Description								
ALL	All files are considered for processing regardless of forms codes.								
fc ₁ /fc ₂ /fc ₃	fc _i is 2 alphanumeric characters or **, indicating null forms codes. Up to three forms codes or ** can be specified; files with the specified forms code(s) are considered for processing.								
fc ₁ -fc ₂	fc _i is 2 alphanumeric characters or **. The 6-bit display code value of fc ₁ must be less than or equal to the 6-bit display code value of fc ₂ . fc ₂ can be ** only if fc ₁ is also **. Files in the range fc ₁ through fc ₂ are considered for processing; if fc ₁ equals **, files with no forms code are (also) considered for processing.								
	A subsequent FC entry replaces any FC entry previously made. Default is to consider all files for processing regardless of forms codes.								
FM=familyname	Name of family to be processed. Default is to process all queued files on all devices in all families. With DFTERM, this parameter specifies the family of devices on which the inactive dayfile resides or on which the new dayfile will reside if the active dayfile is terminated. If an active dayfile is terminated, default is the same family as that in which the dayfile being terminated resides.								
FN=filename	File name of dump or load file. Default is FN=QFILES.								
FS=x	File size in PRUs. The value can be entered as a single file size (FS=100) or as a range of file sizes (FS=10-400). $x \leq 777777_8$. Default is to consider all files for processing regardless of file size.								

Parameter	Description										
FT=t	Type of dayfile to be terminated by the DFTERM utility.										
	<table border="1"> <thead> <tr> <th>t</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ACCOUNT</td> <td>Account dayfile.</td> </tr> <tr> <td>DAYFILE</td> <td>System dayfile.</td> </tr> <tr> <td>ERRLOG</td> <td>Error log.</td> </tr> <tr> <td>MAINLOG</td> <td>Binary maintenance log.</td> </tr> </tbody> </table> <p>This parameter causes the FM and DN parameters to be updated to reflect the current family and device number of the dayfile specified by FT. Default is FT=DAYFILE.</p>	t	Description	ACCOUNT	Account dayfile.	DAYFILE	System dayfile.	ERRLOG	Error log.	MAINLOG	Binary maintenance log.
t	Description										
ACCOUNT	Account dayfile.										
DAYFILE	System dayfile.										
ERRLOG	Error log.										
MAINLOG	Binary maintenance log.										
FU=familyname	Name of family under which the queued files to be processed were created. Queued files created by users of one family can reside on nonremovable devices of another family. The FM parameter specifies the family of devices that will be searched for queued files created by users validated in the family specified by the FU parameter. Default is to process queued files created by users in all families.										
ID=id	The 1- or 2-digit octal number (0 through 67 ₈) specifying that only local batch-destined queued files assigned to that identifier are to be processed. This parameter does not deselect remote batch-destined files. If two identifiers, separated by a hyphen, are specified (for example, ID=6-30), any queued file with an identifier in the specified range (including the end points) is processed. Default is to process all queued files regardless of identifier.										
JSN=jsn	The 4-character job sequence name of queued files to be processed, dumped, loaded, or moved. At most, five job sequence names can be specified. If a job sequence name is specified more than once, it is removed from the list. Default is to process all queued files regardless of job sequence name.										
L=outfile	The 1- to 7-character name of the file to receive output. Default is L=OUTPUT. Specify L=0 if no output file is needed.										
LA=level	The 1- to 7-character name that specifies the lower limit of the range of access levels to process. If this parameter is specified, the UA parameter must also be specified. If neither LA nor UA is specified, the default is that all access levels are selected.										
LD=lid	The 3-character logical identifier of the mainframe with which the file is associated. LD=0 means process all files without LIDs. A subsequent LD entry replaces any LD entry previously made. Default is to process all queued files regardless of logical identifier.										

Parameter	Description																				
LO=opt	List option associated with the output file specified by the L parameter.																				
	<table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>F</td> <td>All routing information for each file.</td> </tr> <tr> <td>S</td> <td>Condensed listing of each selected queued file.</td> </tr> </tbody> </table>	opt	Description	F	All routing information for each file.	S	Condensed listing of each selected queued file.														
opt	Description																				
F	All routing information for each file.																				
S	Condensed listing of each selected queued file.																				
MI=id	The 2-character machine identifier indicating the mainframe on which the queued files to be processed currently reside. If MI=ALL is specified, any queued file residing on any mainframe is processed. The MI entry is cleared if previously specified; that is, specifying two identical MI entries is the same as not specifying any at all. If the MI entry is cleared on the K display, the id is not checked. Default is the machine identifier of the mainframe where the utility is being run.																				
NAL=level	QALTER can change the access level of a queued file. The new access level must be within the origin type limits for the file and within the device limits of the device on which the file resides. This parameter is restricted to users with security administrator privileges.																				
NDC=dt	QALTER can change a print file's disposition type while the file is in the queue. New disposition types that can be selected include the following.																				
	<table border="1"> <thead> <tr> <th>dt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>LP</td> <td>Any printer.</td> </tr> <tr> <td>LR</td> <td>580-12 printer.</td> </tr> <tr> <td>LS</td> <td>580-16 printer.</td> </tr> <tr> <td>LT</td> <td>580-20 printer.</td> </tr> <tr> <td>LX</td> <td>5870 printer.</td> </tr> <tr> <td>LY</td> <td>5970 printer.</td> </tr> <tr> <td>NONE</td> <td>No device code specified.</td> </tr> <tr> <td>PR</td> <td>Any printer.</td> </tr> <tr> <td>P2</td> <td>512 printer.</td> </tr> </tbody> </table>	dt	Description	LP	Any printer.	LR	580-12 printer.	LS	580-16 printer.	LT	580-20 printer.	LX	5870 printer.	LY	5970 printer.	NONE	No device code specified.	PR	Any printer.	P2	512 printer.
dt	Description																				
LP	Any printer.																				
LR	580-12 printer.																				
LS	580-16 printer.																				
LT	580-20 printer.																				
LX	5870 printer.																				
LY	5970 printer.																				
NONE	No device code specified.																				
PR	Any printer.																				
P2	512 printer.																				
NDF=familyname	New destination family name associated with selected output files. This parameter is meaningful only when altering files routed to or changed to remote batch. A subsequent NDF entry replaces any previous NDF entry. Default is no change in destination family name.																				

Parameter	Description																														
NEC=ex	<p>QALTER can change an output file's external characteristics. New external characteristics that can be selected include the following.</p> <p>For print files:</p> <table border="1"> <thead> <tr> <th>ex</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A4¹</td> <td>ASCII graphic 48-character set.</td> </tr> <tr> <td>A6</td> <td>ASCII graphic 63/64-character set.</td> </tr> <tr> <td>A9</td> <td>ASCII graphic 95-character set.</td> </tr> <tr> <td>B4¹</td> <td>CDC graphic 48-character set.</td> </tr> <tr> <td>B6</td> <td>CDC graphic 63/64-character set.</td> </tr> </tbody> </table> <p>For punch files:</p> <table border="1"> <thead> <tr> <th>ex</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AS</td> <td>Punch ASCII.</td> </tr> <tr> <td>PB</td> <td>Punch system binary.</td> </tr> <tr> <td>PH</td> <td>Punch 026 mode.</td> </tr> <tr> <td>P8</td> <td>Punch 80-column binary.</td> </tr> <tr> <td>P9</td> <td>Punch 029 mode.</td> </tr> </tbody> </table> <p>For plot files:</p> <table border="1"> <thead> <tr> <th>ex</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>T6</td> <td>Transparent 6 bit.</td> </tr> <tr> <td>T8</td> <td>Transparent 8 bit.</td> </tr> </tbody> </table>	ex	Description	A4 ¹	ASCII graphic 48-character set.	A6	ASCII graphic 63/64-character set.	A9	ASCII graphic 95-character set.	B4 ¹	CDC graphic 48-character set.	B6	CDC graphic 63/64-character set.	ex	Description	AS	Punch ASCII.	PB	Punch system binary.	PH	Punch 026 mode.	P8	Punch 80-column binary.	P9	Punch 029 mode.	ex	Description	T6	Transparent 6 bit.	T8	Transparent 8 bit.
ex	Description																														
A4 ¹	ASCII graphic 48-character set.																														
A6	ASCII graphic 63/64-character set.																														
A9	ASCII graphic 95-character set.																														
B4 ¹	CDC graphic 48-character set.																														
B6	CDC graphic 63/64-character set.																														
ex	Description																														
AS	Punch ASCII.																														
PB	Punch system binary.																														
PH	Punch 026 mode.																														
P8	Punch 80-column binary.																														
P9	Punch 029 mode.																														
ex	Description																														
T6	Transparent 6 bit.																														
T8	Transparent 8 bit.																														
NF=n	Decimal number of media files to skip. A media file contains all queued files that are dumped by one QDUMP operation. Each of these dumped queued files exists as a record on the media file. If two QDUMP operations are performed, the dumped queued files exist as records on two files. This parameter can be used in conjunction with the SC parameter. Default is 0.																														
NFC=fc	Two alphanumeric characters or ** (null forms code) specifying the new forms code associated with selected output files. The NFC entry is cleared if previously specified; that is, specifying two identical NFC entries is the same as not specifying any at all. Default is no change in forms code.																														

1. Not supported. Provided for NOS/BE compatibility.

Parameter	Description
NID=id	The 1- or 2-digit number (0 through 67 ₈) specifying the new file identifier associated with selected output files. This parameter is meaningful only when altering files routed to or changed to local batch. The NID entry is cleared if previously specified; that is, specifying two identical NID entries is the same as not specifying any at all. Default is no change in file identifier.
NLD=lid	The 3-character logical identifier specifying the new logical identifier associated with selected files. NLD=0 will remove the LID from selected files. A subsequent NLD entry replaces any NLD entry previously made. Default is no change in logical identifier.
NM=filename	The 1- to 5-character suffix for the name of the direct access permanent file on which DFTERM writes the terminated dayfile. DFTERM adds a 2-character prefix indicating the type of dayfile being terminated (AC, DF, ML, or ER). If NM is not specified or the specified name is a duplicate, DFTERM automatically assigns a name for the file as follows: <ul style="list-style-type: none"> • The first two characters indicate the type of dayfile being terminated (AC, DF, ML, or ER). • The third character is a sequence number (A through 9). • The next two characters indicate the month. • The last two characters indicate the day of the month.
NPR=p	The 1- to 4-digit number (0 through 7777 ₈) specifying the new queue priority associated with selected output files. The NPR entry is cleared if previously specified; that is, specifying two identical NPR entries is the same as not specifying any at all. Default is no change in queue priority.
NRC=c	The 1- or 2-digit number (0 through 77 ₈) specifying the new repeat count associated with selected output files. The NRC entry is cleared if previously specified; that is, specifying two identical NRC entries is the same as not specifying any at all. Default is no change in repeat count.
NUN=username	New destination user name associated with selected output files. This parameter is meaningful only when altering any files routed to or changed to remote batch. The NUN entry is cleared if previously specified; that is, specifying two identical NUN entries is the same as not specifying any at all. Default is no change in destination user name.

Parameter	Description																										
OP=opt	<p>Processing option specifying the function to be performed by each utility.</p> <p>With QREC, opt can be any of the following. Default is OP=RI.</p> <table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>DI</td> <td>Selected active queued files are made inactive and the remaining active queued files are ignored.</td> </tr> <tr> <td>PI</td> <td>Selected inactive queued files are purged, and the remaining inactive queued files are ignored.</td> </tr> <tr> <td>RI</td> <td>Selected inactive queued files are activated (requeued), and the remaining inactive queued files are ignored.</td> </tr> <tr> <td>RP</td> <td>Selected inactive queued files are activated (requeued), and the remaining inactive queued files are purged.</td> </tr> </tbody> </table> <p>With QALTER, opt specifies the destination to which the selected queued files should be changed; it can be one of the following. Default is OP=NC.</p> <table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>BC</td> <td>Change to local batch.</td> </tr> <tr> <td>NC</td> <td>Do not change file destination.</td> </tr> <tr> <td>PR</td> <td>Purge files.</td> </tr> <tr> <td>RB</td> <td>Change to remote batch.</td> </tr> </tbody> </table> <p>With DFTERM, opt specifies whether active or inactive dayfiles will be terminated. With QLOAD and QMOVE, opt specifies whether the loaded or moved queued files will be active or inactive. With these utilities, opt can be one of the following. Default is OP=A.</p> <table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Active queued files/dayfiles are specified.</td> </tr> <tr> <td>I</td> <td>Inactive queued files/dayfiles are specified.</td> </tr> </tbody> </table>	opt	Description	DI	Selected active queued files are made inactive and the remaining active queued files are ignored.	PI	Selected inactive queued files are purged, and the remaining inactive queued files are ignored.	RI	Selected inactive queued files are activated (requeued), and the remaining inactive queued files are ignored.	RP	Selected inactive queued files are activated (requeued), and the remaining inactive queued files are purged.	opt	Description	BC	Change to local batch.	NC	Do not change file destination.	PR	Purge files.	RB	Change to remote batch.	opt	Description	A	Active queued files/dayfiles are specified.	I	Inactive queued files/dayfiles are specified.
opt	Description																										
DI	Selected active queued files are made inactive and the remaining active queued files are ignored.																										
PI	Selected inactive queued files are purged, and the remaining inactive queued files are ignored.																										
RI	Selected inactive queued files are activated (requeued), and the remaining inactive queued files are ignored.																										
RP	Selected inactive queued files are activated (requeued), and the remaining inactive queued files are purged.																										
opt	Description																										
BC	Change to local batch.																										
NC	Do not change file destination.																										
PR	Purge files.																										
RB	Change to remote batch.																										
opt	Description																										
A	Active queued files/dayfiles are specified.																										
I	Inactive queued files/dayfiles are specified.																										

Parameter	Description																
PO=opt	Processing option specifying the function to be performed.																
	<table border="1"> <thead> <tr> <th>opt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>E</td> <td> <p>Errored file processing status is toggled. Default conditions for queued file utilities include:</p> <p> QDUMP Dump file.</p> <p> QLOAD Do not load file.</p> <p> QMOVE Do not move file.</p> </td> </tr> <tr> <td>F</td> <td>Entries are made in FOT for inactive families and the files are queued.</td> </tr> <tr> <td>I</td> <td>Activate inactive input queued files when the owner family is not active.</td> </tr> <tr> <td>L</td> <td>Entries are made in the LID table for inactive LIDs and the files are queued.</td> </tr> <tr> <td>N</td> <td>Only command directives are processed. The K display is available to monitor utility progress, but no K-display input is allowed. The N option is implied on all nonsystem origin jobs.</td> </tr> <tr> <td>R</td> <td>The dump or load file is rewound before processing begins.</td> </tr> <tr> <td>U</td> <td>Unconditionally dump, load, recover, or move files (the only restrictions are device security access limits).</td> </tr> </tbody> </table>	opt	Description	E	<p>Errored file processing status is toggled. Default conditions for queued file utilities include:</p> <p> QDUMP Dump file.</p> <p> QLOAD Do not load file.</p> <p> QMOVE Do not move file.</p>	F	Entries are made in FOT for inactive families and the files are queued.	I	Activate inactive input queued files when the owner family is not active.	L	Entries are made in the LID table for inactive LIDs and the files are queued.	N	Only command directives are processed. The K display is available to monitor utility progress, but no K-display input is allowed. The N option is implied on all nonsystem origin jobs.	R	The dump or load file is rewound before processing begins.	U	Unconditionally dump, load, recover, or move files (the only restrictions are device security access limits).
opt	Description																
E	<p>Errored file processing status is toggled. Default conditions for queued file utilities include:</p> <p> QDUMP Dump file.</p> <p> QLOAD Do not load file.</p> <p> QMOVE Do not move file.</p>																
F	Entries are made in FOT for inactive families and the files are queued.																
I	Activate inactive input queued files when the owner family is not active.																
L	Entries are made in the LID table for inactive LIDs and the files are queued.																
N	Only command directives are processed. The K display is available to monitor utility progress, but no K-display input is allowed. The N option is implied on all nonsystem origin jobs.																
R	The dump or load file is rewound before processing begins.																
U	Unconditionally dump, load, recover, or move files (the only restrictions are device security access limits).																

Parameter	Description																																
RB=dt	Queued files destined for remote batch devices can be selected for processing by disposition type (dt).																																
	<table border="1"> <thead> <tr> <th>dt</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ALL</td> <td>Files of all disposition types.</td> </tr> <tr> <td>IN</td> <td>Input files.</td> </tr> <tr> <td>LR</td> <td>580-12 printer files.</td> </tr> <tr> <td>LS</td> <td>580-16 printer files.</td> </tr> <tr> <td>LT</td> <td>580-20 printer files.</td> </tr> <tr> <td>LX</td> <td>5870 printer files.</td> </tr> <tr> <td>LY</td> <td>5970 printer files.</td> </tr> <tr> <td>NONE</td> <td>No files.</td> </tr> <tr> <td>PH,PU</td> <td>Punch files.</td> </tr> <tr> <td>PL</td> <td>Plot files.</td> </tr> <tr> <td>PR</td> <td>Files that can print on any printer.</td> </tr> <tr> <td>PRINT</td> <td>All printer files.</td> </tr> <tr> <td>P2</td> <td>512 printer files.</td> </tr> <tr> <td>SF</td> <td>Installation-defined special files; cannot be used with QALTER or QFTLIST.</td> </tr> <tr> <td>WT</td> <td>Queued files with a wait disposition.</td> </tr> </tbody> </table>	dt	Description	ALL	Files of all disposition types.	IN	Input files.	LR	580-12 printer files.	LS	580-16 printer files.	LT	580-20 printer files.	LX	5870 printer files.	LY	5970 printer files.	NONE	No files.	PH,PU	Punch files.	PL	Plot files.	PR	Files that can print on any printer.	PRINT	All printer files.	P2	512 printer files.	SF	Installation-defined special files; cannot be used with QALTER or QFTLIST.	WT	Queued files with a wait disposition.
dt	Description																																
ALL	Files of all disposition types.																																
IN	Input files.																																
LR	580-12 printer files.																																
LS	580-16 printer files.																																
LT	580-20 printer files.																																
LX	5870 printer files.																																
LY	5970 printer files.																																
NONE	No files.																																
PH,PU	Punch files.																																
PL	Plot files.																																
PR	Files that can print on any printer.																																
PRINT	All printer files.																																
P2	512 printer files.																																
SF	Installation-defined special files; cannot be used with QALTER or QFTLIST.																																
WT	Queued files with a wait disposition.																																

The RB entry can be specified more than once, with each successive entry reversing the previously established condition. By default, all disposition types are selected; so any entry actually clears the automatic selection. For example, RB=PU means remote batch-destined punch files are not to be processed. However, a second RB=PU entry reenables processing of remote batch-destined punch files.

If RB=ALL or RB=NONE is specified, processing of all disposition types is either enabled or disabled, respectively. If all disposition types are currently enabled, entering RB=ALL has no effect. Likewise, entering RB=NONE has no effect if all disposition types are disabled.

Parameter	Description								
SC=n	Decimal number of queued files (and hence records on the dump file) to be skipped before the queue selection begins. Skipping begins at the point on the dump file indicated by the NF parameter. This parameter allows a restart load to begin from the last aborted queue loaded by QLOAD. Default is SC=0.								
TF=familyname	Name of family for remote batch-destined files to be selected. By using the TF parameter with the TUI parameter, files destined for a particular remote batch terminal can be selected. Local batch-destined files are not deselected by this parameter. Default is all families selected.								
TP=t	Type of files to move or dump. Default is TP=ALL.								
	<table border="1"> <thead> <tr> <th>t</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Active files.</td> </tr> <tr> <td>ALL</td> <td>Active and inactive files.</td> </tr> <tr> <td>I</td> <td>Inactive files.</td> </tr> </tbody> </table>	t	Description	A	Active files.	ALL	Active and inactive files.	I	Inactive files.
t	Description								
A	Active files.								
ALL	Active and inactive files.								
I	Inactive files.								
TUI=userindex	The 1- to 6-digit value plus a radix that specifies the destination terminal user index for remote batch destined output files. If one userindex is specified (xxxxxx), only that user index is processed. If a range is specified (xxxxxx-yyyyyy), all indexes in the specified range are processed. If TUI=0, no terminal user index is specified. Default is TUI=0.								
UA=level	The 1- to 7-character name that specifies the upper limit of the range of access levels to process. If this parameter is specified, the LA parameter must also be specified. If neither LA nor UA is specified, the default is that all access levels are selected.								
UI=userindex	User index under which queued files to be processed were created. If two user indexes, separated by a hyphen, are specified (for example, 75-162), all queued files created by all users having indexes within this range (including the end points) are processed. Default is to process all queued files regardless of user index.								

Parameter	Description
UJN=ujn	The 1- to 7-character user job name specified on Job commands associated with queued files being listed or altered. Only one name can be specified. A subsequent UJN entry replaces any UJN entry previously made. This parameter in conjunction with the JSN parameter uniquely describes which files are to be processed. Default is no UJN specified.
UN=username	Destination user name within the family specified by the DF option. UN and DF must be specified together. The utility (QALTER or QFTLIST) will process files belonging to the remote terminal identified by these parameters. The UN entry is cleared if previously specified; that is, specifying two identical UN entries is the same as not specifying any at all.

NOTE

If the currently displayed value for JSN, NFC, NID, NPR, NRC, NUN, or UN is reentered, the value for the specified parameter is cleared.

File Preassignment

When using the LDLIST, QDUMP, and QLOAD utilities, the dump or load file must be preassigned. This can be done by using a LABEL or an ASSIGN command. For example, the following commands can be used.

```
LABEL, QFILES, VSN=DMP1, D=HY.
```

or

```
ASSIGN, MS, QFILES.
```

If no dump file is present, QDUMP will write on a local mass storage file; LDLIST and QLOAD will attempt to read an empty file and terminate.

Active and Inactive Queued Files

A queued file can be active or inactive. All queued files are originally active upon entering a queue. A queued file is deactivated (made inactive) when its entry is removed from the queued file table (QFT) and a corresponding entry is created in the inactive queued file table (IQFT) on the mass storage device where the queued file resides. An inactive queued file is activated when the entry from the IQFT is removed and a corresponding entry is created in the QFT. Inactive queued files are not considered for processing (other than by the queued file utilities), nor do they appear on any DSD display.

Three queued file utilities can deactivate or activate queued files (QLOAD, QMOVE, and QREC). QLOAD and QMOVE each transfer queued files from either tape or disk to a disk device. Upon completion of the transfer, the utility leaves the queued files active or inactive according to the parameters you have specified. QREC does not perform any file transfer but only adjusts table entries as described in the preceding paragraph. These three utilities have uses other than deactivation and activation of queued files (refer to the description of the appropriate utility and parameters).

Queued files are also deactivated when they are recovered during a level 0 deadstart. Such files can be activated automatically by an IPRDECK entry that invokes QREC during a level 0 deadstart. Refer to section 3, Deadstart Decks, for further information on IPRDECK.

Queued File Utility Routines

The queued file utility routines are described next. The parameters available for each of the routines are listed in table 18-1. The K displays (if any) for each utility are shown and contain lists of applicable parameters. Output formats for each utility are also shown.

DFLIST

DFLIST generates a printer listing of all permanent files created by the DFTERM utility.

DFLIST can be initiated only from the system console directly through this command:

```
X.DFLIST.
```

The output generated by DFLIST is the same as that produced by the LIST command available under the DFTERM utility (refer to figure 18-7).

DFTERM

DFTERM terminates an active or inactive dayfile and retains it as a direct access permanent file on user name SYSTEMX for later interrogation or processing. When an active dayfile (that is, the current system dayfile, account dayfile, error log, or binary maintenance log) is terminated, information in the central memory buffer for that dayfile is written to mass storage to be included with the permanent file, and a new active dayfile is started. The new dayfile can reside on the same device, or a new device can be specified.

Terminating an inactive dayfile has no effect on the currently active dayfiles. Inactive dayfiles are not used by the system. Furthermore, the presence of an inactive dayfile in the system is possible only under unusual conditions. For example, if the system is deadstarted and the device that previously contained the account dayfile is turned off, a new account dayfile is started on another device. Two devices in the system now contain account dayfiles. If both devices are turned on when the system is next deadstarted, two account dayfiles are recovered. The most recent account dayfile is made active and is used by the system. The remaining account dayfile is made inactive.

The DFTERM utility can be initiated only from the system console through a K display or directly through the DFTERM command. No batch or terminal input is allowed.

When DFTERM is selected at the system console, the DFTERM K display appears on the left screen. This display contains a list of all applicable options associated with DFTERM, their default values, and a short description of each (refer to figure 18-6).

A list of the permanent dayfiles is displayed on the right screen when you enter this command:

K.LIST.

When this command is processed, the right screen K display is automatically replaced with the DFTERM dayfile list K display (refer to figure 18-7). The information presented in this display is the same as that generated by the DFLIST utility.

If there are more files than can be displayed on one screen, a message appears at the bottom of the screen indicating that more files exist. When the LIST command is entered, a complete list of permanent dayfiles is written to the output file specified by the DFTERM L parameter; thus, a complete printer listing is provided when the DFTERM utility is terminated. If it is necessary to obtain the output listing immediately, enter this command:

K.OUT.

When this command is processed, the initial right screen K display (a list of the queued file utility commands) is automatically returned. However, if the OUT command is not entered, the original right screen K display can be returned by entering the CLEAR command.

After DFTERM processing is completed, an output file indicating the results of this processing is generated if the LIST command has been entered at any time during DFTERM operations. (Refer to figure 18-7, which shows essentially the same format.) If the LIST command was not entered, no output file is generated.

```

K. DFTERM.          jsn
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                        MID=AA   NOS version

                        DFTERM

      ---ACTIVE DAYFILES----      --INACTIVE DAYFILES---
      TYPE      FAMILY   DN      TYPE      FAMILY   DN
      DAYFILE   SYST06   2B
      ACCOUNT   SYST06   2B
      ERRLOG    SYST06   2B
      MAINLOG   SYST06   2B

      OPTION      DESCRIPTION

      FT = DAYFILE  FILE TYPE TO BE TERMINATED
                    (DAYFILE, ACCOUNT, ERRLOG, MAINLOG).
      FM = SYST06   FAMILY NAME OF NEW DEVICE (1-7 CHARACTERS).
      DN = 2B       DEVICE NUMBER OF NEW DEVICE (0-77B).
      OP = A        OPTION (A=ACTIVE or I=INACTIVE).
      NM = DAYF     PERMANENT FILE NAME OVERRIDES AUTOMATIC NAMING
                    (1-5 CHARACTERS).
      L = OUTPUT    OUTPUT FILE (1-7 CHARACTERS).
    
```

Figure 18-6. DFTERM K Display (Left Screen)

K. DFTERM.

jsn

PERMANENT DAYFILE CATALOG

CATALOG OF PERMANENT DAYFILES (377777). 85/01/16 09.25.43

NUM	TYPE	FM/PN	DN	PFNAME	DATE	TIME	LENGTH
1	DAYFILE	SYST06	02	DFA0116	85/01/16	08.00.04	20
2	DAYFILE	SYST06	02	DFB0116	85/01/16	08.00.34	1
3	DAYFILE	SYST06	02	DFC0116	85/01/16	08.01.14	1
4	DAYFILE	SYST06	02	DFD0116	85/01/16	08.01.54	1
5	DAYFILE	SYST06	02	DFE0116	85/01/16	08.02.33	1
6	DAYFILE	SYST06	02	DFF0116	85/01/16	08.03.13	1
7	DAYFILE	SYST06	02	DFG0116	85/01/16	08.03.53	1
8	DAYFILE	SYST06	02	DFH0116	85/01/16	08.04.33	1
9	DAYFILE	SYST06	02	DFI0116	85/01/16	08.05.12	1
10	DAYFILE	SYST06	02	DFJ0116	85/01/16	08.05.52	1
11	DAYFILE	SYST06	02	DFK0116	85/01/16	08.06.45	1
12	DAYFILE	SYST06	02	DFL0116	85/01/16	08.07.34	1
13	DAYFILE	SYST06	02	DFM0116	85/01/16	08.08.13	1
14	DAYFILE	SYST06	02	DFN0116	08/01/16	08.08.30	1

MORE PERMANENT DAYFILES.
 ENTER *OUT* FOR A FULL LIST.

Figure 18-7. DFTERM Dayfile List (Right Screen)

LDLIST

LDLIST generates a printer listing of queued files present on a QDUMP dump tape.

Since LDLIST requires a QDUMP dump tape to be preassigned, it should be entered from DIS or in a batch or terminal job in the form

```
LDLIST, opt1=val1, opt2=val2, . . . , optn=valn.
```

opt_i=val_i are parameters available for LDLIST.

The output generated by LDLIST is the same as that produced by QLOAD (refer to figure 18-22).

QALTER

QALTER displays, lists, and/or alters routing and other information about active queued files. The utility selects files for processing according to criteria you specify. QALTER purges selected files from the system when you specify OP=PR.

You may alter the following information associated with the selected files.

- File destination.
- Access level (requires security administrator privileges).
- Print file disposition type (device code).
- External characteristics.
- Destination LID.
- Destination family (remote batch-destined files).
- Destination user name (remote batch-destined files).
- File identifier (local batch-destined files).
- Forms code.
- Priority of a queued file.
- Repeat count.
- Disposition type.

When QALTER is selected at the system console, the QALTER/QFTLIST K display (3 pages) appears on the left screen. This display contains a list of all applicable parameters associated with QALTER, their default values, and a short description of each (refer to figures 18-8 through 18-10). Follow this procedure to alter active queued files:

1. Enter the selection criteria specifying which files are to be altered (parameters FM, DN, DF, UN, ID, FC, UJN, LD, LA, UA, BC, RB, and JSN).
2. Ensure that the desired files are selected by inspecting the list of files produced by the K-display commands LIST or OUT (refer to figures 18-11 through 18-13). If the file produced by OUT is to be other than OUTPUT or contain a detailed listing of the selected files, specify such with the L and/or LO parameters before entering K.OUT.
3. Enter the new information to be associated with the selected files (parameters OP, NAL, NDC, NEC, NLD, NDF, NUN, NID, NFC, NPR, and NRC). QALTER changes only information specified via parameter entry.
4. Enter K.GO. QALTER makes the requested changes, issues the message *QALTER COMPLETE.*, and resets the K-display parameters to their default values.
5. Repeat steps 1 through 4 to alter additional files, or enter K.STOP. to terminate QALTER activity. After QALTER activity has completed, an output file is generated indicating the changes made to the queued files (see figure 18-13).

Example:

Use the QALTER utility to divert output files from the local printers to a remote terminal. The following command alters the routing of all batch output files with a forms code of AD. These files are routed to a remote terminal with family name FAM1 and user name USRNAM. The BC and RB parameters explicitly disable input file selection.

```
QALTER, PO=N, BC=IN, RB=IN, FC=AD, OP=RB, NDF=FAM1, NUN=USRNAM.
```

The QFTLIST utility can be used in conjunction with QALTER to ensure that the files were properly altered. The following command generates a list of all active queued files belonging to the remote terminal with family name FAM1 and user name USRNAM.

```
QFTLIST, PO=N, LO=F, DF=FAM1, UN=USRNAM, BC=NONE.
```

```

K. QALTER.          jsn
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version

QALTER                                PAGE 1 OF 3

      OPTION                                DESCRIPTION

FM = ALL          FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL          DEVICE NUMBER (0-77B, *ALL*).
DF = SYS06        DESTINATION FAMILY (1-7 CHARACTERS).
UN = DEST01       DESTINATION USER (1-7 CHARACTERS).
L = OUTPUT        FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
LO = S            LIST OPTION (S=SHORT, F=FULL).

ID = OB          ID RANGE OF FILES TO PROCESS (0-67B).
      67B
FC = 11/22/33    FORMS CODE (2 CHARACTERS OR **).
UJN = JOB02       USER JOB NAME (1-7 CHARACTERS).
LD = M06          DESTINATION LID (3 CHARACTERS).
LA = LVL2         LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA = LVL6         UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
    
```

Figure 18-8. QALTER/QFTLIST K Display (Left Screen) (Page 1 of 3)

```

QALTER                                PAGE 2 OF 3

      OPTION                                DESCRIPTION

BC = (DT)         DISPOSITION TYPE FOR LOCAL BATCH DEVICES.
RB = (DT)         DISPOSITION TYPE FOR REMOTE BATCH DEVICES.

                      S I P P P P L L L L W
                      F N U L R 2 R S T X Y T
BC  * * * * - - - - - *
RB  * * - * * * - * * * *
* = OPTION SELECTED.
OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.

JSN = AAAA        JOB SEQUENCE NAME (1-4 CHARACTERS).
      AAAB        JSN WILL BE CLEARED IF PREVIOUSLY ENTERED.
      AAAC
    
```

Figure 18-9. QALTER/QFTLIST K Display (Left Screen) (Page 2 of 3)

QALTER	PAGE 3 OF 3 ²	OPTION
OP = NC		OPTION (BC, RB, NC, PR).
NAL = LVL3		NEW ACCESS LEVEL (1-7 CHARACTERS).
NDC = NONE		NEW DEVICE CODE (2 CHARACTERS OR NONE).
NEC = A9		NEW EXTERNAL CHARACTERISTICS (2 CHARACTERS).
NLD = M06		NEW DESTINATION LID (3 CHARACTERS).
NDF = SYST06		NEW DESTINATION FAMILY (1-7 CHARACTERS).
NUN = SYSTEM		NEW DESTINATION USER (1-7 CHARACTERS).
NID = 67B		NEW FILE IDENTIFIER (0-67B).
NFC = **		NEW FORMS CODE (2 CHARACTERS OR **).
NPR = 7777B		NEW PRIORITY FOR OUTPUT (0-7760B).
NRC = 30D		NEW REPEAT COUNT (0-31D).

Figure 18-10. QALTER/QFTLIST K Display (Left Screen) (Page 3 of 3)

2. Page 3 of 3 of this K display is present for QALTER, but not QFTLIST.

ACTIVE QUEUE FILE LIST									
ORD	JSNS	DT	LID	FAMILY	USERNAM	TUI/ID	FC	EC	LENGTH
1	AAAOS	PL				55		T6	4
2	AAAPB	LS				22		A9	6
3	AAAQR	LP		SYST01	ABC1234	1234	AX	LP	41
4	AAARB	LS				22		A9	5
5	AAASB	LS				22		A9	4
6	AAATR	LP		SYST01	ABC1234	1234	AX	LP	41
7	AAAUR	LP		SYST01	ABC1234	1234	AX	LP	41
10	AAAVB	LP				55		LP	7
11	AAAWR	LP		SYST01	ABC1234	1234	AX	LP	41
12	AAAXB	LS	M01			22		A9	5
13	AAAYR	LP		SYST01	ABC1234	1234	AX	LP	41
16	AABBR	LP	LBK	SYST01	ABC1234	1234	AX	LP	41
17	AABCB	LS				22		A9	3
20	AABDB	LS				22		A9	2

END OF DISPLAY.

Figure 18-11. QALTER/QFTLIST LIST Command Display (Right Screen)

ACTIVE QUEUE LIST	
JSN = AAAD	ORDINAL = 7
SERVICE = SYSTEM	QUEUE = PRINT
DESTINATION	CREATION
FAMILY =	FAMILY = SYST03
LID = M01	LID = M03
USER =	USER =
TUI/ID = 0	USER INDEX = 0
FORMS = AX	USER JOBNAME = DIS
DISPOSITION CODE = LP	DATE = 85/01/16
EXTERNAL CHAR. = LP	LENGTH = 2
INTERNAL CHAR. = DIS	REPEAT = 0
RESIDENCE	
FAMILY = SYST01	INTERRUPT = NO
DEVICE = 2	PRIORITY = 7110
ACCESS LEVEL =	

Figure 18-12. QALTER/QFTLIST LIST=qft Command Display (Right Screen)

CURRENT				DESTINATION				REASON				RESIDENCE													
*****				*****				*****				*****													
QALTER - ALTER/PURGE ACTIVE QUEUED FILES. yy/mm/dd. hh.mm.ss. PAGE 1																									
OPTION = NC NEW FORMS CODE = --N/A-- NEW PRIORITY = 5103 NEW REPEAT COUNT = --N/A--				NEW DESTINATION FAMILY NAME = --N/A-- NEW DESTINATION USER NAME = --N/A-- NEW DISPOSITION CODE = --N/A-- NEW DESTINATION LID = --N/A--				NEW EXTERNAL CHAR = --N/A-- NEW FILE ID = --N/A-- NEW ACCESS LEVEL = --N/A--																	
1.	JSN	= AACG	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LP	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= NO	PRIORITY = 55	ORDINAL = 3	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 13	ACCESS LV=
2.	JSN	= AACI	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LP	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= NO	PRIORITY = 55	ORDINAL = 5	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 5	ACCESS LV=
3.	JSN	= AACJ	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LT	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= NO	PRIORITY = 56	ORDINAL = 6	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 15	ACCESS LV=
4.	JSN	= AACW	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LP	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= YES	PRIORITY = 56	ORDINAL = 23	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 13	ACCESS LV=
5.	JSN	= AACY	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LP	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= YES	PRIORITY = 56	ORDINAL = 25	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 5	ACCESS LV=
6.	JSN	= AACZ	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LP	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= YES	PRIORITY = 56	ORDINAL = 26	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 13	ACCESS LV=
7.	JSN	= AADB	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LT	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= YES	PRIORITY = 56	ORDINAL = 30	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 15	ACCESS LV=
8.	JSN	= AADF	FAMILY = SYS964	FORMS =	FAMILY = SYS964	USR JOBNM= DIS	FAMILY = SYST64	USER = BMS4247	DEVICE = LP	USER = BMS4247	DATE = 83/07/07	DEVICE = 2	QUEUE = PRINT	TUI/ID = 4247	EXT.CHR. = LP	USR INDX = 4247	INTERRUPT= NO	PRIORITY = 57	ORDINAL = 34	DLID =	INT.CHR. = DIS	SLID =	REPEAT = 0	LENGTH = 13	ACCESS LV=

Figure 18-13. QALTER Printer Listing/Output File

QDUMP

QDUMP dumps selected queued files from a single device, a family of devices, or all devices on the system. These queued files can be dumped to a tape or disk. When active queued files are dumped, the QFT is searched to obtain the proper file. When inactive queues are dumped, the IQFT is searched. QDUMP also provides a listing of all files dumped with information about each file processed.

For a secured system, QDUMP determines the maximum range of access levels that can be dumped. If the LA and UA parameters have been used to select access level limits, these limits will be used. If no access level limits were selected, QDUMP uses the device limits determined by taking the lowest lower access limit and the highest upper access limit of all the devices to be processed. If this range of access levels is not within the system access level limits, QDUMP issues a diagnostic message.

When QDUMP is selected at the system console, the initial QDUMP K display (3 pages) appears on the left screen. This display contains a list of all applicable options associated with QDUMP, their default values, and a short description of each (refer to figures 18-14 through 18-16).

If the FM parameter is specified, indicating dumping of queued files from a specified family of devices, the initial QDUMP left screen K display is replaced by the secondary QDUMP K display (refer to figure 18-17). This display closely resembles the initial QDUMP display but restricts QDUMP processing to the family of devices specified. Each device in the specified family containing inactive queued files is indicated. In addition, any parameter entered for the QDUMP utility is reflected in this display. Parameters not entered remain at default values. Only the first screen of the secondary QDUMP K display is shown, since the second and third screens have the same format as the corresponding screens of the initial QDUMP K display.

After QDUMP processing has been completed and K-display interaction has been terminated by

K. STOP.

a QDUMP output file is generated, indicating all files dumped. Figure 18-18 is an example of this output.

```

K. QDUMP.          jsn
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                    MID=AA   NOS version

QDUMP                                     PAGE 1 OF 3

INACTIVE QUEUES RESIDE ON FOLLOWING FAMILIES.
    SYST06   PACKV2   FEATURE
              (ERRORED FILE DUMPING SET)

    OPTION                                DESCRIPTION

MI = 06          MACHINE ID (1-2 CHARACTERS, *ALL*).
FM = ALL        FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL        DEVICE NUMBER (0-77B, *ALL*).
DA = 85/01/11   DUMP DATE RANGE (YYMMDD, *ALL*).
              85/01/16

L = OUTPUT      FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
FN = QFILES     DUMP FILE NAME (1-7 CHARACTERS).
NF = OD         NUMBER OF MEDIA FILES TO SKIP.
TP = ALL        TYPE (A=ACTIVE, I=INACTIVE, *ALL*).

```

Figure 18-14. Initial QDUMP K Display (Left Screen) (Page 1 of 3)

```

QDUMP                                     PAGE 2 OF 3

    OPTION                                DESCRIPTION

FU = ALL        FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = OB        USER INDEX RANGE (0-377777B).
              377777B
TF = ALL        FAMILY FOR TERMINAL (1-7 CHARACTERS).
TUI = OB       DESTINATION TERMINAL UI RANGE (0-377777B).
              377777B
FS = OB        FILE SIZE RANGE IN PRU-S (0-777777B).
              777777B
ID = OB        ID RANGE OF FILES TO PROCESS (0-67B).
              67B
FC = ALL        FORMS CODE (2 CHARACTERS OR **).
LA = LVL2      LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA = LVL6      UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
LD = M42       DESTINATION LID (3 CHARACTERS).

```

Figure 18-15. Initial QDUMP K Display (Left Screen) (Page 2 of 3)

QDUMP		PAGE 3 OF 3
OPTION	DESCRIPTION	
BC = (DT)	DISPOSITION TYPE FOR LOCAL BATCH DEVICES.	
RB = (DT)	DISPOSITION TYPE FOR REMOTE BATCH DEVICES.	
	S I P P P P L L L L L W F N U L R 2 R S T X Y T BC * * * * - - - - - * RB * * - * * * - * * * * * * = OPTION SELECTED. OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.	
JSN = AAAA	JOB SEQUENCE NAME (1-4 CHARACTERS).	
AAAB	JSN WILL BE CLEARED IF PREVIOUSLY ENTERED.	

Figure 18-16. Initial QDUMP K Display (Left Screen) (Page 3 of 3)

K. QDUMP.	jsn
hh.mm.ss. yy/mm/dd.	CDC NETWORK OPERATING SYSTEM
	MID=AA NOS version
QDUMP	
PAGE 1 OF 3	
INACTIVE QUEUES RESIDE ON FOLLOWING DEVICES IN FAMILY - FEATURE	
40	
(ERRORED FILE DUMPING SET)	
OPTION	DESCRIPTION
MI = 06	MACHINE ID (1-2 CHARACTERS, *ALL*).
FM = ALL	FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = 40B	DEVICE NUMBER (0-77B, *ALL*).
DA = 85/01/11	DUMP DATE RANGE (YYMMDD, *ALL*).
85/01/16	
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
FN = QFILES	DUMP FILE NAME (1-7 CHARACTERS).
NF = OD	NUMBER OF MEDIA FILES TO SKIP.
TP = ALL	TYPE (A=ACTIVE, I=INACTIVE, *ALL*).

Figure 18-17. Secondary QDUMP K Display (Left Screen)

QDUMP PROCESSOR.		yy/mm/dd. hh.mm.ss.		PAGE 1						
*** DUMPED QUEUE FILE DISPOSITION ***										
NO.	JSNS	MID	FAMILY	DN DT DS	DATE	LENGTH ACCESS	(PRUS)	LEVEL	DISPOSITION	ERRORS
1.	AABAB	64	SYST64	2 PR BC	83/07/07.	1			ACTIVE	
2.	AABBB	64	SYST64	2 LS BC	83/07/07.	12			ACTIVE	
3.	AABDB	64	SYST64	2 LT BC	83/07/07.	13			ACTIVE	
4.	AABGT	64	SYST64	2 LT BC	83/07/07.	14			ACTIVE	
5.	AABJS	64	SYST64	2 PR BC	83/07/07.	15			ACTIVE	
6.	AABOB	64	SYST64	2 PR BC	83/07/07.	17			ACTIVE	
7.	AABCB	64	SYST64	3 PR RB	83/07/07.	13			ACTIVE	
8.	AABE2	64	SYST64	3 PR RB	83/07/07.	13			ACTIVE	
9.	AABFT	64	SYST64	3 PR RB	83/07/07.	5			ACTIVE	
10.	AABHR	64	SYST64	3 LT RB	83/07/07.	15			ACTIVE	
11.	AABIS	64	SYST64	3 LT BC	83/07/07.	1			ACTIVE	
12.	AABK2	64	SYST64	3 LT BC	83/07/07.	1			ACTIVE	
13.	AABLB	64	SYST64	3 PR BC	83/07/07.	1			ACTIVE	
14.	AABPB	64	SYST64	3 PR BC	83/07/07.	20			ACTIVE	

Figure 18-18. QDUMP Output File

QDUMP does not purge files upon completion. If you want to purge files, the following sequence of commands is recommended.

Command	Description
QREC,OP=DI.	Dequeues the files.
QDUMP,TP=I.	Dumps the inactive files. The dump file (default name = QFILES) should be preassigned before issuing this command.
QREC,OP=PI.	Purges the files.

QFTLIST

QFTLIST displays and/or lists routing and other information about active queued files. Its operation is similar to that of QALTER, except file alteration or purging is not allowed. When QFTLIST is selected at the system console, the QFTLIST K display appears on the left screen. This display contains a list of all applicable parameters associated with QFTLIST, their default values, and a short description of each (refer to figures 18-8 through 18-10). The output generated by QFTLIST LIST and OUT K-display commands (both printed and right screen K display) is similar or identical to that produced by QALTER.

Example 1:

Use the QFTLIST utility to determine what actions are necessary to get queued files processed by the RBF or BIO Subsystems. The following QFTLIST entry generates a detailed list of queued files, from which forms code, id, and other destination routing information can be determined. This information indicates what action must be taken to have the files processed.

```
X.QFTLIST (PO=N, LO=F)
```

Example 2:

Use the QFTLIST utility to determine what queued files are routed to a particular device. The following QFTLIST entry generates a list of all print files with forms code of AA that are routed to a 580-20 (LT) printer. Input file selection is explicitly disabled.

```
X.QFTLIST (PO=N, FC=AA, RB=IN, BC=NONE, BC=LT)
```

QLIST

QLIST lists inactive queued files, which may include all inactive queued files in the system or a selected subset based on parameters specified when the utility is called.

The QLIST utility is not restricted to use from the system console. However, from the system console, QLIST must be called directly with the following command. No K-display interaction is available.

```
X.QLIST(opt1=val1, opt2=val2, . . . , optn=valn)
```

The applicable parameters associated with QLIST are the same as those used with the QREC utility, except for the OP parameter. The OP parameter is not valid with QLIST.

If all current inactive queued files are to be listed, no parameters need be specified. In this case, the format of the QLIST call is:

```
X.QLIST.
```

The output generated by QLIST is the same as that produced by the LIST command available under the QREC utility (refer to figure 18-31).

QLOAD

QLOAD processes the dump files generated by QDUMP or other utilities using the same format. QLOAD can selectively load the queued files from these dump files. QLOAD can also list the contents of a dump file without loading any files.

For secured systems, the access level for each file selected to be loaded must be within the file's origin type limits and the file must be assigned to an appropriate mass storage device. If a selected file cannot be loaded, QLOAD issues a diagnostic message and skips that file.

When QLOAD is selected at the system console, the QLOAD K display (3 pages) appears on the left screen. This display contains a list of all applicable parameters associated with QLOAD, their default values, and a short description of each (refer to figures 18-19 through 18-21).

After QLOAD processing has been completed and K-display interaction has been terminated by

```
K . STOP .
```

a QLOAD output file is generated, indicating all queued files which were contained on the dump file, and whether they were loaded or not. Figure 18-22 contains an example of this output.

Example:

Use the QDUMP and QLOAD utilities to dump queued files from one machine and load them onto another. The following commands dump all punch files from machine AA. The QREC utility is used to dequeue the files and then purge them, because QDUMP does not purge files.

```
QREC , PO=N, I=DIRCTVS, OP=DI .
REWIND, DIRCTVS .
LABEL, QFILES, VSN=QUEUE, MT .
QDUMP, PO=N, I=DIRCTVS, TP=I .
REWIND, DIRCTVS .
QREC , PO=N, OP=PI, I=DIRCTVS .
```

File DIRCTVS is the input directive file that contains the following directives.

```
BC=NONE, BC=PH .
RB=NONE, RB=PH .
```

The following commands load the files dumped by the previous QDUMP onto another machine.

```
LABEL, QFILES, VSN=QUEUE, MT .
QLOAD, PO=N, MI=AA, OP=A, DF=ALL .
```

```

K. QLOAD.          jsn
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version

QLOAD                                     PAGE 1 OF 3

                      (ERRORED FILE DUMPING SET)

      OPTION                                DESCRIPTION

MI = 06                MACHINE ID (1-2 CHARACTERS, *ALL*).
DF = ALL               DESTINATION FAMILY (1-7 CHARACTERS).
DD = ALL               DESTINATION DEVICE (0-77B, *ALL*).
DA = 85/01/11         LOAD DATE RANGE (YYMMDD, *ALL*).
                      85/01/16

L = OUTPUT            FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
FN = QFILES           LOAD FILE NAME (1-7 CHARACTERS).
NF = OD               NUMBER OF MEDIA FILES TO SKIP.
SC = OD               NUMBER OF QUEUED FILES TO SKIP.
OP = A                LOAD OPTION (A=LOAD/ACTIVATE,
                      I=LOAD/INACTIVATE).
    
```

Figure 18-19. QLOAD K Display (Left Screen) (Page 1 of 3)

```

QLOAD                                     PAGE 2 OF 3

      OPTION                                DESCRIPTION

FU = ALL              FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = OB               USER INDEX RANGE (0-377777B).
                      377777B
TF = ALL              FAMILY FOR TERMINAL (1-7 CHARACTERS).
TUI = OB              DESTINATION TERMINAL UI RANGE (0-377777B).
                      377777B
FS = OB               FILE SIZE RANGE IN PRU-S (0-777777B).
                      777777B
ID = OB               ID RANGE OF FILES TO PROCESS (0-67B).
                      67B
FC = ALL              FORMS CODE (2 CHARACTERS OR **).
LA = LVL2             LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA = LVL6             UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
LD = M42              DESTINATION LID (3 CHARACTERS).
    
```

Figure 18-20. QLOAD K Display (Left Screen) (Page 2 of 3)

QLOAD		PAGE 3 OF 3
OPTION	DESCRIPTION	
BC = (DT)	DISPOSITION TYPE FOR LOCAL BATCH DEVICES.	
RB = (DT)	DISPOSITION TYPE FOR REMOTE BATCH DEVICES.	
	S I P P P P L L L L W F N U L R 2 R S T X Y T BC * * * * - - - - - * RB * * - * * * - * * * * * * = OPTION SELECTED. OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.	
JSN = AAAA	JOB SEQUENCE NAME (1-4 CHARACTERS).	
AAAB	JSN WILL BE CLEARED IF PREVIOUSLY ENTERED.	

Figure 18-21. QLOAD K Display (Left Screen) (Page 3 of 3)

QLOAD/LDLIST PROCESSOR.		yy/mm/dd. hh.mm.ss.		PAGE 1 .			
NO.	JSNS	MID	FAMILY	DN DT DS	DATE	LENGTH ACCESS (PRUS) LEVEL	DISPOSITION ERRORS
1.	AABAB	64	SYST64	3 PR BC	83/07/07.	1	ACTIVE
2.	AABBB	64	SYST64	2 LS BC	83/07/07.	12	ACTIVE
3.	AABDB	64	SYST64	3 LT BC	83/07/07.	13	ACTIVE
4.	AABGT	64	SYST64	2 LT BC	83/07/07.	14	ACTIVE
5.	AABJS	64	SYST64	3 PR BC	83/07/07.	15	ACTIVE
6.	AABOB	64	SYST64	2 PR BC	83/07/07.	17	ACTIVE
7.	AABCB	64	SYST64	3 PR RB	83/07/07.	13	ACTIVE
8.	AABE2	64	SYST64	2 PR RB	83/07/07.	13	ACTIVE
9.	AABFT	64	SYST64	3 PR RB	83/07/07.	5	ACTIVE
10.	AABHR	64	SYST64	2 LT RB	83/07/07.	15	ACTIVE
11.	AABIS	64	SYST64	3 LT BC	83/07/07.	1	ACTIVE
12.	AABK2	64	SYST64	2 LT BC	83/07/07.	1	ACTIVE
13.	AABLB	64	SYST64	3 PR BC	83/07/07.	1	ACTIVE
14.	AABPB	64	SYST64	2 PR BC	83/07/07.	20	ACTIVE

Figure 18-22. QLOAD Output File

QMOVE

QMOVE moves queued files from one disk device to another. It also produces a listing of all files moved with information about each file processed.

For secured systems, the destination device for each file selected to be moved must have access level limits that are appropriate to accept the file. If the file is being reactivated as well as moved, its access level must be within the file's origin type limits. If a selected file cannot be moved, **QMOVE** issues a diagnostic message and skips that file.

When **QMOVE** is selected at the system console, the initial **QMOVE K** display (3 pages) appears on the left screen. This display contains a list of all applicable parameters associated with **QMOVE**, their default values, and a short description of each (refer to figures 18-23 through 18-25).

If the **FM** parameter is specified, indicating the loading of queued files to a specified family of devices, the initial **QMOVE** left screen **K** display is replaced by the secondary **QMOVE K** display (refer to figure 18-26). This display closely resembles the initial **QMOVE** display but restricts **QMOVE** processing to the family of devices specified. Each device in the specified family containing inactive queued files is indicated. In addition, any parameter entered for the **QMOVE** utility is reflected in this display. Parameters not entered remain at default values. Only the first screen of the secondary **QMOVE K** display is shown, since the second and third screens have the same format as the corresponding screens of the initial **QMOVE K** display.

After **QMOVE** processing has been completed and **K**-display interaction has been terminated by

K . STOP .

a **QMOVE** output file is generated, indicating all files moved. Figure 18-27 contains an example of this output.

Example 1:

If there are inactive queued files residing on a removable device, you can activate these files but must first move them to a nonremovable device. This example shows how to use the QMOVE utility to move queued files from the removable device with family name SYSTEM and device number 45 to the nonremovable device with family name FAM1 and device number 5. The following entries are made from DSD.

```
X.QFSP.                               jsn is the job sequence name of QFSP.
K, jsn.
K.QMOVE.
K.FM=SYSTEM, DN=45.
K.DF=FAM1, DD=5.                       The moved files are active.
K.OP=A.
K.GO, STOP.
```

Example 2:

This example exhibits the relationships among the ID, TF/TUI, and FC parameters. (Refer to the descriptions of these parameters for the types of files affected by them.) The following command moves all local batch-destined output files with ID=6 and all remote batch-destined output files to device 3 in family SYS. The BC and RB parameters explicitly disable input file processing.

```
QMOVE, PO=N, DF=SYS, DD=3, ID=6, BC=IN, RB=IN.
```

The following command performs the same function as the previous command except that no remote batch files are moved.

```
QMOVE, PO=N, DF=SYS, DD=3, ID=6, BC=IN, RB=NONE.
```

```

K. QMOVE.          jsn
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version

QMOVE                                     PAGE 1 OF 3

INACTIVE QUEUES RESIDE ON FOLLOWING FAMILIES.
  SYST06   PACKV2   FEATURE
              (ERRORED FILE MOVING NOT SET)

OPTION                                DESCRIPTION

MI = 06          MACHINE ID (1-2 CHARACTERS, *ALL*).
FM = ALL         FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL         DEVICE NUMBER (0-77B, *ALL*).
DF = FEATURE     DESTINATION FAMILY (1-7 CHARACTERS).
DD = 40B         DESTINATION DEVICE (0-77B).

DA = 85/01/11    DATE RANGE (YYMMDD, *ALL*).
      85/01/16
L = OUTPUT       FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
TP = ALL         TYPE (A=ACTIVE, I=INACTIVE, *ALL*).
OP = A           OPTION (A=LEAVE ACTIVE, I=LEAVE INACTIVE).
  
```

Figure 18-23. Initial QMOVE K Display (Left Screen) (Page 1 of 3)

```

QMOVE                                     PAGE 2 OF 3
OPTION                                DESCRIPTION

FU = ALL         FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = OB          USER INDEX RANGE (0-377777B).
      377777B
TF = ALL         FAMILY FOR TERMINAL (1-7 CHARACTERS).
TUI = OB         DESTINATION TERMINAL UI RANGE (0-377777B).
      377777B
FS = OB          FILE SIZE RANGE IN PRU-S (0-777777B).
      777777B
ID = OB          ID RANGE OF FILES TO PROCESS (0-67B).
      67B
FC = ALL         FORMS CODE (2 CHARACTERS OR **).
LA = LVL2        LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA = LVL6        UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
LD = M42         DESTINATION LID (3 CHARACTERS).
  
```

Figure 18-24. Initial QMOVE K Display (Left Screen) (Page 2 of 3)

QMOVE		PAGE 3 OF 3
OPTION	DESCRIPTION	
BC = (DT)	DISPOSITION TYPE FOR LOCAL BATCH DEVICES.	
RB = (DT)	DISPOSITION TYPE FOR REMOTE BATCH DEVICES.	
	S I P P P P L L L L L W F N U L R 2 R S T X Y T BC * * * * - - - - - * RB * * - * * * - * * * * * * = OPTION SELECTED. OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.	
JSN = AAAA	JOB SEQUENCE NAME (1-4 CHARACTERS).	
AAAB	JSN WILL BE CLEARED IF PREVIOUSLY ENTERED.	

Figure 18-25. Initial QMOVE K Display (Left Screen) (Page 3 of 3)

K. QMOVE.		jsn
hh.mm.ss. yy/mm/dd.		CDC NETWORK OPERATING SYSTEM
		MID=AA NOS version
QMOVE		PAGE 1 OF 3
INACTIVE QUEUES RESIDE ON FOLLOWING DEVICES IN FAMILY - FEATURE		
40		
(ERRORED FILE MOVING NOT SET)		
OPTION	DESCRIPTION	
MI = 06	MACHINE ID (1-2 CHARACTERS, *ALL*).	
FM = FEATURE	FAMILY FOR DEVICES (1-7 CHARACTERS).	
DN = 40B	DEVICE NUMBER (0-77B, *ALL*).	
DF = FEATURE	DESTINATION FAMILY (1-7 CHARACTERS).	
DD = 40B	DESTINATION DEVICE (0-77B).	
DA = 85/01/11	DATE RANGE (YYMMDD, *ALL*).	
	85/01/16	
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).	
TP = ALL	TYPE (A=ACTIVE, I=INACTIVE, *ALL*).	
OP = A	OPTION (A=LEAVE ACTIVE, I=LEAVE INACTIVE).	

Figure 18-26. Secondary QMOVE K Display (Left Screen)

```

QMOVE PROCESSOR.                yy/mm/dd. hh.mm.ss.  PAGE  1
*** MOVED QUEUE FILE DISPOSITION ***
ACTIVE FILES MOVED TO FAMILY SYST64  DEVICE  2 AND LEFT ACTIVE
                                LENGTH ACCESS
NO.  JSNS MID FAMILY  DN DT DS  DATE      (PRUS) LEVEL  DISPOSITION ERRORS
1.  AABCB  64 SYST64   3 PR RB 83/07/07.    13
2.  AABE2  64 SYST64   3 PR RB 83/07/07.    13
3.  AABFT  64 SYST64   3 PR RB 83/07/07.     5
4.  AABHR  64 SYST64   3 LT RB 83/07/07.    15
5.  AABIS  64 SYST64   3 LT BC 83/07/07.     1
6.  AABK2  64 SYST64   3 LT BC 83/07/07.     1
7.  AABLB  64 SYST64   3 PR BC 83/07/07.     1
8.  AABPB  64 SYST64   3 PR BC 83/07/07.    20
9.  AABQB  64 SYST64   3 PR BC 83/07/07.     2
10. AABSB  64 SYST64   3 LT BC 83/07/07.    13
11. AABUS  64 SYST64   3 PR BC 83/07/07.    15
12. AABWB  64 SYST64   3 PR RB 83/07/07.    13
13. AABYT  64 SYST64   3 PR RB 83/07/07.     5
14. AACAS  64 SYST64   3 LT BC 83/07/07.     2
15. AACCB  64 SYST64   3 PR BC 83/07/07.     2

```

Figure 18-27. QMOVE Output File

QREC

QREC deactivates or activates selected queued files and purges selected inactive queued files.

For secured systems, the access level for any file selected to be reactivated must be within the file's origin type limits. If a selected file cannot be processed, QREC issues a diagnostic message and skips that file.

When QREC is selected at the system console, the initial QREC K display (3 pages) appears on the left screen. This display contains a list of all applicable parameters associated with QREC, their default values, and a short description of each (refer to figures 18-28 through 18-30).

A complete list of current inactive queued files is displayed on the right screen when you enter this command:

```
K.LIST.
```

When this command is processed, the right screen K display is automatically replaced with the inactive queues K display (refer to figure 18-31). The information presented in this display is the same as that generated by the QLIST utility. The list of files displayed is also written to the output file specified by the QREC L parameter.

If the FM parameter is specified, indicating processing of queued files for a specified family of devices, the initial QREC left screen K display is replaced by the secondary QREC K display (refer to figure 18-32). This display closely resembles the initial QREC display but restricts QREC processing to the family of devices specified. Each device in the specified family containing inactive queued files is indicated. In addition, any parameter entered for the QREC utility is reflected in this display. Parameters not entered remain at default values. Only the first screen of the secondary QREC K display is shown, since the second and third screens have the same format as the corresponding screens of the initial QREC K display.

After QREC processing has been completed and K-display interaction has been terminated by

```
K.STOP.
```

a QREC output file is generated, indicating the disposition of all queued files processed. Figure 18-33 is an example of this output.

Example 1:

If the QFT is becoming full, you can enter QREC from DSD to dequeue active queued files and free up QFT space. This QREC entry dequeues all queued files:

```
X.QREC (OP=DI , PO=N)
```

If RBF is not active, you are still able to dequeue remote batch files by using QREC. This QREC entry dequeues all remote batch files:

```
X.QREC (OP=DI , BC=NONE , PO=N)
```

If these are queued files that previously had been dequeued, you can activate these files with QREC. This QREC entry requeues all inactive queued files:

```
X.QREC (OP=RI , PO=NU)
```

Example 2:

Purge all queued files (active and inactive) on a particular device by using QREC either with or without an input directive file. The first method involves two calls to QREC. The first QREC call dequeues (makes inactive) all active queued files on the device with family name SYSTEM and device number 1, and the second QREC call purges all inactive (and hence all) queued files on that same device.

```
QREC , PO=N , OP=DI , FM=SYSTEM , DN=1 .  
QREC , PO=N , OP=PI , FM=SYSTEM , DN=1 .
```

The second method involves only one command call to QREC to perform the same function as the first method:

```
QREC , I=DIR .
```

File DIR is the input directive file that contains the following directives.

```
OP=DI , FM=SYSTEM , DN=1 .  
GO .  
OP=PI , FM=SYSTEM , DN=1 .  
GO , STOP .
```

Example 3:

This example exhibits the relationships among the ID, TF/TUI, and FC parameters. It is important to observe the following restrictions regarding these parameters and QREC.

- ID pertains only to local batch-destined output files.
- TF or TUI pertains only to remote batch-destined output files.
- FC pertains to all output files.
- When you specify the OP parameter, all queued files are processed unless the ID, TF, TUI, or FC parameter restricts output file selection or the BC or RB parameter restricts input or output file selection.

The following command activates all local batch-destined output files with ID=6 and FC=AD, all remote batch-destined files with TF=SYSTAA,TUI=3751 and FC=AD, and all input files.

```
QREC,PO=N,OP=RI, ID=6, TF=SYSTAA, TUI=3751, FC=AD.
```

The following command performs the same functions as the previous command except that no input files are processed.

```
QREC,PO=N,OP=RI, ID=6, TF=SYSTAA, TUI=3751, FC=AD, BC=IN, RB=IN.
```

```

K. QREC.          jsn
  hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                        MID=AA   NOS version

QREC                                     PAGE 1 OF 3

INACTIVE QUEUES RESIDE ON FOLLOWING FAMILIES.
  SYST06   PACKV2   FEATURE

      OPTION                                DESCRIPTION

MI = 06          MACHINE ID (1-2 CHARACTERS, *ALL*).
FM = ALL         FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL         DEVICE NUMBER (0-77B, *ALL*).
DA = 85/01/11   QUEUED DATE RANGE (YYMMDD, *ALL*).
                85/01/16
L = OUTPUT      FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
OP = RI         PROCESS OPTION.
                DI = MAKE SELECTED ACTIVE FILES INACTIVE.
                PI = PURGE SELECTED INACTIVE FILES.
                RI = REQUEUE SELECTED INACTIVE FILES.
                RP = REQUEUE SELECTED INACTIVE FILES AND
                    PURGE REMAINING INACTIVE FILES.

```

Figure 18-28. Initial QREC K Display (Left Screen) (Page 1 of 3)

QREC	OPTION	DESCRIPTION
		PAGE 2 OF 3
FU =	ALL	FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI =	OB 377777B	USER INDEX RANGE (0-377777B).
TF =	ALL	FAMILY FOR TERMINAL (1-7 CHARACTERS).
TUI =	OB 377777B	DESTINATION TERMINAL UI RANGE (0-377777B).
FS =	OB 777777B	FILE SIZE RANGE IN PRU-S (0-777777B).
ID =	OB 67B	ID RANGE OF FILES TO PROCESS (0-67B).
FC =	ALL	FORMS CODE (2 CHARACTERS OR **).
LA =	LVL2	LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA =	LVL6	UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
LD =	M42	DESTINATION LID (3 CHARACTERS).

Figure 18-29. Initial QREC K Display (Left Screen) (Page 2 of 3)

QREC	OPTION	DESCRIPTION
		PAGE 3 OF 3
BC =	(DT)	DISPOSITION TYPE FOR LOCAL BATCH DEVICES.
RB =	(DT)	DISPOSITION TYPE FOR REMOTE BATCH DEVICES.
		S I P P P P L L L L W F N U L R 2 R S T X Y T BC * * * * - - - - - * RB * * - * * * - * * * * * * = OPTION SELECTED. OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.
JSN =	AAAA	JOB SEQUENCE NAME (1-4 CHARACTERS).
	AAAB	JSN WILL BE CLEARED IF PREVIOUSLY ENTERED.

Figure 18-30. Initial QREC K Display (Left Screen) (Page 3 of 3)

```

*** INACTIVE QUEUES LIST ***

```

NO.	JSNS	MID	FAMILY	DN	DT	DS	DATE	LENGTH (PRUS)	ACCESS LEVEL
1.	AABEB	64	SYST64	2	LS	BC	83/08/08.	5	
2.	AABFR	64	SYST64	2	PR	RB	83/08/08.	41	
3.	AABG1	64	SYST64	2	PR	RB	83/08/08.	41	
4.	AABH2	64	SYST64	2	LS	BC	83/08/08.	5	
5.	AABIR	64	SYST64	2	PR	RB	83/08/08.	41	
6.	AABJ2	64	SYST64	2	PR	BC	83/08/08.	7	
7.	AABKR	64	SYST64	2	PR	RB	83/08/08.	41	
8.	AABLR	64	SYST64	2	PR	RB	83/08/08.	41	
9.	AABM2	64	SYST64	2	LS	BC	83/08/08.	4	
10.	AABN2	64	SYST64	2	LS	BC	83/08/08.	5	
11.	AABO1	64	SYST64	2	PR	RB	83/08/08.	41	
12.	AABPB	64	SYST64	2	LS	BC	83/08/08.	6	
13.	AABDB	64	SYST64	2	LS	BC	83/08/08.	2	

END OF DISPLAY.

Figure 18-31. Inactive Queues K Display (Right Screen)

```

K. QREC.                jsn
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM
                      MID=AA   NOS version

QREC                                PAGE 1 OF 3

INACTIVE QUEUES RESIDE ON FOLLOWING DEVICES IN FAMILY - FEATURE
40

OPTION                            DESCRIPTION

MI = 06                            MACHINE ID (1-2 CHARACTERS, *ALL*).
FM = FEATURE                        FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = 40B                            DEVICE NUMBER (0-77B, *ALL*).
DA = 85/01/11                       QUEUED DATE RANGE (YYMMDD, *ALL*).
      85/01/16
L = OUTPUT                          FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
OP = RI                              PROCESS OPTION.
      DI = MAKE SELECTED ACTIVE FILES INACTIVE.
      PI = PURGE SELECTED INACTIVE FILES.
      RI = REQUEUE SELECTED INACTIVE FILES.
      RP = REQUEUE SELECTED INACTIVE FILES AND
           PURGE REMAINING INACTIVE FILES.

```

Figure 18-32. Secondary QREC K Display (Left Screen)

QREC/QLIST PROCESSOR.				yy/mm/dd. hh.mm.ss.			PAGE	1
*** QUEUE DISPOSITION ***								
NO.	JSNS	MID	FAMILY	DN	DT	DS	DATE	LENGTH ACCESS (PRUS) LEVEL DISPOSITION
1.	AABAB	64	SYST64	2	PR	BC	83/07/07.	1 ACTIVATE
2.	AABBB	64	SYST64	2	LS	BC	83/07/07.	12 ACTIVATE
3.	AABE2	64	SYST64	2	PR	RB	83/07/07.	13 ACTIVATE
4.	AABDB	64	SYST64	2	LT	BC	83/07/07.	13 ACTIVATE
5.	AABFT	64	SYST64	2	PR	RB	83/07/07.	5 ACTIVATE
6.	AABHR	64	SYST64	2	LT	RB	83/07/07.	15 ACTIVATE
7.	AABGT	64	SYST64	2	LT	BC	83/07/07.	14 ACTIVATE
8.	AABIS	64	SYST64	2	LT	BC	83/07/07.	1 ACTIVATE
9.	AABK2	64	SYST64	2	LT	BC	83/07/07.	1 ACTIVATE
10.	AABJS	64	SYST64	2	PR	BC	83/07/07.	15 ACTIVATE
11.	AABLB	64	SYST64	2	PR	BC	83/07/07.	1 ACTIVATE
12.	AABPB	64	SYST64	2	PR	BC	83/07/07.	20 ACTIVATE
13.	AABOB	64	SYST64	2	PR	BC	83/07/07.	17 ACTIVATE
14.	AABQB	64	SYST64	2	PR	BC	83/07/07.	2 ACTIVATE
15.	AABSB	64	SYST64	2	LT	BC	83/07/07.	13 ACTIVATE
16.	AABRB	64	SYST64	2	LS	BC	83/07/07.	12 ACTIVATE
17.	AABUS	64	SYST64	2	PR	BC	83/07/07.	15 ACTIVATE
18.	AABTT	64	SYST64	2	LT	BC	83/07/07.	14 ACTIVATE
19.	AABWB	64	SYST64	2	PR	RB	83/07/07.	13 ACTIVATE
20.	AABVB	64	SYST64	2	PR	BC	83/07/07.	17 ACTIVATE
21.	AABYT	64	SYST64	2	PR	RB	83/07/07.	5 ACTIVATE
22.	AABX2	64	SYST64	2	PR	RB	83/07/07.	13 ACTIVATE
23.	AACAS	64	SYST64	2	LT	BC	83/07/07.	2 ACTIVATE
24.	AABZR	64	SYST64	2	LT	RB	83/07/07.	15 ACTIVATE
25.	AACCB	64	SYST64	2	PR	BC	83/07/07.	2 ACTIVATE
26.	AACB2	64	SYST64	2	LT	BC	83/07/07.	2 ACTIVATE
27.	AACDB	64	SYST64	2	PR	BC	83/07/07.	20 ACTIVATE
28.	AACB	64	SYST64	2	PR	RB	83/07/07.	13 ACTIVATE

Figure 18-33. QREC Output File

Dayfile Dumping Utilities

The dayfile dumping utilities (AFD, DFD, ELD, GETLOG, MAINLOG) write all or selected parts of the account dayfile, system dayfile, error log, or binary maintenance log to a file. AFD, DFD, and ELD produce a listing (output) file. GETLOG writes the error log, binary maintenance log, and an image of the equipment status table to files. GETLOG is used to provide input to the Hardware Performance Analyzer (HPA). MAINLOG produces its output in a binary format. A job using any of these utilities must be of system origin, or you must have system origin privileges and the system must be in debug mode. The following format is used to call the dayfile dumping utilities.

```
utility,L=listfile,FR=string,OP=op,PD=pd,PL=pl,I=infile,
B=binfile,E=estfile.
```

or

```
utility,listfile,string,op,pd,pl,infile,binfile,estfile.
```

Parameter	Description												
utility	Utility called to perform the dayfile dump.												
	<table border="1"> <thead> <tr> <th>Utility</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>AFD</td> <td>Dump account dayfile.</td> </tr> <tr> <td>DFD</td> <td>Dump system dayfile.</td> </tr> <tr> <td>ELD</td> <td>Dump error log.</td> </tr> <tr> <td>GETLOG</td> <td>Dump error log, binary maintenance log, and EST image.</td> </tr> <tr> <td>MAINLOG</td> <td>Dump binary maintenance log.</td> </tr> </tbody> </table>	Utility	Description	AFD	Dump account dayfile.	DFD	Dump system dayfile.	ELD	Dump error log.	GETLOG	Dump error log, binary maintenance log, and EST image.	MAINLOG	Dump binary maintenance log.
Utility	Description												
AFD	Dump account dayfile.												
DFD	Dump system dayfile.												
ELD	Dump error log.												
GETLOG	Dump error log, binary maintenance log, and EST image.												
MAINLOG	Dump binary maintenance log.												
L=listfile	<p>Output file containing the dayfile dump produced by AFD, DFD, or ELD. For GETLOG, listfile is the file containing the error log. The utility also places diagnostic messages in this file when various error conditions occur. These messages begin with NOTICE*** to distinguish them from the lines of the dayfile being processed; they are described in the NOS Version 2 Operations Handbook.</p> <p>The default name of listfile is OUTPUT for AFD, DFD, and ELD. The default name is ERR for GETLOG. If L=0 is specified, no file is written. L=listfile is not used with MAINLOG. The utility paginates listfile if it is OUTPUT or if print density and page length are specified.</p>												

Parameter	Description														
FR=string	<p>Search string for selective dayfile dumping. The utility searches the dayfile for this string in the starting position of the field specified by the OP=op parameter.</p> <p>The string specified must begin with the first character in the field to be searched. The time field begins with a blank.</p> <p>If the string contains characters other than letters and numbers (such as blanks), it must be enclosed by \$ delimiters. A \$ within the string must be entered twice (\$\$) so it is not interpreted as a delimiter. Interactive commands entered in the dayfile are preceded by a \$ unless you were under the batch subsystem when you issued the command. To search for an interactive command, the \$ preceding the command is replaced with two \$s, and the command is enclosed by \$ delimiters (for example, \$\$OLD\$).</p> <p>If the string is found, the portion of the dayfile following the last occurrence of the specified string is output. If the string is not found, an informative message and the entire dayfile is output for jobs where the output file is connected to a terminal; where the output file is not connected to a terminal, the job receives only the informative message.</p>														
OP=op	<p>Dump option. If neither FR=string nor the OP dump option is specified, default is OP=F for nonterminal output files. For terminal output files, the default is OP=I. If FR is specified but OP is not, default is OP=M.</p>														
	<table border="1"> <thead> <tr> <th>op</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>F</td> <td>Full dayfile dump is taken.</td> </tr> <tr> <td>I</td> <td>Incremental dump is taken. The dayfile is dumped starting from the point of the last dayfile dump. AFD does not process the I option when executed within a job with system origin privileges that is not system origin.</td> </tr> <tr> <td>J</td> <td>The job sequence name field in the dayfile is searched for the string specified by FR=string. The dump begins from that point.</td> </tr> <tr> <td>M</td> <td>The message field in the dayfile is searched for the string specified by FR=string. The dump begins from that point. OP=M cannot be used with MAINLOG.</td> </tr> <tr> <td>P</td> <td>Incremental dump is taken. The dayfile is dumped starting from the point of the last dayfile dump with this job name.</td> </tr> <tr> <td>T</td> <td>The time field in the dayfile is searched for the string specified by FR=string. The dump begins from that point.</td> </tr> </tbody> </table>	op	Description	F	Full dayfile dump is taken.	I	Incremental dump is taken. The dayfile is dumped starting from the point of the last dayfile dump. AFD does not process the I option when executed within a job with system origin privileges that is not system origin.	J	The job sequence name field in the dayfile is searched for the string specified by FR=string. The dump begins from that point.	M	The message field in the dayfile is searched for the string specified by FR=string. The dump begins from that point. OP=M cannot be used with MAINLOG.	P	Incremental dump is taken. The dayfile is dumped starting from the point of the last dayfile dump with this job name.	T	The time field in the dayfile is searched for the string specified by FR=string. The dump begins from that point.
op	Description														
F	Full dayfile dump is taken.														
I	Incremental dump is taken. The dayfile is dumped starting from the point of the last dayfile dump. AFD does not process the I option when executed within a job with system origin privileges that is not system origin.														
J	The job sequence name field in the dayfile is searched for the string specified by FR=string. The dump begins from that point.														
M	The message field in the dayfile is searched for the string specified by FR=string. The dump begins from that point. OP=M cannot be used with MAINLOG.														
P	Incremental dump is taken. The dayfile is dumped starting from the point of the last dayfile dump with this job name.														
T	The time field in the dayfile is searched for the string specified by FR=string. The dump begins from that point.														

Parameter	Description										
PD=pd	Print density in pd lines per inch (3, 4, 6, or 8). Default is PD=6. PD cannot be used with MAINLOG.										
PL=pl	Page length in pl lines per page. Default is based on the following print densities. PL cannot be used with MAINLOG.										
	<table border="1"> <thead> <tr> <th>pd</th> <th>Default pl</th> </tr> </thead> <tbody> <tr> <td>3</td> <td>30</td> </tr> <tr> <td>4</td> <td>40</td> </tr> <tr> <td>6</td> <td>60</td> </tr> <tr> <td>8</td> <td>80</td> </tr> </tbody> </table>	pd	Default pl	3	30	4	40	6	60	8	80
pd	Default pl										
3	30										
4	40										
6	60										
8	80										
I=infile	An attached, terminated dayfile is to be used for input. For GETLOG, infile is an attached terminated binary maintenance log. If omitted, the utility uses the active dayfile for input.										
B=binfile	File on which GETLOG and MAINLOG write the binary maintenance log. These utilities write binfile with W-type records having C-type blocking (refer to the CYBER Record Manager Basic Access Methods Reference Manual for further information on record types and record blocking). GETLOG and MAINLOG assume B=BML if this parameter is omitted. If B=0 is specified, no file is written.										
E=estfile	File on which GETLOG writes an image of the equipment status table. GETLOG assumes E=EST if this parameter is omitted. If E=0 is specified, no file is written.										

Directives	19-3
*CM Directive	19-3
*MS Directive	19-4
*AD Directive	19-5
*SC Directive	19-7
*FL Directive	19-7
*/ Directive	19-8
*DELETE Directive	19-8
*FILE Directive	19-9
*IGNORE Directive	19-9
*PPSYN Directive	19-9
Record Types	19-10

After the system is loaded, you can use the SYSEDIT command to modify the system library. A job containing a SYSEDIT command must either be a system origin job, or you must be validated for system origin privileges and the system must be in debug mode. Refer to the DEBUG command in section 5 for a description of debug mode. In addition, to prevent unauthorized modification of the operating system on a secured system, SYSEDIT will be allowed only from jobs with security administrator privileges.

NOTE

Using SYSEDIT in a production environment may cause unpredictable results and is not recommended. The system should be idle to ensure predictable results.

The format of the command is:

SYSEDIT, p_1, p_2, \dots, p_n . directive₁, directive₂, ..., directive_n

p_i is either a keyword or a keyword equated to a value.

P_i	Description
B=repfile	Replacement records are on file repfile.
B	Same as B=LGO.
B=0	No replacement file.
B omitted	Same as B=LGO.
C	Checkpoint of the system follows SYSEDIT. This parameter is required if changes made to the system library are to be recovered during a level 3 recovery deadstart.
C omitted	No checkpoint is performed unless the system was generated using the alternate system library residency feature or the PROBE utility was enabled. In these cases, the checkpoint automatically follows SYSEDIT.
I=infile	Directive input is on file infile.
I	Same as I=INPUT.
I=0	No directive input.
I omitted	Same as I=INPUT.
L=outfile	List output is on file outfile.
L	Same as L=OUTPUT.

P_i	Description
L=0	No list output.
L omitted	Same as L=0.
NA	SYSEEDIT does not abort if the system file is busy. Execution continues when the system file is no longer busy.
R=n	Restoration is made to copy n of the system. The system copy number is printed on the output listing.
R	Restoration is made to initial deadstart system.
R=0	No system file restoration.
R omitted	Same as R=0.
Z	Directive input follows the terminator on the command. The I option is ignored.
Z omitted	The system uses the I option for directive input.

NOTE

If the Z option is specified, a terminator must not be placed after the input directives.

Directives

Discussed next are input directives to SYSEEDIT. The directives *CM, *MS, *AD, *SC, *FL, and */ can be placed on LIBDECK. A list of valid record types follows the directive descriptions.

*CM Directive

The *CM directive specifies the system library routines that are to reside in central memory resident (CMR). These routines will reside in the resident peripheral library (RPL) and resident central library (RCL).

The format of the directive is:

```
*CM, ty1/rec1, ty2/rec2, . . . , tyn/recn
```

Parameter Description

ty_i/rec_i Record type and record name of the routine.

Central memory is the most accessible of all storage devices for PP programs. SYSEEDIT automatically includes those programs in CMR that must reside in CMR with the exception of TRACER CPD overlays. If TRACER CPD is used, a *CM directive must be included for PP routines 3CR and 3CS.

SYSEEDIT automatically includes the following programs in CMR:

- Mass storage drivers.
- Programs specified in the internal tables of SLL (SYSEEDIT service routine), such as 1DD and 0DF. These tables are modified at deadstart to exclude programs not required by the configuration.
- Some DSD overlays, such as the 9Cx overlays. A site may make additional DSD overlays reside in CMR.

In addition, it is suggested that high usage PP programs (such as 1MT, 1MS, 4MA through 4MF, and PFM) reside in CMR. However, any addition to central memory allows less space for user jobs. REL type records cannot reside in CMR.

Like any other job in NOS, SYSEEDIT has a field length restriction of 376000₈ central memory words. You cannot use more than approximately 326000₈ central memory words (376000₈ - 50000₈ words for SYSEEDIT's program FL) when defining records to be central memory resident.

***MS Directive**

The *MS directive identifies which system library routines will reside on the system device. It is the default residence for routines; any routine not specified as *CM will automatically reside on the system device, even if *MS is not entered in the LIBDECK.

The format of the directive is:

```
*MS, ty1/rec1, ty2/rec2, . . . , tyn/recn
```

Parameter	Description
------------------	--------------------

ty _i /rec _i	Record type and record name of the routine.
-----------------------------------	---

*AD Directive

The *AD directive places copies of specific system library routines on particular mass storage devices so that they will be accessed from a device other than the system device. For example, it is advantageous to:

- Place frequently used system library routines in extended memory for faster access than is possible from the system device. Notice that CPU programs residing in extended memory can be accessed faster than CPU programs residing in central memory.
- Place PP routines on a mass storage device that has better latency than does the system device.

Once a routine is on an alternate system device, it is accessed from that device instead of from the system device (*MS LIBDECK entries), except in the following cases of error recovery:

- If there is an unrecoverable error for a PP routine on an alternate system device, all of the PP routines will subsequently be accessed from the system device instead of from any of the alternate system devices.
- If there is an unrecoverable error for a CPU, ABS, or OVL routine on an alternate system device, only that routine will subsequently be accessed from the system device.
- If there is an unrecoverable error for an REL CPU routine, the alternate system device must be turned off by the operator to prevent further access to the routine. In this case, the system copy will be used for backup. The alternate system device is not used for loads.

The devices that are to be used as alternate system devices are specified with the ASR entry in the EQPDECK (refer to section 3, Deadstart Decks); the routines that are to reside on each alternate system device are specified in LIBDECK.

The format of the directive is:

```
*AD,xx,ty1/rec1,ty2/rec1,...,tyn/recn
```

Parameter	Description
xx	<p>The 1- or 2-digit EST ordinal or 2-character equipment type of the equipment to be used as an alternate system device. The equipment can be any nonremovable mass storage device (including extended memory) except for a system device. The EST ordinal specified cannot be larger than 37₈.</p> <p>If the equipment is not specified in the EQPDECK with the ASR entry, the *AD entry is ignored without an error indication when the system is loaded.</p>

Parameter	Description										
ty_i	Record type of routine: <table border="1"> <thead> <tr> <th>Record Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ABS</td> <td>CPU multiple entry point overlay.</td> </tr> <tr> <td>OVL</td> <td>CPU overlay.</td> </tr> <tr> <td>PP</td> <td>PP absolute.</td> </tr> <tr> <td>REL</td> <td>Relocatable CPU routine.</td> </tr> </tbody> </table> <p>jIf a record type other than ABS, OVL, PP, or REL is specified, the run is aborted.</p>	Record Type	Description	ABS	CPU multiple entry point overlay.	OVL	CPU overlay.	PP	PP absolute.	REL	Relocatable CPU routine.
Record Type	Description										
ABS	CPU multiple entry point overlay.										
OVL	CPU overlay.										
PP	PP absolute.										
REL	Relocatable CPU routine.										
rec_i	Record name of routine. A routine is allowed on only one alternate system device.										

Additional qualifications:

- Once a routine is placed on an alternate system device, SYSEEDIT may be used to prohibit access to the routine; however, the space for that routine is not released until LIBDECK is modified and the system is reloaded.
- If extended memory is an alternate system device, all ABS, OVL, or REL routines residing there will be loaded from extended memory directly to the load address.
- If DDP is available, PP programs residing in extended memory will be loaded using DDP, and CPU programs will be loaded using the CPU access to extended memory.

***SC Directive**

The *SC directive specifies the commands in a certain program that are to be processed in product set format rather than in NOS format (refer to the NOS Version 2 Reference Set, Volume 3).

The format of the directive is:

```
*SC, ty1/rec1, ty2/rec2, . . . , tyn/recn
```

Parameter	Description
ty _i /rec _i	Record type and record name of the routine to be processed in product set format.

***FL Directive**

The * FL directive specifies the field length that routines to be loaded require to begin execution.

The format of the directive is:

```
*FL, ty1/rec1-fl1, ty2/rec2-fl2, . . . , tyn/recn-fln
```

Parameter	Description
ty _i /rec _i	Record type and record name of the routine.
fl _i	Field length divided by 100 ₈ required by the routine.

The actual field length obtained is subject to the rules governing RFL= and MFL= entry points, since the specified fl_i field is placed in the library directory. The system uses this information to determine field length in the following manner:

1. If bit 11 is not set, an RFL= entry point is indicated. The field length is set to the value in the entry.
2. If bit 11 is set (indicating a value of 400000₈), an MFL= entry point is indicated. The field length is determined in one of two ways:
 - a. If bit 10 is also set, the field length is set to the maximum of the value of the last RFL command and the value in the entry after masking off these upper 2 bits.
 - b. If bit 10 is not set, the field length is set to the maximum of the existing field length and the value in the entry after masking off these upper 2 bits.

***/ Directive**

The */ directive specifies comment lines that are listed on the output file. Other than being listed on the output file, comment lines are ignored. They can occur any place in the directives file or on LIBDECK.

The format of the directive is:

```
*/ comment
```

Parameter	Description
comment	A comment line can contain any valid characters and be used for any purpose.

***DELETE Directive**

The *DELETE directive deletes a record from the system. It cannot, however, delete a user library (ULIB type record).

The format of the directive is:

```
*DELETE, ty1/rec1, ty2/rec2, . . . , tyn/recn
```

or

```
*D, ty1/rec1, . . . , tyn/recn
```

Parameter	Description
ty _i /rec _i	Record type and record name to be deleted from the system. ty _i must not be ULIB.

***FILE Directive**

The ***FILE** directive declares an additional file containing records to be added to the system or to logically replace records on the system.

The format of the directive is:

```
*FILE, filename
```

or

```
*FILE, filename, NR
```

Parameter	Description
filename	Name of local file containing addition or replacement records to be placed on the system. File filename is rewound before processing if NR is omitted.
NR	Optional parameter that inhibits rewinding of file filename before processing.

***IGNORE Directive**

The ***IGNORE** directive specifies that records on a replacement file are to be ignored. If no ***FILE** directive precedes an ***IGNORE** directive, SYSEDIT ignores the records named on this directive on the replacement file specified by the B parameter. If one or more ***FILE** directives precede an ***IGNORE** directive, SYSEDIT ignores the records on the file specified in the most recent ***FILE** directive.

The format of the directive is:

```
*IGNORE, ty1/rec1, ty2/rec2, . . . , tyn/recn
```

Parameter	Description
ty _i /rec _i	Record type and record name to be ignored on the current replacement file.

***PPSYN Directive**

The ***PPSYN** directive specifies one or more names to be synonymous with the name of an existing PP routine.

The format of the directive is:

```
*PPSYN, name/name1, name2, . . . , namen
```

Parameter	Description
name	Name of existing PP routine.
name _i	Additional (synonymous) name for name.

Record Types

The following record types may be specified in SYSEEDIT directives.

Type	Description
ABS	Multiple entry point overlay.
CAP	Fast dynamic load capsule.
OPL	Modify old program library deck.
OPLC	Modify old program library common deck.
OPLD	Modify old program library directory.
OVL	Central processor overlay.
PP	Peripheral processor program.
PPU	Peripheral processor unit program.
PROC	Procedure record.
REL	Relocatable central processor program.
TEXT	Unrecognizable as a program.
ULIB	User library.

The system library contains a maximum of 62 ULIB type records. If more records are added, only the first 62 can be assessed.

System File Initialization

ISF 20-1

Deadstart Sequencing 20-4

This section describes the system mechanisms used to initialize the fast-attach files used by MODVAL, PROFILE, and RESEX.

ISF

The ISF command initializes the fast-attach system files VALIDUs, PROFILa, RSXDId, and RSXVId.¹ A fast-attach file is a special direct-access file under user index 377777₈ that is initialized with the E parameter on the ISF command and released with the R parameter on the ISF command. However, in order to release a fast-attach file, an idle family situation must be present. That is, the job containing the ISF,R=filename command must be the only job in the family (family count is zero, and the direct-access file count equals the number of fast-attach files). Therefore, the DSD command IDLEFAMILY must be used to clear the system of all jobs. All current jobs and subsystems (such as maintenance jobs, MAGNET, and IAF) must be allowed to complete or be terminated. Refer to the family status display (E,F) in the NOS 2 Operations Handbook to determine when the above conditions have been met.

The fast-attach file mechanism should be used by special system jobs for files that are to be retained as permanent files but have a high enough access rate to make permanent file ATTACHs excessively time-consuming. When a permanent file is activated as a fast-attach file, an entry in the system FNT is made that retains the basic data normally kept in the catalog entry and system sector of the file (interlocks and file name). This eliminates the catalog search and system sector read normally necessary to attach a permanent file.

If the file is a shared (global) fast-attach file for a multimainframe network, additional information is also maintained in the fast-attach table on the link device. The type of file determines whether it is entered as global fast attach or local fast attach. This criterion is kept internal to ISF. Basically, VALIDUs and PROFILa are entered as global fast-attach files if they reside on a shared device. The resource files are always entered as local fast attach. A limit of 77₈ exists on the number of files that can be entered as global fast attach in a multimainframe environment.

Because of the special nature of fast-attach files, a job containing an ISF command must be a system origin job. Processing the command causes a search of the system permanent file catalog (UI=377777₈) for files with the predefined names previously listed. They are defined in a table internal to ISF.

The format of the ISF command is:

```
SF,option,FM=familyname,SJ=filename,SP=filename.
```

1. The resource files are generated and maintained separately for each machine id in a multimainframe or single mainframe system by appending the machine id to the file name (for example, RSXVId becomes RSXVAB on the machine with an id of AB).

Parameter	Description						
option	Specifies whether to initialize or release system files.						
	<table border="1"> <thead> <tr> <th>option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>E=filename</td> <td>System file that is initialized. If E=0 or no filename is specified (neither E nor R appear), all files defined in the ISF table are initialized (refer to table 20-1).</td> </tr> <tr> <td>R=filename</td> <td>Currently active system file that is released from fast-attach status. If R=0, all of the files in the ISF table for the specified family that are currently active are released. When this parameter is specified, an idle family situation (family count is zero, and the direct-access file count equals the number of fast-attach files) must first be created with the DSD command IDLEFAMILY. When the family is idle, the IDLEFAMILY command must be entered again so that the system will accept the ISF command. To release fast-attach system files, you must use this format of the ISF command: X.ISF(R=filename) Since initialization and release are mutually exclusive, E and R cannot appear on the same command.</td> </tr> </tbody> </table>	option	Description	E=filename	System file that is initialized. If E=0 or no filename is specified (neither E nor R appear), all files defined in the ISF table are initialized (refer to table 20-1).	R=filename	Currently active system file that is released from fast-attach status. If R=0, all of the files in the ISF table for the specified family that are currently active are released. When this parameter is specified, an idle family situation (family count is zero, and the direct-access file count equals the number of fast-attach files) must first be created with the DSD command IDLEFAMILY. When the family is idle, the IDLEFAMILY command must be entered again so that the system will accept the ISF command. To release fast-attach system files, you must use this format of the ISF command: X.ISF(R=filename) Since initialization and release are mutually exclusive, E and R cannot appear on the same command.
option	Description						
E=filename	System file that is initialized. If E=0 or no filename is specified (neither E nor R appear), all files defined in the ISF table are initialized (refer to table 20-1).						
R=filename	Currently active system file that is released from fast-attach status. If R=0, all of the files in the ISF table for the specified family that are currently active are released. When this parameter is specified, an idle family situation (family count is zero, and the direct-access file count equals the number of fast-attach files) must first be created with the DSD command IDLEFAMILY. When the family is idle, the IDLEFAMILY command must be entered again so that the system will accept the ISF command. To release fast-attach system files, you must use this format of the ISF command: X.ISF(R=filename) Since initialization and release are mutually exclusive, E and R cannot appear on the same command.						
FM=familyname	Family of devices. If FM is not specified, the calling job's current family is used. The calling job's family will be restored upon exit from ISF.						
SJ=filename	Job file that ISF submits as a system origin job. The file must be an indirect-access permanent file stored under the system user index (377777 ₈). If SJ is specified without =filename, ISF assumes SJ=SYSJOB. If SJ=0 is specified, no job is submitted.						
SP=filename	Procedure file that ISF calls with system origin. The file must be an indirect-access permanent file stored under the system user index. If SP is specified without =filename, ISF assumes SP=SYSPROC. If SP=0 is specified, no procedure is called.						

ISF is automatically executed at each deadstart (refer to Deadstart Sequencing later in this section). This enables the use of SYSJOB and/or SYSPROC to SYSEdit local modifications into the system.

ISF can also be entered as a command from the console with the DSD X command or from any system origin job.

The matrix in table 20-1 shows how the initialize (E) and release (R) parameters affect individual fast-attach files.

Table 20-1. Initialize and Release Parameters of Fast-Attach Files

Name of Fast-Attach System File	Initialize (E) Parameter	Release (R) Parameter
VALIDUs PROFILa	Make global fast-attach file.	Return from fast-attach status to normal direct access.
RSXDid RSXVid	Make local fast-attach file. If the file does not exist in the system catalog (UI=377777 ₈), ISF creates the file and makes it a fast-attach file. The file is created with the backup requirement set to none (BR=N). This prevents PFDUMP from dumping the file. If either RSXDid or RSXVid is specified, ISF initializes both files.	Return from fast-attach status to normal direct access. If either RSXDid or RSXVid is specified, ISF returns both files to normal direct access.

One use of the R parameter is to release fast-attach files activated on a device to be initialized or unloaded. Device initialization is not initiated as long as any direct-access files are active on the device (an activated fast-attach file is treated the same as an active direct-access file). Until these files are released, the system will reply to an attempted device initialization with the error message:

ACTIVE FILES ON DEVICE

Similarly, a device cannot be unloaded until all its fast-attach files are released.

Deadstart Sequencing

During a level 0 deadstart, a job is queued for input using the service class DSSC. This job executes the PP routine CMS, whose normal functions include mass storage initialization and recovery. After completing these functions, CMS checks its service class. If the service class is DSSC (indicating a deadstart sequencing call), CMS places one of the following ISF commands in its command buffer and causes the system to execute it.

Command	Description
---------	-------------

ISF,FM=0,SJ.	Used for recovery deadstarts.
--------------	-------------------------------

ISF,FM=0,SJ,SP.	Used for other deadstarts.
-----------------	----------------------------

The CMS deadstart sequencing job begins execution only after job processing has been enabled by the DSD AUTO or MAINTENANCE command. The job scheduler prohibits scheduling of all other jobs until the CMS job completes, thereby ensuring that all system files in the default family are initialized and that a SYSPROC procedure (if any) is executed before normal job processing begins (refer to ISF earlier in this section).

NOTE

Since the deadstart job (including SYSPROC commands) must complete before normal job processing begins, avoid using SYSPROC for unnecessary or time-consuming tasks. SYSJOB may be more appropriate for some tasks. Also, SYSPROC cannot use tapes since MAGNET is not available. If SYSPROC uses removable disk packs, they must be mounted and ready.

Tracer Utility	21-1
TRACER Commands	21-4
ACPD Command	21-4
ENDCPD Command	21-5
ICPD Command	21-6
Output File Format	21-7
Summary File Format	21-17
Data Items Reported by TRACER	21-18
Fast Loop Items	21-21
Fast Loop Samples	21-21
PPs Active	21-21
Move Request Pending	21-21
No PP Available	21-22
No CPP Available	21-22
EM Transfer in Progress	21-22
MTR Cycle Time	21-22
Monitor Mode — CPU0/CPU1	21-22
Scheduler Active	21-23
Channel Reserved	21-23
Channel Active	21-24
Channel Requested	21-25
Requests Pending	21-25
Buffered I/O Lists	21-25
Buffered I/O Channel Busy	21-25
CPU Usage	21-26
Subsystem CPU Usage	21-27
Medium Loop Items	21-27
Medium Loop Samples	21-27
CPS in W Status	21-27
CPS in X Status	21-27
CPS in I Status	21-27
Same Move Request	21-27
FL Available	21-27
User EM Available	21-28
Noninteractive Jobs	21-28
Detached Jobs	21-28
Online Jobs	21-28
Preinitial Job Step	21-28
Executing	21-28
Scheduler Rollout	21-28
SCP Rollin	21-29
SCP Rollout	21-29
Timed/Event Rollout	21-29
Interactive Rollout	21-29
Disabled Rollout	21-29
Suspended Rollout	21-29
Rollout File Error	21-29
EJT Entries in Use	21-29
FL at Control Points	21-29

FL at Pseudo-control Points	21-30
FL in Rollout Queue	21-30
EM Memory at Control Points	21-30
EM Memory at Pseudo-control Points	21-30
EM in Rollout Queue	21-30
Subsystem FL	21-30
Slow Loop Items	21-31
Slow Loop Samples	21-31
IAF Users	21-31
IAF Pots Allocated	21-31
IAF Pots in Use	21-31
Queue Files Assigned	21-31
Input Files	21-31
Print Files	21-31
Punch Files	21-31
Other Queue Files	21-31
QFT Entries in Use	21-32
Tape Drives in Use	21-32
Tracks Available	21-32
CPPs Active	21-32
ISHARED Table Changes	21-32
ISHARED Device Updates	21-32
ISHARED Seek Time	21-32
ISHARED Updating Time	21-32
Statistical Summary	21-33
MTR Maximum Time	21-33
Direct Moves	21-33
Worst Case MTR Cycle Time	21-33
Missed Clock Updates	21-33
EM Moves	21-33
Central Memory Moves	21-33
Rollouts to MS	21-33
Pseudo-Rollouts	21-33
PCP Job Rollouts	21-33
Secondary Rollouts	21-33
Total Sectors Rolled	21-34
Secondary Sec Rolled to MS	21-34
Rollouts/User Limits	21-34
Time Slices	21-34
PP Priority Exchanges	21-34
No Comm Buffer Avail	21-34
No PCP Available	21-34
EJT Scans	21-34
Schedulable Jobs	21-34
Jobs Preempted	21-34
Jobs Scheduled	21-34
Scheduled No Constraints	21-35
Insufficient CM Scans	21-35
Insufficient EM Scans	21-35
No Control Point Scans	21-35
 PROBE Utility	 21-36

The TRACER and PROBE utilities described in this section provide data for statistical analysis of the system. The data is used to determine where problems occur and where improvements in design might be made, and to perform system tuning. The TRACER utility monitors the system's activity and gathers data periodically. The PROBE utility traps and measures particular internal events in the system. Both utilities capture valuable data that may not be obtainable any other way.

Tracer Utility

The TRACER utility monitors these conditions:

- Channel activity by channel.
- Channel reserved.
- Channel requested.
- Requests pending.
- Number of active PPs.
- Number of active concurrent PPs.
- Buffered input/output list parameters.
- Buffered input/output channel busy.
- CPU use (idle, system, subsystem, system-related activity, or user activity).
- Subsystem CPU use.
- Storage moves pending.
- PP saturation.
- Extended memory transfer in progress.
- MTR cycle time.
- CPU0 or CPU1 is or is not in monitor mode.
- Same storage move request is pending.
- Control points in automatic recall (I) status.
- Control points in periodic or automatic recall (X) status.
- Control points in waiting (W) status.
- Amount of available memory.
- Amount of memory at control points by service class.
- Amount of memory in queue by service class.
- Amount of memory at control points by subsystem.
- Number of noninteractive jobs.

- Number of detached interactive jobs.
- Number of online jobs.
- Number of preinitial jobs by service class.
- Number of executing jobs by service class.
- Number of jobs rolled out by job scheduler for each service class.
- Number of jobs rolled in or out by system control point processing for each service class.
- Number of jobs in a timed/event rollout queue by service class.
- Number of jobs rolled out by interactive input/output processing for each service class.
- Number of disabled jobs rolled out by service class.
- Number of suspended jobs rolled out by service class.
- Number of jobs with rollout file errors by service class.
- Number of EST entries in use.
- Number of FNT entries in use.
- Number of EJT entries in use.
- Number of queued files assigned to jobs at control points.
- Number of input files by service class.
- Number of print and punch files by service class.
- Number of other queued files not assigned to jobs at control points.
- Number of QFT entries in use.
- Number of FOT entries in use.
- Number of control points in use.
- Number of pseudo-control points in use.
- Number of IAF active users.
- Number of IAF pots available.
- Number of IAF pots in use.
- Number of tape drives in use.
- Number of tracks available by mass storage device.
- Number of segment table reads.
- Number of missed clock updates.
- Number of extended memory moves.
- Number of central memory moves.
- Number of all rollouts.

- Number of rollouts to secondary rollout devices.
- Number of all sectors rolled.
- Number of sectors rolled out to secondary rollout devices.
- Number of rollouts/user limits.
- Number of time slices.
- Number of PP priority exchanges.
- Number of times communication buffer not available.
- Number of EJT scans.
- Number of schedulable jobs.
- Number of jobs preempted.
- Number of jobs scheduled.
- Number of jobs scheduled with no constraints.
- Number of insufficient CM scans.
- Number of insufficient EM scans.
- Number of no control point scans.
- Number of ISHARED table changes.
- Number of ISHARED device updates.
- Amount of ISHARED seek time.
- Amount of ISHARED updating time.

The TRACER utility includes the following programs:

Program	Description
ICPD	A CPU program that initiates system monitoring by CPD.
CPD	A PP program that monitors any of the system activities just listed. CPD is dedicated to a PP while it is monitoring system activity. Data is written to a direct access permanent file for future analysis.
ACPD	A postprocessor program that generates an output report from the direct access permanent file written by CPD.
ENDCPD	A CPU program that terminates system monitoring by CPD.

TRACER Commands

TRACER commands are described next.

ACPD Command

ACPD reads the sample data file produced by CPD and generates reports, in both user-readable and machine-readable formats, for further analysis. The sample data file must be attached before ACPD is called. If ICPD is called with the M=A or M=M parameters, the sample data file can be accessed while CPD is still active.

ACPD assumes a continuity of the sample data file. Therefore, the uncollected information during the time gap separating two consecutive files is assumed to be present, although the information is not reported. As a result, if the consecutive files on the sample data file are not in chronological order, ACPD terminates and issues an error message.

Format:

ACPD, P₁, P₂ . . . P_n.

P_i	Description
FN=datafile	Name of sample file. Default is SAMPLE. This file is not rewound before or after processing.
L=outfile	Name of output file generated by ACPD. Default is OUTPUT.
S=sumfile	Name of machine-readable summary file generated by ACPD. If this parameter is omitted, no summary file is generated. If S is specified without sumfile, SUMMARY is assumed.
IC=nn	Select the report interval by specifying the number of CPD sample file records. This allows selection of report intervals of less than 1 minute. You can specify the report interval by using either the IC or IN parameter, but not both.
IN=nn	Time span of report interval, in minutes. Default value is 6 minutes. ACPD generates a report for each report interval. You can specify the report interval by using either the IN or IC parameter, but not both.
LO=Z	Report data items with zero values. Default is to suppress data items with zero values.
N=nn	Number of files on the sample data file to be analyzed and reported. Default is only one file processed. If N is not equivalenced, all files are processed until EOI is reached on the sample data file. The sample data file is not rewound before processing.
BT=hhmmss	Beginning time in the form hour minute second. If BT=hhmmss is specified, only data collected after this time on the date specified by BD=yymmdd is reported. If BT is omitted or no time is specified, BT=0 is assumed.

p_i	Description
ET=hhmmss	Ending time in the form hour minute second. If ET=hhmmss is specified, only data collected before this time on the date specified by ED=yymmdd is reported. If ET is omitted or no time is specified, ET=0 is assumed.
BD=yymmdd	Beginning date in the form year month day. If BD=yymmdd is specified, only data collected on or after this date is reported. If BD is omitted or no date is specified, the beginning date is the date of the current file on the sample data file.
ED=yymmdd	Ending date in the form year month day. If ED=yymmdd is specified, only data collected on or before this date is reported. If ED is omitted or no date is specified, but ET=hhmmss is specified, the ending date is the same as the beginning date. If neither ED nor ET is specified, ACPD terminates when the number of files specified by the N parameter are processed, or end-of-information is encountered, whichever happens first.

ACPD begins by processing the command parameters. If the beginning time (BT) and beginning date (BD) are specified, ACPD first locates the correct file and then processing begins. Processing continues until the ending time (ET) and ending date (ED) are encountered, the number of files specified by the N parameter have been processed, or end-of-information is encountered, whichever happens first.

If the BT and BD parameters are not specified, ACPD starts at the current position of the sample data file. The sample data file is not rewound before processing starts.

ENDCPD Command

The ENDCPD command terminates all CPD data gathering.

Format:

```
ENDCPD.
```

ICPD Command

ICPD defines a mass storage file to which CPD will write statistical data and then initiates system monitoring by CPD.

NOTE

If you want the statistical data to include only data from the start of this TRACER run rather than an accumulation of data from the last deadstart, enter the command PROBE(OP=C,L=O) before entering the ICPD command so that the statistical data area of CMR will be cleared prior to the reporting done by CPD.

Format:

ICPD, P₁, P₂, . . . P_n.

P_i	Description								
FL=fl	Fast loop sampling frequency, in milliseconds, during which items such as PPs active and move request pending are sampled. Default is 5 milliseconds.								
ML=ml	Medium loop sampling frequency, in milliseconds, during which items such as control points in I, X, and W status and field length available are sampled. Default is 100 milliseconds.								
SL=sl	Slow loop sampling frequency, in milliseconds, during which items such as IAF users and tape drives in use are sampled. Default is 1000 milliseconds.								
FW=fw	Snapshot loop sampling frequency, in seconds. Default is 5 seconds.								
FN=filename	Name of sample data file. Default is SAMPLE. ICPD will attempt to attach a direct access file by this name. If no file exists, it will be defined. If a file is found, ICPD will skip to EOI and write an EOF. CPD will then start writing data after the EOF.								
M=mode	Permanent file mode for sample data file. Default is M=WRITE. mode can have one of these values:								
	<table border="1"> <thead> <tr> <th>mode</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>WRITE or W</td> <td>Sample data file attached in write mode.</td> </tr> <tr> <td>APPEND or A</td> <td>Sample data file attached in append mode.</td> </tr> <tr> <td>MODIFY or M</td> <td>Sample data file attached in modify mode.</td> </tr> </tbody> </table>	mode	Description	WRITE or W	Sample data file attached in write mode.	APPEND or A	Sample data file attached in append mode.	MODIFY or M	Sample data file attached in modify mode.
mode	Description								
WRITE or W	Sample data file attached in write mode.								
APPEND or A	Sample data file attached in append mode.								
MODIFY or M	Sample data file attached in modify mode.								

NOTE

If the sample data file is attached in write mode, the file cannot be accessed until ENDCPD is run. If the sample data file is to be accessed while data is being collected, append or modify mode must be specified. In this situation, the file may be attached in read/allow modify (RM) mode. (Attaching the file in write mode rather than in modify or append mode expends less overhead when interlocking and writing the data file.)

If a loop time is set to 0 (zero), no samples for that loop will be taken. If the data block sample time is set to 0 (zero), the data file will be written only when the sampling interval terminates.

All numeric data should lie within the range 0 through 4095 (0 through 7777₈).

Output File Format

The first three pages of the output report produced by ACPD contain the header block information. Next the data items are reported for fast, medium, and slow loop samples. The report ends with the snapshot data items.

Figure 21-1 is an example showing the format of the output report. The example has been simplified and condensed to reduce the amount of output. Also, supporting text has been added to the example.

Data items monitored at successive time intervals are listed in the same row. For each data item, the average, standard deviation, and percentage are listed in successive rows. Up to 10 intervals can be listed per page in successive vertical columns. If the output file contains more than 10 columns per row, the output report lists the first 10 columns for all rows of data items and then resumes listing subsequent intervals following the snapshot data items.

The SUBTOTAL column contains the values of the data items for the time spanned by the preceding intervals on the current page; that is, the time spanned by the preceding intervals is considered one interval. The TOTAL column appears after the last interval reported and contains the statistical values of the data items for the entire run. The SUBTOTAL column is not listed if the subtotal data and total data are identical and would appear on the same page. In this case, only the TOTAL column is listed. The SUBTOTAL and TOTAL columns are not reported for the snapshot data items. The *MAX* and *MIN* columns appear at the end of the report and contain the maximum and minimum interval values of the data items for the entire run. The maximum and minimum interval values on each page are indicated by brackets and parentheses, respectively.

The average is not reported for data items that have a weighting factor of 1, and the percentage is not reported for data items that have a weighting factor of 100, since this information is redundant.

A C P D - VER 9.1		CDC NETWORK OPERATING SYSTEM.	NOS 2	yy/mm/dd. hh.mm.ss. PAGE 1	
START DATE	86/05/25.			} Start of sampling interval	
START TIME	23.50.13.				
DATA FILE NAME	SAMPLE			} ACPD command parameters (FN, IC, IN)	
REPORT INTERVAL (MINUTES)	2				
CPD VERSION	9.1			} ICPD command parameters (FL, ML, SL, PW)	
FAST LOOP INTERVAL (MSECS)	10				
MEDIUM LOOP INTERVAL (MSECS)	100				
SLOW LOOP INTERVAL (MSECS)	1000				
SNAPSHOT LOOP INTERVAL (SECS)	60				
NUMBER OF CPUS	1			} Hardware configuration at beginning of sampling interval	
NUMBER OF MPUS	20				
NUMBER OF CPFS	10				
MEMORY SIZE / 100B	7777B				
USER EN / 1000B	200B				
MAGNETIC TAPE DEVICES	4			} Software configuration at beginning of sampling interval	
NUMBER OF EST ENTRIES	76B				
NUMBER OF FWT ENTRIES	23B				
NUMBER OF EJT ENTRIES	1000B				
NUMBER OF QFT ENTRIES	1000B				
NUMBER OF PQT ENTRIES	10B				
NUMBER OF CONTROL POINTS	33B				
NUMBER OF PCP-S	20B				
CHSIZE / 100B	1632B				
LINKAGE NUMBER	0				
RECOVERY LEVEL	0				
NUMBER OF IAP TERMINALS	400B				
MACHINE ID	32				
CPU USAGE CATEGORIES (CTOUL)	5B				} System statistics; the symbols are assembly constants used in NOS
CHANNEL TABLE LENGTH (CDALL)	55B				
LENGTH OF AN EST ENTRY (ESTE)	2B				
LENGTH OF AN FWT ENTRY (FWTE)	2B				
LENGTH OF AN EJT ENTRY (EJTE)	4B				
LENGTH OF A QFT ENTRY (QFTE)	4B				
NUMBER OF SERVICE CLASSES (MSC)	17B				
NUMBER OF JOB STATUSES (OJLS)	16B				
NUMBER OF CONNECTION STATUSES (MCS)	3B				
NUMBER OF FILE TYPES (MFT)	21B				
MAXIMUM ORIGIN TYPES (MOT)	6B				
NUMBER OF MASS STORAGE DEVICES	11B			} System delay parameters at beginning of sampling interval	
LENGTH OF MSAL TABLE (MIRS)	14B				
CPU RECALL DELAY	30B				
PP AUTO RECALL DELAY	1750B				
CPU PRIORITY INCREMENTING FREQUENCY	200B				
MEMORY PAB SIZE	400B			} System buffers at beginning of sampling interval	
INPUT FILE SCHEDULING INTERVAL	2B				
SCHEDULER CYCLE INTERVAL	1B				
HIGH SPEED DISK BUFFERS	0B				
EN/PP BUFFERS	1B				

Figure 21-1. Example of TRACER Output (Continued)

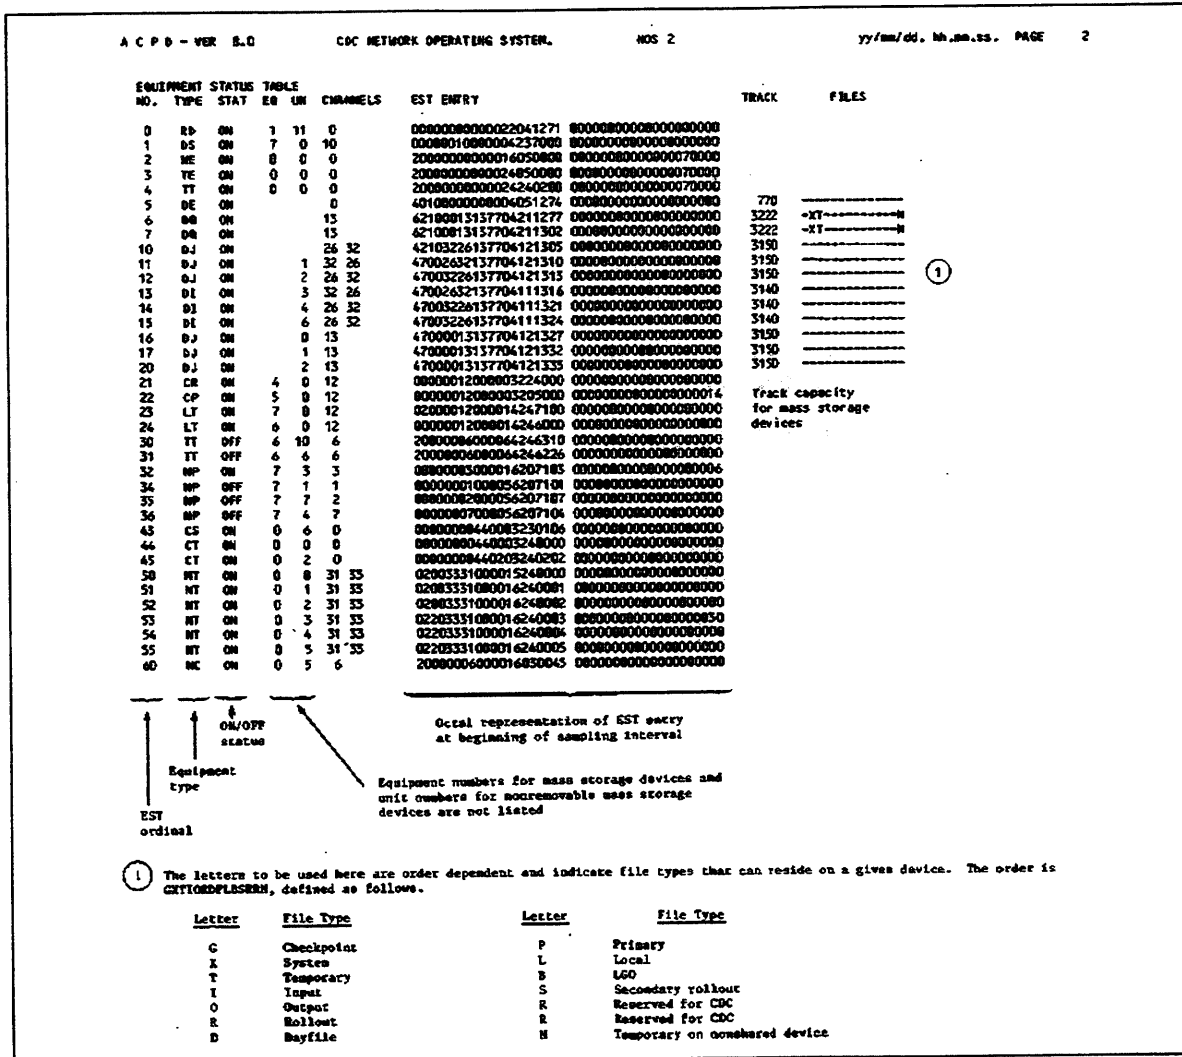


Figure 21-1. Example of TRACER Output (Continued)

A C P D - VER 9.4		COC NETWORK OPERATING SYSTEM.					NOS 2		yy/mm/dd. hh:mm:ss. PAGE 33				
SYSTEM CONTROL INFORMATION													
SERVICE CLASS	QUEUE	PRIORITIES					SERVICE LIMITS						
		IL	LP	UP	MF	TP	CP	CT	CN	NJ	TO		
SY	IN		7770	7776	1		6770	10	20	7777	341		
	EX	4000	2000	7800	1	7000	3777	7777	4004	1	DI		
	OT		7000	7776	1		3777	7777	0	0	0		
							30	20	4	20			
BC	IN		10	4000	1		3770	20	200	7777	341		
	EX	2000	1000	4000	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
RB	IN		10	4000	1		3770	20	200	7777	341		
	EX	2000	1000	4000	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
TS	IN		7000	7770	1		6770	10	10	7777	341		
	EX	3770	3700	7000	1	4000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
DI	IN		10	4000	1		3770	20	20	7777	341		
	EX	2000	1000	4000	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
NS	IN		7770	7776	1		7770	20	200	7777	341		
	EX	7772	7770	7776	1	7772	3777	7777	4004	14	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							74	20	4	20			
SS	IN		7770	7776	1		7770	10	20	7777	341		
	EX	7772	7770	7776	1	7772	3777	7777	4004	6	DI		
	OT		7400	7776	1		3777	7777	0	0	0		
							70	20	4	20			
PA	IN		1	10	1		1	10	20	7777	341		
	EX	4	1	10	1	10	3777	7777	4004	0	DI		
	OT		7000	7776	1		3777	7777	0	0	0		
							2	20	4	20			
CT	IN		7770	7776	1		6770	10	200	7777	341		
	EX	4004	4000	7000	1	7000	3777	7777	4004	2	DI		
	OT		7000	7776	1		3777	7777	0	0	0		
							30	20	4	20			
IB	IN		10	4000	1		3770	10	200	7777	341		
	EX	2000	1000	4004	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
II	IN		10	4000	1		3770	10	200	7777	341		
	EX	2000	1000	4004	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
I2	IN		10	4000	1		3770	10	200	7777	341		
	EX	2000	1000	4004	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			
I3	IN		10	4000	1		3770	10	200	7777	341		
	EX	2000	1000	4004	1	2000	3777	7777	4004	0	DI		
	OT		1	7000	1		3777	7777	0	0	0		
							30	20	4	20			

This page shows the status of the job control area at the beginning of the sampling interval. The values are set IPDECK entries QUEUE and SERVICE or by the DSD commands QUEUE and SERVICE.

Figure 21-1. Example of TRACER Output (Continued)

ACPD - VER 3.0		CDC NETWORK OPERATING SYSTEM				NOS 2		yy/mm/dd. hh.mm.ss. PAGE 4				
2	MINUS INTERVAL	25.52.15	25.54.15	25.56.15	25.58.12	00.00.12	00.02.13	00.04.12	00.06.13	00.08.13	00.10.13	HP SUBTOTAL
3	FAST LOOP SAMPLES	10480	10519	10525	10517	10528	10524	10516	10507	10517	10517	94433
	PLUS ACTIVE	(4.236)	6.272	4.907	7.186	[7.788]	4.352	4.504	4.250	6.154	5.539	
	AV	(0.014)	0.139	0.118	0.191	(0.365)	0.238	0.054	0.021	0.143	1.281	
	SD	(21.188)	31.359	24.933	35.951	[38.588]	22.748	22.519	21.250	30.770	27.693	
	PC	(0.080)	0.080	0.080	0.080	(0.124)	0.048	0.000	0.000	0.067	0.035	
	NO PPU AVAILABLE	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
	SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
	PC	(0.048)	0.152	0.105	0.095	0.209	0.162	(0.209)	0.086	0.057	0.125	
	ER TRANSFER IN PROG	0.000	0.000	0.000	0.000	0.001	0.000	(0.001)	0.000	0.000	0.001	
	SD	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	0.000	0.001	
	PC	(0.048)	0.152	0.105	0.095	0.209	0.162	(0.209)	0.086	0.057	0.125	
	NRR CYCLE TIME	0.371	0.393	0.372	(0.436)	0.412	0.372	0.370	0.364	0.433	0.392	
	SD	0.010	0.007	0.008	(0.000)	0.013	0.007	0.001	0.004	0.023	0.029	
	PC	(0.013)	0.002	0.012	0.009	0.003	(0.002)	0.010	0.002	0.005	0.020	
	MONITOR NODE - CPU 0	(7.887)	12.358	8.688	(13.605)	10.961	8.200	9.878	8.576	9.186	9.895	
	SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
	PC	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
	MONITOR NODE - CPU 1	(0.001)	0.001	0.006	0.001	0.002	0.019	0.001	0.001	(0.069)	0.057	
	SD	0.001	0.001	0.006	0.001	0.002	0.019	0.001	0.001	(0.069)	0.057	
	PC	(0.001)	0.001	0.006	0.001	0.002	0.019	0.001	0.001	(0.069)	0.057	
	SCHEDULER ACTIVE	5.114	(5.067)	7.335	15.915	(19.136)	7.051	5.696	5.805	12.110	9.248	
	SD	0.000	0.000	(0.000)	0.000	0.054	0.000	0.001	0.000	(0.068)	0.054	
	PC	0.105	0.076	0.076	0.095	3.776	0.095	0.095	(0.058)	6.898	1.250	
	CH 12	0.010	0.262	0.012	(0.005)	0.053	0.059	0.016	0.005	(0.275)	0.168	
	SD	0.010	0.262	0.012	(0.005)	0.053	0.059	0.016	0.005	(0.275)	0.168	
	PC	7.781	(59.808)	11.135	16.810	12.914	19.718	25.217	(7.434)	33.625	19.583	

- 2 The interval ending time is shown at the top of each report column. The report interval is specified by the IC or IN parameter of the ACPD command.
- 3 The sampling frequency for fast loop items is specified by the PL parameter of the ICPD command. Refer to table 21-1 for all fast loop items that TRACER reports. Fast loop items are described following the table.
- 4 The statistics reported for each data item are the average (AV), standard deviation (SD), and percentage (PC) for the report interval.
- 5 The SUBTOTAL column contains a summation of all data gathered from all intervals of each page. This column is not listed if the subtotal data and the total data are identical and would appear on the same page.
- 6 This example has 13 columns of information for each row of data items. However, only 10 columns can be shown per page, so the remaining columns are listed at the end of the report.
- 7 The values enclosed in brackets are the maximum interval values for each row of data items on each page. Refer to note 14.
- 8 The values enclosed in parentheses are the minimum interval values for each row of data items on each page. Refer to note 15.

Figure 21-1. Example of TRACER Output (Continued)

9	MEMOR LOOP SAMPLES		1170	1175	1175	1175	1175	1175	1175	10566
	CPS IN W STATUS	AV	(0.225)	0.339	0.689	1.868	(2.494)	1.699	2.085	1.282
		SD	0.863	0.744	(0.328)	0.113	0.295	(0.802)	0.122	0.134
		PC	(0.805)	1.209	3.175	6.671	(8.907)	6.067	7.444	4.579
	CPS IN X STATUS	AV	(5.787)	6.452	5.894	(5.802)	6.878	5.986	6.088	5.827
		SD	0.020	0.014	(0.003)	0.044	0.101	0.009	0.182	(0.190)
		PC	(20.669)	23.043	21.052	(24.293)	21.705	21.099	21.714	20.811
	CPS IN J STATUS	AV	(8.067)	0.130	0.083	(0.185)	0.177	0.091	0.095	0.075
		SD	(0.001)	0.016	0.015	0.004	0.005	0.011	0.006	0.002
		PC	(8.258)	8.445	8.795	(8.643)	8.432	8.325	8.340	8.268
	SAME PHONE REQUEST	SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		PC	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	FL AVAILABLE	AV	(434.78)	42348	33658	26758	24718	30178	(24128)	33758
		SD	(1.227)	7.478	22.167	82.904	(126.290)	5.241	60.961	9.916
		PC	(61.750)	59.760	48.185	39.794	36.213	42.016	(34.957)	48.459
	USER IDN AVAILABLE	AV	00	00	00	00	00	00	00	00
		SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		PC	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	NON INTERACTIVE JOBS	AV	(1.022)	2.896	2.668	4.263	(6.591)	2.603	2.351	2.191
		SD	0.136	(0.023)	0.456	(0.619)	0.139	0.462	0.258	0.033
		PC	(0.177)	0.364	8.463	0.757	(1.444)	0.452	0.488	0.380
	DETACHED JOBS	AV	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		PC	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	ON-LINE JOBS	AV	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
		SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		PC	0.174	0.174	0.174	0.174	0.174	0.174	0.174	0.174
	PRE-INITIAL JOB STEP	AV	(1.000)	1.000	1.001	1.029	1.000	1.000	1.000	1.001
		SD	(8.080)	0.080	0.008	0.029	0.080	0.000	0.000	0.001
		PC	(0.174)	0.174	0.174	0.179	0.174	0.174	0.174	(0.233)
	BC	AV	(0.025)	1.000	0.684	1.166	(1.665)	0.094	0.236	0.092
		SD	8.025	(0.000)	0.125	0.166	0.335	0.094	0.236	0.029
		PC	(0.004)	0.174	8.119	0.202	(0.289)	8.016	0.041	0.016
	DB	AV	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		PC	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	TS	AV	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		PC	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

9 The sampling frequency for medium loop items is specified by the UL parameter of the ICPD command. Refer to table 21-1 for all medium loop items that TRACER reports. Medium loop items are described following the table.

Figure 21-1. Example of TRACER Output (Continued)

⑩

```

*****
SLOW LOOP SAMPLES
IAF USERS
IAF POTS ALLOCATED
IAF POTS IN USE
QUEUE FILES ASSIGNED
INPUT FILES
PRINT FILES
    
```

	119	119	119	119	119	119	119	119	119	1071
IAF USERS	AV 1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.781	0.781	0.781	0.781	0.781	0.781	0.781	0.781	0.781	0.781
IAF POTS ALLOCATED	AV 455.000	455.000	455.000	455.000	455.000	455.000	455.000	455.000	455.000	455.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 193.000	193.000	193.000	193.000	193.000	193.000	193.000	193.000	193.000	193.432
IAF POTS IN USE	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 42.418	42.418	42.421	42.419	42.421	42.418	42.418	42.418	42.418	42.526
QUEUE FILES ASSIGNED	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.002	0.000	0.002	0.012	0.000	0.000	0.000	0.000	0.004
INPUT FILES										
SY	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
BC	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MB	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
TS	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
PRINT FILES										
SY	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.004
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
BC	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MB	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
TS	AV 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	SB 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	PC 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

⑩ The sampling frequency for slow loop items is specified by the SL parameter of the IOPD command. Refer to table 21-1 for all slow loop items that TRACER reports. Slow loop items are described following the table.

Figure 21-1. Example of TRACER Output (Continued)




Figure 21-1. Example of TRACER Output (Continued)

A C P D - VER 9.4 (31) CY180-860 SN302. OPEN SHOP. NOS 750C2/24R1/3A. Date/Time	
6 MINS INTERVAL	13.02.05
	INTERVAL
MTR MXN TIME	436
WORST CASE MTR CYCLE TIME	4037
MISSED CLOCK UPDATES	0
EM MOVES	0
CM MOVES	49
DIRECT MOVES	41
ROLLOUTS TO MS	53
PSEUDO-ROLLOUTS	0
PCP JOB ROLLOUTS	0
SECONDARY ROLLOUTS	30
SECTORS ROLLED TO MS	9228
SECONDARY SEC ROLLED	3099
ROLLOUTS/USER LIMITS	0
TIME SLICES	28
PP PRIORITY EXCHANGES	0
NO COMM BUFFER AVAIL	0
NO PCP AVAILABLE	0
EJT SCANS	437
SCHEDULABLE JOBS	79
JOBS PREEMPTED	0
JOBS SCHEDULED	54
SCHEDULED NO CONSTRAINTS	0
INSUFFICIENT CM SCANS	0
INSUFFICIENT EM SCANS	0
NO CONTROL POINT SCANS	0

Figure 21-1. Example of TRACER Output (Continued)

Tracer Utility

A C P D - VER 9.4	(31) CY180-860 SN302. OPEN SHOP. NOS 750C2/24R1/3A.	Date/Time
6 MINS INTERVAL	13.02.05	
	INTERVAL	
PRG MD FUNCTIONS TIME	0000	
	0000	
	0245	
	7164	
	0240	
STORAGE MOVE TIME	0000	
	0000	
	0013	
	0701	
	6700	
PSEUDO-PP TIME	0000	
	0000	
	0000	
	0000	
	0000	
CPU 0 IDLE TIME	0000	
	0017	
	5253	
	4501	
	1200	
CPU 1 IDLE TIME	0000	
	0000	
	0000	
	0000	
	0000	
SNAP SHOT OF RTCL	0040	
	0370	
	0000	
	0075	
	0106	
SNAP SHOT OF PDTL	0000	
	0000	
	2403	
	1015	
	0205	

Figure 21-1. Example of TRACER Output (Continued)

A C P D - VER 3.0		CDC NETWORK OPERATING SYSTEM		NOS 2		yy/mm/dd. hh.mm.ss. PAGE 30	
2 WING INTERVAL		00:18 HR	03/05/25	TO	03/05/26		
		(13) TOTAL	(14) *MAX*	(15) *MIN*			
(12) FAST LOOP SAMPLES		94633					
PPUS ACTIVE	AV	5.539	7.708	4.256			
	SD	1.284	0.343	0.016			
	PC	27.695	38.538	21.180			
MOVE REQUEST PENDING	SD	0.001	0.001	0.000			
	PC	0.035	0.124	0.000			
NO PPU AVAILABLE	SD	0.000	0.000	0.000			
	PC	0.000	0.000	0.000			
EN TRANSFER IN PRG	SD	0.001	0.001	0.000			
	PC	0.125	0.299	0.048			
MTR CYCLE TIME	AV	0.392	0.436	0.366			
	SD	0.029	0.023	0.000			
MONITOR NODE - CPU 0	SD	0.020	0.013	0.002			
	PC	9.895	15.605	7.887			
MONITOR NODE - CPU 1	SD	0.000	0.000	0.000			
	PC	0.000	0.000	0.000			
SCHEDULER ACTIVE	SD	0.057	0.069	0.001			
	PC	9.248	19.136	5.067			
CHANNEL RESERVED							
CH 12	SD	0.054	0.068	0.000			
	PC	1.250	6.898	0.038			
CH 13	SD	0.168	0.275	0.003			
	PC	19.345	39.808	7.434			

(12) These fast loop items are the remaining columns from page 4 of the example. Refer to note (6).

(13) The TOTAL column contains the total samples, average, standard deviation, and percentage for the entire ACPD run.

(14) The *MAX* column contains the maximum interval value for each row of preceding intervals for the entire ACPD run. Refer to note (7).

(15) The *MIN* column contains the minimum interval value for each row of preceding intervals for the entire ACPD run. Refer to note (8).

Figure 21-1. Example of TRACER Output

Summary File Format

The machine-readable summary file has two types of records, the header block record and the data block record. An *EOR* separates consecutive records.

The header block record contains the header data in an unpacked format.

Each data block record is divided into two equal length parts. The first part contains the average values of the data items. The second part contains the corresponding standard deviations of the data items.

Total and subtotal data and interval percentage data does not appear on the summary file. The loop sample times and snapshot data items have 0 standard deviations to simplify the summary file format.

Data Items Reported by TRACER

The data items reported by TRACER are described in the following paragraphs and are summarized in tables 21-1 through 21-3. TRACER increments each item's counter when the item is in a given state, and periodically writes the contents of the counters to the data file for future processing by ACPD. When and how the counter for a given item is incremented is also discussed.

Suggestions are given to help you determine system performance. Performance may be affected by several areas within the system; therefore, a full analysis should be done prior to making any changes.

Table 21-1. Fast Loop Data Items

Fast Loop Items	Weighting Factor
Fast loop samples	None
PPs active	Number of PPs available
Move request pending	1 ¹
No PP available	1
No CPP available	1
EM transfer in progress	1
MTR cycle time	100 ²
Monitor mode — CPU0	1
Monitor mode — CPU1	1
Scheduler active	1
Channel reserved	1
Channel active	1
Channel requested	1
Requests pending	100
Buffered I/O lists	Number of I/O buffers
Buffered I/O channel busy	1
CPU usage — CPU0	1
CPU usage — CPU1	1
Subsystem CPU usage	1

1. The average value will not be reported if the weighting factor is 1.

2. The percentage value will not be reported if the weighting factor is 100.

Table 21-2. Medium Loop Data Items

Medium Loop Items	Weighting Factor
Medium loop samples	None
CPS in W status	Number of control points
CPS in X status	Number of control points
CPS in I status	Number of control points
Same move request	1
FL available	Available field length (machine size — CMR size)
User EM available	Available user extended memory field length
Noninteractive jobs	Number of EJT entries
Detached jobs	Number of EJT entries
Online jobs	Number of EJT entries
Preinitial job step	Number of EJT entries
Executing	Number of EJT entries
Scheduler rollout	Number of EJT entries
SCP rollin	Number of EJT entries
SCP rollout	Number of EJT entries
Timed/event rollout	Number of EJT entries
Interactive rollout	Number of EJT entries
Disabled rollout	Number of EJT entries
Suspended rollout	Number of EJT entries
Rollout file error	Number of EJT entries
EJT entries in use	Number of EJT entries
FL at control points	Available field length
FL at pseudo-control points	Available field length
FL in rollout queue	Available field length
EM memory at control point	User extended memory size/1000 ₈
EM memory at pseudo-control point	User extended memory size/1000 ₈
EM in rollout queue	User extended memory size/1000 ₈
Subsystem FL	Available field length

Table 21-3. Slow Loop Data Items

Slow Loop Items	Weighting Factor
Slow loop samples	None
IAF users	Number of terminals defined
IAF pots allocated	100
IAF pots in use	IAF pots allocated
Queue files assigned	Number of QFT entries
Input files	Number of QFT entries
Print files	Number of QFT entries
Punch files	Number of QFT entries
Other queue files	Number of QFT entries
QFT entries in use	Number of QFT entries
Tape drives in use	Number of available tape drives
Tracks available	Maximum number of tracks for monitored device
CPPs active	Number of CPPs available
ISHARED table changes	None
ISHARED device updates	None
ISHARED seek time	None
ISHARED updating time	None

Fast Loop Items

The following paragraphs describe fast loop items. Fast loop items are continuously changing, so they should be sampled frequently.

Fast Loop Samples

TRACER increments the fast loop samples counter each time the fast loop items are sampled.

PPs Active

TRACER increments the PP active counter for each active PP it finds when scanning the PP communication area. Disabled PPs (those turned off at deadstart time) are not considered active. An active PP is one that has a nonzero input register. MTR (PP0) and DSD (PP1) are always counted as active PPs.

If this counter's percentage exceeds 80 percent and the no PP available counter is greater than 20 percent, you may need more PPs. Check the channel active counters for disks. If the disk channels range from 35 to 40 percent, you may not need additional PPs; instead, look at what is causing your disk channels to be so busy. If you can reduce their activity to below 30 percent, the percentages should also drop for the PPs active counter and the no PP available counter.

Examine which PPs are dedicated. If the percentage of floating PPs is low (about 20 percent of the total PPs configured), you may need more PPs.

Move Request Pending

TRACER increments the move request pending counter whenever it determines, from word CMCL of CMR, that a storage move request is outstanding.

If this counter's percentage ranges from 30 to 40 percent, the system is degrading. The following items could cause excessive storage movements.

- Are subsystems at the highest or lowest control points? If not, NOS may be moving them up or down in memory frequently in order to satisfy other user job memory requests. If a subsystem like NAM is being moved, its users may experience slow response time since a subsystem cannot execute while it is moving in memory.
- Try to stabilize the field lengths of your subsystems. For example, if TAF is not built correctly for your site's use, it may roll out and roll in quite often. Depending on your CM constraints, you may want TAF to roll out less frequently. You should evaluate each subsystem and place it at the correct control point.
- Are some user jobs changing their field length excessively? For example, jobs that open and close many files repeatedly cause CMM to adjust the user's field length, which causes other job movement within memory. This may save on field length but takes resources from the system, which increases system overhead. This type of job should be changed to be more efficient.

No PP Available

TRACER increments the no PP available counter each time it determines, from word PPAL of CMR, that there are no PPs available.

The percentage for this counter should be examined with the percentage for the PP active counter. If the PP active percentage is low (less than 50 percent) but you have counts of no PPs available, it means that you have peaks in your work load and this may be acceptable. However, if the PP active percentage is high (from 80 to 90 percent) and you have counts of no PPs available, you may need more PPs.

No CPP Available

TRACER increments the no CPP available counter each time it determines, from word PPAL of CMR, that there are no CPPs available.

EM Transfer in Progress

TRACER increments the extended memory transfer in progress counter whenever the S/C register (maintenance register for models 865 and 875) indicates an extended memory transfer is active. This data is available for CYBER 170 Computer Systems except models 176, 815, 825, 835, 845, and 855.

MTR Cycle Time

TRACER increments the MTR cycle time counter by the cycle time for the last MTR cycle. This item indicates how fast MTR is completing one complete scan of all PP output registers and processing those functions present.

The MTR cycle time should not exceed 0.5 seconds. If it does, MTR could miss updating the real time clock. Refer to MTR MXN TIME reported in the statistical data area from CMR (figure 21-1) for the worst case.

Monitor Mode — CPU0/CPU1

TRACER increments the monitor mode counter for CPU0/CPU1 whenever the S/C register (maintenance register for models 865 and 875) indicates that the CPU0/CPU1 is in monitor mode. This data is available for CYBER 170 Computer Systems except models 815, 825, 835, 845, and 855.

During monitor mode, the CPU is executing the code of CPUMTR, which performs NOS overhead. During program mode, CPUMTR is executing to perform a function that a user job requested and only CPUMTR has the capability or security to do.

When the monitor mode counter reaches 20 percent for each CPU on a dual CPU configuration or 30 percent for a single CPU machine, you should examine the data provided by PROBE. PROBE reports each monitor request for both monitor mode and program mode. Examine the requests to determine what functions were called in excess. As a basis of comparison, use a PROBE report for which system performance was good and the monitor mode statistics were normal.

Scheduler Active

TRACER increments the scheduler active counter whenever it determines that the job scheduler (1SJ) is active. The job scheduler is considered active when the scheduler active flag in word JSCL is set.

The job scheduler is called to process input queues and rollout queues. The scheduler searches all of the corresponding queues, starting with the last EJT or QFT entry scheduled to a control point. It schedules the highest priority jobs into execution. Queue length and decision processing determines the amount of time the scheduler spends. Limit the size of the queues to the amount shown in the EJT entries in use and the QFT entries in use statistics.

If the scheduler active counter reaches 70 to 90 percent, you may notice problems. The reason is that 1SJ is called to process EJT entries or QFT entries, not both. For example, if 1SJ spends all of its time processing EJT entries, no input queue jobs are processed. In this case, interactive users entering the system for the first time will never get started. Also, excessive rollin/rollout because of central memory constraints or CM/CP time slice expirations could cause the scheduler to be busy. Examine the SERVICE parameters; you might have to do some tuning.

Channel Reserved

TRACER maintains a channel reserved counter for each available channel or concurrent channel and increments a counter whenever its channel is logically reserved by the operating system. TRACER samples the channel status table for this data.

A PP program may have a channel reserved without actually using it to transfer data. This can cause a problem only if the channel has a high reserve percentage (from 80 to 100 percent) and the corresponding activity for the channel is low (less than 20 percent). The exceptions are those channels that are dedicated to PPs such as DSD, PIP, and 1HP.

Channel Active

TRACER maintains a channel active counter for each available channel and increments a counter whenever it detects that the channel is not inactive, as determined by an IJM PP instruction.

System performance is affected most by disk channel activity. A percentage of from 10 to 30 percent is considered normal. When this percentage reaches 35 to 40 percent, your system may experience performance degradation. High channel activity could be caused by the following factors. If you can scale down these problems, you may decrease channel activity and improve system performance; otherwise, you may need more disk controllers.

- Do all the disk channels show similar activity levels? If not, examine the attributes you have described for each disk unit. For example, what units contain permanent files, rollout files, a copy of the system, and temporary files? You may have to redistribute the work load by changing attributes. Determine which files are busy and the devices on which they reside using the PFCAT utility. You may find that altering device masks and/or moving files will lower the channel activity. You are looking for an even percentage of work performed by each of the disk channels.

The system selects file residency based on several factors. If the file is a permanent file, the system uses the device masks you have set for your family. If it is a temporary file or rollout file, the system looks at the attributes set for each device and selects the best candidate based on the following.

- First, the system looks at the number of tracks available on each potential device. If any of the devices have fewer tracks available than the low space threshold set by the THRESHOLD command, those devices will no longer be candidates for selection unless all of the other potential devices are also below their threshold levels.
- Second, the system picks (from the remaining candidates) the device with the least activity.
- Are there an excessive number of rollouts? User jobs may be reaching the CM or CP time slice limits. Check your SERVICE parameters and examine the number of time slices reported in the statistical data area from CMR. The system may be rolling jobs in and out of the control points if your SERVICE limits are too small. By adjusting the limits, you should improve system performance.
- Are large user jobs or subsystems rolling in or out? This can be determined by watching the DSD rollout status display (R). One large job can cause problems if it uses too many resources. Redesigning or rescheduling the job may help.
- Examine the PROBE utility output. PROBE reports the PP programs that were loaded (program name, residency, and number of loads). Any program that is loading more than once per second, should be made central memory resident. If you are short of central memory, place these programs in extended memory.

Channel Requested

TRACER maintains a channel requested counter for each available channel or concurrent channel and increments a counter whenever there is an outstanding request for that channel. TRACER uses the channel status table to determine the channel requested status.

This counter tells you how often PP programs have requested a channel but were unable to reserve it because another PP program had it reserved. A normal count is from 20 to 50 percent. The exceptions are those channels that are dedicated to PPs such as DSD, PIP, and 1HP.

This can be used for debugging purposes when developing your own PP programs that require channels. Be sure that you follow the rules governing channel dialog.

Requests Pending

TRACER maintains a request pending counter for each mass storage device available and increments a counter by the number of outstanding requests on that device when the sampling occurs. The MST provides this information.

Ensure that the number of requests pending is evenly distributed between all units. If not, you may have to redistribute files (permanent, system, temporary, and rollout files) for the devices to obtain a more even distribution. Adding more disk units to a configuration will also aid in reducing the number of requests outstanding.

Buffered I/O Lists

TRACER maintains a counter for each type of buffered I/O buffer list (Empty, Data Written, Read, Write). The number of buffers assigned to each list is incremented based on the buffered I/O data tables.

An average of 25 percent of the buffers on the Empty and Data Written lists is normal. The Write list should contain from 2 to 3 percent of the buffers, but a higher percentage is acceptable. The Read list is normal at from 65 to 75 percent. If there seems to be a shortage of Empty and Data Written buffers, you can add more extended memory for I/O buffer space.

Buffered I/O Channel Busy

TRACER maintains a PP driver busy counter for each buffered I/O channel and increments a counter whenever a data transfer associated with that channel is in progress. For 885-4x disks, the data transfer occurs through the low-speed port to ESM rather than the channel; only control and status information transfers over the channel. TRACER uses the buffered I/O channel control table to determine the PP driver busy status.

The normal percentage is from 15 to 20 percent for a model 760 with two channels. This percentage may be higher for machines with slower CPUs.

CPU Usage

TRACER maintains a set of CPU usage counters for each CPU. There is a counter for each type of CPU use.

Counter	CPU Use
IDLE	CPU is not currently being used.
PR MD FTN	CPU is being used by CPUMTR for program-mode monitor functions.
STOR MOVE	CPU is being used by CPUMTR for storage move.
FL VER	CPU is being used by CPUMTR for field length verification.
PSEUDO-PP	CPU is being used by CPUMTR to process buffered I/O requests in the pseudo-PP.
SUBSYSTEM	CPU is being used by a subsystem.
SYS ORG	CPU is being used by a system origin job (subsystems are not considered system origin).
USER	CPU is being used by a user program.
EI	CPU is being used by EI, the environmental interface. This is a segment of native-mode CYBER 180 code which processes illegal CYBER 170-mode instructions, and which performs communication between NOS and NOS/VE.
VE SYSTEM	CPU is being used by NOS/VE for a system task.
VE USER	CPU is being used by NOS/VE for a user job.

Each time the CPU status is sampled (in words CPAL and CPAL+1 in CMR), TRACER increments one of the CPU usage counters. It determines which counter to update by investigating the control point area to which the CPU is assigned.

Tune your system to get the most CPU usage for your users. Following are suggestions on how to do this.

- If you run a mixture of jobs, such as interactive and batch jobs, set the limit of executing batch jobs to a number that minimally impacts interactive users with respect to central memory and control points. You may have to assign batch jobs a higher queue priority so they essentially will be locked into the control points. Since interactive jobs tend to use the CPU for short durations, the batch jobs will be there using up whatever CPU time is left from the interactive jobs.
- On machines with big memories, it is usually advantageous to assign most of the memory to the users instead of using UEM as a rollout device. This helps in two ways: 1) rolling jobs to disk usually takes less CPU overhead than rolling jobs to UEM; 2) with more memory for the users, NOS performs less storage moves.

However, if you are rolling jobs to 887, 895, or 9853 disks and your system is CPU-saturated, you may want to rollout jobs to UEM and use the ENABLE,DDP ROLLOUT PATH entry in the IPRDECK. This will make rollout processing slower but should reduce system CPU usage.

- If you run site-developed subsystems, examine the possibility of performing more of the work as disk activity rather than manipulating the data in central memory.

Subsystem CPU Usage

Subsystem CPU usage is a further breakdown of subsystems of the SUBSYSTEM CPU usage data. TRACER maintains a CPU usage counter for every subsystem. Whenever the SUBSYSTEM CPU usage counter is incremented, TRACER also increments the appropriate subsystem CPU usage counter.

Medium Loop Items

The following paragraphs describe medium loop items.

Medium Loop Samples

TRACER increments the medium loop samples counter each time the medium loop items are sampled.

CPS in W Status

TRACER increments the control points in waiting (W) status counter whenever a control point is found in W status (determined by STSW word of control point area). TRACER scans all control points during each medium loop cycle.

CPS in X Status

TRACER increments the control points in periodic or automatic recall (X) status counter whenever it finds a control point in X status (determined by STSW word of control point area). TRACER scans all control points during each medium loop cycle.

CPS in I Status

TRACER increments the control points in automatic recall (I) status counter whenever a control point is found in I status (determined by STSW word of control point area). TRACER scans all control points during each medium loop cycle.

Same Move Request

TRACER increments the same move request counter every time the move request pending for the current medium loop cycle is the same as the previous medium loop cycle. Word CMCL of CMR is used to determine this status.

Use this counter to determine if NOS is having trouble moving jobs in central memory. A value of from 5 to 10 percent is not unreasonable. A higher percentage indicates that you probably are short of memory. Examine the FL available counter.

FL Available

TRACER increments the FL available counter by the amount of available FL during the current medium loop cycle. Word ACML of CMR determines the available FL.

Compare this counter with the average FL in the rollout queue (FL in the rollout queue divided by the average number of jobs in the queue). On a busy system, the average amount of FL available should be less than that required for the average job in the rollout queue. You may need to consider any jobs that are in the rollout queue but are not schedulable, such as NAM supervisory applications (CS, NS, and NVF). On a system that is not heavily used, the FL available counter is of little value since there should always be memory available.

User EM Available

TRACER increments the user extended memory available counter by the amount of available user extended memory FL during the current medium loop cycle. Word AECL of CMR determines the available user extended memory FL.

Compare this number to the amount of extended memory at control points and rollout queues to determine if you have enough user extended memory.

Noninteractive Jobs

TRACER maintains a count of noninteractive jobs by monitoring the number of EJT entries with a connection status of NICS.

Detached Jobs

TRACER maintains a count of detached interactive jobs by monitoring the number of EJT entries with a connection status of DTCS.

Online Jobs

TRACER maintains a count of online jobs by monitoring the number of EJT entries with a connection status of OLCS.

Preinitial Job Step

TRACER maintains a count of preinitial jobs for each service class by monitoring the number of EJT entries with a job status of PRJS. A job is in this state after it has been assigned to an EJT entry but before it has been assigned to a control point for the first time.

Executing

TRACER maintains a count of executing jobs for each service class by monitoring the number of EJT entries with a job status of EXJS.

Scheduler Rollout

TRACER maintains a count of jobs rolled out by the job scheduler for each service class by monitoring the number of EJT entries with a job status of ROJS.

SCP Rollin

TRACER maintains a count of jobs rolled in by system control point processing for each service class by monitoring the number of EJT entries with a job status of SIJS.

SCP Rollout

TRACER maintains a count of jobs rolled out by system control point processing for each service class by monitoring the number of EJT entries with a job status of SOJS.

Timed / Event Rollout

TRACER maintains a count of jobs in a timed/event rollout queue for each service class by monitoring the number of EJT entries with a job status of TOJS.

Interactive Rollout

TRACER maintains a count of jobs rolled out by interactive input/output processing for each service class by monitoring the number of EJT entries with a job status of IOJS.

Disabled Rollout

TRACER maintains a count of disabled jobs rolled out for each service class by monitoring the number of EJT entries with a job status of DOJS.

Suspended Rollout

TRACER maintains a count of suspended jobs rolled out for each service class by monitoring the number of EJT entries with a job status of SUJS.

Rollout File Error

TRACER maintains a count of jobs with rollout file errors for each service class by monitoring the number of EJT entries with a job status of ERJS.

EJT Entries in Use

TRACER maintains a count of the number of EJT entries in use for all jobs during the specified time interval. Each EJT entry is four central memory words long. If you are short of memory, you can gain some memory by decreasing the number of EJT entries.

FL at Control Points

TRACER maintains an FL at control point counter for each service class. When an active control point is found, the service class is determined and the appropriate counter is incremented by the amount of FL assigned to that control point (determined by FLSW of control point area). TRACER scans all control points during the medium loop cycle.

Use this count to determine which service class of jobs is using the most memory. For example, if you have poor interactive job response and notice that the batch service class is using the most memory, it may help to restrict batch jobs to a smaller central memory time slice using the CM parameter of the SERVICE entry in the IPRDECK.

FL at Pseudo-control Points

TRACER maintains an FL at pseudo-control point counter for each service class. When an active pseudo-control point is found, the service class is determined and the appropriate counter is incremented by the amount of FL assigned to that pseudo-control point (determined by FLSW of pseudo-control point area). TRACER scans all pseudo-control points during the medium loop cycle.

FL in Rollout Queue

TRACER maintains an FL in rollout queue counter for each service class. When a rollout file is found in the queue, and it is not assigned to a control point, the amount of FL it will require (determined from the EJT) when it is rolled into a control point is added to the appropriate counter. TRACER scans all EJT entries during the medium loop cycle.

Use this count in conjunction with the FL available counter to determine if you have a central memory shortage.

EM Memory at Control Points

TRACER maintains an extended memory at control point counter for each service class. When an active control point is found, the service class is determined and the appropriate counter is incremented by the amount of extended memory assigned to that control point (determined by ECSW of control point area). TRACER scans all control points during the medium loop cycle.

Use this count in conjunction with the user EM available counter and the EM in rollout queue counter to determine if you have an extended memory shortage.

EM Memory at Pseudo-control Points

Use this count in conjunction with the user EM available counter and the EM in rollout queue counter to determine if you have an extended memory shortage.

EM in Rollout Queue

TRACER maintains an extended memory in rollout queue counter for each service class. When a rollout file is found in the queue, and it is not assigned to a control point, the amount of extended memory FL it will require when rolled in (determined from the EJT) is added to the appropriate counter. TRACER scans all EJT entries during the medium loop cycle.

Use this count in conjunction with the user EM available counter and the EM memory at CP counter to determine if you have an extended memory shortage.

Subsystem FL

TRACER maintains a subsystem FL counter for every possible subsystem. When a subsystem is found at a control point, the amount of FL assigned to that subsystem is added to the appropriate counter. TRACER scans all control points during the medium loop cycle.

Slow Loop Items

The following paragraphs describe slow loop items.

Slow Loop Samples

TRACER increments the slow loop samples counter each time the slow loop items are sampled.

IAF Users

TRACER increments the IAF users counter by the number of users connected to IAF during the slow loop cycle. IAF must be active for this data to be collected. Word VANL of IAF FL determines the number of users.

IAF Pots Allocated

TRACER increments the pots allocated counter by the number of pots that are currently available for use, whether they are being used or not. Word VPAL of IAF FL determines the number of pots allocated.

IAF Pots in Use

TRACER increments the pots in use counter by the number of pots currently assigned to a connection. Word VPUL of IAF FL determines the number of pots in use.

Queue Files Assigned

TRACER maintains a count of the number of queued files assigned to control points by monitoring QFT entries.

Input Files

TRACER maintains a count of the number of input files for each service class by monitoring QFT entries.

Print Files

TRACER maintains a count of the number of print files for each service class by monitoring QFT entries.

Punch Files

TRACER maintains a count of the number of punch files for each service class by monitoring QFT entries.

Other Queue Files

TRACER maintains a count of the number of other queued files not assigned to control points for each service class by monitoring QFT entries. This includes any queued file that is not an input, print, or punch file.

QFT Entries in Use

TRACER maintains a count of the number of QFT entries in use for all jobs during the specified time interval. Each QFT entry is four central memory words long. If you are short of memory, you can gain some memory by decreasing the number of QFT entries.

Tape Drives in Use

TRACER maintains two tape drives in use counters; one for 7-track and one for 9-track drives. A tape drive is considered to be in use if an EST entry indicates it is logically turned on, and it is assigned to a job. TRACER increments the appropriate counter for each drive found in use.

Tracks Available

TRACER maintains a tracks available counter for each mass storage equipment and adds the number of available tracks (tracks not currently assigned to a file) for a device to the appropriate counter for each mass storage device found in the EST. TRACER obtains this information from TDGL of the MST.

CPPs Active

TRACER increments the CPP active counter for each active CPP it finds when scanning the PP communication area. Disabled CPPs (those turned off at deadstart time) are not considered active. An active CPP is one that has a nonzero input register.

ISHARED Table Changes

TRACER maintains a count of the number of monitor function calls that force a table update in an ISHARED environment.

ISHARED Device Updates

TRACER maintains a count of the number of times the device-resident tables are updated in an ISHARED environment.

ISHARED Seek Time

TRACER monitors the ISHARED seek time by recording the total time spent reading the first sector during a table update sequence.

ISHARED Updating Time

TRACER monitors the ISHARED updating time by recording the total time spent updating tables and the maximum time and minimum time to update the tables for both single and simultaneous operations.

Statistical Summary

The following paragraphs describe various statistics collected by TRACER.

MTR Maximum Time

TRACER reports the maximum time, in microseconds, that MTR waits for a CPU exchange to occur after being initiated by a monitor exchange jump (MXN).

Direct Moves

Tracer reports the number of direct storage moves of central memory or extended memory.

Worst Case MTR Cycle Time

TRACER reports the maximum time, in milliseconds, that it takes MTR to make a complete cycle; processing all PP output register requests and performing its other system functions.

Missed Clock Updates

TRACER reports the number of times MTR was not able to update the clock.

EM Moves

TRACER reports the number of storage moves of extended memory performed by CPUMTR.

Central Memory Moves

TRACER reports the number of storage moves of central memory performed by CPUMTR.

Rollouts to MS

TRACER reports the total number of jobs rolled out to mass storage by the job rollout routine (1RO).

Pseudo-Rollouts

TRACER reports the number of jobs rolled to a pseudo-control point.

PCP Job Rollouts

TRACER reports the number of pseudo-control point jobs rolled out to mass storage.

Secondary Rollouts

TRACER reports the number of jobs rolled out to secondary rollout devices by the job rollout routine (1RO).

Total Sectors Rolled

TRACER reports the total number of sectors of mass storage used by the job rollout routine (1RO) for rollouts.

Secondary Sec Rolled to MS

TRACER reports the number of sectors of mass storage used by the job rollout routine (1RO) for rollouts to secondary rollout devices.

Rollouts/User Limits

TRACER reports the number of jobs rolled out due to a time limit or SRU limit detected by the job scheduler (1SJ).

Time Slices

TRACER reports the number of jobs whose scheduling is set to the lower bound for its service class due to an expired CM time slice.

PP Priority Exchanges

TRACER reports the number of priority exchange requests issued by PPs and CPPs. A priority exchange is issued after three unsuccessful attempts to perform a monitor exchange jump (MXN).

No Comm Buffer Avail

TRACER reports the number of times CPUMTR could not use a communications buffer because all buffers were full.

No PCP Available

TRACER reports the number of times a pseudo-control point was not available.

EJT Scans

TRACER reports the number of times the job scheduler (1SJ) scans the executing job table (EJT) to schedule a job to a control point.

Schedulable Jobs

TRACER reports the total number of schedulable jobs. This is determined by adding the number of schedulable jobs in the EJT on each EJT scan.

Jobs Preempted

TRACER reports the number of jobs that were rolled out so that the job scheduler (1SJ) could schedule a higher priority job.

Jobs Scheduled

TRACER reports the number of jobs scheduled to a control point.

Scheduled No Constraints

TRACER reports the number of jobs scheduled with no service class, central memory, or extended memory constraints imposed.

Insufficient CM Scans

TRACER reports the number of EJT scans for which a schedulable job was found, but could not be scheduled due to insufficient central memory.

Insufficient EM Scans

TRACER reports the number of EJT scans for which a schedulable job was found, but could not be scheduled due to insufficient extended memory.

No Control Point Scans

TRACER reports the number of EJT scans for which a schedulable job was found, but could not be scheduled because all control points were busy.

PROBE Utility

The PROBE utility measures the following:

- Number of times a PP or CPP routine was loaded.
- Number of CIO RA+1 requests by function number.
- Number of PP requests to CPUMTR by function number.
- Number of MTR requests to CPUMTR by function number.
- Statistical data accumulated in CMR includes items such as number of sectors rolled, number of rollouts, and number of pseudo-rollouts.
- Input/output statistics for 885-42 and 895 Disk Storage Subsystems showing the number of sectors transferred and control buffer tables transferred.

PROBE data gathering is selectable at deadstart time by an IPRDECK entry. SYSEDIT resets the PROBE data tables to zeros.

The PROBE utility generates a report from the data collected by the system. PROBE analyzes data either from system tables or from a binary file containing data from a previous PROBE run. An IPRDECK entry ENABLE,PROBE must be specified at deadstart time to allow the system to collect the data. (Refer to section 3, Deadstart Decks, for more information.)

Format:

PROBE, P₁, P₂, . . . , P_n.

P_i	Description
B=readfile	Binary file to be read. Default is STATS.
L=outfile	Report file. Default is OUTPUT.
L=0	No report is to be generated.
LO=opt	Sort option for PP program load information. Default is F.
	opt Description
	A Sort data in alphabetic order.
	F Sort data by frequency of loads.
	R Sort data by location and frequency of loads.

P_i	Description
OP=opt	Processing option. Default is P.
	opt Description
C	Perform R option functions and clear system tables after they are read.
P	Generate report from binary file specified by the B parameter. This binary file has been created by a previous PROBE run with OP=C or OP=R specified.
R	Read system tables, and write binary file and report file as specified.
P=writefile	Binary file to be written. Default is STATS.
P=0	No binary file is to be written.
R	Rewind binary files before and after operation. Default is no rewind.

NOTE

The file names must be unique. If you specify OP=C or OP=R, then you may not specify P=0.

Figure 21-2 is an example of the PROBE output file format.

When looking at the PP program load information reported on the PROBE output, check to see if it meets these suggested performance requirements:

- PP programs called at least once every second should be central memory resident.
- PP programs called at least once every 2 seconds are good candidates for alternate residency using extended memory.
- PP programs called every 3 seconds or less should be disk resident.

You can change a PP program's residency using the SYSEEDIT command (refer to section 19, SYSEEDIT).

PROBE VERSION 1.0.		yy/mm/dd.	hh.mm.ss.	PAGE	1
PROBE VERSION 1.0		yy/mm/dd.	hh.mm.ss.		
START OF SAMPLE INTERVAL		yy/mm/dd.	hh.mm.ss.		
END OF SAMPLE INTERVAL		yy/mm/dd.	hh.mm.ss.		
MONITOR REQUESTS					
	PROGRAM MODE	MONITOR MODE			
	CPU0	CPU0	CPU1		
CHGM (21)	0	0	0		
HNGM (22)	0	0	0		
(23)	0	0	0		
AFAM (24)	0	7177	0		
DLKM (25)	0	0	0		
DTKM (26)	131214	189955	0		
RTCM (27)	129969	129969	0		
STBM (30)	0	121555	0		
VMSM (31)	9108	9108	0		
ACTM (32)	1133	20059	0		
BFMM (33)	0	0	0		
CKSM (34)	435	435	0		
CSTM (35)	34405	34670	0		
ECSM (36)	0	0	0		
FIOM (37)	0	0	0		
RDCM (40)	3170	3170	0		
ABTM (41)	0	76	0		
BIOM (42)	0	0	0		
BMIM (43)	0	0	0		
CCAM (44)	0	77321	0		
CEFM (45)	0	832	0		
CPRM (46)	0	0	0		
DCPM (47)	0	25696	0		
DEQM (50)	0	41425	0		
DFMM (51)	0	36827	0		
DPPM (52)	0	1108876	0		
EATM (53)	0	48694	0		
JACM (54)	0	59846	0		
LDAM (55)	0	844946	0		
LMSM (56)	0	0	0		
MTRM (57)	0	2495	0		
PLFM (60)	0	76738	0		
RCLM (61)	0	447037	0		
RCPM (62)	0	1110	0		
RECM (63)	0	46638	0		
REQM (64)	0	41426	0		
RJSM (65)	0	675	0		
RLMM (66)	0	30666	0		
ROCM (67)	0	15092	0		
RPNM (70)	0	0	0		
RPPM (71)	0	106311	0		
RSJM (72)	0	15481	0		
SCDM (73)	0	62387	0		
SFBM (74)	0	70802	0		
SJCM (75)	0	22237	0		
SPLM (76)	0	1358228	0		
TDAM (77)	0	17566	0		
TGPM (100)	0	11532	0		
TIOM (101)	0	21228	0		
TSEM (102)	0	37099	0		
UADM (103)	0	583735	0		
UTEM (104)	0	1473766	0		

Figure 21-2. Example of PROBE Output (Continued)

PROBE VERSION 1.0.				yy/mm/dd. hh.mm.ss.	PAGE	2
VFLM(105)	0	0	0			
VFPM(106)	0	14518	0			
VSAM(107)	0	13180	0			
(110)	0	0	0			
SUBTOTAL		7230584	0			
TOTAL	309434		7230584			
MTR REQUESTS						
		CPU0	CPU1			
ARTF(1)	27487		0			
IARF(2)	0		0			
CSLF(3)	23220		0			
RCLF(4)	6134559		0			
MFLF(5)	66297		0			
SCSF(6)	0		0			
SMSF(7)	0		0			
CMSF(10)	0		0			
PRQF(11)	41638		0			
ACSF(12)	93425		0			
PCXF(13)	0		0			
ARMF(14)	0		0			
----(15)	0		0			
MFEF(16)	0		0			
SUBTOTAL	6386626		0			
TOTAL	6386626					
PROGRAM MODE						
		CPU0				
MSTF(1)	23220					
PDMF(2)	0					
PMRF(3)	0					
MECF(4)	0					
TOTAL	23220					
MTR PERFORMANCE PARAMETERS						
WORST CASE *MXN* TIME =				751		
WORST CASE CYCLE TIME =				6		
MISSED CLOCK UPDATES =				2		
CPUMTR PERFORMANCE PARAMETERS						
EM STORAGE MOVES =				0		
CM STORAGE MOVES =				24173		
DIRECT MOVES =				20540		
COMMUNICATIONS BUFFER NOT AVAILABLE =				0		
PP PRIORITY EXCHANGE REQUESTS =				0		
PSEUDO-CONTROL POINT NOT AVAILABLE				351		
LRO PERFORMANCE PARAMETERS						
NUMBER OF ROLLOUTS =				18393		
PSEUDO-ROLLOUTS =				4000		
PSEUDO-CONTROL POINT ROLLOUTS =				1000		
ROLLOUTS TO SECONDARY DEVICES =				0		
NUMBER OF SECTORS ROLLED =				2918049		
SECTORS ROLLED TO SECONDARY DEVICES =				0		

Figure 21-2. Example of PROBE Output (Continued)

PROBE VERSION 1.0.		yy/mm/dd. hh.mm.ss. PAGE 3	
1SJ PERFORMANCE PARAMETERS			
EJT SCANS =			76193
SCHEDULABLE JOBS=			21917
JOBS PREEMPTED =			95
JOBS SCHEDULED =			17866
JOBS SCHEDULED NO CONSTRAINTS =			0
EJT SCANS WITH INSUFFICIENT CM =			130
EJT SCANS WITH INSUFFICIENT EM =			0
EJT SCANS WITH NO CONTROL POINT =			2
ROLLOUTS FOR RESOURCE LIMITS =			1
TIME SLICES =			460
CIO REQUESTS			
FUNCTION	CALLS	FUNCTION	CALLS
RPHR(000)	3015	READ(010)	154978
WRITE(014)	83438	READSKP(020)	4907
WRITER(024)	53751	WRITEF(034)	5237
BKSP(040)	465	REWIND(050)	55857
****(54)	128	UNLOAD(060)	459
RETURN(070)	30345	OPEN(100)	2200
OPEN(104)	2643	EVICT(114)	1653
OPEN(120)	2602	CLOSE(130)	1356
OPEN(140)	753	OPEN(144)	4885
CLOSE(150)	582	OPEN(160)	543
CLOSE(170)	5339	CLOSE(174)	465
READCW(200)	15476	WRITECW(204)	7361
READLS(210)	1235	REWRITE(214)	13111
REWRITER(224)	11907	RPHRLS(230)	25137
SKIPF(240)	9784	READNS(250)	400
READEI(600)	5893	SKIPB(640)	4386
TOTAL CIO FUNCTIONS PROCESSED		510291	

Figure 21-2. Example of PROBE Output (Continued)

PPU PROGRAM LOADS					
NAME	LOC	LOADS	NAME	LOC	LOADS
CPM	CMR	605181	PIP	DSK	0
1MS	CMR	460275	PNC	DSK	0
4MB	CMR	243588	SBP	DSK	0
4MD	CMR	169670	SLL	DSK	0
4ME	CMR	155651	SMP	DSK	0
QAC	CMR	83410	TMG	DSK	0
3Q3	CMR	83372	VEJ	DSK	0
1SJ	CMR	73838	VER	DSK	0
3AE	CMR	70341	WRM	DSK	0
LFM	CMR	69611	XIS	DSK	0
1AJ	CMR	47141	X26	DSK	0
3LB	CMR	46964	OCI	DSK	0
TCS	CMR	46406	OIP	DSK	0
3QU	DSK	46381	OMC	DSK	0
LDR	CMR	43575	OMD	DSK	0
1MA	CMR	39939	OSD	CMR	0
0BF	CMR	37181	OSI	DSK	0
3QS	CMR	36988	OTD	DSK	0
4MC	CMR	31223	1FA	DSK	0
3NW	CMR	29968	1HP	DSK	0
1IO	CMR	27405	1IS	DSK	0
NDR	CMR	26489	1KB	DSK	0
RHH	DSK	25371	1LC	DSK	0
1MT	CMR	24435	1LT	CMR	0
3ME	CMR	24435	1MF	DSK	0
ODF	CMR	24317	1MR	DSK	0
3MF	CMR	23257	1PC	DSK	0
3CD	CMR	21678	1PD	DSK	0
3AF	CMR	19990	1PL	DSK	0
3MJ	CMR	19709	1PR	DSK	0
QFM	DSK	19637	1RM	DSK	0
RPV	CMR	18598	1RP	DSK	0
1RO	CMR	18282	1RU	DSK	0
1RI	CMR	18262	1TM	DSK	0
1MI	CMR	14294	1TN	DSK	0
3LD	CMR	12765	1TP	DSK	0
SFM	CMR	12584	1TS	DSK	0
3SX	CMR	10536	1VP	DSK	0
STD	DSK	9223	2DC	DSK	0
4SD	DSK	9222	2DD	DSK	0
3AD	DSK	8614	2IE	DSK	0
2MA	DSK	8501	2IM	DSK	0
TLX	CMR	8444	2IN	DSK	0
1DD	CMR	7995	2IP	DSK	0
4MF	CMR	7600	2IQ	DSK	0
3RP	CMR	7357	2KA	DSK	0
3RH	CMR	7342	2KB	DSK	0
PFM	CMR	6961	2KC	DSK	0
3FA	CMR	5122	2LD	DSK	0
3LC	CMR	5085	2MB	DSK	0
3AI	CMR	4722	2ME	DSK	0

Figure 21-2. Example of PROBE Output (Continued)

PROBE VERSION 1.0.				yy/mm/dd. hh.mm.ss.		PAGE	5
3LE	CMR	4024	2MF	DSK	0		
3PC	CMR	3960	2MN	DSK	0		
3QD	DSK	3959	2MR	DSK	0		
LDQ	CMR	3790	2MZ	DSK	0		
3PD	CMR	3619	2NT	DSK	0		
3QA	DSK	2800	2NU	DSK	0		
3CR	CMR	2749	2PT	DSK	0		
3CS	CMR	2748	2PV	DSK	0		
1TO	CMR	2296	2PX	DSK	0		
2TO	CMR	2203	2RP	DSK	0		
3PE	CMR	2080	2RU	DSK	0		
0AV	CMR	1887	2SB	DSK	0		
1TA	DSK	1867	2SC	DSK	0		
3CB	CMR	1822	2SE	DSK	0		
1CK	DSK	1699	2SF	DSK	0		
3PR	CMR	1682	2SG	DSK	0		
3MQ	CMR	1670	2SI	DSK	0		
3CC	DSK	1643	2TM	DSK	0		
3SV	DSK	1598	2TN	DSK	0		
3PN	CMR	1489	3AG	DSK	0		
ORP	CMR	1325	3AH	CMR	0		
LDD	CMR	1301	3BB	DSK	0		
3TC	DSK	1200	3BC	DSK	0		
0FA	CMR	1162	3BD	DSK	0		
DSF	DSK	1161	3BE	DSK	0		
3DB	DSK	1120	3BF	DSK	0		
0DQ	DSK	1086	3CA	DSK	0		
3RG	DSK	1029	3CV	DSK	0		
2NW	DSK	995	3FB	DSK	0		
3MI	CMR	981	3IC	DSK	0		
0AU	DSK	947	3ID	DSK	0		
3AK	DSK	941	3IE	DSK	0		
3AA	DSK	934	3IM	DSK	0		
3PQ	CMR	923	3IN	DSK	0		
2NR	DSK	922	3IO	DSK	0		
3AJ	DSK	918	3IP	DSK	0		
3QW	DSK	904	3IQ	DSK	0		
3QX	DSK	903	3KA	DSK	0		
3QO	DSK	903	3KE	DSK	0		
3PH	CMR	791	3KC	DSK	0		
CVL	DSK	747	3LA	DSK	0		
3QC	DSK	685	3MA	DSK	0		
SFP	CMR	605	3MB	DSK	0		
3AB	DSK	596	3MD	DSK	0		
1DS	CMR	521	3MK	CMR	0		
3DA	DSK	491	3MM	DSK	0		
1ML	DSK	490	3MR	CMR	0		
3TE	DSK	480	3MU	DSK	0		
3PI	DSK	479	3PM	DSK	0		
1DL	CMR	454	3PP	DSK	0		
CMS	DSK	429	3QR	DSK	0		
4DB	DSK	426	3QI	DSK	0		
4DD	DSK	426	3RF	DSK	0		
3SA	DSK	425	3RU	DSK	0		
3MO	CMR	383	3SU	DSK	0		
3NM	DSK	382	3SZ	DSK	0		
3ML	DSK	354	3TA	DSK	0		
3MH	DSK	350	4DF	DSK	0		

Figure 21-2. Example of PROBE Output (Continued)

PROBE VERSION 1.0.			yy/mm/dd. hh.mm.ss. PAGE 6		
3PF	CMR	348	4DI	DSK	0
2MD	DSK	338	4DJ	DSK	0
4MA	CMR	313	4DK	DSK	0
0VJ	DSK	310	4DZ	DSK	0
0VU	DSK	310	4IM	DSK	0
3PK	DSK	245	4IO	DSK	0
3MC	CMR	237	4IP	DSK	0
3AL	DSK	235	4IQ	DSK	0
2MC	DSK	216	4KB	DSK	0
9AA	DSK	215	4KC	DSK	0
SSH	DSK	210	4KD	DSK	0
3MG	DSK	193	4MG	DSK	0
3TB	DSK	188	4NB	DSK	0
PFU	DSK	177	4NC	DSK	0
9EA	DSK	168	4NM	DSK	0
0RF	DSK	163	4RA	DSK	0
2SD	DSK	141	4RB	DSK	0
3MT	DSK	138	4RC	DSK	0
2NS	DSK	136	4RD	DSK	0
3MP	CMR	133	4RE	DSK	0
3PB	DSK	129	4RF	DSK	0
3PJ	DSK	126	4RG	DSK	0
3AM	DSK	91	5BA	DSK	0
9BA	DSK	83	5BB	DSK	0
3MN	DSK	82	5BC	DSK	0
0RT	DSK	74	5BD	DSK	0
3MS	DSK	67	5CU	DSK	0
2NQ	DSK	63	5CV	DSK	0
T76	DSK	58	5D1	DSK	0
6DI	CMR	58	5D2	DSK	0
0PT	CMR	51	5D3	DSK	0
QAP	DSK	48	5FF	DSK	0
NLD	DSK	46	5FH	DSK	0
3DC	DSK	36	5IA	DSK	0
3FA	DSK	35	5IC	DSK	0
3Q2	DSK	32	5ID	DSK	0
DIS	DSK	29	5IG	DSK	0
3BA	DSK	24	5IH	DSK	0
3BG	DSK	24	5LC	DSK	0
3IA	DSK	24	5LL	DSK	0
3IB	DSK	24	5MA	DSK	0
5IE	DSK	24	5MB	DSK	0
3QB	DSK	22	5MC	DSK	0
1CD	DSK	21	5MD	DSK	0
3PL	DSK	19	5ME	DSK	0
3AC	DSK	15	5MG	DSK	0
2DA	DSK	13	5MR	DSK	0
2SA	DSK	13	5MS	DSK	0
7CI	CMR	13	5MT	DSK	0
7DI	CMR	13	5MU	DSK	0
7EI	CMR	13	5MW	DSK	0
7EL	CMR	13	5SE	DSK	0
7EM	CMR	13	5SU	DSK	0
7EN	CMR	13	5SV	DSK	0
7EO	CMR	13	5S1	DSK	0
7EP	CMR	13	5S2	DSK	0
7SI	CMR	13	5XA	DSK	0
4MH	DSK	12	5XB	DSK	0

Figure 21-2. Example of PROBE Output (Continued)

PROBE VERSION 1.0.			yy/mm/dd. hh.mm.ss. PAGE 7		
IMS	DSK	10	5XC	DSK	0
3PG	DSK	10	5XD	DSK	0
3QY	DSK	9	5XE	DSK	0
3QZ	DSK	9	5XF	DSK	0
2SH	DSK	8	5XG	DSK	0
3CE	DSK	8	541	DSK	0
0PI	DSK	7	55X	DSK	0
1MB	DSK	7	56X	DSK	0
1MP	CMR	7	57X	DSK	0
2DB	DSK	6	58F	DSK	0
3DD	DSK	6	58H	DSK	0
4DA	DSK	6	58X	DSK	0
4DC	DSK	6	7ER	CMR	0
4DG	DSK	6	7FI	CMR	0
026	DSK	5	7GI	CMR	0
9GA	DSK	5	7HI	CMR	0
3RI	DSK	4	7II	CMR	0
2NX	DSK	3	7JI	CMR	0
2NY	DSK	3	7KI	CMR	0
3AN	DSK	3	7LI	CMR	0
3PO	DSK	3	7SE	CMR	0
3QT	DSK	3	7WI	CMR	0
3QV	DSK	3	8XA	DSK	0
7EQ	CMR	3	8XB	DSK	0
7ES	CMR	3	8XC	DSK	0
0ST	DSK	2	9CA	CMR	0
3SY	DSK	2	9CB	CMR	0
0TI	DSK	1	9CC	CMR	0
0TJ	DSK	1	9CD	CMR	0
3CK	DSK	1	9CE	CMR	0
3RJ	DSK	1	9CF	CMR	0
3SW	DSK	1	9CG	CMR	0
3TD	DSK	1	9CH	CMR	0
ADC	DSK	0	9CI	CMR	0
BAT	DSK	0	9CJ	CMR	0
CPD	DSK	0	9CK	CMR	0
CUX	DSK	0	9CL	CMR	0
DDF	DSK	0	9CM	CMR	0
DOG	DSK	0	9CN	CMR	0
DS1	DSK	0	9CO	CMR	0
ELM	DSK	0	9CP	CMR	0
EYE	DSK	0	9CQ	CMR	0
FDP	DSK	0	9CR	CMR	0
HPM	DSK	0	9CS	CMR	0
LIF	DSK	0	9JN	DSK	0
MDD	DSK	0	9VA	DSK	0
MLD	DSK	0	9WA	DSK	0
MP3	DSK	0			
PIM	DSK	0			
NUMBER OF ASR PROGRAMS =			0		
NUMBER OF CMR PROGRAMS =			123		
NUMBER OF DSK PROGRAMS =			313		
NUMBER OF PLD ENTRIES =			436		

Figure 21-2. Example of PROBE Output (Continued)

SECTORS		CBT-S	READS	WRITES
I/O STATISTICS FOR DB006:				
TRANSFERRED		TRANSFERRED		
1		1-40B	0	1
2-3		41B-140B	0	23
4-7		141B-340B	0	0
10B-17B		341B-740B	0	0
20B-37B		741B-1740B	0	0
40B-77B		1741B-3740B	0	0
100B-177B		3741B-7740B	0	0
200B-MORE		7741B-MORE	0	0
TOTALS			0	24
I/O STATISTICS FOR DB007:				
TRANSFERRED		TRANSFERRED		
1		1-40B	19	1
2-3		41B-140B	4	1
4-7		141B-340B	0	0
10B-17B		341B-740B	0	2
20B-37B		741B-1740B	0	0
40B-77B		1741B-3740B	0	0
100B-177B		3741B-7740B	0	0
200B-MORE		7741B-MORE	0	0
TOTALS			23	4
I/O STATISTICS FOR DB010:				
TRANSFERRED		TRANSFERRED		
1		1-40B	0	1
2-3		41B-140B	0	0
4-7		141B-340B	0	0
10B-17B		341B-740B	0	3
20B-37B		741B-1740B	0	0
40B-77B		1741B-3740B	0	0
100B-177B		3741B-7740B	0	0
200B-MORE		7741B-MORE	0	0
TOTALS			0	4

Figure 21-2. Example of PROBE Output

Cartridge Tape Subsystem (CTS) and Automated Cartridge Subsystem (ACS) 22

Cartridge Tape Subsystem	22-1
CTS Components	22-1
CTS Configuration	22-2
Dual-State CTS Configurations	22-4
Automated Cartridge Subsystem	22-6
ACS Components	22-6
ACS Configurations	22-9
Communication Paths	22-10
EQPDECK and IPRDECK Considerations	22-10
Example of Minimum ACS Configuration	22-11
Example of ACS Configuration with Two Library Storage Modules	22-11
Automated Tape Facility Subsystem (ATF)	22-13
Workstation Commands	22-13
Cross-Coupled Controllers for CTS and ACS	22-13
Requirements for CTS and ACS	22-16
Constraints for CTS and ACS	22-16
Performance Considerations for CTS and ACS	22-17
Maintenance for CTS and ACS	22-17

Cartridge Tape Subsystem (CTS) and Automated Cartridge Subsystem (ACS) 22

Both the Control Data 5680 Cartridge Tape Subsystem (CTS) and the Control Data 5744 Automated Cartridge Subsystem (ACS) consist of hardware and software that allow reading and writing of IBM 3480-compatible cartridge tapes.

In CTS, operators load and unload the cartridges on the tape units. ACS has additional hardware and software that robotically handles the cartridges. Once the cartridges are loaded into an ACS storage module, the ACS robot can perform all the conventional tape handling operations without operator assistance: tape retrieval and mounting, dismounting, and storage of tape volumes. The operator can also manually add or remove cartridges to and from the ACS storage module.

Use of CTS and ACS is transparent to users; user requests for files stored on cartridge tapes are identical to requests for files on other types of tapes, except users specify a different density (CE for CTS tapes and AE for ACS tapes) and/or device type (CT for CTS tapes and AT for ACS tapes).

CTS and ACS operate on Control Data CYBER 180 models 835 through 995 and require the Control Data CDCNET communications network.

The next subsections describe:

- CTS components and configurations
- ACS components and configurations
- Cross-coupling controllers for CTS and ACS
- Configuration requirements and constraints for CTS and ACS
- Performance considerations for CTS and ACS
- Maintenance for CTS and ACS

Cartridge Tape Subsystem

The Cartridge Tape Subsystem (CTS) reads and writes IBM 3480-compatible cartridge tapes. CTS provides increased recording capacity and speed over that of reel tape units and is fully compatible with the optional Tape Management System (TMS).

CTS Components

The basic components of CTS are:

- CYBER Channel Couplers (CCCs) that connect the mainframe(s) to Federal Information Processing Standard (FIPS) channels on cartridge tape controllers.
- One or more 5680 cartridge tape controllers (also called control units).
- One or more 5682 cartridge tape cabinets (also called cartridge drives).
- Two or four tape units (also called tape transports) in each 5682 cartridge tape cabinet.
- Operational and control software under NOS.

NOS and other operating systems can share CTS; each operating system, however, must have a separate, dedicated controller and tape unit(s).

You can purchase additional hardware to cross-couple CTS controllers; refer to Cross-Coupled Controllers for CTS and ACS. You can also purchase hardware and software to upgrade CTS to ACS.

Table 22-1 lists and describes the functions of the CTS hardware components as well as the maximum CTS configuration.

Table 22-1. CTS Hardware

Hardware Component	Description
Cartridge	Plastic housing enclosing magnetic tape on which data is stored. This cartridge is IBM 3480-compatible; it has a density of 38,000 cpi and a storage capacity of approximately 130 megabytes when recorded in I format and 220 megabytes when recorded in LI format.
5682-1x cartridge tape cabinet	Cabinet containing cartridge tape units. The 5682-12 cabinet contains two tape units; the 5682-14 cabinet contains four tape units. The tape units can be operated concurrently. A maximum of 16 tape units (four cabinets) can be configured for CTS.
5680-1x cartridge tape controller	<p>Controller for the tape units in the 5682-1x cabinet. There are four types of controllers: 5680-11, 5680-12, 5680-13, and 5680-14. The last digit indicates how many FIPS channels in the controller can be used.</p> <p>A controller can simultaneously position tapes on four tape units, but can transfer data to or from only one tape at a time.</p> <p>A controller contains a 1-megabyte data buffer to maximize transfer rates. This buffer is divided into four 256-kilobyte unit buffers so that I/O can be performed simultaneously on four units without reusing a buffer. Attempts to simultaneously perform I/O on more than four units results in performance degradation.</p> <p>A controller can be connected to a maximum of two tape cabinets. A maximum of eight controllers can be configured for CTS.</p>

CTS Configuration

The minimum CTS n (as shown in figure 22-1) consists of:

- One NIO channel to the CCC
- One cartridge tape controller
- One cartridge tape cabinet containing two tape units

After you install the CTS hardware, you must modify the EQPDECK and IPRDECK to reflect the CTS configuration:

- Create EST entries for the CTS tape units in the EQPDECK.
- Decide whether CTS is your site's default tape device type. If it is, enter in the IPRDECK:

TDTY=CT.

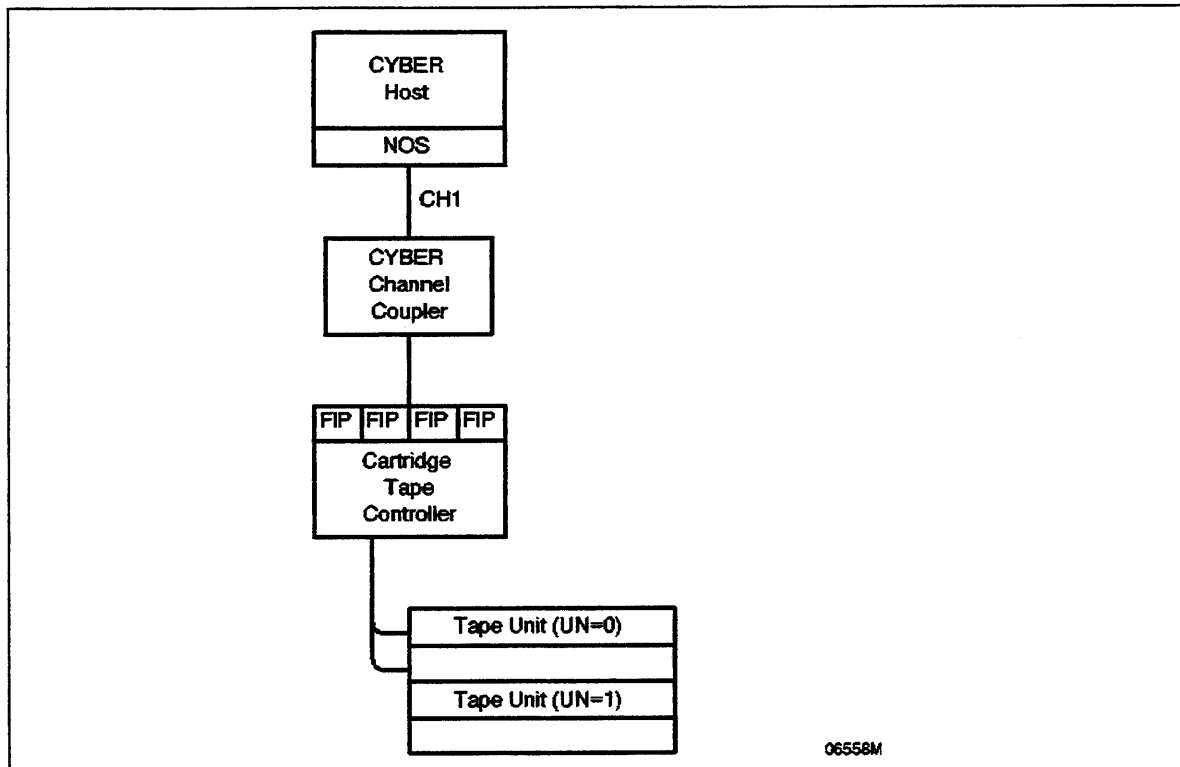


Figure 22-1. CTS Minimum Configuration

EQPDECK and IPRDECK entries are described in section 3, Deadstart Decks.

For the minimum configuration shown in figure 22-1, you could enter the following EQPDECK entry to define the two tape units:

```
EQ60=CT-2, ST=ON, CH=1, UN=0.
```

For the preceding entry, NOS generates the EST entries:

```
EQ60=CT, ST=ON, CH=1, UN=0.
EQ61=CT, ST=ON, CH=1, UN=1.
```

Dual-State CTS Configurations

For dual-state mainframes, you can configure a fixed number of units to each operating system, or you can configure the units so that you can change unit assignments from one operating system to the other. Do the latter with caution, because it is possible to accidentally assign one unit to both operating systems.

An example of a fixed number of CTS units assigned to dual-state partners is shown in figure 22-2. The following EQPDECK entry defines the two units to which NOS has exclusive access; do not define the tape units that NOS/VE (or another mainframe) uses.

```
EQ62=CT-2, ST=ON, CH=1, UN=2.
```

For the preceding entries, NOS generates the EST entries:

```
EQ62=CT, ST=ON, CH=1, UN=2.
EQ63=CT, ST=ON, CH=1, UN=3.
```

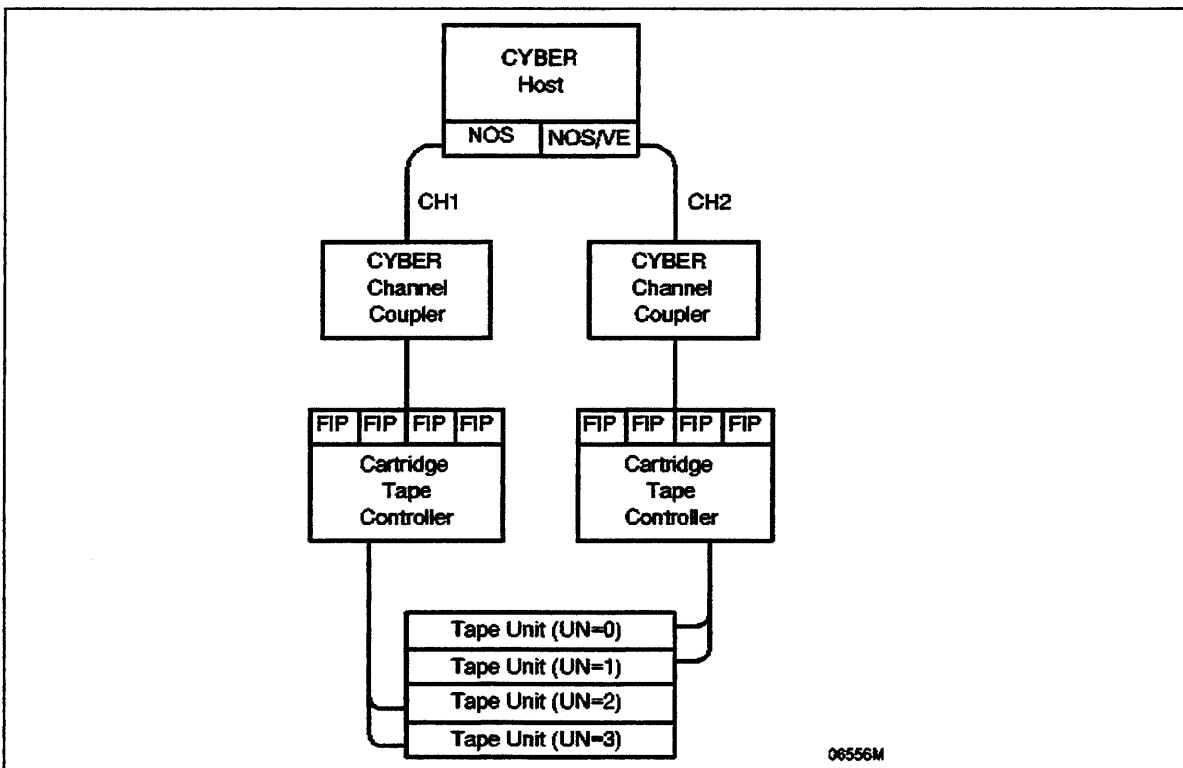


Figure 22-2. Dual-State Fixed Configuration for CTS

An example of configuring all CTS units to both NOS and NOS/VE is shown in figure 22-3. The following EQPDECK entries define the four units, assign a status of OFF to units 0 and 1, and assign status of ON to units 2 and 3:

EQ60=CT,ST=DOWN,CH=1,UN=0.
EQ62=CT,ST=ON,CH=1,UN=2.

For the preceding entries, NOS generates the EST entries:

EQ60=CT,ST=DOWN,CH=1,UN=0.
EQ61=CT,ST=DOWN,CH=1,UN=1.
EQ62=CT,ST=ON,CH=1,UN=2.
EQ63=CT,ST=ON,CH=1,UN=3.

Similarly, all four units must be defined on NOS/VE, but the status of units 0 and 1 must be ON and the status of units 2 and 3 must be OFF.

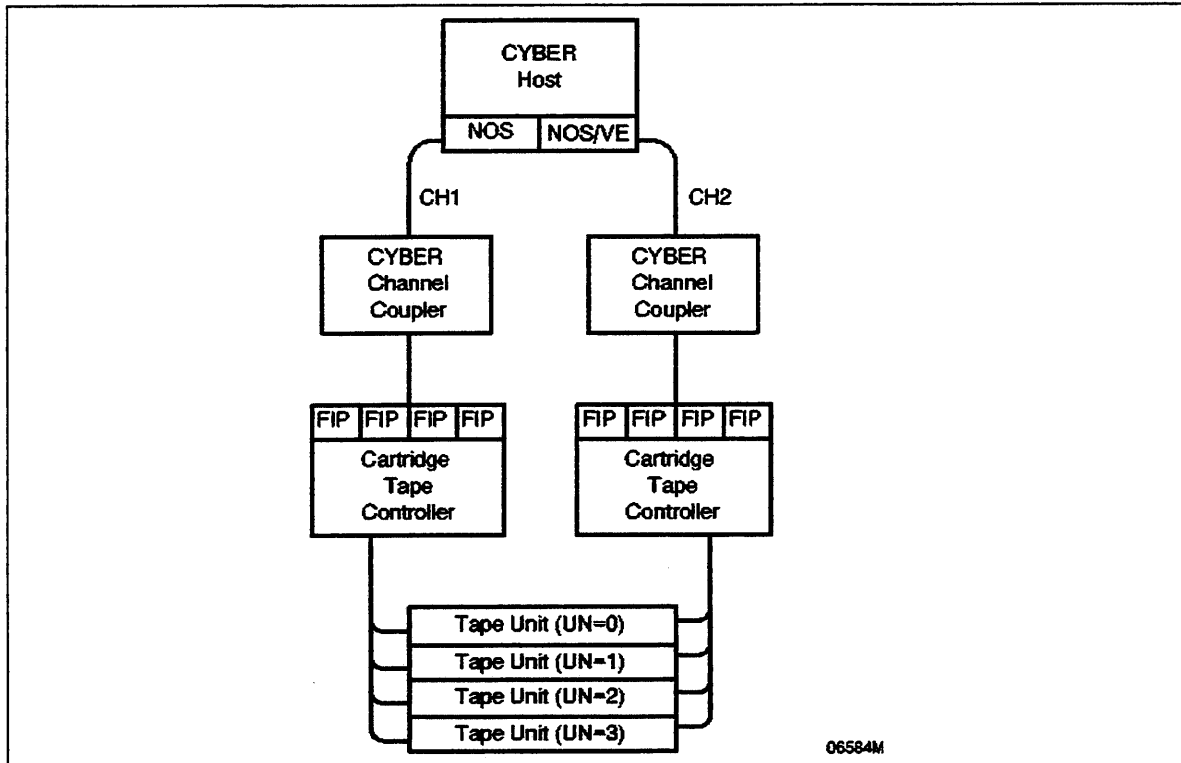


Figure 22-3. Dual-State Flexible Configuration for CTS

As workload requirements change, you can logically move CTS units between NOS and NOS/VE. To move a unit from NOS to NOS/VE, enter the DOWN,EQ=nn DSD command on NOS and the CHANGE_ELEMENT_STATE subcommand of the Logical Configuration Utility on NOS/VE. Conversely, to move unit from NOS/VE to NOS, enter the CHANGE_ELEMENT_STATE subcommand on NOS/VE and the ON,EQ=nn DSD command on NOS.

NOTE

Be sure that each unit is always assigned to only one operating system. Otherwise, tapes may be incorrectly read and written.

For an example of cross-coupled controllers, refer to Cross-Coupled Controllers for CTS and ACS.

Automated Cartridge Subsystem

The Automated Cartridge Subsystem (ACS) is an integrated hardware and software package for robotically handling IBM 3480-compatible cartridge tapes. ACS provides for the automated retrieval, mounting, dismounting, and storage of tape volumes and is capable of unattended operations. Operators, of course, can manually add and remove cartridges from ACS and can enter a variety of UNIX¹ operator commands and utility functions at the ACS workstation.

To use ACS, your site must install and activate the Tape Management System (TMS). TMS provides the tape security for labeled and unlabeled cartridge tapes. To verify that the correct tape is mounted and that the user has permission to access the tape, NOS compares the cartridge's external VSN (the bar-coded identification on the edge of the cartridge) and internal VSN with the VSNs specified on the request; if the tape is unlabeled, only the external VSNs are compared.

TMS scratch tapes in ACS must have matching external and internal VSNs; the internal VSN should be written in ASCII mode and have blank volume and file accessibility fields (the BLANK utility's default values).

To read ACS tapes without using TMS, initiate a system origin job and specify TO=F on the LABEL command.

ACS Components

The basic components of ACS are:

- ACS hardware manufactured by Storage Technology Corporation (StorageTek). Its hallmark is a robotic unit that reliably manipulates tape cartridges within silo-shaped storage modules called library storage modules (LSM).
- Sun Workstation² and associated software.
- CYBER Channel Couplers (CCC) that connect the mainframe(s) to Federal Information Processing Standard (FIPS) channels on cartridge tape controllers.
- One or more 5680 cartridge tape controllers (also called control units).
- One or more 5682 cartridge tape cabinets (also called cartridge drives).
- Two or four tape units (also called tape transports) in each 5682 cartridge tape cabinet.
- Operational and control software under NOS.

You can purchase additional hardware to cross-couple ACS controllers; refer to Cross-Coupled Controllers for CTS and ACS.

NOS and other operating systems can share ACS; each operating system must have a separate, dedicated controller and tape units.

Table 22-2 lists and describes the functions of the ACS hardware components as well as the maximum ACS configuration.

1. UNIX is a registered trademark of UNIX System Laboratories, Inc.

2. Sun Workstation is a trademark of Sun Microsystems, Inc.

Table 22-2. ACS Hardware

Hardware Component	Description
Cartridge	Plastic housing enclosing magnetic tape on which data is stored. This cartridge is IBM 3480-compatible; it has a density of 38,000 cpi and a storage capacity of 130 megabytes when recorded in I format and 220 megabytes when recorded in LI format. Each cartridge has a unique external volume serial number (VSN) appearing as a bar-coded identification on the edge of the cartridge. Up to 6000 cartridges can be stored in a library storage module (LSM). Because NOS uses software to enforce the read-only option (PO=R), all cartridges should be write enabled (white dot is not visible on the thumbwheel) before they are stored in an LSM.
5682-2x cartridge tape cabinet	Cabinet containing cartridge tape units. The 5682-22 cabinet contains two tape units; the 5682-24 cabinet contains four tape units. The tape units can be operated concurrently. A maximum of 16 tape units (four cabinets) can be connected to an LSM.
5680-1x cartridge tape controller	<p>Controller for the tape units in the 5682-2x (or 5682-1x) cabinet. There are four types of controllers: 5680-11, 5680-12, 5680-13, and 5680-14. The last digit indicates how many FIPS channels the controller has.</p> <p>A controller can simultaneously position tapes on four tape units, but can transfer data to or from only one tape at a time.</p> <p>A controller contains a 1 megabyte data buffer to maximize transfer rates. This buffer is divided into four 256-kilobyte unit buffers so that I/O can be performed simultaneously on four units without reusing a buffer. Attempts to simultaneously perform I/O on more than four units results in performance degradation.</p> <p>A controller can be connected to a maximum of two cabinets. A maximum of eight controllers can be configured per LSM.</p>

(Continued on next page)

Table 22-2. ACS Hardware

Hardware Component	Description
<i>(Continued from previous page)</i>	
5744-10 library storage module (LSM)	Twelve-sided cylindrical hardware unit enclosing the cartridge storage area and the robot for moving the cartridges. Each LSM has storage cells for approximately 6,000 cartridges (storing approximately 1,200 gigabytes of data). Up to 16 tape units can be connected to an LSM. Up to 16 LSMs can be connected in ACS.
5744-11 library control unit (LCU)	Microprocessor that controls the robot's movements. There is one library control unit per library storage module.
5744-30 library management unit (LMU)	Microprocessor that manages all the library storage modules in ACS. The library management unit receives mount requests from the workstation and selects a library storage module. Sending messages through the library control unit, it instructs the robot to select a cartridge, validate its bar-coded identification label, retrieve it, mount it in a tape unit, and dismount it when the operation is complete. The library management unit communicates with the library storage modules through a local area network (LAN). There is only one library management unit in ACS, and it is attached to all library storage modules via the library control units.
5744-40 cartridge library server workstation (CLSW)	Sun Workstation providing the interface between the CYBER host (via an Ethernet ³ link) and the library storage module. The workstation services ACS and is connected to the library management unit through an RS423 interface. We recommend at least two RS423 connections to provide redundancy. The library server software resides on the workstation and automates the mounting and dismounting of cartridges.

3. Ethernet is a trademark of Xerox Corporation.

ACS Configurations

The minimum ACS configuration (as shown in figure 22-4) is:

- One Sun Workstation with the library server software
- One cartridge library with one library management unit, one library control unit, and one library storage module
- One NIO channel to the CCC
- One cartridge tape controller
- Two cartridge tape units

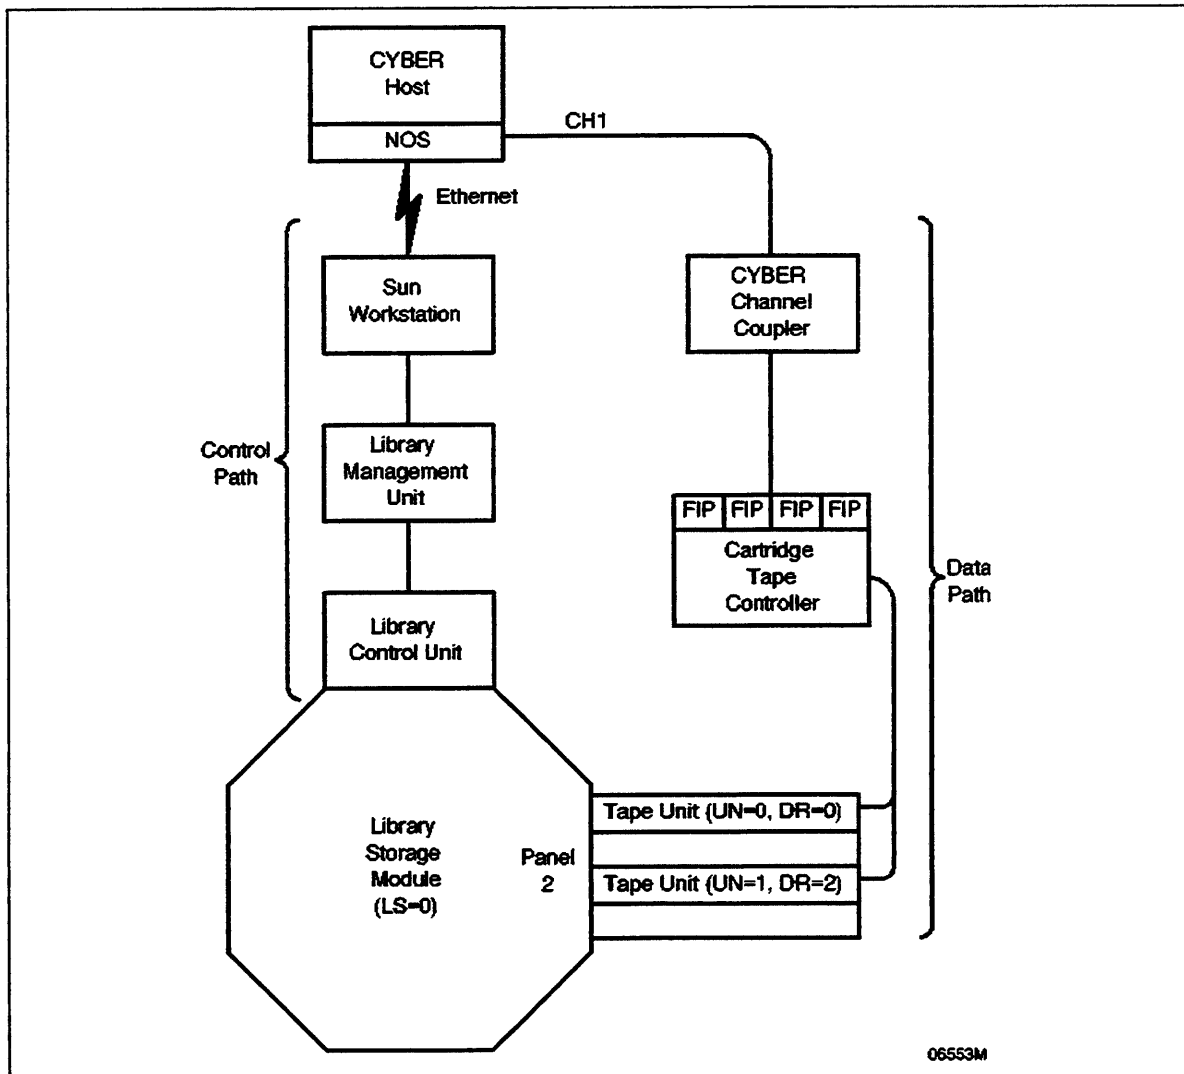


Figure 22-4. ACS Minimum Configuration

Communication Paths

There are two communication paths from the mainframe to a cartridge unit, the control path and the data path. These two paths are completely independent; the control path has no access to the data path, and the data path has no access to the control path.

The control path is managed by the ACS library server and is used to mount and dismount a cartridge from a unit. The ACS library server has no access to the data on a cartridge.

The data path is managed by NOS and is used to verify the tape label (if applicable), to access data on the volume, and to unload the unit. The same data path can access both tape units connected to ACS and tape units connected to CTS.

A tape unit has two addresses, one for the data path and one for the control path. The magnetic tape equipment EST entry defines the data path. The EST entry also defines part of the control path; it specifies the exact location of the tape unit in the library storage module by specifying the addresses (indexes) of the library storage module, panel, and tape unit. The CDCNET configuration file defines the remainder of the control path.

EQPDECK and IPRDECK Considerations

In the EQPDECK, you must create EST entries for the ACS tape units; the entries specify the exact location of the tape units by using the indexes associated with the library storage module, the panel in the module, and the tape unit in the panel. A site can have from 1 to 16 library storage modules; the library storage module indexes are assigned consecutively from 0 to 17₈. Each library storage module has 12 panels; the panel indexes are assigned consecutively from 0 to 13₈. One tape cabinet can be connected to a panel, and a maximum of four panels can have tape cabinets. Within the cabinet, the index of the tape unit (drive) is determined by the position of the unit in the cabinet. The uppermost unit has the index zero (0), the one beneath has the index one (1), and so on.

In the IPRDECK you must:

- Enable the NAM subsystem. Before you enable it, add ATF to the NAM startup file.
- Decide whether ACS is your site's default device type. If it is, enter:

```
TDTY=AT.
```

- Decide whether permanent files that reside on ACS as alternate tape storage are automatically staged to disk when requested by a user. To automatically stage the files, enter:

```
ENABLE,ACS TAPE PF STAGING.
```

You can also enable file staging with a DSD command. The default is to not stage the files.

The EQPDECK and IPRDECK are described in section 3, Deadstart Decks.

Example of Minimum ACS Configuration

For the minimum configuration shown in figure 22-4, you enter EQPDECK entries and an IPRDECK entry.

The EQPDECK entries, the EST entries for the two tape units, could be:

```
EQ60=AT, ST=ON, CH=1, UN=0, LS=0, PA=2, DR=0.
EQ61=AT, ST=ON, CH=1, UN=1, LS=0, PA=2, DR=2.
```

The IPRDECK entry could specify 9-track tapes as the default tape device type:

```
TDTY=NT.
```

Example of ACS Configuration with Two Library Storage Modules

An ACS with two library storage modules is shown in figure 22-5. You can define the EST ordinals for the tape units in one library storage module (LS=0) with the EQPDECK entries:

```
EQ60=AT, ST=ON, CH=1, UN=0, LS=0, PA=2, DR=0.
EQ61=AT, ST=ON, CH=1, UN=1, LS=0, PA=2, DR=2.
```

Define the EST ordinals for the tape units in the other library storage module (LS=1) with the EQPDECK entries:

```
EQ62=AT, ST=ON, CH=1, UN=2, LS=1, PA=3, DR=0.
EQ63=AT, ST=ON, CH=1, UN=3, LS=1, PA=3, DR=2.
```

The IPRDECK entry could specify 9-track tapes as the default tape device type:

```
TDTY=NT.
```

For an example of cross-coupled controllers, refer to Cross-Coupled Controllers for CTS and ACS.

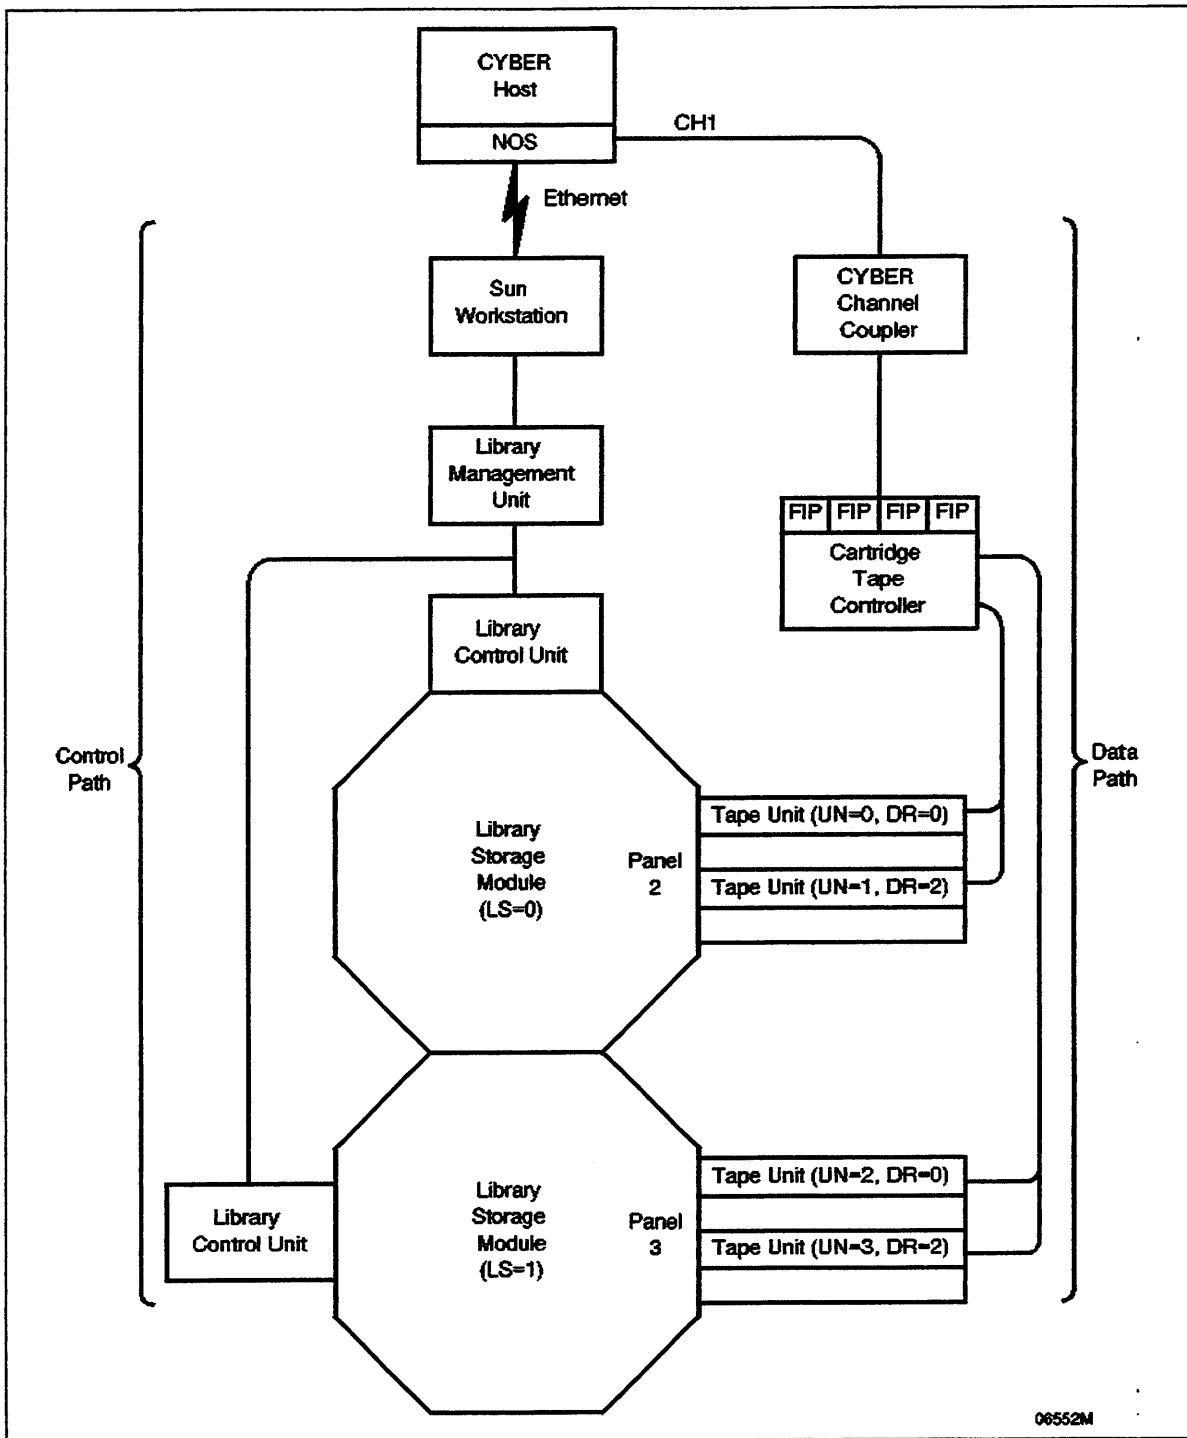


Figure 22-5. ACS Configuration with Two LSMs

Automated Tape Facility Subsystem (ATF)

The Automated Tape Facility Subsystem (ATF) is a NAM application that enables communication between the Sun Workstation software and NOS. The NAM startup file activates ATF, and ATF monitors both the network and NOS for requests. To configure ATF, refer to the CDCNET Configuration Guide.

Workstation Commands

Although ACS normally runs unattended, the Sun Workstation provides a variety of UNIX operator commands and utility functions. These commands and utilities enable the workstation operator to perform the following tasks:

- Start and stop operations of the workstation.
- Allocate and deallocate library resources.
- Recover library resources if a workstation failure occurs.
- Display the status of all library components.
- Locate a volume within the library.
- Obtain information about cartridges.
- Redistribute scratch volumes.
- Mount and dismount cartridges.
- Enter and eject cartridges from the library.
- Audit the physical contents of storage cells.

For information about the Sun Workstation commands and utility functions, refer to the 5744 Automated Cartridge Subsystem (ACS) Administration Guide.

Cross-Coupled Controllers for CTS and ACS

In addition to the hardware components previously discussed, you can purchase 5683-33 cables to cross-couple two tape controllers in either CTS or ACS. Cross-coupling allows:

- Two cartridges to be read or written simultaneously.
- One controller to access all the tape units when the other channel or controller is being repaired or maintained.

When you cross-couple controllers, the system automatically selects one controller to access all even-numbered tape units and the other controller to access all odd-numbered tape units. Consequently, an even-numbered tape unit and an odd-numbered tape unit can be accessed concurrently.

If one of the cross-coupled controllers is faulty, change the status of its channel to off and physically turn the controller's switch to offline. Once the faulty controller has processed the offline status, the other controller can access all tape units.

CTS Example:

A pair of cross-coupled controllers for CTS is shown in figure 22-6. The following example creates EST entries for the tape units shown in figure 22-6:

EQ60=CT-4, ST=ON, CH=1/2, UN=0.

For the preceding EST entry, NOS generates the EST entries:

EQ60=CT, ST=ON, CH=1/2, UN=0.

EQ61=CT, ST=ON, CH=1/2, UN=1.

EQ62=CT, ST=ON, CH=1/2, UN=2.

EQ63=CT, ST=ON, CH=1/2, UN=3.

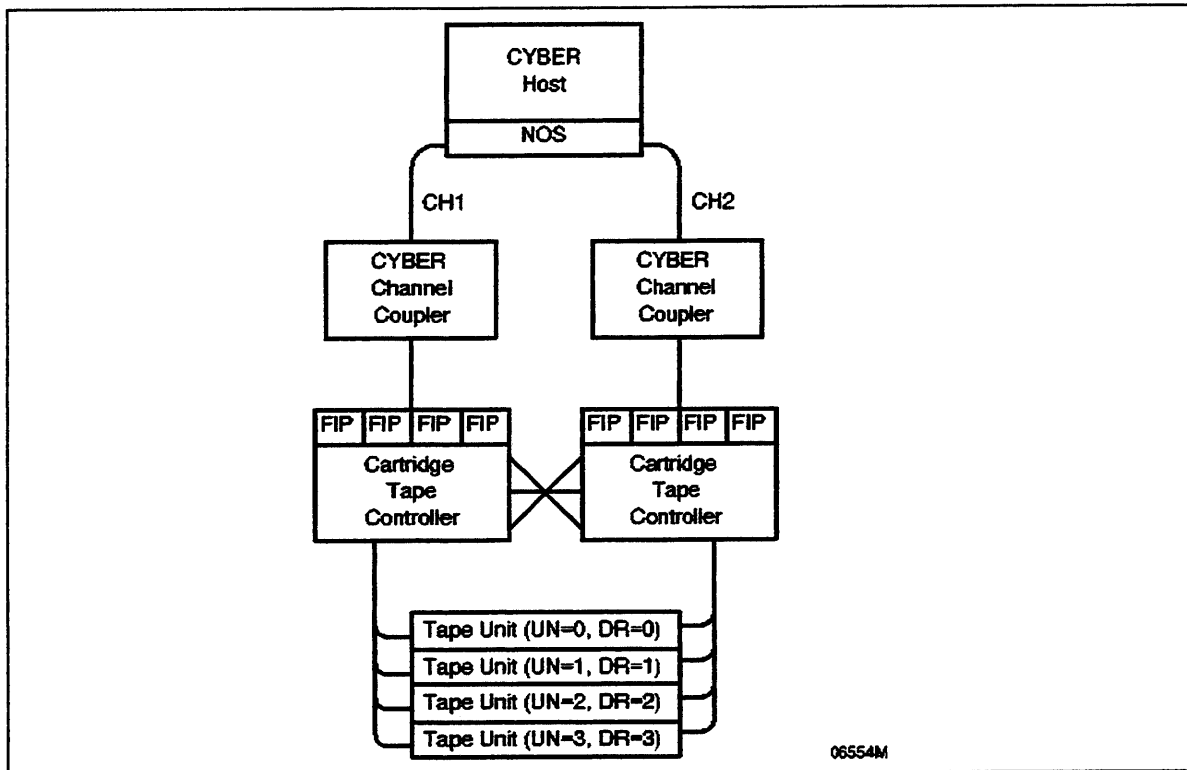


Figure 22-6. CTS Cross-Coupled Controllers

ACS Example:

A pair of cross-coupled controllers for ACS is shown in figure 22-7. The following example creates EST entries for the tape units shown in figure 22-7:

EQ60=AT, ST=ON, CH=1/2, UN=0, LS=1, PA=2, DR=0.
 EQ61=AT, ST=ON, CH=1/2, UN=1, LS=1, PA=2, DR=1.
 EQ62=AT, ST=ON, CH=1/2, UN=2, LS=1, PA=2, DR=2.
 EQ63=AT, ST=ON, CH=1/2, UN=3, LS=1, PA=2, DR=3.

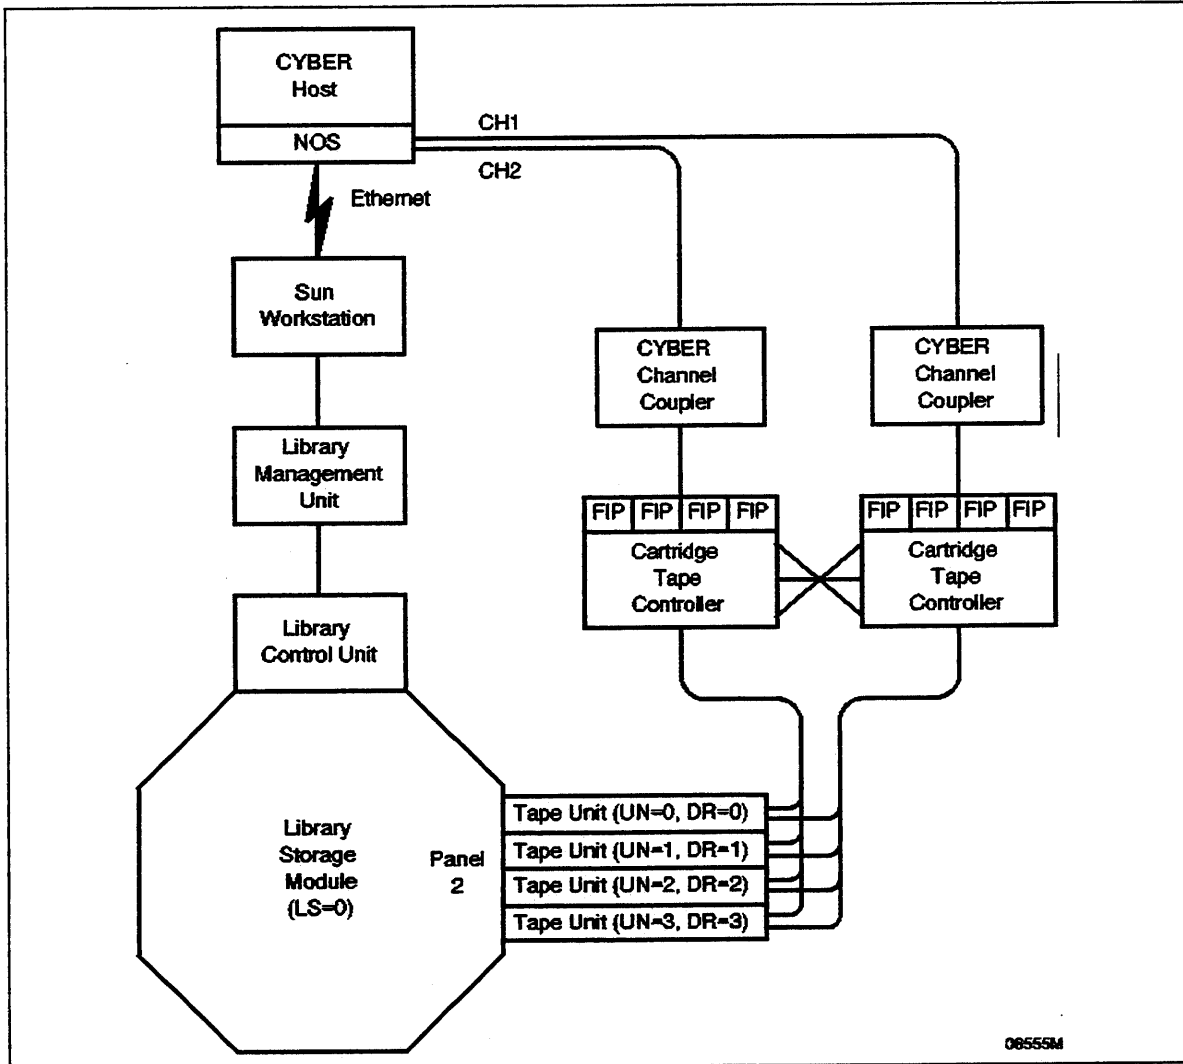


Figure 22-7. ACS Cross-Coupled Controllers

Requirements for CTS and ACS

CTS and ACS configurations have the following requirements:

- A CYBER 180 mainframe that has 4X speed peripheral processors (PPs). All models 835 and higher meet this requirement.
- The CDCNET communications network.
- Two PPs per channel connection when F, L, or LI format tapes are used. Otherwise, one PP per channel connection when only I format tapes are used.
- One NIO channel for each CCC connection.
- A CCC can connect to only one controller.
- A CCC must be a C model with field change order (FCO) #51479 installed; the microcode must be at least at level MB468-DO4.

ACS configurations have the following additional requirements:

- Install and enable the Tape Management System (TMS). Refer to the NOS Installation Handbook and the Tape Management System (TMS) Site Operations Manual.
- Configure the Automated Tape Facility Subsystem (ATF); refer to the CDCNET Configuration Guide. ATF is described earlier in this section.
- Install the Transmission Control Protocol/Internet Protocol (TCP/IP) and configure TCP and UDP. Refer to the CDCNET Configuration Guide.
- Create an entry for each CYBER system in the */etc/hosts* file on the Sun Workstation. Refer to the UNIX documentation accompanying the Sun Workstation.
- Install the library server software on the Sun Workstation. For the appropriate software levels, refer to the NOS Software Release Bulletin. For information on installing library server software, refer to the 5744 Automated Cartridge Subsystem (ACS) Installation and Service manual.
- Modify the NOS EQPDECK and IPRDECK as described earlier in EQPDECK and IPRDECK Considerations.

Constraints for CTS and ACS

NOS places the following constraints on both CTS and ACS unless noted otherwise:

- The maximum number of tape channels is four; this includes reel and cartridge tapes.
- Cartridge tapes cannot be used as a deadstart device nor as an Express Deadstart Dump (EDD) tape.
- NOS supports only binary mode on cartridge tapes. You can, however, access coded mode tapes in binary mode and use the FCOPY command to make the necessary character conversion.
- If NOS is run in secured mode (OPSECM is nonzero), all ACS tape units should be defined with the same range of access levels.

Performance Considerations for CTS and ACS

The following items affect performance for both CTS and ACS:

- A cartridge tape controller has four partitions; each partition controls data transfers for one tape unit. If a controller controls more than four tape units and receives more than four tape requests, the requests are time-shared. For optimum performance, configure no more than four tape units per controller.
- You can connect up to four channels to the same controller to provide either redundant access or multiple mainframe access. The maximum I/O throughput is achieved, however, with a single channel connection.
- If a configuration has multiple controllers, cross-couple the controllers for optimum performance.
- If multiple CYBER mainframes share one CTS or ACS configuration, each mainframe must have its own dedicated tape unit(s). For optimum performance, configure a separate controller for each mainframe.
- Multiple operating systems can share CTS and ACS. For optimal performance, each operating system should have its own separate, dedicated controller and tape unit(s).

Maintenance for CTS and ACS

The controller and tape units have extensive self-testing capabilities at power-up and load time. For a controller, a customer engineer can also execute inline diagnostics with the Detached Diagnostic Device. Before doing preventive or corrective maintenance on a controller or tape unit, change the status of its channel to off and physically turn the controller's switch to offline. Refer to the 5744 Automated Cartridge Subsystem (ACS) Maintenance Manual.

Character Sets

A

Character Set Anomolies	A-2
Character Set Tables	A-2
Interactive Jobs	A-3
Batch Jobs	A-3
Jobs Using Line Printers	A-4
Jobs Using Magnetic Tape	A-17

A character set is composed of graphic and control characters. A code set is a set of codes used to represent each character within a character set.

A graphic character may be displayed at a terminal or printed by a line printer. Examples are the characters A through Z and the digits 0 through 9. A control character initiates, modifies, or stops a control operation. An example is the backspace character that moves the terminal carriage or cursor back one space. Although a control character is not a graphic character, a terminal may produce a graphic representation when it receives a control character.

All references within this manual to the ASCII character set or the ASCII code set refer to the character set and code set defined in the American National Standard Code for Information Interchange (ASCII, ANSI Standard X3.4-1977). References in this manual to the ASCII character set do not necessarily refer to the ASCII code set.

NOS supports the following character sets.

- CDC graphic 64- (or 63-) character set.
- ASCII 128-character set.
- ASCII graphic 64- (or 63-) character set.
- ASCII graphic 95-character set.

Each installation selects either the 64-character set or the 63-character set. The differences between the two are described next, under Character Set Anomalies. Any reference in this appendix to the 64-character set implies either the 63- or 64-character set, unless otherwise stated.

NOS supports the following code sets.

- 6-bit display code.
- 6/12-bit display code.
- 7-bit ASCII code.

Display code is a set of 6-bit codes from 00_8 to 77_8 .

The 6/12-bit display code is a combination of 6-bit codes and 12-bit codes. The 6-bit codes are 00_8 through 77_8 , excluding 74_8 and 76_8 . (Refer to Character Set Anomalies, next, for the interpretation of the 00_8 and 63_8 codes.) The 12-bit codes begin with either 74_8 or 76_8 and are followed by a 6-bit code. Thus, 74_8 and 76_8 are considered escape codes and are never used as 6-bit codes within the 6/12-bit display code set. The 12-bit codes are 7401_8 , 7402_8 , 7404_8 , 7407_8 , and 7601_8 through 7677_8 . All other 12-bit codes ($74xx_8$ and 7600_8) are undefined.

The 7-bit ASCII code (as defined by ANSI Standard X3.4-1977) is right-justified in a 12-bit byte. Assuming that the bits are numbered from the right starting with 0, bits 0 through 6 contain the ASCII code, bits 7 through 10 contain zeros, and bit 11 distinguishes the 0000_8 code from the end-of-line byte. The 7-bit codes are 0001_8 through 0177_8 and 4000_8 .

Character Set Anomalies

NOS interprets two codes differently when the installation selects the 63-character set rather than the 64-character set. In tables A-1, A-2, and A3, the codes for the colon and percent graphic characters in the 64-character set are unshaded; the codes for the colon and percent graphic characters in the 63-character set are shaded.

If an installation uses the 63-character set, the colon graphic character is always represented by a 63₈ code, and the 00₈ code is undefined. However, if the installation uses the 64-character set, output of 6/12-bit display codes 7404₈ or 6-bit display code 00₈ produces a colon. In ASCII mode for interactive jobs, a colon can be input only as a 7404₈ 6/12-bit display code.

When using either the 63- or 64-character set, the use of undefined 6/12-bit display codes in output files produces unpredictable results and should be avoided.

On input, NOS recognizes alternate 029 punch codes of 11-0 for the right bracket (]) and 12-0 for the left bracket ([). The alternate codes support the COBOL sign overpunch convention and are not recommended for other uses. Refer to the COBOL 5 Reference Manual.

Also, two 00₈ codes may be confused with an end-of-line byte and should be avoided (refer to the NOS 2 Reference Set, Volume 3 for further explanation).

Translation of 7-bit ASCII to 6-bit display code causes character set folding from the 128-character ASCII set to the 63- or 64-character ASCII subset. The following special character substitutions occur:

7-Bit ASCII		6-Bit Display Code		7-Bit ASCII	
Code	Character	Code	Character	Code	Character
0140	~	74	@	0100	@
0173	{	61	[0133	[
0174		75	\	0134	\
0175	}	62]	0135]
0176	~	76	^	0136	^

Character Set Tables

This appendix includes character set tables for interactive jobs, batch jobs, and jobs involving magnetic tapes. Table A-1 is for interactive jobs, and table A-2 is for batch jobs. Table A-3 is a conversion table used to cross-reference 7-bit ASCII codes and 6/12-bit display codes and to convert ASCII codes from octal to hexadecimal.

Tables A-4, A-5, and A-6 list the magnetic tape codes and their display code equivalents.

The character set tables are designed so that you can find the character represented by a code (such as in a dump) or find the code that represents a character. To find the character represented by a code, look up the code in the column listing the appropriate code set and then find the character on that line in the column listing the appropriate character set. To find the code that represents a character, you first look up the character and then find the code on the same line in the appropriate column.

Interactive Jobs

Table A-1 shows the character sets and code sets available to you at an ASCII code terminal. When in NORMAL mode (specified by the NORMAL command), NOS displays the ASCII graphic 64-character set and interprets all input and output as display code. When in ASCII mode (specified by the ASCII command), NOS displays the ASCII 128-character set and interprets all input and output as 6/12-bit display code.

To determine the octal or hexadecimal ASCII code for a character, refer to table A-3. (Certain terminal definition commands require specification of an ASCII code.)

On output, the US code is reserved for network use and defined as an end-of-line. Use of this character, except in transparent mode, causes incorrect formatting and possible loss of output characters.

Batch Jobs

Table A-2 lists the CDC graphic 64-character set, the ASCII graphic 64-character set, and the ASCII graphic 95-character sets. It also lists the code sets and card punch codes (O26 and O29) that represent the characters.

The 64-character sets use display code as their code set; the 95-character set uses 7-bit ASCII code. The 95-character set is composed of all the characters in the ASCII 128-character set that can be printed at a line printer (refer to Jobs Using Line Printers, next). Only 7-bit ASCII code files can be printed using the ASCII graphic 95-character set. To print a 6/12-bit display code file (usually created by an interactive job in ASCII mode), you must convert the file to 7-bit ASCII code. To do this, you enter the FCOPY command (refer to the NOS 2 Reference Set, Volume 3). The 95-character set is represented by 7-bit ASCII codes 0040₈ through 0176₈.

Jobs Using Line Printers

The batch character set printed depends on the print train used on the line printer to which the file is sent (refer to the ROUTE command in the NOS 2 Reference Set, Volume 3). The following are the print trains corresponding to each of the batch character sets.

Character Set	Print Train	PSU Printer Print Band
CDC graphic 64-character set	596-1	—
ASCII graphic 64-character set	596-5	530-1
ASCII graphic 95-character set	596-6	530-2

The characters of the default 596-1 print train are listed in the table A-2 column labeled CDC Graphic (64 Character); the 596-5 print train characters are listed in the table A-2 column labeled ASCII Graphic (64 Character); and the 596-6 print train characters are listed in the table A-2 column labeled ASCII Graphic (95 Character).

If a transmission error occurs when printing a line, the system prints the line again. The CDC graphic print train prints a concatenation symbol (. .) in the first printable column of the repeated listing of the line. The ASCII print trains print an underline (_) instead of the concatenation symbol.

If an unprintable character exists in a line (that is, a 7-bit ASCII code outside the range 0040₈ through 0176₈), the number sign (#) appears in the first printable column of a print line, and a space replaces the unprintable character.

To route and correctly print a 6/12-bit display code file on a line printer with the ASCII graphic 95-character set, you must convert the 6/12-bit display code file to a 7-bit ASCII code file with the FCOPY command (refer to the NOS 2 Reference Set, Volume 3). The resulting 7-bit ASCII file can be routed to a line printer (refer to the ROUTE command in the NOS 2 Reference Set, Volume 3) but cannot be output at an interactive terminal.

Table A-1. Character Sets for Interactive Jobs

ASCII Graphic (64-Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code
: colon		00 ^{1,2}		
A	A	01	01	0101
B	B	02	02	0102
C	C	03	03	0103
D	D	04	04	0104
E	E	05	05	0105
F	F	06	06	0106
G	G	07	07	0107
H	H	10	10	0110
I	I	11	11	0111
J	J	12	12	0112
K	K	13	13	0113
L	L	14	14	0114
M	M	15	15	0115
N	N	16	16	0116
O	O	17	17	0117
P	P	20	20	0120
Q	Q	21	21	0121
R	R	22	22	0122
S	S	23	23	0123
T	T	24	24	0124
U	U	25	25	0125
V	V	26	26	0126
W	W	27	27	0127
X	X	30	30	0130
Y	Y	31	31	0131
Z	Z	32	32	0132
0	0	33	33	0060
1	1	34	34	0061
2	2	35	35	0062
3	3	36	36	0063
4	4	37	37	0064

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

2. Display code 00 is undefined at sites using the 63-character set.

(Continued on next page)

Table A-1. Character Sets for Interactive Jobs

ASCII Graphic (64-Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code
<i>(Continued from previous page)</i>				
5	5	40	40	0065
6	6	41	41	0066
7	7	42	42	0067
8	8	43	43	0070
9	9	44	44	0071
+ plus	+ plus	45	45	0053
- dash	- dash	46	46	0055
* asterisk	* asterisk	47	47	0052
/ slant	/ slant	50	50	0057
(opening parenthesis	(opening parenthesis	51	51	0050
) closing parenthesis) closing parenthesis	52	52	0051
\$ dollar sign	\$ dollar sign	53	53	0044
= equal	= equal	54	54	0075
space	space	55	55	0040
, comma	, comma	56	56	0054
. period	. period	57	57	0056
# number sign	# number sign	60	60	0043
[opening bracket	[opening bracket	61	61	0133
] closing bracket] closing bracket	62	62	0135
% percent sign	% percent sign	63	63 ¹	0045
: colon	: colon	63	63	0072
" quote	" quote	64	64	0042
_ underline	_ underline	65	65	0137
! exclamation point	! exclamation point	66	66	0041
& ampersand	& ampersand	67	67	0046
' apostrophe	' apostrophe	70	70	0047
? question mark	? question mark	71	71	0077
< less than	< less than	72	72	0074
> greater than	> greater than	73	73	0076
@ commercial at	@ commercial at	74 ¹	7401	0100
\ reverse slant	\ reverse slant	75	75	0134
^ circumflex	^ circumflex	76	7402	0136
; semicolon	; semicolon	77	77	0073

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

(Continued on next page)

Table A-1. Character Sets for Interactive Jobs

ASCII Graphic (64-Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code
<i>(Continued from previous page)</i>				
	: colon		7404	0072
	% percent		7404	0045
	‘ grave accent	74 ¹	7407	0140
	a		7601	0141
	b		7602	0142
	c		7603	0143
	d		7604	0144
	e		7605	0145
	f		7606	0146
	g		7607	0147
	h		7610	0150
	i		7611	0151
	j		7612	0152
	k		7613	0153
	l		7614	0154
	m		7615	0155
	n		7616	0156
	o		7617	0157
	p		7620	0160
	q		7621	0161
	r		7622	0162
	s		7623	0163
	t		7624	0164
	u		7625	0165
	v		7626	0166
	w		7627	0167
	x		7630	0170
	y		7631	0171
	z		7632	0172
	{ opening brace	61 ¹	7633	0173
	vertical line	75 ¹	7634	0174
	} closing brace	62 ¹	7635	0175
	~ tilde	76 ¹	7636	0176
	DEL		7637	0177

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

(Continued on next page)

Table A-1. Character Sets for Interactive Jobs

ASCII Graphic (64-Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code
<i>(Continued from previous page)</i>				
	NUL		7640	4000
	SOH		7641	0001
	STX		7642	0002
	ETX		7643	0003
	EOT		7644	0004
	ENQ		7645	0005
	ACK		7646	0006
	BEL		7647	0007
	BS		7650	0010
	HT		7651	0011
	LF		7652	0012
	VT		7653	0013
	FF		7654	0014
	CR		7655	0015
	SO		7656	0016
	SI		7657	0017
	DLE		7660	0020
	DC1		7661	0021
	DC2		7662	0022
	DC3		7663	0023
	DC4		7664	0024
	NAK		7665	0025
	SYN		7666	0026
	ETB		7667	0027
	CAN		7670	0030
	EM		7671	0031
	SUB		7672	0032
	ESC		7673	0033
	FS		7674	0034
	GS		7675	0035
	RS		7676	0036
	US ¹		7677	0037

1. Reserved for network use. Refer to Character Set Tables in this appendix.

Table A-2. Character Sets for Batch Jobs

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Dis- play Code	6/12- Bit Dis- play Code	7-Bit ASCII Code	Punch Code 026	Punch Code 029
: colon ^{1,2}	: colon ¹		00 ¹			8-2	8-2
A	A	A	01	01	0101	12-1	12-1
B	B	B	02	02	0102	12-2	12-2
C	C	C	03	03	0103	12-3	12-3
D	D	D	04	04	0104	12-4	12-4
E	E	E	05	05	0105	12-5	12-5
F	F	F	06	06	0106	12-6	12-6
G	G	G	07	07	0107	12-7	12-7
H	H	H	10	10	0110	12-8	12-8
I	I	I	11	11	0111	12-9	12-9
J	J	J	12	12	0112	11-1	11-1
K	K	K	13	13	0113	11-2	11-2
L	L	L	14	14	0114	11-3	11-3
M	M	M	15	15	0115	11-4	11-4
N	N	N	16	16	0116	11-5	11-5
O	O	O	17	17	0117	11-6	11-6
P	P	P	20	20	0120	11-7	11-7
Q	Q	Q	21	21	0121	11-8	11-8
R	R	R	22	22	0122	11-9	11-9
S	S	S	23	23	0123	0-2	0-2
T	T	T	24	24	0124	0-3	0-3
U	U	U	25	25	0125	0-4	0-4
V	V	V	26	26	0126	0-5	0-5
W	W	W	27	27	0127	0-6	0-6
X	X	X	30	30	0130	0-7	0-7
Y	Y	Y	31	31	0131	0-8	0-8
Z	Z	Z	32	32	0132	0-9	0-9
0	0	0	33	33	0060	0	0
1	1	1	34	34	0061	1	1
2	2	2	35	35	0062	2	2
3	3	3	36	36	0063	3	3
4	4	4	37	37	0064	4	4

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

2. Display code 00 is undefined at sites using the 63-character set.

(Continued on next page)

Table A-2. Character Sets for Batch Jobs

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Dis- play Code	6/12- Bit Dis- play Code	7-Bit ASCII Code	Punch Code 026	Punch Code 029
<i>(Continued from previous page)</i>							
5	5	5	40	40	0065	5	5
6	6	6	41	41	0066	6	6
7	7	7	42	42	0067	7	7
8	8	8	43	43	0070	8	8
9	9	9	44	44	0071	9	9
+ plus	+ plus	+ plus	45	45	0053	12	12-8-6
- dash	- dash	-dash	46	46	0055	11	11
* asterisk	* asterisk	* asterisk	47	47	0052	11-8-4	11-8-4
/ slant	/ slant	/ slant	50	50	0057	0-1	0-1
(opening parenthesis	(opening parenthesis	(opening parenthesis	51	51	0050	0-8-4	12-8-5
) closing parenthesis) closing parenthesis) closing parenthesis	52	52	0051	12-8-4	11-8-5
\$ dollar sign	\$ dollar sign	\$ dollar sign	53	53	0044	11-8-3	11-8-3
= equal	= equal	= equal	54	54	0075	8-3	8-6
space	space	space	55	55	0040	no punch	no punch
, comma	, comma	, comma	56	56	0054	0-8-3	0-8-3
. period	. period	. period	57	57	0056	12-8-3	12-8-3
≡ equivalence	# number sign	# number sign	60	60	0043	0-8-6	8-3
[opening bracket	[opening bracket	[opening bracket	61	61	0133	8-7	12-8-2 ¹
] closing bracket] closing bracket] closing bracket	62	62	0135	0-8-2	11-8-2 ¹
% percent sign ¹	% percent sign ¹	% percent sign ¹	63 ¹	63 ¹	0045	8-6	0-8-4
: colon	: colon	: colon	63	63	0072	8-2	8-2
≠ not equal	" quote	"quote	64	64	0042	8-4	8-7
.. concaten- ation	— underline	— underline	65	65	0137	0-8-5	0-8-5
∨ logical OR	! exclam- ation point	! exclam- ation point	66	66	0041	11-0	12-8-7
∧ logical AND	& ampersand	& ampersand	67	67	0046	0-8-7	12

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

(Continued on next page)

Table A-2. Character Sets for Batch Jobs

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Dis- play Code	6/12- Bit Dis- play Code	7-Bit ASCII Code	Punch Code 026	Punch Code 029
<i>(Continued from previous page)</i>							
↑ superscript	' apostrophe	' apostrophe	70	70	0047	11-8-5	8-5
↓ subscript	? question mark	? question mark	71	71	0077	11-8-6	0-8-7
< less than	< less than	< less than	72	72	0074	12-0	12-8-4
> greater than	> greater than	> greater than	73	73	0076	11-8-7	0-8-6
≤ less or equal	@ commercial at		74			8-5	8-4
≥ greater or equal	\ reverse slant	\ reverse slant	75	75	0134	12-8-5	0-8-2
– logical NOT	^ circumflex		76			12-8-6	11-8-7
; semicolon	; semicolon	; semicolon	77	77	0073	12-8-7	11-8-6
		@ commercial at	74 ¹	7401	0100		
		^ circumflex	76 ¹	7402	0136		
		: colon ¹		7404 ¹	0072		
		% percent sign		7404	0045		
		' grave accent	74 ¹	7407	0140		
		a		7601	0141		
		b		7602	0142		
		c		7603	0143		
		d		7604	0144		
		e		7605	0145		
		f		7606	0146		
		g		7607	0147		

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

(Continued on next page)

Table A-2. Character Sets for Batch Jobs

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Dis- play Code	6/12- Bit Dis- play Code	7-Bit ASCII Code	Punch Code 026	Punch Code 029
<i>(Continued from previous page)</i>							
		h		7610	0150		
		i		7611	0151		
		j		7612	0152		
		k		7613	0153		
		l		7614	0154		
		m		7615	0155		
		n		7616	0156		
		o		7617	0157		
		p		7620	0160		
		q		7621	0161		
		r		7622	0162		
		s		7623	0163		
		t		7624	0164		
		u		7625	0165		
		v		7626	0166		
		w		7627	0167		
		x		7630	0170		
		y		7631	0171		
		z		7632	0172		
		{ opening brace	61 ¹	7633	0173		
		vertical line	75 ¹	7634	0174		
		} closing brace	62 ¹	7635	0175		
		~ tilde	76 ¹	7636	0176		

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

Table A-3. ASCII to 6/12-Bit Display Code Conversion

ASCII Character (128 Character)	7-Bit ASCII Code Octal	7-Bit ASCII Code Hexadecimal	6/12-Bit Display Code
NUL	4000	00	7640
SOH	0001	01	7641
STX	0002	02	7642
ETX	0003	03	7643
EOT	0004	04	7644
ENQ	0005	05	7645
ACK	0006	06	7646
BEL	0007	07	7647
BS	0010	08	7650
HT	0011	09	7651
LF	0012	0A	7652
VT	0013	0B	7653
FF	0014	0C	7654
CR	0015	0D	7655
SO	0016	0E	7656
SI	0017	0F	7657
DLE	0020	10	7660
DC1	0021	11	7661
DC2	0022	12	7662
DC3	0023	13	7663
DC4	0024	14	7664
NAK	0025	15	7665
SYN	0026	16	7666
ETB	0027	17	7667
CAN	0030	18	7670
EM	0031	19	7671
SUB	0032	1A	7672
ESC	0033	1B	7673
FS	0034	1C	7674
GS	0035	1D	7675
RS	0036	1E	7676
US ¹	0037	1F	7677 ¹
space	0040	20	55
! exclamation point	0041	21	66
" quote	0042	22	64
# number sign	0043	23	60
\$ dollar sign	0044	24	53
% percent sign ²	0045	25	63 ²

1. Reserved for network use. Refer to Character Set Tables in this appendix.

2. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies in this appendix.

(Continued on next page)

Table A-3. ASCII to 6/12-Bit Display Code Conversion

ASCII Character (128 Character)	7-Bit ASCII Code Octal	7-Bit ASCII Code Hexadecimal	6/12-Bit Display Code
<i>(Continued from previous page)</i>			
% percent sign	0045	25	7404
& ampersand	0046	26	67
' apostrophe	0047	27	70
(opening parenthesis	0050	28	51
) closing parenthesis	0051	29	52
* asterisk	0052	2A	47
+ plus	0053	2B	45
, comma	0054	2C	56
- dash	0055	2D	46
. period	0056	2E	57
/ slant	0057	2F	50
0	0060	30	33
1	0061	31	34
2	0062	32	35
3	0063	33	36
4	0064	34	37
5	0065	35	40
6	0066	36	41
7	0067	37	42
8	0070	38	43
9	0071	39	44
: colon¹	0072	3A	7404 ¹
: colon	0072	3A	63
; semicolon	0073	3B	77
< less than	0074	3C	72
= equal	0075	3D	54
> greater than	0076	3E	73
? question mark	0077	3F	71
@ commercial at	0100	40	7401
A	0101	41	01
B	0102	42	02
C	0103	43	03
D	0104	44	04
E	0105	45	05
F	0106	46	06
G	0107	47	07
H	0110	48	10

1. The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies in this appendix.

(Continued on next page)

Table A-3. ASCII to 6/12-Bit Display Code Conversion

ASCII Character (128 Character)	7-Bit ASCII Code Octal	7-Bit ASCII Code Hexadecimal	6/12-Bit Display Code
<i>(Continued from previous page)</i>			
I	0111	49	11
J	0112	4A	12
K	0113	4B	13
L	0114	4C	14
M	0115	4D	15
N	0116	4E	16
O	0117	4F	17
P	0120	50	20
Q	0121	51	21
R	0122	52	22
S	0123	53	23
T	0124	54	24
U	0125	55	25
V	0126	56	26
W	0127	57	27
X	0130	58	30
Y	0131	59	31
Z	0132	5A	32
[opening bracket	0133	5B	61
\ reverse slant	0134	5C	75
] closing bracket	0135	5D	62
^ circumflex	0136	5E	7402
_ underline	0137	5F	65
‘ grave accent	0140	60	7407
a	0141	61	7601
b	0142	62	7602
c	0143	63	7603
d	0144	64	7604
e	0145	65	7605
f	0146	66	7606
g	0147	67	7607
h	0150	68	7610
i	0151	69	7611
j	0152	6A	7612
k	0153	6B	7613
l	0154	6C	7614
m	0155	6D	7615
n	0156	6E	7616
o	0157	6F	7617

(Continued on next page)

Table A-3. ASCII to 6/12-Bit Display Code Conversion

ASCII Character (128 Character)	7-Bit ASCII Code Octal	7-Bit ASCII Code Hexadecimal	6/12-Bit Display Code
<i>(Continued from previous page)</i>			
p	0160	70	7620
q	0161	71	7621
r	0162	72	7622
s	0163	73	7623
t	0164	74	7624
u	0165	75	7625
v	0166	76	7626
w	0167	77	7627
x	0170	78	7630
y	0171	79	7631
z	0172	7A	7632
{ opening brace	0173	7B	7633
vertical line	0174	7C	7634
} closing brace	0175	7D	7635
~ tilde	0176	7E	7636
DEL	0177	7F	7637

Jobs Using Magnetic Tape

Coded data to be copied from disk to magnetic tape is assumed to be represented in display code. NOS converts the data to external BCD code when writing a coded 7-track tape and to ASCII or EBCDIC code (as specified on the tape assignment command) when writing a coded 9-track tape.

Because only 63 characters can be represented in 7-track even parity, one of the 64 display codes is lost in conversion to and from external BCD code. Figure A-1 shows the differences in conversion depending on the character set (63 or 64) that the system uses. The ASCII character for the specified character code is shown in parentheses. The output arrow shows how the 6-bit display code changes when it is written on tape in external BCD. The input arrow shows how the external BCD code changes when the tape is read and converted to 6-bit display code.

<u>63-Character Set</u>		
<u>6-Bit Display Code</u>		<u>6-Bit Display Code</u>
00		00
33 (0)	Output →	12 (0)
63 (:)		12 (0)
		Input →
		33 (0)
		33 (0)
<u>64-Character Set</u>		
<u>6-Bit Display Code</u>		<u>6-Bit Display Code</u>
00 (:)		12 (0)
33 (0)	Output →	12 (0)
63 (Z)		16 (Z)
		Input →
		33 (0)
		63 (Z)

Figure A-1. Conversion Differences

If a lowercase ASCII or EBCDIC code is read from a 9-track coded tape, it is converted to its uppercase 6-bit display code equivalent. To read or write lowercase ASCII or EBCDIC characters, you must assign the tape in binary mode and use FCOPY to read or write the tape.

Tables A-4 and A-5 show the character set conversion for 9-track tapes. Table A-4 lists the conversions to and from the 7-bit ASCII character code and 6-bit display code. Table A-5 lists the conversions between the EBCDIC character code and the 6-bit display code. Table A-6 shows the character set conversions between external BCD and 6-bit display code for 7-track tapes.

Table A-4. Nine-Track ASCII Coded Tape Conversion

7-Bit ASCII Code (Hex)	7-Bit ASCII Char¹	7-Bit ASCII Code (Hex)	7-Bit ASCII Char²	6-Bit Display Code Char	6-Bit Display Code Code (Octal)
20	space	00	NUL	space	55
21	!	7D	}	!	66
22	"	02	STX	"	64
23	#	03	ETX	#	60
24	\$	04	EOT	\$	53
25	%	05	ENQ	%	63
25	%	05	ENQ	space³	55
26	&	06	ACK	&	67
27	'	07	BEL	'	70
28	(08	BS	(51
29)	09	HT)	52
2A	*	0A	LF	*	47
2B	+	0B	VT	+	45
2C	,	0C	FF	,	56
2D	-	0D	CR	-	46
2E	.	0E	SO	.	57
2F	/	0F	SI	/	50
30	0	10	DLE	0	33
31	1	11	DC1	1	34
32	2	12	DC2	2	35
33	3	13	DC3	3	36
34	4	14	DC4	4	37
35	5	15	NAK	5	40
36	6	16	SYN	6	41
37	7	17	ETB	7	42
38	8	18	CAN	8	43
39	9	19	EM	9	44
3A	:	1A	SUB ⁴	:	00
3A	:	1A	SUB	:	63
3B	;	1B	ESC	;	77
3C	<	7B	{	<	72
3D	=	1D	GS	=	54

1. When these characters are copied from/or to a tape, the characters remain the same but the codes change from one code set to the other.

2. These characters do not exist in 6-bit display code. Therefore, when the characters are copied from a tape, each 7-bit ASCII character is changed to an alternate 6-bit display code character. The corresponding codes are also changed. Example: When the system copies a lowercase a, 61₁₆, from tape, it writes an uppercase A, 01₈.

3. A 6-bit display code space always translates to a 7-bit ASCII space.

4. 6-Bit display code 00 is undefined at sites using the 63-character set.

(Continued on next page)

Table A-4. Nine-Track ASCII Coded Tape Conversion

7-Bit ASCII Code (Hex)	7-Bit ASCII Char¹	7-Bit ASCII Code (Hex)	7-Bit ASCII Char²	6-Bit Display Code Char	6-Bit Display Code Code (Octal)
<i>(Continued from previous page)</i>					
3E	>	1E	RS	>	73
3F	?	1F	US	?	71
40	@	60	'	@	74
41	A	61	a	A	01
42	B	62	b	B	02
43	C	63	c	C	03
44	D	64	d	D	04
45	E	65	e	E	05
46	F	66	f	F	06
47	G	67	g	G	07
48	H	68	h	H	10
49	I	69	i	I	11
4A	J	6A	j	J	12
4B	K	6B	k	K	13
4C	L	6C	l	L	14
4D	M	6D	m	M	15
4E	N	6E	n	N	16
4F	O	6F	o	O	17
50	P	70	p	P	20
51	Q	71	q	Q	21
52	R	72	r	R	22
53	S	73	s	S	23
54	T	74	t	T	24
55	U	75	u	U	25
56	V	76	v	V	26
57	W	77	w	W	27
58	X	78	x	X	30
59	Y	79	y	Y	31
5A	Z	7A	z	Z	32
5B	[1C	FS	[61
5C	\	7C		\	75
5D]	01	SOH]	62
5E	^	7E	~	^	76
5F	_	7F	DEL	_	65

1. When these characters are copied from/or to a tape, the characters remain the same but the codes change from one code set to the other.

2. These characters do not exist in 6-bit display code. Therefore, when the characters are copied from a tape, each 7-bit ASCII character is changed to an alternate 6-bit display code character. The corresponding codes are also changed. Example: When the system copies a lowercase a, 61₁₆, from tape, it writes an uppercase A, 01₈.

Table A-5. Nine-Track EBCDIC Coded Tape Conversion

EBCDIC Code (Hex)	EBCDIC Char¹	EBCDIC Code (Hex)	EBCDIC Char²	6-Bit Display Code Char	6-Bit Display Code Code (Octal)
40	space	00	NUL	space	55
4A	ø	1C	IFS	[61
4B	.	0E	SO	.	57
4C	<	C0	{	<	72
4D	(16	BS	(51
4E	+	0B	VT	+	45
4F		D0	}	!	66
50	&	2E	ACK	&	67
5A	!	01	SOH]	62
5B	\$	37	EOT	\$	53
5C	*	25	LF	*	47
5D)	05	HT)	52
5E	;	27	ESC	;	77
5F	-	A1	-	^	76
60	-	0D	CR	-	46
61	/	0F	SI	/	50
6B	,	0C	FF	,	56
6C	%	2D	ENQ	%	63
6C	%	2D	ENQ	space³	55
6D	-	07	DEL	-	65
6E	>	1E	IRS	>	73
6F	?	1F	IUS	?	71
7A	:	3F	SUB ⁴	:	00
7A	:	3F	SUB	:	63

1. When these characters are copied from/or to a tape, the characters remain the same (except EBCDIC codes 4A, 4F, 5A, and 5F) but the codes change from one code set to the other.

2. These characters do not exist in 6-bit display code. Therefore, when the characters are copied from a tape, each EBCDIC character is changed to an alternate 6-bit display code character. The corresponding codes are also changed. Example: When the system copies a lowercase a, 81₁₆, from tape, it writes an uppercase A, 01₈.

3. All EBCDIC codes not listed translate to 6-bit display code 55₈ (space). A 6-bit display code space always translates to an EBCDIC space.

4. 6-Bit display code 00 is undefined at sites using the 63-character set.

(Continued on next page)

Table A-5. Nine-Track EBCDIC Coded Tape Conversion

EBCDIC Code (Hex)	EBCDIC Char¹	EBCDIC Code (Hex)	EBCDIC Char²	6-Bit Display Code Char	6-Bit Display Code (Octal)
<i>(Continued from previous page)</i>					
7B	#	03	ETX	#	60
7C	@	79	\	@	74
7D	'	2F	BEL	'	70
7E	=	1D	IGS	=	54
7F	"	02	STX	"	64
C1	A	81	a	A	01
C2	B	82	b	B	02
C3	C	83	c	C	03
C4	D	84	d	D	04
C5	E	85	e	E	05
C6	F	86	f	F	06
C7	G	87	g	G	07
C8	H	88	h	H	10
C9	I	89	i	I	11
D1	J	91	j	J	12
D2	K	92	k	K	13
D3	L	93	l	L	14
D4	M	94	m	M	15
D5	N	95	n	N	16
D6	O	96	o	O	17
D7	P	97	p	P	20
D8	Q	98	q	Q	21
D9	R	99	r	R	22
E0	\	6A	l	\	75
E2	S	A2	s	S	23
E3	T	A3	t	T	24
E4	U	A4	u	U	25
E5	V	A5	v	V	26
E6	W	A6	w	W	27
E7	X	A7	x	X	30

1. When these characters are copied from/or to a tape, the characters remain the same (except EBCDIC codes 4A, 4F, 5A, and 5F) but the codes change from one code set to the other.

2. These characters do not exist in 6-bit display code. Therefore, when the characters are copied from a tape, each EBCDIC character is changed to an alternate 6-bit display code character. The corresponding codes are also changed. Example: When the system copies a lowercase a, 81_{16} , from tape, it writes an uppercase A, 01_8 .

(Continued on next page)

Table A-5. Nine-Track EBCDIC Coded Tape Conversion

EBCDIC Code (Hex)	EBCDIC Char¹	EBCDIC Code (Hex)	EBCDIC Char²	6-Bit Display Code Char	6-Bit Display Code (Octal)
<i>(Continued from previous page)</i>					
E8	Y	A8	y	Y	31
E9	Z	A9	z	Z	32
F0	0	10	DLE	0	33
F1	1	11	DC1	1	34
F2	2	12	DC2	2	35
F3	3	13	TM	3	36
F4	4	3C	DC4	4	37
F5	5	3D	NAK	5	40
F6	6	32	SYN	6	41
F7	7	26	ETB	7	42
F8	8	18	CAN	8	43
F9	9	19	EM	9	44

1. When these characters are copied from/or to a tape, the characters remain the same (except EBCDIC codes 4A, 4F, 5A, and 5F) but the codes change from one code set to the other.

2. These characters do not exist in 6-bit display code. Therefore, when the characters are copied from a tape, each EBCDIC character is changed to an alternate 6-bit display code character. The corresponding codes are also changed. Example: When the system copies a lowercase a, 81_{16} , from tape, it writes an uppercase A, 01_8 .

Table A-6. Seven-Track Coded Tape Conversions

External BCD	ASCII Character	6-Bit Display Code (Octal)
01	1	34
02	2	35
03	3	36
04	4	37
05	5	40
06	6	41
07	7	42
10	8	43
11	9	44
12 ¹	0	33
13	=	54
14	"	64
15	@	74
16 ¹	%	63
17	[61
20	space	55
21	/	50
22	S	23
23	T	24
24	U	25
25	V	26
26	W	27
27	X	30
30	Y	31
31	Z	32
32]	62
33	,	56
34	(51
35	-	65
36	#	60
37	&	67

1. As explained previously in this appendix, conversion of these codes depends on whether the tape is being read or written.

(Continued on next page)

Table A-6. Seven-Track Coded Tape Conversions

External BCD	ASCII Character	6-Bit Display Code (Octal)
<i>(Continued from previous page)</i>		
40	-	46
41	J	12
42	K	13
43	L	14
44	M	15
45	N	16
46	O	17
47	P	20
50	Q	21
51	R	22
52	!	66
53	\$	53
54	*	47
55	,	70
56	?	71
57	>	73
60	+	45
61	A	01
62	B	02
63	C	03
64	D	04
65	E	05
66	F	06
67	G	07
70	H	10
71	I	11
72	<	72
73	.	57
74)	52
75	\	75
76	^	76
77	;	77



A

Access Category

See File Access Category and System Access Categories.

Access Code

A hardware/software security code assigned to each NAD on the network. A NAD may communicate only with other NADs having matching codes.

Access Level

A property of each file, job, and equipment on a secured system that is used to indicate the sensitivity of information in the file or job, or the sensitivity of information that can be processed by the equipment. On a secured system, there are up to eight access levels corresponding to increasing levels of sensitivity; you are authorized to access some or all of those levels. Refer also to Equipment Access Levels, File Access Level, Job Access Level, and System Access Levels.

Access Level Limits

See Job Access Level Limits.

Account Dayfile

A dayfile that provides a history of system usage over the life of the account. It provides information necessary for accurate billing and system usage analysis.

ACN

See Application Connection Number.

ACS

See Automated Cartridge Subsystem.

AFD Utility

A dayfile dumping utility that dumps all or selected parts of the account dayfile to produce a listing.

AIP Trace Utility

See Application Interface Program Trace Utility.

Allocation Summary Table (AST)

A table that contains information used by the allocation algorithm to select the cartridges on which a file will reside in an MSE environment.

Allocation Unit (AU)

In an MSE environment, an allocation unit is the smallest allocatable portion of a cartridge. Each AU consists of 14 data strips; each cartridge has 1931 AUs.

Alternate Storage

The storage of permanent file data on external media other than mass storage such as MSE, tape, or optical disk alternate storage. When a file resides on alternate storage, the file's permanent file catalog (PFC) entry and permit data still reside on disk, but the disk space occupied by the file data can be released.

Alternate Storage Address (ASA)

The pointers in a file's PFC entry that point to the location of the file data on MSE, tape, or optical disk alternate storage.

Application

See Application Program.

Application Connection Number (ACN)

A number assigned by the NAM program to identify a particular logical connection within an application.

Application Interface Program (AIP) Trace Utility

A utility that produces a trace file of the messages transferred between IAF and NAM. The information contained in this trace can be useful in tracking network problems and in debugging application programs.

Application Program

A program resident in a host computer that provides an information storage, retrieval, and/or processing service to a remote user via the data communication network and the Network Access Method. Application programs use the system control point feature of NOS to communicate with the Network Access Method.

In the context of network software, an application program is not an interactive job, but rather a terminal servicing facility that provides terminal users with a specific processing capability such as remote job entry from batch terminals, transaction processing, entry and execution of interactive jobs, and so forth. For example, the standard Control Data Interactive Facility IAF makes terminal input and output appear the same to an executing program as file input and output; IAF is a network application program, but the executing program using IAF is an interactive job.

Archive Files

A dump of permanent files accumulated on disk that are dumped as a whole or in part to a backup tape (or other type of backup medium) to protect the files from loss in case of a device malfunction or to free a device for temporary use during preventive maintenance.

ASCII

American National Standard Code for Information Interchange. The standard character set and code used for information interchange between systems. It is a 7-bit code representing a prescribed set of 128 characters.

AST

See Allocation Summary Table.

ATF

See Automated Tape Facility.

Attach

The process of making a direct access permanent file accessible to a job by specifying the proper permanent file identification and passwords.

AU

See Allocation Unit.

Auto Recall

The act of a program releasing control of the CPU until a requested function is complete. Refer to Recall.

Automated Cartridge Subsystem (ACS)

The Control Data 5744 Automated Cartridge Subsystem (ACS) consists of hardware and software that allow reading and writing of ACS tapes (IBM 3480-compatible cartridge tapes). The ACS hardware includes a robot that automatically loads and unloads the cartridge tapes. Contrast with Cartridge Tape Subsystem.

Automated Tape Facility (ATF)

The Automated Tape Facility (ATF) is a NAM application that enables communication between the ACS library server (Sun Workstation) software and NOS. ATF must be active before users can access the Automated Cartridge Subsystem.

Auxiliary Device

A disk device that is not part of a permanent file family. Auxiliary devices can contain direct or indirect access permanent files.

B**Backup Tape**

See Archive Files.

Batch Job

The instructions and data that are submitted as a complete unit without further user intervention. The job can be punched on cards or created and submitted from a terminal.

Beginning-of-Information (BOI)

The start of the first programmer record in a file is known as the beginning-of-information. System information, such as tape labels on sequential files or indexes, does not affect the beginning-of-information.

Binary File

A noneditable file that contains a precompiled program.

BOI

BOI

See Beginning-of-Information.

Breakpoint

A designated location in a program where, if reached during program execution, a break or suspension in execution occurs.

Buffer

An intermediate storage area used to compensate for a difference in rates of data flow, or times of event occurrences, when transmitting data between central memory and an external device during input/output operations.

Byte

A group of bits. Unless prefixed (for example, a 6-bit byte), the term means 8-bit groups. When used for encoding character data, a byte represents a single character.

C

Cache

A high-speed memory that resides in the central processor.

Cartridge Alternate Storage

MSE or optical disk used as an alternate storage medium for permanent files. See Alternate Storage.

Cartridge Tape Subsystem (CTS)

The Control Data 5680 Cartridge Tape Subsystem (CTS) consists of hardware and software that allow reading and writing of CTS tapes (IBM 3480-compatible cartridge tapes). An operator manually loads and unloads the tapes. Contrast with Automated Cartridge Subsystem.

Cassette

The magnetic tape device in an NPU used for bootstrap loading of offline diagnostics and (in remote NPUs) the bootstrap load/dump operation.

Catalog Image Record (CIR)

A record written at the beginning of the archive file on which the permanent files are dumped for each incremental dump. When a file is loaded, this CIR information is placed in the permanent file catalog of the device being loaded.

Catalog Track

A track on a user's master device containing the catalog entries that define and specify the location of each permanent file created by that user. Users are assigned by groups to catalog tracks according to user index and number of catalog tracks on the master device.

Catenet

A collection of network solutions connected using gateways. Abbreviated from concatenated network.

CCITT

See Consultative Committee of International Telephone and Telegraph (CCITT).

CCP

See Communications Control Program.

CDCNET

See CONTROL DATA® Distributed Communications Network.

CDCNET Operator (COP)

The administrative operator who resides at either the system console or a terminal and controls CDCNET network elements (such as lines, terminals, trunks, and DIs associated with CDCNET).

Central Memory Resident (CMR)

The low address area of central memory reserved for tables, pointers, and subroutines necessary for operation of the operating system. It is never accessible to a user's central processor program. The monitor allocates the remainder of central memory to jobs as they are selected on a priority basis for execution.

Channel Number

The number of the data channel on which a peripheral device controller can be accessed.

Character

Unless otherwise specified, references to characters in this manual are to 7-bit ASCII code.

Charge Number

An alphanumeric identifier the installation uses to allocate charges to individual users for system usage.

Checkpoint

The process of writing a copy of your job's central memory, the system information used for job control, and the names and contents of all assigned files that are identified in a CHECKPT request to a magnetic tape or disk.

CIO

See Combined Input/Output

CIR

See Catalog Image Record.

CLA

See Communications Line Adapter.

CMR

See Central Memory Resident.

CMRDECK

CMRDECK

The central memory resident deck that resides on the deadstart file. It defines central memory, table sizes, and configuration information not oriented to equipment.

Coldstart

Procedure used to deadstart if the tape or disk controller has not yet been loaded with controlware, or the controlware is not running.

Combined Input/Output (CIO)

A system routine that performs NOS I/O.

Common Testing and Initialization (CTI)

Common deadstart process that resides on the deadstart file and the maintenance system library.

Communication Line

A complete communication circuit between a terminal and its network processing unit.

Communications Control Program (CCP)

A portion of the network software that resides in a 255x series network processing unit. This software can include routines such as the terminal interface program.

Communications Line Adapter (CLA)

Hardware that provides the interface between NPUs and modems.

Communications Supervisor (CS)

A portion of the network software written as an application program; the CS coordinates the network-oriented activities of the host computer and of the lines and terminals logically linked to it.

COMPASS

COMPrehensive ASsembly System. The standard assembly language used with CYBER 180, CYBER 170, CYBER 70, and 6000 Computer Systems. Also, the command used to assemble a program written in the COMPASS assembly language.

Connection Number

A number assigned to an IAF terminal by the system when the terminal is logged in and an entry is made for the job in the executing job table.

Connection Status

A job attribute kept in the job's executing job table (EJT) entry. The system uses it to determine the job's relationship with IAF.

Consultative Committee of International Telephone and Telegraph (CCITT)

(Actually, the Comite Consultif International Telephonique et Telegraphique). An organization chartered by the United Nations to develop and publish international standards for the communications industry.

CONTROL DATA® Distributed Communications Network (CDCNET)

A catenet system using Control Data Network Architecture, Control Data hardware, and Control Data software.

Control Point

The portion of central memory that is assigned to a job. When a job is allocated a portion of central memory, it becomes eligible for assignment to the central processor for execution.

Control Point Number

The number of the control point to which a job is assigned while the job resides in central memory. The actual number of control points is an installation parameter. Before the job can execute, each central processor program must be assigned to a control point.

Controller

Hardware device that connects channels to peripheral devices. For example, a tape controller might connect up to eight tape units to one channel.

Controlling NOP

An NPU operator (NOP) who is allowed to change the status of network elements (lines, logical links, terminals, and trunks) connected to an NPU.

Controlware

A special type of software that resides in a peripheral controller. The controlware defines the functional characteristics of the controller.

COP

See CDCNET Operator.

CS

See Communications Supervisor.

CTI

See Common Testing and Initialization.

CTS

See Cartridge Tape Subsystem.

D

DAS

See Disk Array Subsystem.

DAT

See Device Access Table.

Data Channel

One of the 9 to 24 channels (12-bit) by which information passes between the peripheral processors and peripheral devices. Refer to Channel Number.

Data Recording Drive (DRD)

A component of the 7991 Storage Module. The DRD reads data from and writes data to the cartridges.

Dayfile

A chronological file created during job execution that forms a permanent accounting and job history record. Dayfile messages are generated by operator action or when some commands are processed. A copy of the dayfile is printed with the output for each job. The user must explicitly request it in an interactive job.

DDP

The distributive data path.

Deadstart

The process of initializing the system by loading the operating system library programs and any of the product set from magnetic tape or disk. Deadstart recovery is reinitialization after system failure.

Deadstart Sequencing

The execution of a selected set of commands before normal system job scheduling is enabled.

Debug Log File Processor (DLFP)

A processor that analyzes trace files produced by the application interface program trace utility.

Destage Dump Tapes

Tapes that contain files that were destaged from disk using the PFDUMP utility with the DT parameter specified.

Destaging

The process of creating an alternate storage copy of a file on MSE, tape, or optical disk alternate storage. Files are destaged to tape or optical disk alternate storage when the site executes PFDUMP with the DT parameter specified. Files are destaged to MSE when the site executes SSMOVE with the appropriate parameters specified.

Detached Job

An interactive service class job removed from control of the Interactive Subsystem. It may or may not continue to execute, depending on the presence of commands in the command buffer or an active job step. Control is regained by recovering the EJT entry for the job.

Device Access Table (DAT)

A table that contains the logical description (family name/pack name and device number) of each disk device (shared or nonshared) that is accessible by any machine in a linked shared device multmainframe complex.

Device Index Table (DIT)

A table that is used to determine device usage in an independent shared device multimainframe complex.

Device Interface (DI)

CDCNET hardware for open system interconnection. The device interface houses processor boards in configurations that permit a network of various other data processing equipment.

Device Mask

An 8-bit quantity that identifies the group of users who have the particular device as their master device; that is, it identifies the device that contains their file catalogs, all their indirect access files, and possibly some or all of their direct access files.

DFD Utility

A dayfile dumping utility that dumps all or selected parts of the system dayfile to produce a listing.

DFLIST Utility

A utility that generates a printer listing of all permanent files created by the DFTERM utility.

DFTERM Utility

A utility that terminates an active or inactive dayfile and retains it as a direct access permanent file for later use.

DI

See Device Interface.

Diagnostic Operator (DOP)

An NPU operator who resides at a terminal and has permission only to check status and test the network.

Direct Access File

A NOS permanent file that can be attached to the your job. All changes to this file are made on the file itself rather than a temporary copy of the file (compare with Indirect Access File).

DIS (Job Display)

A system peripheral processor program similar to the system display (DSD) that provides communication between a job in central memory and the operator at the console, and permits the operator to control execution of the program through the console keyboard.

Disabled Job

An interactive service class job temporarily rolled out due to user break processing or encountering the end of its command stream. The job scheduler ignores disabled jobs.

Disk Array Subsystem (DAS)

A disk storage subsystem. The 5830 Disk Array Subsystem (DAS) consists of multiple disks that are combined to form one logical device in a single cabinet. The disks can be combined to increase storage capacity and operational speed as well as to do parity checking. The only devices that can be included in DAS are 5832 Solid State Disk (SSD) units, 5833 Sabre units, and 5838 Elite units. Refer also to Solid State Disk, Sabre, and Elite.

Display Code

A 6-bit character code set that represents alphanumeric and special characters.

Disposition Code

A 2-character mnemonic indicating the destination queue and format for processing a file named on a ROUTE function.

DIT

See Device Index Table.

DLFP

See Debug Log File Processor.

DMPNAD

See Dump NAD Memory.

DOP

See Diagnostic Operator.

Downline

The direction of output flow, from host to terminal.

DRD

See Data Recording Drive.

DSD (System Display)

The operating system program that provides communication between the operator and the system by accepting control information typed on the console keyboard and by displaying information pertinent to all jobs known to the system. DSD is permanently assigned to peripheral processor 1.

DSDI Utility

A deadstart dump interpreter utility that is called by a batch command to convert selected portions of the binary information on an express deadstart dump file into reports to be listed.

Dump NAD Memory (DMPNAD)

A utility that reads the NAD memory and formats the data into an output file.

E**ECS**

See Extended Core Storage.

EDD File

See Express Deadstart Dump File.

EDD Utility

See Express Deadstart Dump Utility.

EJT

See Executing Job Table.

EJT Ordinal

An index into the executing job table (EJT). It uniquely identifies an EJT entry. The acceptable range is from 0001 to 4095.

ELD Utility

A dayfile dumping utility that dumps all or selected parts of the error log dayfile to produce a listing.

Elite

A disk unit that can be configured only as part of the 5830 Disk Array Subsystem (DAS). The Elite disk has a formatted capacity of 1759 megabytes and can be configured two different ways: as a single unit, or as a single unit with another unit used for parity.

End-of-Chain Flag (EOC)

In an MSE environment, EOC is a flag in the 7990 catalog that identifies the last volume in a chain of allocated AUs.

End-of-File (EOF)

A boundary within a sequential file, but not necessarily the end of a file that can be referenced by name. The actual end of a named file is defined by EOI. For labeled tape, EOF and EOI (denoted by the EOF1 label) are the same. For multifile tape files, EOF and EOI do not correspond. In the product set manuals, an end-of-file is also referred to as an end-of-partition.

End-of-Information (EOI)

The end of data on a file. Information appearing after this point is not considered part of file data. In card decks, a card with a 6/7/8/9 multiple punch in column 1. On mass storage devices, the position of the last written data. On labeled tape, it is the EOF1 label. CYBER Record Manager defines end-of-information in terms of file residency and organization.

End-of-Record (EOR)

An indicator that marks the end of a logical record. Also referred to as end-of-section.

Entry Point

Entry Point

A location within a program or procedure that can be referenced from other programs. Each entry point has a unique name with which it is associated.

EOC

See End-of-Chain Flag.

EOF

See End-of-File.

EOI

See End-of-Information.

EOR

See End-of-Record.

EQPDECK

The equipment description deck used during deadstart to define the system equipment configuration.

Equipment Access Levels

A range of access levels specified for each equipment on a secured system. In order for a file to be stored or output on a given equipment, the file's access level must be within the equipment access levels for that equipment.

Equipment Number

A number from 0 to 7 that identifies the setting on a peripheral device controller.

Equipment Status Table (EST)

A central memory resident table listing all the defined equipments, parameters affecting their operation, and the status of the equipments.

ESM

See Extended Semiconductor Memory.

EST

See Equipment Status Table.

EST Ordinal

The number designating the position of an entry within the equipment status table (EST) established at each installation. Devices are identified in operator commands by EST ordinals. The EST ordinal is sometimes referred to as equipment number.

Ethernet

A baseband local area network protocol developed by the XEROX, Intel, and Digital Equipment Corporations. CDCNET is an Ethernet-compatible network.

Exchange Package

A table that contains information used during job execution. It is printed as part of the output when a job aborts.

Executing Job Table (EJT)

A central memory resident table that contains a 4-word entry for all executing jobs including interactive service class jobs. It is used to control jobs that are executing at a control point and jobs that are rolled out. Every executing job in the system has an EJT entry.

Execution

An input job is in execution after it is selected by the operating system and assigned to a control point. A job remains in execution until terminated, but it can be temporarily swapped or rolled out by the operating system.

Express Deadstart Dump (EDD) File

A file that is generated on magnetic tape by the express deadstart dump utility. This file contains a dump of memory, executing exchange packages, hardware registers, and controller memory.

Express Deadstart Dump (EDD) Utility

A utility that may be run at deadstart time after a system malfunction has occurred. It generates the express deadstart dump file on magnetic tape.

Extended Core Storage (ECS)

A type of extended memory that is an option available for 6000 Computer Systems, CYBER 70 Computer Systems, CYBER 170 Computer Systems (except model 176), and CYBER 180 Computer Systems. The maximum size of ECS is two million words. See Extended Memory.

Extended Memory (EM)

An additional portion of memory that is available as an option. This memory can be used for program and data storage, but not for program execution. Special hardware instructions exist for transferring data between central memory and extended memory. Extended memory consists of either extended core storage (ECS), extended semiconductor memory (ESM), large central memory extended (LCME), or unified extended memory (UEM).

Extended Memory File Space

The portion of extended memory that is defined as a mass storage device.

Extended Memory I/O Buffers

The portion of extended memory that is used for buffering I/O data to and from disks. Disks requiring I/O buffers are: 819, 885-42, 887, 895, and 9853.

Extended Semiconductor Memory (ESM)

A type of extended memory that is an option available for 6000 Computer Systems, CYBER 70 Computer Systems, CYBER 170 Computer Systems (except model 176), and CYBER 180 Computer Systems. The maximum size of ESM is 16 million words. See Extended Memory.

F

Facility Interface Program (FIP)

A program consisting of routines and buffers that are loaded into each application program's field length. This program is the interface between the application program and RHF.

Family Device

A mass storage permanent file device associated with a specific system. A family may consist of from 1 to 63 logical devices. Normally, a system runs with one family of permanent file devices available. However, additional families may be introduced during normal operation. This enables users associated with the additional families to access their permanent files via the alternate family.

Family Name

Name of the permanent file storage device or set of devices on which all of a user's permanent files are stored. When a user requests a permanent file, the system looks for it on this family (group) of devices. Usually a system has only one family of permanent file devices, but it is possible to have alternate families in the system. At login, the user may have to specify which family he/she is using. A user gets a family name from his/her employer, instructor, or computer center personnel.

Family Ordinal Table (FOT)

A table that maintains the relationship between family ordinals and family names.

FCT

See File and Cartridge Table.

FET

See File Environment Table.

Field Length

The area in central memory allocated to a particular job; the only part of central memory that a job can directly access. Also the number of central memory words required to process a job.

File

A collection of information referred to by a file name (from 1 to 7 alphanumeric characters). You can create a file at the terminal or retrieve a file from permanent file storage for use during a terminal session.

File Access Category

A property of a permanent file used by the creator of the file on a secured system to restrict access of the file to a particular group of users. A secured system supports up to 32 access categories, and a user is authorized to use some, all, or none of those categories. Refer also to System Access Categories.

File Access Level

A property of each file on a secured system used to indicate the sensitivity of information contained on the file. A file is assigned the current job access level by default when it is

created or stored; the file creator may specify any access level for that file that is within the set of access levels valid for the job, the system, the file creator, and (for interactive jobs) the communication line to the host mainframe. If a user accesses a file on a secured system, that user must be validated for the access level of the file. Refer also to Access Level, Job Access Level, and Job Access Level Limits.

File and Cartridge Table (FCT)

Table that has an entry for each cubicle assigned to the subfamily from a given SM.

File Category

Each permanent file is assigned a category of private, semiprivate, or public.

File Count

A maximum number of permanent files allowed each user.

File Environment Table (FET)

A table within a program's field length through which the program communicates with operating system input/output routines. One FET exists for each file in use by the program.

File Name Table (FNT)

A system-managed table that contains the local file name, the file type, and other job control information. All active files in the system have an FNT entry.

File Status Table (FST)

A system-managed table that contains information pertaining to the file's location in mass storage and other job control information. Each active file in the system has an FST entry. See also File Name Table.

FIP

See Facility Interface Program.

First Level Peripheral Processor (FLPP)

The processor that is connected directly to the CYBER 170 Model 176 mainframe and operates synchronously with the mainframe.

FLPP

See First Level Peripheral Processor.

FNT

See File Name Table.

Forms Code

An attribute of output files and output devices. The user can specify special forms required for output, and then mount the special forms and use the FORM command to let the system process his/her output.

FOT

FOT

See Family Ordinal Table.

FST

See File Status Table.

Function Processor

A system CPU or PP program that the user can call by placing a request in location RA+1. Function processors perform input/output, local and permanent file manipulations, and so on.

H

Hardware Initialization and Verification Software (HIVS)

The software package that assists CTI during deadstart and provides deadstart confidence-level testing (HVS).

Hardware Verification Sequence (HVS)

HVS is a member of HIVS. It tests the ability of memory to hold patterns of data and execute instructions. The user can choose to test central memory, extended memory, PP memory, and central processor memory.

Head-of-Chain Flag (HOC)

In an MSE environment, HOC is a flag in the 7990 catalog that identifies the first volume in a chain of allocated AUs.

Header

A word or set of words at the beginning of a block, record, file, or buffer that contains control information for that unit of data.

HIVS

See Hardware Initialization and Verification Software.

HOC

See Head-of-Chain Flag.

HOP

See Host Operator.

Host

The computer that controls a network, executes the application programs, and processes network messages.

Host Operator (HOP)

The administrative operator who resides at the system console, initiates NAM, and controls NPUs and network related host elements.

HVS

See Hardware Verification Sequence.

I**IAF**

See Interactive Facility.

Inactive Queued File Table (IQFT)

A table of file entries that has been removed from the queued file table. An IQFT file is on each mass storage device on which one or more inactive queued files reside.

Incremental Dump

An incremental dump copies those permanent files modified after a specified date. Each incremental dump writes a catalog image record at the beginning of the archive file on which the permanent files are dumped.

Incremental Load

An incremental load builds up an accumulation of the most recently modified versions of the files extracted from the archive files for loading. A series of archive files is read in reverse order of creation. The CIR is read and checked against the archive files. If a file matches an entry on the CIR, that file is a candidate for loading.

Indirect Access File

A NOS permanent file that you access by making a temporary copy of the file (GET or OLD command). You create or alter it by saving or substituting the contents of an existing temporary file (REPLACE or SAVE command).

Input File

The system-defined file that contains the entire job the user submits for processing. It is also known as the job file.

Input/Output Unit (IOU)

A collection of all PPs, PP channels, and related hardware for models 865 and 875 and CYBER 180-class machines.

Interactive Facility (IAF)

An application that provides a terminal operator with interactive processing capability. The Interactive Facility makes terminal input/output and file input/output appear the same to an executing program.

Interactive Transfer Facility (ITF)

A network application that allows the user to connect an interactive terminal to a remote CYBER 200 computer system linked to a host mainframe by a loosely coupled network.

IOU

See Input/Output Unit.

IQFT

IQFT

See Inactive Queued File Table.

ITF

See Interactive Transfer Facility.

J

Job Access Level

On a secured system, each job has an access level. This is the default access level that is assigned to files that are created or stored in the job. A job's initial access level is the lower access level limit for the job. The job's access level is automatically raised to the access level of any file from which information is read. The user can also change the job access level. Refer also to Job Access Level Limits.

Job Access Level Limits

An upper limit and a lower limit that determine the range of access levels that are valid for a particular job on a secured system. All files used in a given job must have an access level within the job's access level limits.

Job Sequence Name (JSN)

The unique, system-defined name assigned to every executing job or queued file. The JSN is a string of four alphabetic characters.

Job Status

A job attribute kept in the job's executing job table (EJT) entry. It is used by the system to determine if a job is rolled in or rolled out. If the job is rolled out, job status indicates why it was rolled out.

JSN

See Job Sequence Name.

L

LAN

See Local Area Network.

Large Central Memory Extended (LCME)

A type of extended memory that is an option available for model 176. Refer to Extended Memory.

LCF

See Local Configuration File.

LCME

See Large Central Memory Extended.

LCN

See Loosely Coupled Network.

LDLIST Utility

A utility that generates a printer listing of queued files present on a dump tape produced by the QDUMP utility.

LFG

See Load File Generator.

LID

See Logical Identifier.

LISTPPM Utility

A PIP dump analyzer program that converts all available PIP dump binary records on the PIP memory dump file into a report to be listed in byte format.

Load File Generator (LFG)

A utility program that reformats communications control program files for subsequent use by the network supervisor of NAM to load network processing units.

Load Point

Metallic strip marking the beginning of the recordable portion of a magnetic tape. Data, including labels, is written after the load point. A rewind positions a single file volume to the load point.

Load Sequence

A sequence of load operations that encompasses all of the loader's processing from the time that nothing is loaded until the time execution begins. It includes initialization, specification of specified loader requests, and completion of load.

Local Area Network (LAN)

A privately owned network that interconnects data processing equipment to provide high-speed communications. It allows users and services to exchange messages and share resources.

Local Batch Job

A batch job submitted at the central computer site through a card reader or terminal.

Local Configuration File (LCF)

A file in the host computer system containing information on the logical makeup of the communication elements of the host. The file contains a list of the application programs available for execution in the host computer, and the users that can access it. This is a NOS direct access permanent file.

Local File

Local File

Any file that is currently associated with a job. Local files include all temporary files and attached direct access files.

Local File Name

The file name assigned to a file while it is local (assigned) to a job. The name is contained in the local file name table.

Local NAD

A 380-170 NAD connected to the host mainframe using a channel and configured in the EST.

Local NPU

An NPU that is connected to the host via a coupler. A local NPU always contains a host interface program for processing block protocol transfers across the host/local NPU interface.

Logical Identifier (LID)

A 3-character alphanumeric string used to identify a particular mainframe. LIDs are identified by the user's site.

Loosely Coupled Network (LCN)

A network of physically connected computer systems. The LCN environment allows jobs, data files, and messages to be transmitted from one computer system to another.

M

Machine Identification (MID)

The identifier that associates a specific machine with its access to a shared device.

Machine Recovery Table (MRT)

A table that provides the information needed to recover the mass storage space and interlocks of a machine that shares a mass storage device.

Machine Recovery Utility (MREC)

A utility that clears interlocks held by the machine to be deadstarted that have not been cleared by CPUMTR. It also recovers mass storage space on a shared device that is currently not accessible because of a machine interruption.

Macro

A sequence of source statements that is saved and then assembled whenever needed through a macro call.

Mainframe Device Interface (MDI)

A device interface that is configured to connect a CYBER mainframe to Ethernet.

Mainframe to Terminal Interface (MTI)

A terminal interface that is configured to connect a CYBER mainframe to a terminal for support of terminal-to-network communications.

MAINLOG Utility

A dayfile dumping utility that dumps all or selected parts of the binary maintenance log to produce an output file in binary format.

Maintenance Logging Transfer Utility (MLTF)

A utility that controls logging NAD error logs into the binary maintenance log.

Maintenance Register

A hardware register used in error detection, logging, and recovery procedures. Maintenance registers are used on models 865 and 875 instead of status/control registers. Refer to Status/Control (S/C) Register.

MAP

The Matrix Algorithm Processor.

Mass Storage Device

An extended memory or disk unit that has defined logical attributes such as family, file residency, and so on.

Mass Storage Extended Subsystem (MSE)

MSE is the product consisting of the 7990 hardware, the channel interface, the diagnostics, and the operational software. MSE stores data on the 7990 and moves it to disk upon request for access by an authorized user.

Mass Storage Table (MST)

A table that contains an entry for each logical device in the configuration of mass storage devices currently available to the system.

Master Device

A disk device that contains the user's permanent file catalog entries; all of the user's indirect access files; and all, part, or none of the user's direct access files.

MCT

See Memory Control Table.

MDI

See Mainframe Device Interface.

MID

See Machine Identification.

Memory Control Table

A central memory table used in allocating central memory and extended memory to user jobs.

MLIA

MLIA

See Multiplex Loop Interface Adapter.

Monitor

The system routine that coordinates and controls all activities of the computer system. It occupies peripheral processor 0 and part of central memory. It schedules the use of the central processor and the other peripheral processors.

MREC

See Machine Recovery Utility.

MRT

See Machine Recovery Table.

MSE

See Mass Storage Extended Subsystem.

MST

See Mass Storage Table.

MTI

See Mainframe to Terminal Interface.

Multimainframe Operation

An operation that provides mechanisms by which more than one computer can share mass storage devices.

Multiplex Loop Interface Adapter (MLIA)

The hardware portion of the multiplex subsystem that controls the multiplex loops (input and output) as well as the interface between the NPU and the multiplex subsystem.

Multispindle Device

A logical mass storage device that includes from two to eight disk units.

Multiterminal Job

A job that does one specific task for many terminals while being scheduled into the system only once.

N

NAD

See Network Access Device.

NAM

See Network Access Method.

NCF

See Network Configuration File.

NCTF

See Network Description File.

NDA

See NPU Dump Analyzer.

NDI

See Network Device Interface.

NDL Processor

See Network Definition Language Processor.

NDR

See Network Driver.

Negative Field Length (NFL)

Central memory assigned to a control point that physically precedes the job's reference address (RA).

NETLOG

A program that uses the network configuration file to determine which remote NADs should be logged.

NETOU

See Network Operator Utility.

Network

An interconnected set of network elements consisting of a host and one or more NPUs and terminals.

Network Access Device (NAD)

The primary element in a loosely coupled network. Each NAD connects a computer system to the network.

Network Access Method (NAM)

A software package that provides a generalized method of using a communications network for switching, buffering, queuing, and transmitting data. NAM is a set of interface routines used by a terminal servicing facility for shared access to a network of terminals and other applications, so that the facility program does not need to support the physical structures and protocols of a private communication network.

Network Configuration File (NCF)

A network definition file in the host computer containing information on the network elements and permissible linkages between them. The status of the elements described in

this file is modified by the NPU operator in the course of managing the network. This is a NOS direct access permanent file.

Network Definition Language (NDL) Processor

The network software module that processes an NDL program as an offline batch job to create the network definition files and other NDL program output.

Network Description File (NCTF)

A file that must be present if the Transaction Facility is used. The file is prepared by the site analyst.

Network Device Interface (NDI)

A device interface that is configured to transfer data between networks; for example, LAN to LAN, LAN to PDN, PDN to PDN, or LAN to communication lines.

Network Driver (NDR)

A program that executes in a dedicated peripheral processor unit. It communicates with the network access devices using a host computer data channel, and is the interface between RHF and the communication network.

Network Invocation Number (NIN)

A 1- to 3-digit decimal number. NIN is incremented by 1 every time NAM is brought up.

Network Load File (NLF)

An output file generated by the load file generator utility for use by the network supervisor.

Network Operator Utility (NETOU)

A group of programs residing in a host computer and in a mainframe device interface connected to the mainframe that allow a network operator to access, monitor, control, and configure a CDCNET network from the host console or a remote terminal. NETOU allows commands from network operators to be sent through the CDCNET network to specific device interfaces or all of the DIs in the network.

Network Processing Unit (NPU)

The collection of hardware and software that switches, buffers, and transmits data between terminals and host computers.

Network Supervisor (NS)

A portion of the network software written as a NAM application program. NS dumps and loads NPUs upon request.

Network Terminal

A terminal that communicates with the operating system through the network.

Network Validation Facility (NVF)

A portion of the network software, written as a NAM application program. The network validation facility performs application validation and all connection validation processing and supports login dialog with the terminal user.

NFL

See Negative Field Length.

NIN

See Network Invocation Number.

NLF

See Network Load File.

Nonincremental Load

A nonincremental load does no CIR checking and uses only parameter options specified on the PFLOAD call, if any, to select candidates for loading.

NOP

See NPU Operator.

NPU

See Network Processing Unit.

NPU Dump Analyzer (NDA)

A utility program that produces a readable printout from the NPU dump files.

NPU Operator (NOP)

The administrative operator who resides at a terminal and controls NPUs.

NS

See Network Supervisor.

NVF

See Network Validation Facility.

O**Object Code**

The machine language version of a program that has been translated (compiled) from source code written in a higher-level language.

Operating System

The set of system programs that controls the execution of computer programs and provides scheduling, error detection, input/output control, accounting, compilation, storage assignment, and other related services.

Optical Disk Alternate Storage

Optical disk used as an alternate storage medium for permanent files. The optical disk units may be housed in a jukebox. See Alternate Storage.

Origin Type

Origin Type

A job attribute that indicates how a job entered the system. The four origin types are interactive origin, batch origin, remote batch origin, and system origin.

Output File

The system-defined file that contains the output from job processing. It is also known as the print or punch file.

P

PACKER Utility

A utility that provides the capability to manage the holes within the indirect access permanent file chain on a permanent file device.

Paging (Screen)

The process of filling a CRT display with data and holding additional data for subsequent displays. Changing the page display is an operator-controlled function if the page-wait option is selected.

Parity

In writing data, an extra bit is either set or cleared in each byte so that every byte has either an odd number of set bits (odd parity) or an even number of set bits (even parity). Parity is checked on a read for error detection and possible recovery.

Partial Dump

A partial dump copies permanent files according to any specified options, except those defining a full or incremental dump.

Password

A name or word the user enters during login to provide extra security for his/her user name. A unique password ensures that no one else can log into the system with someone else's user name and access that user's files. A user's password is given to him/her by that user's employer, instructor, or computer center personnel.

PCP

See Pseudo-control Point.

PDN

See Public Data Network.

Peripheral Interface Package (PIP)

The interface package between the PPU of the CYBER computer and the network application.

Peripheral Processor (PP)

The hardware unit within the host computer that performs physical input and output through the computer's data channels.

Peripheral Processor Unit (PPU)

First level peripheral processor. A PPU is contained in the mainframe in a multimainframe environment and operates synchronously with the mainframe. Sometimes referred to as FLPP.

Permanent File

A mass storage file that is cataloged by the system so that its location and identification are always known to the system. Permanent files cannot be destroyed accidentally during normal system operation. They are protected by the system from unauthorized access according to privacy controls specified when they are created.

Permanent File Catalog Entry (PFC)

A 16-word entry that the system maintains and uses to determine the file name, owner, identification, disk pointers, alternate storage pointers, and other attributes of a permanent file.

Permanent File Family

The permanent files that reside on the family devices of a specific system.

Permanent File Manager (PFM)

PFM identifies the master device and catalog track information when a user submits a job.

Permanent File Supervisor (PFS)

The PFS processes parameters in utility commands and loads the correct processing overlays.

Permanent File Transfer Facility (PTF)

PTF is an application program initiator started by the user using an MFLINK command. It is responsible for initiating and completing (with the help of its servicing application, PTFS on another host) a permanent file transfer.

Permanent File Transfer Facility Servicer (PTFS)

PTFS is an application program servicer started by RHF or NAM when requested by a PTF on another host. The PTFS application assists the PTF application in completing the file transfer by performing those permanent file functions requested by the user and then transferring the file between PTF and PTFS.

Permanent File Utility (PFU)

A utility that manages the catalogs, permits, data allocation on a device, and the data transfer between the device and the overlay.

Permission Mode

A mode of operation that a user is allowed for a particular permanent file, such as write, modify, append, read, and so forth.

PFATC Utility

A utility that produces a cataloged directory of file information derived from an archive file previously created by the PFDUMP utility.

PFC

PFC

See Permanent File Catalog.

PFCAT Utility

A utility that produces a cataloged directory of file information derived from catalog tracks on a master device.

PFCOPY Utility

A utility that extracts files from an archive file and copies them to one or more files at a control point.

PFDUMP Utility

A utility that dumps permanent files to an archive file. Dumps can be reloaded by the PFLOAD utility and can be accessed by the PFATC and PFCOPY utilities for cataloging and copying.

PFLOAD Utility

A utility that loads archived files produced by the PFDUMP utility back into the permanent file system. The load can reestablish the permanent file system exactly as it was at the time of the dump, or can load only a desired subset of files on the archive file.

PFM

See Permanent File Manager.

PFREL Utility

A utility that releases disk space for permanent files that have copies on alternate storage.

PFS

See Permanent File Supervisor.

PFU

See Permanent File Utility.

Physical Identifier (PID)

The unique 3-character identifier of a specific host.

Physical Record Unit (PRU)

The amount of information transmitted by a single physical operation of a specified device. For mass storage files, a PRU is 64 central memory words (640 characters); for magnetic tape files, the size of the PRU depends upon the tape format. A PRU that is not full of user data is called a short PRU; a PRU that has a level terminator but no user data is called a zero-length PRU.

PICB

See Program Initiation Control Block.

PID

See Physical Identifier.

PIP

See Peripheral Interface Package.

PP

See Peripheral Processor.

PPS

The Peripheral Processor Subsystem.

PPU

See Peripheral Processor Unit.

Preserved File

A mass storage file that is recovered on all levels of system deadstart. Preserved files include permanent files, queued files, and system dayfiles.

Primary File

A temporary file created with the OLD, NEW, LIB, (interactive jobs only), or PRIMARY command. The primary file is assumed to be the file on which most system operations are performed unless another file is specified. There can be only one primary file associated with your job.

Primary VSN

The volume serial number (VSN) of a single volume destage dump tape or the VSN of the first of volume a multivolume set of destage dump tapes. The VSN must be two alphanumeric characters followed by a four-digit decimal number; the number must be in the range 0000 through 4095.

Privileged Analyst

A user with AW=CPAM validation (refer to the NOS Version 2 Administration Handbook for information on MODVAL validation). Such a user can read system status information (such as the system dayfile, account file, and error log) using a nonsystem-origin job if PRIVILEGED ANALYST MODE is enabled (via IPRDECK or DSD command). Note that PRIVILEGED ANALYST MODE cannot be enabled on a secured system.

PROBE Utility

A utility that traps and measures particular interval events in the system. PROBE generates a report from the data collected by the system.

Procedure

A user-defined set of instructions that can be referenced by name. The instructions consist of procedure directives and system commands.

Program Initiation Control Block (PICB)

A sequence of commands that initiates NPU load and dump operations for a specific NPU variant. Several PICBs may exist on the network load file. Each PICB is a separate record with a unique NPU variant name as its record name.

Programmable Format Control

Spacing and format control for 580 line printers provided by the use of software and a microprocessor instead of a carriage control format tape.

Project Number

An alphanumeric identifier that may be required at a user's installation for accounting and billing to a specific project. If it is required, the project number is entered during the login procedure. It is given to the user by personnel at his/her installation.

Protocol

A set of conventions or rules that must be used to achieve complete communication between entities in a network. A protocol can be a set of predefined coding sequences such as the control byte envelopes added to or removed from data exchanged with a terminal; a set of data addressing and division methods, such as the block mechanism used between an application program and NAM; or a set of procedures used to control communication, such as the supervisory message sequences used between an application program and NAM.

PRU

See Physical Record Unit.

Pseudo A Register

A software register used by DSD to function channels and to manipulate peripheral hardware devices from the operator's console.

Pseudo-control Point (PCP)

The portion of central memory that is assigned to a job that has been pseudo-rolled. Pseudo-control points are the same as control points except that a job at a pseudo-control point cannot have any activity (such as PP assignment or CPU assignment).

Pseudo-control Point Number

The number of the pseudo-control point to which a job is assigned while the job resides in central memory. The actual number of pseudo-control points is an installation parameter.

Pseudo-rollout

The removal of jobs from a control point to a pseudo-control point so that the control point and central memory can be assigned to another job. A job is rolled out to a pseudo-control point when its control point is needed by a higher priority job.

Pseudo-rollin

The return of jobs from a pseudo-control point to a control point.

PTF

See Permanent File Transfer Facility.

PTFS

See Permanent File Transfer Facility Servicer.

Public Auxiliary Device

An auxiliary device that is available for access by all validated users knowing the correct pack name. Additional validation is required to create or replace files on an auxiliary device.

Public Data Network (PDN)

A commercial packet-switching network that supports the interface described in the CCITT protocol X.25.

Q**QALTER Utility**

A utility that displays, lists, and/or alters routing and other information about active queued files. It selects files for processing according to specified criteria. QALTER can also purge selected files from the system.

QDUMP Utility

A utility that dumps selected queued files from a single device, a family of devices, or all devices on the system. These queued files can be dumped either to a tape or disk. QDUMP also provides a listing of all files dumped with information about each file processed.

QFSP

See Queue File Supervisor Program.

QFT

See Queued File Table.

QFTLIST Utility

A utility that displays and/or lists routing and other information about active queued files. Its operation is similar to that of QALTER, except file alteration or purging is not allowed.

QLIST Utility

A utility that lists inactive queued files, which may include all inactive queued files in the system or a selected subset based on options specified when the utility is called.

QLOAD Utility

A utility that processes the dump files generated by QDUMP or other utilities using the same format. QLOAD can selectively load the queued files from these dump files. QLOAD can also list the contents of a dump file without loading any files.

QMOVE Utility

A utility that moves queued files from one disk device to another. It also produces a listing of all files moved with information about each file processed.

QREC Utility

A utility that deactivates or activates selected queued files and purges selected inactive queued files.

QTF

See Queue File Transfer Facility.

QTFS

See Queue File Transfer Facility Servicer.

Queue File Supervisor Program (QFSP)

A program that provides control for the queue file utilities.

Queue File Transfer Facility (QTF)

QTF is an application program initiator that periodically scans the I/O queues searching for files to transfer. When it finds a file to transfer, it initiates and completes the queue file transfer with the help of its servicing application, QTFS, on another host.

Queue File Transfer Facility Servicer (QTFS)

QTFS is an application program servicer started by RHF or NAM when requested by a QTF on another host. The QTFS application assists the QTF application in completing the transfer by receiving the queue file and placing it in the I/O queue.

Queue Priority

An attribute associated with input and output files. If all other factors are equal, queue priority is used to select the best file for processing.

Queued File

An input, print, plot, or punch file that has an entry in the QFT, is not assigned to an EJT entry, and is waiting to be selected for processing.

Queued File Table (QFT)

A central memory resident table that contains a 4-word entry for all active input and output queue files.

R

Random Access

An access method by which any record in a file can be accessed at any time. Random access applies only to mass storage files with an organization other than sequential. Refer to Sequential Access.

RCFGEN

See RHF Configuration File Generation.

RCL

See Resident Central Library.

RDF

See Release Data File.

Recall

The state of a program when it has released control of the central processor until a fixed time has elapsed (periodic recall) or until a requested function is completed (auto recall). Recall is a system action request as well as an optional parameter of some file action requests.

Record

A unit of information. In CYBER Record Manager and its language processors, a record is a unit of information produced by a single read or write request.

Eight different record types exist within CRM. The user defines the structure and characteristics of records within a file by declaring a record format.

Regulation Level

A number that indicates to NAM the existence of a logical link and indicates what types of information exchange are possible on that logical link.

Release Data File (RDF)

A file created by PFDUMP that identifies those 7990-resident files that are pointed to by PFC entries at the time of the dump.

Remote Batch Job

A job submitted from a remote batch terminal.

Remote Host Facility (RHF)

A central processor program that executes at a system control point. It performs data buffering and switching, and is the intermediary between application programs and the network.

Remote NAD

Any 380 NAD accessible to a local NAD using a loosely coupled network trunk.

Remote NPU

A network processing unit linked to a host computer through other network processing units.

Removable Device

A disk storage device that can be physically detached from the disk drive.

Resident Central Library (RCL)

An area in central memory resident that central library routines specified by the *CM directive reside.

Resident Peripheral Library (RPL)

An area in central memory resident that peripheral library routines specified by the *CM directive reside.

RHF

See Remote Host Facility.

RHF Configuration File Generation (RCFGEN)

A utility that reads configuration definition statements to create a permanent file that RHF uses for the network description and access.

Rollout

The removal of jobs from central memory to mass storage before execution is complete, so the control point and central memory can be assigned to another job. A job is rolled out when it is waiting for an external event, when its control point and/or central memory is needed by a higher priority job, or when it exceeds its central memory time slice.

Rollout File

A file containing a job (and system information) that has been temporarily removed from the main processing area of the system.

RPL

See Resident Peripheral Library.

S

Sabre

The 5833 Sabre is a disk unit that can be configured only as part of the 5830 Disk Array Subsystem (DAS). The Sabre disk has a formatted capacity of 1039 megabytes. It can be configured in four ways:

- As a single unit
- As a single unit with another unit used for parity
- As a pair of units driven in parallel
- As a pair of units driven in parallel with another unit used for parity

Parallel units increase the device's size and speed.

SC

See Service Class.

S/C Register

See Status/Control Register.

Scheduling Priority

An attribute associated with an executing job available for job scheduling. Scheduling priority is used to select the best executing service class job for processing.

SCOPE 2 Station Facility (SSF)

A NOS subsystem that allows a NOS user to submit jobs (including batch jobs that require interactive I/O) to a linked SCOPE 2 system. The submitted job uses standard SCOPE 2 commands to access NOS files stored on the originating NOS system.

SCP

See Subcontrol Point.

Screen Management Facility (SMF)

A subsystem that alters the performance characteristics of the Full Screen Editor (FSE). The absence or presence of SMF is not detectable by the FSE user. Performance can be optimized by disabling SMF for small mainframes and interactive work loads, and by enabling SMF for large configurations and heavy work loads.

SCRSIM

See Status/Control Register Simulator.

SDF

See System Deadstart File.

SECDED

See Single Error Correction Double Error Detection.

Secondary Mask

An 8-bit quantity used to identify groups of users who can place direct access files on a particular device.

Secondary VSN

The VSN of a single volume destage dump tape or the VSN of the first volume of a multivolume set of destage dump tapes having a sequence number (the last four characters of the VSN) in the range 5000 through 9095. The existence of a secondary VSN assumes the existence of a primary VSN having the same two-character prefix and a sequence number of 5000 less. The primary and secondary VSN tapes or multivolume sets of tapes are assumed to contain identical data, although the individual volumes of multivolume sets cannot be assumed to be identical.

Secured System

A system in which a mandatory security mechanism has been enabled during deadstart. A secured system protects information by enforcing restrictions based on access levels and access categories, and restricts many sensitive system functions to security administrators.

Security Administrator

A secured system prevents users and operators from performing certain functions that could result in the unauthorized disclosure or modification of information. These functions can only be performed by someone who is designated a security administrator. A security administrator is always authorized to access the highest level of information stored on the

system. This person performs functions in the areas of installation, user validation, system operation, and system maintenance.

Security Unlock Status

This status of the system console applies only to a secured system and must be set by a security administrator. The console must be in security unlock status in order for the security administrator to perform certain functions that are restricted on a secured system.

Sequential Access

A method in which only the record located at the current file position can be accessed. Refer to Random Access.

Sequential (SQ) File

A file in which records are accessed in the order in which they occur. Any file can be accessed sequentially.

Service Class (SC)

An attribute associated with a queued file or executing job. The service class determines how the system services the job.

SFS

See Special File Supervisor.

Single Error Correction Double Error Detection (SECDED)

A hardware technique that detects and corrects single bit errors in memory. Double bit errors are detected by not corrected.

SM

See Storage Module.

SM Map

See Storage Module Map.

SMF

See Screen Management Facility.

Solid State Disk (SSD)

The 5832 Solid State Disk (SSD) is a disk unit that can be configured only as part of the 5830 Disk Array Subsystem (DAS). The SSD has a formatted capacity of 165 megabytes. It can be configured in DAS as single units or as pairs of units driven in parallel. Parallel units increase the device's size and speed.

Source Code

Code input to the computer for later translation into executable machine language instructions (object code).

Special File Supervisor (SFS)

A program that provides routines, table management, data manipulation, and I/O processing for special system jobs.

SQ File

See Sequential File.

SRU

See System Resource Unit.

SSALTER Utility

A utility that displays the current 7990 hardware configuration and allows an analyst to change the status of certain elements in the 7990 configuration.

SSBLD Utility

A utility that processes statements consisting of 7990 component-oriented mnemonics that define the logical mapping within the 7990 configuration. SSBLD generates a direct-access permanent file which SSEXEC uses as its unit device table.

SSDEBUG Utility

A utility that allows an analyst to update appropriate entries in the SM maps and/or 7990 catalogs and thereby resolve inconsistencies reported by the SSVAl utility. SSDEBUG can also be used to copy data from selected 7990 files or cartridges to disk.

SSD

See Solid State Disk.

SSDEF Utility

A utility that creates the system files (SM maps and 7990 catalogs) that are necessary for MSE processing.

SSEXEC Program

The main processing program that controls MSE activities, such as destaging files from disk to the 7990, purging unneeded 7990 files, labeling or relabeling cartridges, updating SM maps, and updating 7990 catalogs.

SSF

See SCOPE 2 Station Facility.

SSLABEL Utility

A utility that manages cartridge assignment and cubicle allocation in a storage module.

SSMOVE Utility

A utility that manages disk and 7990 residence. SSMOVE determines which files to leave on disk, which files to release from disk and move to 7990, and which files should reside both on disk and on 7990.

SSSLV Program

A program that runs on each slave mainframe and communicates with the SSEXEC program to retrieve files from the 7990 in response to ATTACH requests by jobs running on the slave mainframes.

SSUSE Utility

A utility that reads data in the 7990 catalogs and SM maps and produces reports on the availability of space on 7990 cartridges and the allocation of cubicle space within an SM.

SSVAL Utility

A utility that either performs release processing or reports on problems with the current MSE system files.

Staging

The process of restoring file data to disk residence from a copy on alternate storage. Staging is initiated when a user executes an APPEND, ATTACH, GET, or OLD command for a file that is not currently disk resident.

Status

Information relating to the current state of a device, line, and so forth. Service messages are the principal carriers of status information. Statistics are a special subclass of status.

Status/Control Register Simulator (SCRSIM)

A program that enables the user to set status/control register bits in order to aid in the testing of error logging and error recovery procedures.

Status/Control (S/C) Register

A hardware register used in error detection, logging, and recovery procedures. This register is present on all CYBER 170 Computer Systems. For models 865 and 875, the S/C register is replaced by a maintenance register. Refer to Maintenance Register.

Step Mode

A protected or debugging mode for the operating system monitor. The keyboard spacebar must be pressed to process each PP request.

Stimulator

A collection of central memory and peripheral processor programs that enters a hypothetical work load into the system to analyze the effects of such a load on response time and system reliability.

Storage Module (SM)

An MSE hardware unit that houses up to 312 usable data cartridges, a cartridge accessor unit that picks cartridges from and puts cartridges in their cubicles, and one or two data recording drives.

Storage Module (SM) Map

A direct access permanent file that contains information indicating the cartridges that reside in the SM.

Subcontrol Point (SCP)

A division of a central memory control point. You can set up a control point to contain two or more programs; one of the programs is the executive, and monitors the other programs executing at the subcontrol points.

Subfamily

Each permanent file family consists of eight subfamilies, subfamily 0 through subfamily 7. The lower 3 bits of the user index identify the subfamily to which a user belongs.

Suspended Job

An interactive job placed in a inactive state. Processing is stopped immediately and recovery information is copied to the rollout file. Processing is resumed as if no interruption took place, if the job's EJT entry is recovered.

System Access Categories

On a secured system, a set of access categories are set during level 0 deadstart. This set may consist of some, all, or none of the 32 possible access categories. While the system is running in security mode, you may only use access categories that are within the set of system access categories.

System Access Levels

On a secured system, a range of access levels is set during level 0 deadstart. This range may contain some or all of the eight possible access levels. While the system is running, users may only use access levels that are within the range of system access levels.

System Deadstart File (SDF)

A file that is a copy of the deadstart tape that resides on a disk storage deadstart device. When the system is deadstarted from disk, this file is read to generate copies of the running system.

System Library (SYSTEM)

The collection of tables and object language programs that reside in central memory or on mass storage and are necessary for running the operating system and its product set.

System Origin Job

A job entered at the system console.

System Resource Unit (SRU)

A unit of measurement of system usage. The number of SRUs includes the central processor time, memory usage, and input/ output resources used for a given job.

T**TAF**

See Transaction Facility.

Tape Alternate Storage

Magnetic tape used as an alternate storage medium for permanent files. See Alternate Storage.

TCU

See Trunk Control Unit.

TDI

See Terminal Device Interface.

Temporary File

A file associated with a job that is not a permanent file. Temporary files no longer exist when the user logs off the system or releases the files.

Terminal Device Interface (TDI)

A device interface that is configured to support terminal-to-network communications.

Timed/Event Rollout

A condition in which an executing job has been temporarily removed from central memory but will be rolled back into central memory when a specified event (such as a file is no longer busy) or a specified time period has elapsed.

TRACER Utility

A utility that monitors the system's activity and gathers data periodically for statistical analysis of the system.

Track Link

An address of the next track that is a logical continuation of a file.

Track Reservation Table (TRT)

A table that describes the physical layout of data on a device and is the key to allocating information on the device.

Transaction Facility (TAF)

An application program that provides the transaction terminal with access to a data base. A terminal using TAF can enter, retrieve, and modify information in the data base.

TRT

See Track Reservation Table.

Trunk

The communication line connecting two network processing units.

Trunk Control Unit (TCU)

The hardware part of a network access device (NAD) that interfaces with a network trunk.

U

UDT

See Unit Device Table.

UEM

See Unified Extended Memory.

Unified Extended Memory (UEM)

A type of extended memory that is available as an option for CYBER 180-class machines and models 865 and 875. UEM differs from other types of extended memory in that it is a portion of central memory and not a separate memory unit. See Extended Memory.

Unit Device Table (UDT)

A table that defines the logical mapping of the components within a 7990 configuration.

Unit Number

The setting of a hardware device. The unit number is used when more than one hardware unit can be connected to a controller.

Unsecured System

A system in which the multilevel security mechanism has not been enabled during deadstart. The restrictions based on access levels and access categories are not enforced on an unsecured system.

Upline

The direction of input flow from terminal to host.

User Break 1 Sequence

The character or sequence of characters that causes an executing program to be interrupted (also called the interruption sequence).

User Break 2 Sequence

The character or sequence of characters that causes an executing program to be terminated (also called the termination sequence).

User Index

A unique 17-bit identifier that is associated with each user name. The user index is used by the permanent file manager to identify the device and catalog track for the user's permanent files.

User Job Name (UJN)

A 1- to 7-character alphanumeric name you specify to replace the system defined JSN for a queued file or executing job.

User Name

A name given to the user by his/her employer, instructor, or computer center personnel. A user name has certain resources and privileges assigned to it. When logging in to the system, a user specifies his/her user name as identification, so that it knows that this person is an authorized user and what resources this user is entitled to use. A user name also represents a specific catalog in the permanent file system. All files a user makes permanent are associated with that user's name and this catalog.

V

Validation File

A file that contains validation information for all users (user names, passwords, resources allowed, and so on).

Volume Serial Number (VSN)

A from 1- to 6-character identifier that identifies the volume of magnetic tape to the system.

VSN

See Volume Serial Number.

W

Word

A group of bits (or 6-bit characters) between boundaries imposed by the computer system. A word is 60 bits in length. The bits are numbered 59 through 0 starting from the left. A word is also composed of five 12-bit bytes, numbered 0 through 4 from the left.

Write Mode

A mode that allows a user to write, modify, append, read, execute, or purge the file (modify permission applies only to direct access files).

Z

Zero-Length PRU

A PRU that contains system information but no user data. Under NOS, a zero-length PRU defines EOF.

7990

A hardware product that is a large capacity mass storage device. The 7990 is a cost-effective extension to the disk file storage system and an alternative to conventional magnetic tape storage.

7990 Catalog

A disk-resident direct access permanent file that contains information describing which AUs of each cartridge assigned to a particular subfamily are allocated to 7990 files and which AUs are available for allocation.

Scope 2 Station Facility

C

Equipment Configuration C-1

SSF File Transfers C-1

 Staged File Transfers C-2

 Spooled File Transfers C-3

Error Logout C-3

Scope 2 Station Facility

C

The SCOPE 2 Station Facility (SSF) is a NOS software subsystem that allows a NOS user to submit jobs and transfer files to a linked SCOPE 2 system residing in a 7000 Computer Systems mainframe. With SSF, a NOS user can route jobs to a linked SCOPE 2 system for processing. While processing the job, the SCOPE 2 system can access permanent files and tape files at the NOS host.

SSF provides these services:

- Provides NOS users with local batch capabilities at a linked SCOPE 2 system.
- Optionally serves as the system operator station for SCOPE 2.
- Provides the NOS user with remote batch capabilities, through RBF, at a linked SCOPE 2 system.
- Allows simultaneous transfer of multiple files.
- Optionally furnishes the deadstart file for SCOPE 2.
- Provides a means of obtaining absolute SCOPE 2 dumps.
- Allows a job executing on a SCOPE 2 system to use SCOPE 2 commands (GETPF, SAVEPF, and PURGE) to access permanent files at a linked NOS system.
- Allows a job executing on a SCOPE 2 system to access magnetic tape files at a linked NOS system.

Equipment Configuration

To establish a NOS-SCOPE 2 link using the SSF Subsystem, the NOS and SCOPE 2 mainframes must be attached to a loosely coupled network (LCN) by a 6683 and 7683 Satellite Coupler, respectively. An equipment status table (EST) device type entry of CC describes the coupler link. The EST table appears on the E,A display.

SSF File Transfers

SSF file transfers between NOS and SCOPE 2 are of two types:

- Staged file transfers.
- Spooled file transfers.

Staged file transfers process one file at a time and handle all tape file and permanent file transfers. Spooled file transfers handle multiple I/O file transfers. Normally, operator intervention is not required for either type of transfer.

SSF initiates a special type of job, called a spunoff task (SPOT) job, to perform file transfers. Staged file SPOT jobs differ from spooled file SPOT jobs as described in this appendix. The primary functions of any SPOT job are performing user validation and queuing tasks required to send a file across the link. All SPOT jobs are executed on the NOS system.

Installation parameters or operator commands can be used to limit the number of file transfer activities that can be active concurrently. Independent limits can be set on the following activities:

- Purging NOS files.
- Reading tape files from NOS to SCOPE 2.
- Writing tape files from SCOPE 2 to NOS.
- Reading permanent files from NOS to SCOPE 2.
- Writing permanent files from SCOPE 2 to NOS.
- Reading input files from NOS terminals to SCOPE 2.
- Writing output files from SCOPE 2 to NOS terminals.

The operator commands used to set limits are described in section 8, K-Display Utilities.

Staged File Transfers

Staged file transfers are essentially the same for both tape and permanent file transfers. In response to an SSF query, SCOPE 2 indicates to SSF that an executing SCOPE 2 job requires access to a NOS file. The SCOPE 2 request can specify either a NOS to SCOPE 2 file transfer, also called a prestaging transfer; or a SCOPE 2 to NOS (poststaging) transfer. In either case, SSF normally responds by creating a SPOT job to perform the file transfer. The SPOT job name is of the form xyyy, where x is the last character of the PID of the SCOPE 2 mainframe, and yyy is the JCB ordinal of the SCOPE 2 job. SSF places the newly created SPOT job into the NOS input queue, and the job executes on NOS in the same manner as any other job.

During execution, the job checks that the user has the appropriate file access permission. If so, the SPOT job initiates the file transfer.

For staged file transfers, each SPOT job terminates upon completion, and the job's dayfile (containing all processing and error messages) is sent to SCOPE 2 to be included in the dayfile of the job that requested the file transfer. SSF initiates a new SPOT job for each staged file transfer request it receives.

SSF periodically queries SCOPE 2 about file staging activities that should be dropped. For example, the SCOPE 2 system operator may have dropped a job for which staging was active. If any file staging activities are to be dropped, SSF requests information to identify the staging activity and then drops the associated SPOT job. No operator intervention is required.

Spooled File Transfers

Spooled file transfers use a single SPOT job to perform all input/output file transfers between SSF and SCOPE 2. As soon as communications are established between SSF and SCOPE 2, SSF creates a SPOT job to handle input/output file transfers. The job is called xSTA, where x is the PID of the SCOPE 2 mainframe. This SPOT job can be swapped out during periods of low activity, but the job is not terminated until the SSF-SCOPE 2 link is dropped.

An input file transfer begins when the spooling SPOT job selects a file from the NOS input queue that is destined for the SCOPE 2 system. The SPOT job queries SCOPE 2 for a system file table (SFT) entry for the input file to be transferred. When an entry is assigned, the spooling SPOT job transfers the input file to SCOPE 2. Upon completion of the transfer, the SPOT job requests that the input file be removed from the NOS input queue.

For each output file transferred from SCOPE 2, the spooling SPOT job transfers the output file from the SCOPE 2 output queue to a local NOS file. When the transfer is complete, the SPOT job routes the local file to the NOS output queue (with the appropriate disposition code, forms code, or terminal ID). The spooling SPOT job then requests the removal of the output file from the SCOPE 2 output queue.

Error Logout

A loss of communications between SSF and SCOPE 2 results in an error logout condition. Error logout occurs as a result of a STOP command entered by the system operator, or a software or hardware error. When SSF detects an interruption in communications, it logs out of the SCOPE 2 system after:

- All partially transmitted files are rewound on the sending end of the link and are deleted at the receiving end.
- All SPOT jobs are dropped.

No files are lost as a result of the communications break. After logging out of the SCOPE 2 system, SSF periodically queries SCOPE 2 to determine if the link has been restored.

Status/Control Register Simulator

D

Using the Simulator	D-2
Console Operation	D-2
Batch Input	D-4
Simulator Commands	D-5
Holding Register Commands	D-5
Cycle Commands	D-7
+ And Stop Commands	D-7

Error logging on a CYBER 170 Computer System enables you to detect and log errors identified in the status/control (S/C) register. The status/control register simulator (SCRSIM) allows you to set S/C register bits in order to aid in the testing of error logging and error recovery procedures.

SCRSIM does not run on CYBER 180-class machines. Models 865 and 875 use maintenance registers instead of S/C registers. Throughout this appendix all references to S/C registers also apply to the model 865 and 875 maintenance registers.

SCRSIM runs on CYBER 170 Computer Systems using the S/C register on channel 16 and, if more than 10 PPs are available on the system, the S/C register on channel 36. On CYBER 70 Computer Systems, SCRSIM uses the interlock register on channel 15.

With the aid of a K display, you can specify commands to set and clear bits, set bytes, and set lines and areas in holding registers. This allows both S/C registers to be set up completely. The contents of the holding register can then be transferred to the S/C registers (64 or 128 bits are transferred to the interlock register of a CYBER 70 Computer System).

The bits set through this simulator are logged in the error log if an error bit is set, thus aiding in testing and software checkout. (Refer to the appropriate hardware reference manual for a complete description of the significance of each S/C register bit.) The simulator job dayfile lists all simulator commands entered. This error logging does not occur in a CYBER 70 Computer System, however, unless an ENABLE,SCRSIM IPRDECK entry has been made.

NOTE

Be careful when using the simulator. Improper use may result in serious system malfunctions.

Refer to the NOS Version 2 Operations Handbook for descriptions of messages produced by SCRSIM.

Using the Simulator

Error logging is always enabled on a CYBER 170 Computer System except models 815, 825, 835, 845, and 855. Error logging is enabled on a CYBER 70 Computer System only if the ENABLE,SCRSIM IPRDECK entry is present.

Console Operation

The simulator is called from the console by entering:

```
X.SCRSIM.
```

The simulator K display (refer to figure D-1) appears on the left screen after entering:

```
K,jsn.
```

Entry	Description
jsn	The job sequence name of SCRSIM noted on the B,O display.

This K display shows the contents of the temporary holding registers, as well as a central memory buffer. The buffer contains:

- A history of all error status bits since the last level 0 deadstart. If an error status bit has been set in the S/C register, it remains set in the buffer, even though it may have been cleared in the actual S/C register. This history may be useful in diagnosing system malfunctions.
- All other bits in the buffer reflect actual values in the S/C register at the time the last error bit was set. Each time an error bit is set, the entire buffer is updated.

Unless the simulator is running on a CYBER 170 Computer System with more than 10 PPs, the message

```
CHANNEL 36 NOT AVAILABLE
```

also appears. This indicates that no channel 36 S/C register is present on the machine, and thus, no simulation need be done for it.

By entering

```
KK.
```

the simulator commands K display (refer to figure D-2) appears on the right screen. This display gives a brief description of the commands available.

Commands can be entered on the K display by entering:

```
K.command.
```

Entry	Description
command	One of the commands shown in figure D-2.


```

TEMPORARY HOLDING REGISTER CONTENTS
CHANNEL 16 REGISTER - LINES 0-3

BITS 59- 0                                BYTE 00
000000000000 000000000000 000000000000 000000000000 000000000000
    0000          0000          0000          0000          0000

BITS 119- 60                              BYTE 05
000000000000 000000000000 000000000000 000000000000 000000000000
    0000          0000          0000          0000          0000

BITS 179-120                              BYTE 10
000000000000 000000000000 000000000000 000000000000 000000000000
    0000          0000          0000          0000          0000

BITS 203-180                              BYTE 15
                                000000000000 000000000000
                                0000          0000

CHANNEL 36 NOT AVAILABLE

ACTUAL S/C REGISTER ERROR BUFFER

CHANNEL 16 REGISTER - WORDS 113-116
000000000000 000000000000 000000000000 000000000000 000000000000
000000000000 000000000000 000000000000 000000000000 000000000000
000000000000 000000000000 000000000000 000000000000 000000000000
                                000000000000 000000000000

```

Figure D-1. Simulator K Display (Left Screen)

COMMAND	SIMULATOR COMMANDS DESCRIPTION
AREA, A, M, Y.	SET M BITS FROM A TO OCTAL VALUE Y
BYTE, XX, YYYY.	SET BYTE XX TO OCTAL VALUE YYYY
CLEAR, A, B, . . . , Z.	CLEAR BITS A, B, . . . , Z
CYCLE, X, T, R.	SET BIT X EVERY 16*T MS. R TIMES
END.	END CYCLE COMMAND BEFORE R REACHED
LINE, X, Y.	SET LINE X TO OCTAL VALUE Y
READ.	READ S/C REGISTER INTO HOLDING REGISTER
SET, A, B, . . . , Z.	SET BITS A, B, . . . , Z
+	CHANGE REGISTER BEING USED AND K DISPLAY
GO.	ENTER HOLDING REGISTERS IN S/C REGISTERS
STOP.	END THE SIMULATOR
ALL BIT, BYTE, AND LINE NUMBERS ASSUMED DECIMAL. TIME VALUES ASSUMED DECIMAL Y AND YYYY VALUES MUST BE OCTAL.	

Figure D-2. Simulator Commands K Display (Right Screen)

Batch Input

The simulator may also be called from batch input by using the SCRSIM command. The input file must have a record containing the commands to be processed, one command per card. The system must be in debug mode and the user must be validated for system origin privileges.

Simulator Commands

The simulator commands are described next. You must enter the entire command keyword and only one command may be entered at a time. Each command, except +, must end with a terminator. In all cases, a null argument is assumed to be zero.

Holding Register Commands

The following commands, except GO., affect only the holding register currently displayed on the left screen. These commands are used to set up the entire 204 bits in the holding registers. GO. transfers the holding register contents to the actual S/C register. (On a CYBER 70 Computer System, the channel 16 S/C register is simulated by the interlock register. GO. transfers the first 64 of 128 bits of the holding register to the interlock register.) The current contents of the holding register is displayed in binary and octal on the left screen (refer to figure D-1). The contents of the actual S/C register are also displayed in binary on the left screen.

Command	Description								
AREA,a,m,y.	Set m bits in the holding register, from bit a to bit a+m-1, to the octal value y.								
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>a</td> <td>Starting bit number; from 0 to 203. a is assumed to be decimal, but a postradix of D or B may also be included.</td> </tr> <tr> <td>m</td> <td>Number of bits to be set. m is assumed to be decimal, but a postradix of D or B may also be included.</td> </tr> <tr> <td>y</td> <td>Value to which the bits are to be set. y may be up to m bits of octal value.</td> </tr> </tbody> </table>	Parameter	Description	a	Starting bit number; from 0 to 203. a is assumed to be decimal, but a postradix of D or B may also be included.	m	Number of bits to be set. m is assumed to be decimal, but a postradix of D or B may also be included.	y	Value to which the bits are to be set. y may be up to m bits of octal value.
Parameter	Description								
a	Starting bit number; from 0 to 203. a is assumed to be decimal, but a postradix of D or B may also be included.								
m	Number of bits to be set. m is assumed to be decimal, but a postradix of D or B may also be included.								
y	Value to which the bits are to be set. y may be up to m bits of octal value.								
BYTE,xx,yyyy.	Set byte xx in the holding register to the octal value yyyy.								
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>xx</td> <td>Byte number; from 0 to 16. xx is assumed to be decimal, but a postradix of D or B may be included.</td> </tr> <tr> <td>yyyy</td> <td>Value to which byte xx is to be set. yyyy may be up to 12 bits of octal value.</td> </tr> </tbody> </table>	Parameter	Description	xx	Byte number; from 0 to 16. xx is assumed to be decimal, but a postradix of D or B may be included.	yyyy	Value to which byte xx is to be set. yyyy may be up to 12 bits of octal value.		
Parameter	Description								
xx	Byte number; from 0 to 16. xx is assumed to be decimal, but a postradix of D or B may be included.								
yyyy	Value to which byte xx is to be set. yyyy may be up to 12 bits of octal value.								
CLEAR,a ₁ ,a ₂ ,...,a _n .	Clear bits a ₁ ,a ₂ ,...,a _n in the holding register. If more than 30 bit numbers are entered, only the first 30 are processed. All others are ignored.								
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>a₁</td> <td>Bit number to be cleared; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.</td> </tr> </tbody> </table>	Parameter	Description	a ₁	Bit number to be cleared; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.				
Parameter	Description								
a ₁	Bit number to be cleared; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.								

Command	Description								
LINE,x,y.	Set line x of the holding register to the octal value y.								
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>x</td> <td>Line number of the holding register shown on the left display screen (refer to figure D-1); from 0 to 3. Line 0 is positioned at the top and line 3 is positioned at the bottom.</td> </tr> <tr> <td>y</td> <td>Value to which line x is to be set. y may be up to 60 bits of octal value.</td> </tr> <tr> <td>READ.</td> <td>Transfer the contents of the actual S/C register to the holding register.</td> </tr> </tbody> </table>	Parameter	Description	x	Line number of the holding register shown on the left display screen (refer to figure D-1); from 0 to 3. Line 0 is positioned at the top and line 3 is positioned at the bottom.	y	Value to which line x is to be set. y may be up to 60 bits of octal value.	READ.	Transfer the contents of the actual S/C register to the holding register.
Parameter	Description								
x	Line number of the holding register shown on the left display screen (refer to figure D-1); from 0 to 3. Line 0 is positioned at the top and line 3 is positioned at the bottom.								
y	Value to which line x is to be set. y may be up to 60 bits of octal value.								
READ.	Transfer the contents of the actual S/C register to the holding register.								
SET,a ₁ ,a ₂ ,...,a _n .	Set bits a ₁ ,a ₂ ,...,a _n in the holding register. If more than 30 bit numbers are entered, only the first 30 are processed. All others are ignored.								
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>a_i</td> <td>Bit number to be set; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.</td> </tr> </tbody> </table>	Parameter	Description	a _i	Bit number to be set; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.				
Parameter	Description								
a _i	Bit number to be set; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.								
GO.	Transfer the contents of the holding register to the actual S/C register or to the interlock register of a CYBER 70 Computer System. No bits in the holding register are changed by this command.								

Cycle Commands

Command	Description								
CYCLE,x,t,r.	Set bit x every t periods of time, a total of r times. This command assumes control of the simulator for the total time period specified. During this time, no command is accepted except END. Bit x is set in the holding register by this command.								
	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>x</td> <td>Bit number to be set; from 0 to 203.</td> </tr> <tr> <td>t</td> <td>Number of periods of time for each cycle; one period is 16 milliseconds. t=32 is approximately 0.5 second.</td> </tr> <tr> <td>r</td> <td>Number of times to set bit x. r may not exceed 4095.</td> </tr> </tbody> </table>	Parameter	Description	x	Bit number to be set; from 0 to 203.	t	Number of periods of time for each cycle; one period is 16 milliseconds. t=32 is approximately 0.5 second.	r	Number of times to set bit x. r may not exceed 4095.
Parameter	Description								
x	Bit number to be set; from 0 to 203.								
t	Number of periods of time for each cycle; one period is 16 milliseconds. t=32 is approximately 0.5 second.								
r	Number of times to set bit x. r may not exceed 4095.								
END.	End CYCLE command processing before r is reached. Control of the simulator is returned to the operator.								

+ And Stop Commands

Command	Description
+	Toggle the K display between the channel 16 and channel 36 register displays and also change the register currently being worked on, if the simulator is being operated on a CYBER 170 Computer System with two S/C registers. If two S/C registers are not present on the machine, no action is taken. The channel 16 register is assumed when the simulator begins. The only holding register that is affected by the holding register and cycle commands is the one currently displayed.
STOP.	End simulator processing.

580 Programmable Format Control

E

Building Programmable Format Control Arrays E-2
Adding Programmable Format Control Arrays E-3

You control the spacing and format on 580 line printers by using carriage control format tapes or programmable format control. The carriage control format tape is punched to indicate particular format channels for each frame. A printer with programmable format control does not use carriage control format tapes; instead, it contains a microprocessor plus memory. Programmable format control arrays are loaded into this memory, performing the same function as the format tape. This appendix describes how to create and load programmable format control arrays. A description of format tapes is included in the NOS Version 2 Reference Set, Volume 3.

Certain 580 printers are not equipped with a carriage control tape; instead, a microprocessor plus memory called programmable format control is used. Instead of a tape controlling the page format, software is used. This software is called a programmable format control array. A programmable format control array consists of numbers from 0 to 12 and 17₈. Each nonzero character represents a channel. A 0 (zero) specifies that no channel is selected. An array is similar to a format tape since each number in a programmable format control array corresponds to a line on the print form. As each line of a page is printed, the next number in the programmable format control array buffer is addressed. A carriage control character in column 1 of the output line, indicating a skip to a particular channel, causes the memory in the programmable format control array buffer to be sequentially addressed until the particular number is found. The paper is spaced a similar number of spaces.

A programmable format control array differs from a format tape because only one channel can be specified per line. A number of channels can be specified per line using a format tape. Also, programmable format control arrays are accessed in pairs, one for 6 lines per inch (lpi) printing and one for 8-lpi printing. The 8-lpi array is usually larger, allowing more lines to be printed on the same size page.

Building Programmable Format Control Arrays

Observe the following rules when building a programmable format control array.

- Enter only valid numbers (from 0 to 12 and 17_8) in the programmable format control array.
- A 1 must be the first number in the programmable format control array, indicating a top-of-form position.
- A 12 should always indicate the last line of the form (bottom of page).
- A 17_8 should appear as the last number in the array, denoting the end of valid numbers for a given array. This number does not correspond to any particular line on the form.
- Maximum length programmable format control arrays (132 for 6 lpi and 176 or 136 for 8 lpi) must include a 9 only at location 132 for 6 lpi and location 176 or 136 for 8 lpi. Improper paper alignment may occur if a 9 is placed elsewhere.

To properly load the appropriate programmable format control buffer for a particular form, it is necessary to assemble data that will contain, when transmitted and stored in the programmable format control array buffer, as many numbers as lines on the form. As stored within the programmable format control array buffer, each number is a 4-bit code used to represent channels (1 through 12) or a null code (no channel selected).

NOTE

Channels 13 and 14 are valid channels but they are not selectable. No programmable format control error occurs when loading these numbers into the programmable format control buffer.

The maximum capacity of the 6-lpi programmable format control buffer is 132 numbers (22-inch form maximum) plus the last line number whereas the 8-lpi programmable format control buffer has a capacity of 176 or 136 numbers (22-inch or 17-inch form maximum) plus the last line number. Fewer than the maximum amount of numbers may be loaded into the programmable format control buffer, but an excessive amount of numbers will cause a programmable format control overflow error.

Adding Programmable Format Control Arrays

Programmable format control arrays must occur in pairs (one 6-lpi array and one 8-lpi array); therefore, when a particular array is specified with the SC option of the ROUTE command, it is possible to switch from 6- to 8-lpi spacing except when using 8.5-inch forms, which are always printed at 8-lpi spacing. Four pairs of arrays are provided with the operating system, two pairs of arrays for 11-inch forms and two pairs of arrays for 8.5-inch forms (refer to table E-1). The numbers are omitted for 6-lpi 8.5-inch forms because this combination is not selectable. The form or paper size is defined in the EQPDECK unit record equipment EST entry.

To add additional programmable format control arrays, BIO must be modified to contain these additional arrays as overlays in QAP. The overlay names and corresponding ROUTE command SC parameter options must be added to PFCO, and the programmable format control overlay table in QAP. The first entry in this table specifies the default programmable control format array. This array is used when the SC parameter is not specified. If an added array is to be the default array, the PFCO entry for it should be placed at the beginning of the table. The value that must be specified with the SC parameter (any 6-bit binary value) is included in this entry unless it is to be the default array. In this case, the value must be zero. The arrays that are also added to QAP must follow the same format, structure, and labeling conventions as the arrays already provided in QAP. To generate the numbers to be loaded into the programmable format control buffer, the DPFC (define programmable format control code) macro is used with the following format. Only the lower 4 bits of each 6 bits are used as programmable format control numbers.

Location	Operation	Variable
	DPFC	a,b,c,d

Variable	Description
----------	-------------

a, b, c, d	Programmable format control numbers for 6 or 8 lpi. These numbers can be any value from 0 to 12 and 17 ₈ .
------------	---

NOTE

In each macro call, all four parameters must be specified (0 denotes no channel is selected). If all four are not required to complete an array (the number of lines on the form is not a multiple of 4), zeros should be used for the remaining parameters.

Table E-1. Released Programmable Format Control Arrays

Location	11 in. ¹	8.5 in. ¹	11 in. ²	8.5 in. ²	11 in. ³	8.5 in. ³	11 in. ⁴	8.5 in. ⁴
1	1		1		1	1	1	1
2	6		6		6	6	6	6
3	0		0		0	0	0	0
4	2		2		0	2	0	2
5	0		0		2	0	2	0
6	0		0		0	0	0	0
7	3		3		0	3	0	3
8	0		0		0	0	0	0
9	0		0		3	0	3	0
10	4		4		0	4	0	4
11	0		0		0	0	0	0
12	0		0		0	0	0	0
13	5		5		4	5	4	5
14	0		0		0	0	0	0
15	0		0		0	0	0	0
16	2		7		0	2	0	7
17	0		0		5	0	5	0
18	0		0		0	0	0	0
19	3		8		0	3	0	8
20	0		0		0	0	0	0
21	0		0		2	0	7	0
22	4		10		0	4	0	10
23	0		0		0	0	0	0
24	0		0		0	0	0	0
25	5		11		3	5	8	11
26	0		0		0	0	0	0
27	0		0		0	0	0	0
28	2		2		0	2	0	2
29	0		0		4	0	10	0
30	0		0		0	0	0	0
31	3		3		0	3	0	3
32	0		0		0	0	0	0
33	0		0		5	0	11	0
34	4		4		0	4	0	4
35	0		0		0	0	0	0

1. 6 lpi default (default is SC=1).
2. 6 lpi alternate (alternate is SC=2).
3. 8 lpi default.
4. 8 lpi alternate.

(Continued on next page)

Table E-1. Released Programmable Format Control Arrays

Location	11 in. ¹	8.5 in. ¹	11 in. ²	8.5 in. ²	11 in. ³	8.5 in. ³	11 in. ⁴	8.5 in. ⁴
<i>(Continued from previous page)</i>								
36	0		0		0	0	0	0
37	5		5		2	5	2	5
38	0		0		0	0	0	0
39	0		0		0	0	0	0
40	2		7		0	2	0	7
41	0		0		3	0	3	0
42	0		0		0	0	0	0
43	3		8		0	3	0	8
44	0		0		0	0	0	0
45	0		0		4	0	4	0
46	4		10		0	4	0	10
47	0		0		0	0	0	0
48	0		0		0	0	0	0
49	5		11		5	5	5	11
50	0		0		0	0	0	0
51	0		0		0	0	0	0
52	2		2		0	2	0	2
53	0		0		2	0	7	0
54	0		0		0	0	0	0
55	3		3		0	3	0	3
56	0		0		0	0	0	0
57	0		0		3	0	8	0
58	4		4		0	4	0	4
59	0		0		0	0	0	0
60	0		0		0	0	0	0
61	5		5		4	5	10	5
62	0		0		0	0	0	0
63	0		0		0	0	0	0
64	12		12		0	2	0	7
65	0		0		5	0	11	0
66	0		0		0	12	0	12
67	1		1		0	0	0	0
68	0		0		0	0	0	0
69	0		0		2	1	2	1
70	2		2		0	0	0	0
71	0		0		0	0	0	0

1. 6 lpi default (default is SC=1).
2. 6 lpi alternate (alternate is SC=2).
3. 8 lpi default.
4. 8 lpi alternate.

(Continued on next page)

Table E-1. Released Programmable Format Control Arrays

Location	11 in. ¹	8.5 in. ¹	11 in. ²	8.5 in. ²	11 in. ³	8.5 in. ³	11 in. ⁴	8.5 in. ⁴
<i>(Continued from previous page)</i>								
72	0		0		0	2	0	2
73	3		3		3	0	3	0
74	0		0		0	0	0	0
75	0		0		0	3	0	3
76	4		4		0	0	0	0
77	0		0		4	0	4	0
78	0		0		0	4	0	4
79	5		5		0	0	0	0
80	0		0		0	0	0	0
81	0		0		5	5	5	5
82	2		7		0	0	0	0
83	0		0		0	0	0	0
84	0		0		0	0	0	7
85	3		8		12	0	12	0
86	0		0		0	0	0	0
87	0		0		0	3	0	8
88	4		10		0	0	0	0
89	0		0		1	0	1	0
90	0		0		0	4	0	10
91	5		11		0	0	0	0
92	0		0		0	0	0	0
93	0		0		2	5	2	11
94	2		2		0	0	0	0
95	0		0		0	0	0	0
96	0		0		0	2	0	2
97	3		3		3	0	3	0
98	0		0		0	0	0	0
99	0		0		0	3	0	3
100	4		4		0	0	0	0
101	0		0		4	0	4	0
102	0		0		0	4	0	4
103	5		5		0	0	0	0
104	0		0		0	0	0	0
105	0		0		5	5	5	5
106	2		7		0	0	0	0
107	0		0		0	0	0	0

1. 6 lpi default (default is SC=1).
2. 6 lpi alternate (alternate is SC=2).
3. 8 lpi default.
4. 8 lpi alternate.

(Continued on next page)

Table E-1. Released Programmable Format Control Arrays

Location	11 in. ¹	8.5 in. ¹	11 in. ²	8.5 in. ²	11 in. ³	8.5 in. ³	11 in. ⁴	8.5 in. ⁴
<i>(Continued from previous page)</i>								
108	0		0		0	2	0	7
109	3		8		2	0	7	0
110	0		0		0	0	0	0
111	0		0		0	3	0	8
112	4		10		0	0	0	0
113	0		0		3	0	8	0
114	0		0		0	4	0	10
115	5		11		0	0	0	0
116	0		0		0	0	0	0
117	0		0		4	5	10	11
118	2		2		0	0	0	0
119	0		0		0	0	0	0
120	0		0		0	2	0	2
121	3		3		5	0	11	0
122	0		0		0	0	0	0
123	0		0		0	3	0	3
124	4		4		0	0	0	0
125	0		0		2	0	2	0
126	0		0		0	4	0	4
127	5		5		0	0	0	0
128	0		0		0	0	0	0
129	0		0		3	5	3	5
130	12		12		0	0	0	0
131	0		0		0	0	0	0
132	9		9		0	2	0	7
133	17 ₈	17 ₈			4	0	4	0
134					0	12	0	12
135					0	0	0	0
136					0	9	0	9
137					5	17 ₈	5	17 ₈
138					0		0	
139					0		0	
140					0		0	
141					2		7	
142					0		0	
143					0		0	

1. 6 lpi default (default is SC=1).
2. 6 lpi alternate (alternate is SC=2).
3. 8 lpi default.
4. 8 lpi alternate.

(Continued on next page)

Table E-1. Released Programmable Format Control Arrays

Location	11 in. ¹	8.5 in. ¹	11 in. ²	8.5 in. ²	11 in. ³	8.5 in. ³	11 in. ⁴	8.5 in. ⁴
<i>(Continued from previous page)</i>								
144					0		0	
145					3		8	
146					0		0	
147					0		0	
148					0		0	
149					4		10	
150					0		0	
151					0		0	
152					0		0	
153					5		11	
154					0		0	
155					0		0	
156					0		0	
157					2		2	
158					0		0	
159					0		0	
160					0		0	
161					3		3	
162					0		0	
163					0		0	
164					0		0	
165					4		4	
166					0		0	
167					0		0	
168					0		0	
169					5		5	
170					0		0	
171					0		0	
172					0		0	
173					12		12	
174					0		0	
175					0		0	
176					9		9	
177					17 ₈		17 ₈	

1. 6 lpi default (default is SC=1).
2. 6 lpi alternate (alternate is SC=2).
3. 8 lpi default.
4. 8 lpi alternate.

533/536, 537, and 585 Printer Electronic Vertical Format Unit (EVFU)

F

EVFU Directives for 533/536 Printers	F-1
Format Control Directives	F-1
EVFU Load File Example	F-3
EVFU Load Image	F-5

533/536, 537, and 585 Printer Electronic Vertical Format Unit (EVFU)

F

Upon startup, the printer support utility (PSU) looks for an EVFU load file which is a public file named EVFULFN under the network user name NETOPS. The EVFU load file contains definitions which describe all 533, 536, 537, and 585 printers to be serviced by PSU. In addition, for 533/536 printers it contains directives that describe the actions to be taken while printing files that have format control characters in column 1 and the page length and print density specified. EVFU information for 537 and 585 printers is supplied by CDCNET, and is documented in the CDCNET Configuration and Site Administration Guide.

PSU supports up to 12 printers with any mix of 533/536, 537, and 585 printers. The printer directives in the released EVFU file specify PRINT01 through PRINT04 as 533/536 printers; PRINT05 and PRINT06 as 537 interactive printers and PRINT07 and PRINT08 as 537 batch printers; and PRINT09 through PRINT12 as 585 printers. If these printer names are not correct for the printers being used, the EVFU file will have to be changed accordingly. If 533/536 printers are used, the 533/536 printer definitions must follow all other printer definitions, which are then followed by the 533/536 EVFU load data. If a mismatch occurs between the EVFU load file printer definition and the connection received from the network, the connection is placed in a HOLD state and a K-display message provides the pertinent information. If the K-display is not assigned when this occurs, a flashing B-display message requests that the K-display be assigned to PSU.

EVFU Directives for 533/536 Printers

EVFU directives follow the printer definitions in the EVFU load file. The presence of a character in column 1 initiates a new EVFU directive. A line with a blank in column 1 is a continuation of the line that started the current EVFU directive. Blank lines are not permitted in the EVFU load file. An EVFU directive terminates with a semicolon or end-of-line. You can add comments to an EVFU directive line following the semicolon.

An EVFU load file contains two types of EVFU directives:

- Format control directives.
- EVFU load directives.

Format Control Directives

The character in column 1 of each line of a print file is a format control character. PSU translates this character into a string of one or more data characters that are output instead of the format control character.

A format control directive consists of a single format control character followed by the associated data for PSU to output in place of the format control character specified in the directive. At least one space must separate the format control character from the data characters. There are two formats for the data characters, and you can intermix the two formats. The first data format is pairs of hexadecimal digits representing ASCII characters. These pairs of digits may be separated from each other by spaces for readability. The second data format is ASCII strings delimited by quotation marks, for example, "data".

You can specify a maximum of six data items in a format control directive. They can be hexadecimal digit pairs or ASCII strings or both, but not more than a total of six items representing ASCII characters. Character data should not extend over line boundaries.

PSU supports the following format control characters:

Character	Description
1	Eject page before printing.
2	Skip to last line of form before printing.
+	Skip zero lines before printing (overprint).
space	Single space.
0	Skip one line before printing (double space).
-	Skip two lines before printing (triple space).
8	Skip to top of page (channel 1) before printing.
7	Skip two lines (channel 2) before printing.
6	Skip three lines (channel 3) before printing.
5	Skip four lines (channel 4) before printing.
4	Skip five lines (channel 5) before printing.
3	Skip one page (channel 6) before printing.
9	Skip seven lines (channel 7) before printing.
X	Skip eight lines (channel 8) before printing.
Y	Skip three lines (channel 9) before printing.
Z	Skip ten lines (channel 10) before printing.
W	Eject page (channel 11) before printing.
U	Eject page (channel 12) before printing.
A	Eject page after printing.
B	Skip to last line of form after printing.
/	Skip zero lines after printing (overprint).
H	Skip to top of page (channel 1) after printing.
G	Skip two lines (channel 2) after printing.
F	Skip three lines (channel 3) after printing.
E	Skip four lines (channel 4) after printing.
D	Skip five lines (channel 5) after printing.
C	Skip one page (channel 6) after printing.
I	Skip seven lines (channel 7) after printing.
J	Skip eight lines (channel 8) after printing.
K	Skip three lines (channel 9) after printing.
L	Skip ten lines (channel 10) after printing.
M	Eject page (channel 11) after printing.
N	Eject page (channel 12) after printing.
Q	Clear auto page eject.
R	Set auto page eject.
S	Set 6-lpi print density.
T	Set 8-lpi print density.

If column 1 of a format control directive contains a blank or any character that has not been defined, PSU performs a line termination and line feed sequence. Thus, the concept of an invalid format control character does not exist. Lines with something other than a defined character in column 1 will have that character replaced by a blank. Following the line feed, processing of the rest of the line will continue as normal.

Note that Q, R, S, and T may appear either as format control characters or as EVFU load characters. However, if both are specified, the format control character takes precedence and the EVFU load character will not be used.

The level number mentioned earlier corresponds to various combinations of page size and print density or spacing codes that can be specified. The following EVFU level numbers are defined for the default EVFU load file:

Level Number	Description
1	Defines a 12-inch form length with a 6-lpi print density.
2	Defines a 12-inch form length with an 8-lpi print density.
3	Defines an 11-inch form length with a 6-lpi print density.
4	Defines an 11-inch form length with an 8-lpi print density.
5	Defines an 8.5-inch form length with a 6-lpi print density.
6	Defines an 8.5-inch form length with an 8-lpi print density.

If you specify a spacing code or request a forms length/density for which no load data exists, PSU will issue a message to the printer, then rewind the output file, and print it using the default EVFU level of 3.

EVFU Load File Example

Figure F-1 provides an example of an EVFU Load File. It shows the required order of directives for a 533/536 printer EVFU load file; printer definitions are first, format control directives are next, and EVFU load directives are last.

In this example, 1B 5C is the trailer for all EVFU data except for the print density selection. For format control directives (skip to channel), the leader is 1B 50 22. For EVFU load directives, the leader is 1B 50 23. However, the printer paper must be positioned at the top of form any time the EVFU is changed because that is how the printer determines the top of form (when the EVFU is changed). Hence, all EVFU load data starts with 0C (form feed). The 1B 63 in the V load directive is a printer reset. A T specifies the total number of lines on the form (T66). An L specifies where to set a stop in a particular line; for example L41;4;5 sets stops at positions 4 and 5 in line 41. Similarly, a C clears a stop. Items following the leader 1B 50 23 and preceding the trailer 1B 5C are separated by semicolons.

```

PRINT09,PC=BATWUFU,DOWN.
PRINT10,PC=BATWUFU,DOWN.
PRINT11,PC=BATWUFU,DOWN.
PRINT12,PC=BATWUFU,DOWN.
PRINT05,PC=INTNVFU,DOWN.
PRINT06,PC=INTNVFU,DOWN.
PRINT07,PC=INTNVFU,DOWN.
PRINT08,PC=INTNVFU,DOWN.
PRINT01,PC=C536INT,DOWN.
PRINT02,PC=C536INT,DOWN.
PRINT03,PC=C536INT,DOWN.
PRINT04,PC=C536INT,DOWN.
8 1B 50 22 "1" 1B 5C
7 1B 50 22 "2" 1B 5C
6 1B 50 22 "3" 1B 5C
5 1B 50 22 "4" 1B 5C
4 1B 50 22 "5" 1B 5C
3 1B 50 22 "6" 1B 5C
H 1B 50 22 "1" 1B 5C
G 1B 50 22 "2" 1B 5C
F 1B 50 22 "3" 1B 5C
E 1B 50 22 "4" 1B 5C
D 1B 50 22 "5" 1B 5C
C 1B 50 22 "6" 1B 5C
Q3 0C 1B 50 23 "C64;2;L66;2" 1B 5C
R3 0C 1B 50 23 "C66;2;L64;2" 1B 5C
S3 0C 1B 50 23 "T66;L64;2;C65;4;C66;5" 1B 5C
    1B 5B 31 32 30 20 47
T4 0C 1B 50 23 "T88;C64;2;7;L65;4;L86;2;7" 1B 5C
    1B 5B 39 30 20 47
V3 0C 1B 63 1B 50 23 "T66;"
    "L1;1;3;4;5;6;7;L3;7;L4;3;L5;4;7;"
    "L6;5;L7;3;7;L9;4;7;L10;3;"
    "L11;5;7;L13;3;4;7;L15;7;"
    "L16;3;5;L17;4;L19;3;7;"
    "L21;4;5;7;L22;3;L23;7;L25;3;4;7;"
    "L26;5;L27;7;L28;3;L29;4;7;"
    "L31;3;5;7;L33;4;7;L34;3;L35;7;"
    "L36;5;L37;3;4;7;L39;7;L40;3;"
    "L41;4;5;7;L43;3;7;L45;4;7;"
    "L46;3;5;L47;7;L49;3;4;7;"
    "L51;5;7;L52;3;L53;4;7;L55;3;7;"
    "L56;5;L57;4;7;L58;3;L59;7;"
    "L61;3;4;5;7;L63;7;L64;2;3;"
    "L67;3;7;L69;4;7;L70;3;"
    "L71;5;7;L73;3;4;7;L75;7;"
    "L76;3;5;L77;4;7;L79;3;7;"
    "L81;4;5;7;L82;3;L83;7;L85;3;4;7;"
    "L86;2;5"
1B 5C
} Printer Definitions
} Format Control
} Directives
}
;CLEAR AUTO PAGE EJECT
;SET AUTO PAGE EJECT
;FF, 66 LPI, BOF
;SELECT 6 LPI
;FF, 88 LPI, BOF
;SELECT 8 LPI
;FF, RESET, LINES/FORM
;LINES 1 - 5
;LINES 6 - 10
;LINES 11 - 15
;LINES 16 - 20
;LINES 21 - 25
;LINES 26 - 30
;LINES 31 - 35
;LINES 36 - 40
;LINES 41 - 45
;LINES 46 - 50
;LINES 51 - 55
;LINES 56 - 60
;LINES 61 - 65
;LINES 66 - 70
;LINES 71 - 75
;LINES 76 - 80
;LINES 81 - 85
;LINES 86 - 88
;TERMINATE
} EVFU
} Load
} Directives

```

Figure F-1. EVFU Load File

EVFU Load Image

Table F-1 shows the released EVFU load image for 533/536 printers. The channel numbers in the format control character descriptions correspond to the channel numbers in table F-1. For example, format control character Y causes the printer to skip three lines before printing. It does this by using channel 9 in the EVFU load image.

Table F-1. Released EVFU Load Image: Print Lines and Channel Numbers

Print Line	1	2	3	4	5	6	7	8	9	10	11	12
1	X	X	X	X	X	X	X	X	X	X	X	X
2												
3	X											
4		X							X			
5	X		X									
6				X								
7	X	X							X			
8						X						
9	X		X				X					
10		X							X			
11	X			X						X		
12												
13	X	X	X						X			
14												
15	X					X						
16		X		X					X			
17	X		X				X					
18												
19	X	X							X			
20												
21	X		X	X							X	
22		X				X			X			
23	X											
24												
25	X	X	X					X	X			
26				X								
27	X											
28		X							X			
29	X		X			X						
30												
31	X	X		X					X	X		
32												
33	X		X				X					
34		X							X			
35	X											
36				X		X						
37	X	X	X						X			
38												

(Continued on next page)

Table F-1. Released EVFU Load Image: Print Lines and Channel Numbers

Print Line	1	2	3	4	5	6	7	8	9	10	11	12
39	X											
40		X						X				
41		X	X	X				X	X			
42												
43		X	X				X		X			
44												
45		X		X								
46			X		X					X		
47		X										
48												
49		X	X	X				X	X		1	1
50							X					
51		X			X					X		
52			X							X		
53		X		X								
54												
55		X	X							X		
56					X							
57		X		X			X	X				
58			X							X		
59		X										
60												
61		X	X	X	X					X	X	
62												
63		X										
64			X				X		X		2	2
65		X		X				X				
66					X						3	3
67		X	X							X		
68												
69		X		X								
70			X							X	4	4

1. Set for 8.5 inch form at 6 lines per inch.
2. Set for 11 inch form at 6 lines per inch.
3. Set for 8.5 inch form at 8 lines per inch.
4. Set for 11 inch form at 8 lines per inch.

(Continued on next page)

Table F-1. Released EVFU Load Image: Print Lines and Channel Numbers

Print Line	1	2	3	4	5	6	7	8	9	10	11	12
<i>(Continued from previous page)</i>												
71		X			X		X			X		
72												
73		X	X	X				X	X			
74												
75		X										
76			X		X				X			
77		X		X								
78							X					
79		X	X						X			
80												
81		X		X	X			X		X		
82			X						X			
83		X										
84												
85		X	X	X			X		X			
86					X						1	1
87		X										
88			X						X			
89		X		X				X				
90												
91		X	X		X				X	X		
92							X					
93		X		X								
94			X						X		2	2
95		X										
96					X							

1. Set for 12 inch form at 6 lines per inch.
2. Set for 12 inch form at 8 lines per inch.

881/883 Pack Reformatting Utility

G

- FORMAT Command G-2
- Input Formats G-4
 - Control Directives G-4
 - Data Statements G-5
- Output Formats G-6
- Accessing Disk Devices G-11
 - Access for Read Operations G-11
 - Access for Read and Write Operations G-12

881/883 Pack Reformatting Utility

G

Each 881 disk pack used in the 844 disk contains factory-recorded flawing information on cylinder 632_g (410), track 0, sectors 0, 1, and 2. Each 883 pack contains this information on cylinder 1466_g (822), track 0, sectors 0, 1, and 2. The following information is included on the cylinders.

- Cylinder 632_g (or 1466_g for 883 packs), track 0, sector 0 contains the factory-recorded manufacturing data. This data consists of the pack serial number and the manufacturing date.
- Cylinder 632_g (or 1466_g), track 0, sector 1 contains the factory map. This map contains a list of all factory-detected flaws, both correctable and uncorrectable.
- Cylinder 632_g (or 1466_g), track 0, sector 2 contains the utility map. This map originally contains all factory-detected uncorrectable flaws. This map is updated by the reformatting utility.

FORMAT is a CPU program that operates in conjunction with FDP, a PP program, to maintain and reformat 881/883 disk packs. It is used to perform the following functions.

- Factory-recorded manufacturing data, factory-recorded flaw data, and utility flaw data can be retrieved from a factory-formatted disk pack.
- Sector and track flaws can be set or cleared on a factory-formatted disk pack.
- Address fields of a previously factory-formatted disk pack can be restored. (This function is used only in the event that addresses on the pack are lost.)

In order to function, FORMAT requires that the factory-recorded data [sectors 0 and 1 of cylinder 632_g (or 1466_g)] be correct and readable. The pack cannot be processed if this data is unreadable. If packs are available that do not contain this factory-recorded information, consult a customer engineer to have this information placed on the packs. Also, the correct level of controlware must be present in order for FORMAT to function. To determine the controlware level and for procedures to install this controlware, refer to the NOS Version 2 Installation Handbook. Since the operating system requires that the utility map contain the physical flaw information in order for automatic logical flawing to be performed, it is important that the utility map be properly maintained.

The operating system automatically sets logical flaws when initializing 844 equipment. This is done by reading the utility map of the 844 units involved, and mapping this physical flaw information into the corresponding logical track addresses. Logical track flaw reservations are then made in the track reservation table (TRT) for the 844 equipment being installed. For example, if the 844 equipment being initialized consists of two physical units (such as a DI-2 configuration), the logical flaws set in the TRT are obtained from the physical flaw information recorded in the utility maps of both units making up the DI-2 configuration. This automatic flawing occurs when an equipment is initialized, regardless of whether the initialization is done during deadstart or online. Automatic flawing also occurs when an X.FLAW request is made from the console.

The operating system allows for setting and clearing flaw information. The SLF and CLF APRDECK entries set or clear logical track reservations in the TRT of the equipment. (Refer to section 3, Deadstart Decks, for information concerning these entries.) If the device is then checkpointed, this flaw information is preserved in the TRT portion of the device label. The SLF entry can be made during deadstart, during online initialization, or by using the FLAW utility (as described in section 8, K-Display Utilities). In any case, the flawing done via these entries is only logical; the flaw information remains only in the TRT and is discarded on subsequent deadstart initialization. This information is also lost during online initializations if it was not possible to recover the equipment. Using SLF does not cause any additional information to be recorded in the utility map; only FORMAT is capable of updating the utility map data. Caution should be used if attempts are made to cancel a logical flaw that was made during automatic flawing, since the physical disk sector is still marked as flawed and attempts to access that sector yield error conditions.

The use of the FORMAT utility for maintaining and reformatting 881/883 disk packs is described next.

FORMAT Command

Processing maintenance operations on an 881/883 type disk pack is initiated by the FORMAT command. This program interfaces with you as required. The format of this command is:

```
FORMAT, P1, P2, . . . , Pn.
```

Each p_i is a keyword or a keyword equated to a value.

P_i	Description
G=m	Relative unit of a multispindle device. This value is checked for validity within the device. For example, if the device is a DI-2 and G=2 is specified, an error results.
I=infile	File on which input directives and data are written.
I	Same as I=INPUT.
L=outfile	Output file on which the information extracted from the disk pack is to be written. (Refer to Output Formats later in this appendix.)
L	Same as L=OUTPUT.

P_i	Description								
MODE=mode	Operational mode for FORMAT.								
	<table border="1"> <thead> <tr> <th>mode</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>ALTER</td> <td>The input file contains directives to control the set or clear flaw operations (refer to Input Formats later in this appendix).</td> </tr> <tr> <td>FETCH</td> <td>The factory-recorded flawing information contained on cylinder 632₈ (or 1466₈), track 0, sectors 0, 1, and 2 is obtained and copied to the output file (and optional output file, if available).</td> </tr> <tr> <td>RESTORE</td> <td>The addresses, flawed sectors, and tracks are restored according to information given in the utility flaw map. If the utility flaw map is not intact, the program aborts.</td> </tr> </tbody> </table>	mode	Description	ALTER	The input file contains directives to control the set or clear flaw operations (refer to Input Formats later in this appendix).	FETCH	The factory-recorded flawing information contained on cylinder 632 ₈ (or 1466 ₈), track 0, sectors 0, 1, and 2 is obtained and copied to the output file (and optional output file, if available).	RESTORE	The addresses, flawed sectors, and tracks are restored according to information given in the utility flaw map. If the utility flaw map is not intact, the program aborts.
mode	Description								
ALTER	The input file contains directives to control the set or clear flaw operations (refer to Input Formats later in this appendix).								
FETCH	The factory-recorded flawing information contained on cylinder 632 ₈ (or 1466 ₈), track 0, sectors 0, 1, and 2 is obtained and copied to the output file (and optional output file, if available).								
RESTORE	The addresses, flawed sectors, and tracks are restored according to information given in the utility flaw map. If the utility flaw map is not intact, the program aborts.								
MODE	Same as MODE=FETCH.								
O=filename	Optional output file to contain the output extracted from the disk pack.								
	<p>NOTE</p> <hr/> <p>If output files other than OUTPUT or optional output files are specified, they should be created prior to the initiation of FORMAT. If they are not, they are destroyed upon completion of FORMAT processing.</p> <hr/>								
P=serialnumber	Pack serial number in decimal of the pack to be processed. If serialnumber does not match the serial number recorded on the disk pack at the factory, processing does not occur.								
P	Same as P=0.								
U=est	EST ordinal of the 844 drive on which the disk pack is mounted. The unit is checked to ensure that it is available for formatting (refer to Accessing Disk Devices later in this appendix).								
V	Specifies that the utility is to verify the addresses recorded on the disk pack. This parameter is valid only if MODE=FETCH or MODE=RESTORE is specified.								

Example:

If all default values are used, the following call is made.

```
FORMAT , I=INPUT , L=OUTPUT , MODE=FETCH , P=0 .
```

At least the U and the P parameter must be correctly specified to initiate processing.

Input Formats

Input to **FORMAT** consists of control directives and data statements. Control directives specify the type of operation to be performed. Data statements specify locations on the pack where the operations are to be performed. A number of data statements may follow each control directive. Control directives and data statements are contained on the input file. This file is accessed only when **MODE=ALTER** has been specified on the **FORMAT** command. The input file (and therefore, control directives and data statements) has no significance when **MODE=FETCH** or **MODE=RESTORE** is specified.

Control Directives

Control directives begin in column 1. The format is:

directive

The following are acceptable directives.

Directive	Description
SET	Declares that the following data statements contain the addresses of flaws to be set and entered in the utility flaw map.
CLEAR	Declares that the following data statements contain the addresses of flaws to be cleared and deleted from the utility flaw map.
FINIS	Declares the end of the input. No information following this directive is processed. This directive is optional.

SET and **CLEAR** directives may be intermixed in the input file. However, all **CLEAR** operations are performed before any **SET** operation. Any attempt to alter the factory map or to set or clear sector flaws in a previously flawed track results in an error.

Data Statements

Data statements begin in column 1. The format is:

`x,cccc,tt,ss`

Parameter	Description						
x	Type of flaw to be set or cleared. Acceptable values are: <table border="1"> <thead> <tr> <th>x</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>S</td> <td>Specifies that the SET or CLEAR directive applies to a sector.</td> </tr> <tr> <td>T</td> <td>Specifies that the SET or CLEAR directive applies to a track.</td> </tr> </tbody> </table>	x	Description	S	Specifies that the SET or CLEAR directive applies to a sector.	T	Specifies that the SET or CLEAR directive applies to a track.
x	Description						
S	Specifies that the SET or CLEAR directive applies to a sector.						
T	Specifies that the SET or CLEAR directive applies to a track.						
cccc	Octal number specifying the cylinder; from 0 to 632 ₈ (or 1466 ₈ for 883 packs).						
tt	Octal number specifying the track; from 0 to 22 ₈ .						
ss	Octal number specifying the sector; from 0 to 27 ₈ . This field is ignored for track flaws (x=T).						

All input data is checked to ensure that the values are within range. Any errors in input result in the termination of the utility before the disk is accessed. Any attempt to alter the factory map, or to set or clear sector flaws in a previously flawed track results in an error.

A maximum of 157 data statements can appear in the input stream.

Output Formats

Output generated by FORMAT is placed on the output file (L=filename on the FORMAT command). This file, for all modes of operation (ALTER, FETCH, and RESTORE), contains:

- A listing of the input stream, if any.
- The pack serial number and date of factory formatting on cylinder 632₈ (or 1466₈), track 0, sector 0.
- A listing of the factory flaw map contained on cylinder 632₈ (or 1466₈), track 0, sector 1.
- A listing of the utility flaw map contained on cylinder 632₈ (or 1466₈), track 0, sector 2.
- A listing of the utility flaw map following any changes resulting from SET or CLEAR directives. This listing appears only when MODE=ALTER is specified on the FORMAT command.
- A listing of the flawed sectors and tracks as read from the disk during address verification. This listing appears only when MODE=FETCH or MODE=RESTORE, and the V parameter are specified on the FORMAT command.

The output generated by FORMAT can be directed to an optional output file (O=filename). This file can then be used as input to another program, or it can be punched or printed.

The following three examples of standard output illustrate a series of reformatting operations performed on the same pack.

Example 1:

A RESTORE operation is performed on an 881 pack. A command similar to the following was entered.

```
FORMAT,U=est,P=819545,MODE=RESTORE.
```

Figure G-1 illustrates the resulting output.

```
DISK PACK REFORMATTING UTILITY                -VERSION 1.1 - 82/01/25.
MODE = RESTORE

        DISK PACK SERIAL NUMBER
        819545

        DATE OF ORIGINAL FACTORY FORMATTING
        74/04/30

        FACTORY FLAW MAP
        (C=CORRECTABLE ERROR,S=SECTOR FLAW,T=TRACK FLAW)
        S,632, 00, 00
        S,632, 00, 01
        S,632, 00, 02
        T,302, 16, 00
        T,362, 01, 00
        T,373, 21, 00
        S,626, 15, 15

        UTILITY FLAW MAP
        (S=SECTOR FLAW,T=TRACK FLAW)
        MAP EMPTY

        PACK FORMATTING COMPLETE, VERIFICATION FOLLOWS
        S,632, 00, 00
        S,632, 00, 02
        S,632, 00, 01

        ADDRESS VERIFICATION COMPLETE
```

Figure G-1. FORMAT Output, MODE=RESTORE

Output Formats

Example 2:

The flaws noted in the factory flaw map from example 1 (refer to figure G-1) are now set in the utility flaw map.

Input similar to the following was entered.

```
FORMAT,U=est,P=819545,MODE=ALTER.  
--EOR--  
SET  
S,626,15,15  
T,302,16,0  
T,362,01,00  
T,373,21,00  
FINIS  
--EOI--
```

Figure G-2 illustrates the resulting output.

```
DISK PACK REFORMATTING UTILITY                -VERSION 1.1 -                82/01/25.
MODE = ALTER
INPUT DATA
SET
S,626, 15, 15
T,302, 16, 00
T,362, 01, 00
T,373, 21, 00
FINIS

DISK PACK SERIAL NUMBER
819545

DATE OF ORIGINAL FACTORY FORMATTING
74/04/30

FACTORY FLAW MAP
(C=CORRECTABLE ERROR, S=SECTOR FLAW, T=TRACK FLAW)
S,632, 00, 00
S,632, 00, 01
S,632, 00, 02
T,302, 16, 00
T,362, 01, 00
T,373, 21, 00
S,626, 15, 15

UTILITY FLAW MAP
(S=SECTOR FLAW, T=TRACK FLAW)
MAP EMPTY

UTILITY FLAW MAP      (ALTERED)
(S=SECTOR FLAW, T=TRACK FLAW)
S,626, 15, 15
T,302, 16, 00
T,362, 01, 00
T,373, 21, 00
```

Figure G-2. FORMAT Output, MODE=ALTER

Output Formats

Example 3:

A FETCH with verification operation is performed to ensure proper reformatting. A command similar to the following was entered.

```
FORMAT,U=est,P=819545,MODE=FETCH,V.
```

Figure G-3 illustrates the resulting output.

```
DISK PACK REFORMATTING UTILITY          -VERSION 1.1 -      82/01/25.
MODE = FETCH
      DISK PACK SERIAL NUMBER
      819545

      DATE OF ORIGINAL FACTORY FORMATTING
      74/04/30

      FACTORY FLAW MAP
      (C=CORRECTABLE ERROR, S=SECTOR FLAW, T=TRACK FLAW)
      S,632, 00, 00
      S,632, 00, 01
      S,632, 00, 02
      T,302, 16, 00
      T,362, 01, 00
      T,373, 21, 00
      S,626, 15, 15

      UTILITY FLAW MAP
      (S=SECTOR FLAW,T=TRACK FLAW)
      S,626, 15, 15
      T,302, 16, 00
      T,362, 01, 00
      T,373, 21, 00

      ADDRESS VERIFICATION FOLLOWS
      T,302, 16, 00
      T,362, 01, 00
      T,373, 21, 00
      S,626, 15, 15
      S,632, 00, 00
      S,632, 00, 02
      S,632, 00, 01

      ADDRESS VERIFICATION COMPLETE
```

Figure G-3. FORMAT Output, MODE=FETCH

Accessing Disk Devices

Special procedures must be used to access the 844 drive used in the reformatting utility. Since certain FORMAT operations (ALTER and RESTORE) can change addresses on the pack, user access to the pack must be restricted.

FORMAT can operate on the pack as follows.

- A read operation (FETCH) obtains formatting information from the pack. The integrity of the pack is maintained.
- Read and write operations (ALTER and RESTORE) can set and clear flaws, and addresses can be rewritten. Users cannot place permanent files on the pack when these operations occur. The integrity of the data on the pack is lost, so a full initialization of the pack must occur before system usage occurs.

Access for Read Operations

Accessing the pack for read operations requires that you specify the U parameter on the FORMAT command with the correct EST ordinal of the device containing the pack. In this case, the device must be a single-spindle device unless you also specify the G parameter. You must also specify the P parameter with the correct pack serial number.

In addition, FORMAT must be called from one of the following.

- A system origin job (from the console).
- A system privileged job (in this case, engineering mode must have been selected on the system console).

Access for Read and Write Operations

In addition to the information specified for read only operations, the following additional steps must be taken to access a device when write operations (ALTER and RESTORE) are to be performed.

1. The pack to be accessed should be mounted on a removable disk device.
2. One of the following conditions is required.
 - The device should not be a shared device. (Refer to section 13, Multiframe Operations.)
 - If the device is shared, a global unload should be set.
3. The device must be declared logically off. Use the OFF DSD command or the OFF parameter in the EQPDECK EQ entry.
4. The disk status display (E,M) must show that the device is not in use.
5. One of the following conditions is required.
 - The disk status display (E,M) must show that the device is unavailable for permanent file access.
 - The following conditions are required.
 - The full initialize status and the format pending status must be set. Use the DSD command, INITIALIZE,FP,est.
 - The direct access file user count should be equal to zero. The family status display (E,F) gives this information.

If all the necessary conditions are satisfied, FORMAT is able to access the pack for reformatting purposes. FORMAT repeatedly checks to ensure that these conditions are satisfied throughout the FORMAT operation.

At the end of the FORMAT process the disk must be initialized (INITIALIZE,AL,est) to copy the flaw information to the TRT and label. Otherwise, the flaw map changes will have no effect on NOS until the next time the disk is recovered. At that time, unreserved tracks marked as flaws in the flaw map will be flawed in the appropriate NOS tables.

NOS/VE Address Formats

H

When doing an analysis of a NOS/VE dump tape, you can specify NOS/VE addresses to the DSDI utility in any of several formats. These formats allow a flexible specification of the address range in which you are interested.

NOTE

To examine NOS/VE memory, the EDD dump file must contain all central memory (not just critical memory), and the FULL option must be specified on the DSDI command.

Address Format	Description
nn.n	A real memory address (RMA).
asid#nn..n	A system virtual address (SVA).
seg#nn..n#exch	A process virtual address (PVA) based on the exchange package that contains the address.
reg#exch	The register and exchange package that contains the address.
nn..n#reg#exch	A PVA formed from the segment field of the specified register pointed to by the specified exchange package. In this format nn..n is the byte offset from the beginning of the segment.
PVA#exch	The address in the pseudo register PVA of the specified exchange package.
exch	The address as the beginning of the specified exchange package.

The number sign (#) character is used in address formats to separate parts of the address parameter. Commas cannot be used because they delimit parameters within DSDI directives.

The symbols used for the various parts of the address formats are described as follows:

Symbol	Description
nn..n	A number (octal, decimal, or hexadecimal as required by the directive) of from 1 to 8 digits.
seg	A hexadecimal number of from 1 to 3 digits specifying a segment number.
asid	A hexadecimal number of from 1 to 4 digits specifying an actual segment identifier.

Symbol	Description												
reg	Any one of the following registers:												
	<table border="1"> <thead> <tr> <th>reg</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>P</td> <td>Specifies the processor P register.</td> </tr> <tr> <td>RA0 to RAF</td> <td>Specifies the corresponding address register 0 to F hexadecimal.</td> </tr> <tr> <td>RN1 to RN15</td> <td>Specifies the top of stack address for the corresponding ring 1 to 15.</td> </tr> <tr> <td>UTP</td> <td>Specifies the address in the untranslatable pointer register.</td> </tr> <tr> <td>TP</td> <td>Specifies the address in the trap pointer register.</td> </tr> </tbody> </table>	reg	Description	P	Specifies the processor P register.	RA0 to RAF	Specifies the corresponding address register 0 to F hexadecimal.	RN1 to RN15	Specifies the top of stack address for the corresponding ring 1 to 15.	UTP	Specifies the address in the untranslatable pointer register.	TP	Specifies the address in the trap pointer register.
reg	Description												
P	Specifies the processor P register.												
RA0 to RAF	Specifies the corresponding address register 0 to F hexadecimal.												
RN1 to RN15	Specifies the top of stack address for the corresponding ring 1 to 15.												
UTP	Specifies the address in the untranslatable pointer register.												
TP	Specifies the address in the trap pointer register.												
exch	Either a hexadecimal number of from 1 to 6 digits specifying an RMA of an exchange package (used to get the register values, segment table address, and segment table length needed to change the real address into a PVA), or one of the following keywords:												
	<table border="1"> <thead> <tr> <th>exch</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>MPS</td> <td>Address pointed to by the monitor process state register.</td> </tr> <tr> <td>JPS</td> <td>Address pointed to by the job process state register.</td> </tr> <tr> <td>RMA</td> <td>Real memory address calculated using the SETRMA DSDI directive.</td> </tr> <tr> <td>PXP</td> <td>Specifies that the current processor exchange package be used to obtain the P register address.</td> </tr> </tbody> </table>	exch	Description	MPS	Address pointed to by the monitor process state register.	JPS	Address pointed to by the job process state register.	RMA	Real memory address calculated using the SETRMA DSDI directive.	PXP	Specifies that the current processor exchange package be used to obtain the P register address.		
exch	Description												
MPS	Address pointed to by the monitor process state register.												
JPS	Address pointed to by the job process state register.												
RMA	Real memory address calculated using the SETRMA DSDI directive.												
PXP	Specifies that the current processor exchange package be used to obtain the P register address.												

You can modify each address by adding or subtracting a hexadecimal number of bytes from the specified address. If the number added or subtracted is entered with no address, the last address specified is used as the base address.

The following illustrates some examples of valid address formats.

Example	Address Format
25E101	nn..n
MPS	exch
RA0#JPS-40	reg#exch
1000#RN1#MPS	nn..n#reg#exch
P#PXP+40	PVA#exch

Management Of Storage Media Defects

I

Detecting Failing Devices	I-1
Disk Idle Mode	I-2
Mass Storage Thresholds	I-2
Dealing With Known Media Defects	I-3
Keeping An Accurate Record of Media Defects	I-3
Preventing The System From Using Flawed Disk Space	I-3
Media Management Tools	I-4
Error Log Messages for 887, 895, and 9853 Disks	I-4
FLAW Utility	I-4
APRDECK Entries	I-5
EQPDECK Entries	I-5
FORMAT Utility for 881/883 Disk Packs	I-5
Media Defect Example	I-6
Releasing Special Files	I-7
Reinitializing A Device	I-7

This appendix describes the process of detecting failing devices that are suspected of having storage media defects and the process of dealing with known storage media defects.

Detecting Failing Devices

This aspect of storage media management is the process of detecting failing mass storage devices and making informed decisions regarding the proper course of action. The action may be to place a failing disk in idle mode, change its threshold values, or set its device state to OFF or DOWN. These actions are discussed in the following paragraphs.

When a mass storage device or channel begins to fail, the system directs new file assignments away from the device and inhibits access to files on the device that are being used by executing jobs.

The operator is alerted to the device failure through the A,OPERATOR display. Also the E,E and E,H displays provide an overview of the mass storage devices that have abnormal conditions present. The E,E display shows that a device error has occurred on a particular device and the current state of that device. The E,H display shows the current mass storage threshold values and the current count associated with each threshold value. Refer to the NOS Version 2 Operations Handbook for additional information about the E,E and E,H displays.

You can use the IDLE command in conjunction with the THRESHOLD command to maintain some degree of flexibility in deciding when to repair the failing device. For more information, refer to the IDLE and THRESHOLD commands in section 5, DSD Commands.

NOTE

The decision to continue using a failing device should be made under the advisement of a customer engineer. If there is a chance of data corruption due to the nature of the failure, it may be best to leave the device OFF or DOWN until it can be repaired.

Disk Idle Mode

When a disk is in idle mode, the system temporarily eliminates access to the device when it starts to fail. Nonsubsystem jobs accessing the device will be rolled out so that the rest of the system and jobs not accessing the device can continue running. A verification algorithm will be used to attempt to isolate the problem to a particular part of the hardware (media, drive, controller, or channel). If the cause of the problem can be isolated to a failing channel or drive, then the system will attempt to DOWN that particular hardware resource.

Mass Storage Thresholds

The initial threshold values for a mass storage device are determined by THRESHOLD EQPDECK entries in the deadstart file. You can use the DSD THRESHOLD command to change the following threshold values for any mass storage device:

- Verification failure threshold
- Restricted activity threshold
- Low space threshold
- Recovered error threshold
- Unrecovered error threshold

These threshold values are described under the DSD THRESHOLD command in section 5.

When the space available on a mass storage device falls below the low space threshold, the following message appears on the A,OPERATOR display:

```
LOW SPACE ON MASS STORAGE DEVICE.
```


Dealing With Known Media Defects

This aspect of storage media management is the process of dealing with known media defects on mass storage devices. This process includes keeping an accurate record of all known defects and preventing the system from using defective areas. You should follow these general steps when you or the system encounter a media defect:

1. Record the media defect in a logbook.
2. Use the FLAW utility to reserve (flaw) the defective area and prevent the system from using that space.
3. Modify or create an APRDECK for the device containing the defect to include a command to flaw the defective area.
4. Modify the EQPDECK entry for the device to include a reference to the APRDECK entry.

This appendix describes general procedures for storage media management and the software tools you will need to perform these procedures. Later in this appendix an example shows the specific steps you should perform given a defined media defect problem.

Keeping An Accurate Record of Media Defects

Your site should maintain a logbook that documents all of the known mass storage media defects. This logbook should be kept by the system console so that an operator can record any media defect messages displayed. The information can then be used to reserve defective areas and thus prevent the system from using them.

Preventing The System From Using Flawed Disk Space

NOS allocates space on a mass storage device in units of logical tracks. For example, when you write a file to disk, NOS assigns a logical track to the file and then writes the file information to that logical track. The logical track remains assigned to the file until you release the file. That is, as long as the logical track is assigned to your file, it cannot be assigned to another file.

To prevent NOS from assigning a defective logical track to a file, you must reserve the track. This is called **flawing a track**. However, you cannot flaw a track while it is in use; that is, if the logical track with the defect is currently assigned to a file, you cannot flaw the track until the file is released.

If a media defect is encountered on a logical track assigned to a type of file that will not be released (such as, the system file, dayfiles, permanent file catalogs, and system checkpoint files) you will have to perform special procedures to explicitly move these files and then flaw the track. For example, you may need to dump and reload permanent files, reinitialize a disk, or deadstart the system. For a discussion of dealing with media defects encountered on areas assigned to these files, refer to **Releasing Special Files and Reinitializing a Device** later in this appendix.

Media Management Tools

You can use the following software tools to help with storage media management:

- Error log messages for 887, 895, and 9853 disks
- FLAW utility
- APRDECK entries
- EQPDECK entries
- FORMAT utility for 881/883 disk packs

Error Log Messages for 887, 895, and 9853 Disks

NOS automatically detects the presence of a media defect on an 887, 895, or 9853 disk. NOS automatically flaws the logical track that contains the media defect when the file containing the logical track is released.

When NOS encounters a media defect on an 887, 895, or 9853 disk, it issues a SEE A,ERROR LOG message to the operator. The operator should then use the A,ERROR LOG command to view the error log screen. The following message describes the media defect encountered:

```
dt est,TKnnnn,MEDIA DEFECT ENCOUNTERED
```

where dt is the device type, est is the EST ordinal of the device, and nnnn is the number of the logical track containing the media defect.

The operator should copy this message to a logbook for future use. This message indicates that NOS has detected a media defect; it does not indicate that the defective area has been automatically flawed. When the file assigned to the logical track containing the defect is released, NOS will then automatically flaw the track.

However, because media defects are rare (they should occur only once per device every two years), NOS stores only one media defect; the last defect encountered. Thus, your operator should always log the media defects encountered so that the information will not be lost if multiple defects occur.

When NOS automatically flaws a logical track, it again issues a SEE A,ERROR LOG message to the operator and displays this message on the error log screen:

```
dt est,TKnnnn,MEDIA DEFECT FLAWED
```

The operator should copy this message to a logbook for future use.

FLAW Utility

If a media defect occurs during normal system operation, use the FLAW utility to reserve (flaw) the defective track. (You can also use the FLAW utility to clear any flaws previously set.) To flaw a defective logical track, you must know the EST ordinal of the device containing the defect and the logical track number of the defect. For more information, refer to FLAW K Display in section 8, K-Display Utilities.

APRDECK Entries

In addition to using the **FLAW** utility to flaw a defective area of mass storage, you should also modify the **APRDECK** (auxiliary mass storage parameter deck) for the device containing the defect.

An **APRDECK** is a text record on the deadstart file that is used when the device is initialized. **APRDECK** entries identify areas of mass storage that are unusable (flawed) and prevent the system from accessing them. The system uses the information in the **APRDECK** entries to build the track reservation table (**TRT**) that resides in central memory resident (**CMR**) and also in the mass storage device label for each device.

By making **APRDECK** entries, you cause the system to flaw the media defects without your intervention during the next deadstart initialization of the device. For more information, refer to **APRDECK** in section 3, **Deadstart Decks**.

EQPDECK Entries

An **EQPDECK** (equipment deck) is a text record on the deadstart file that declares all the devices that comprise your computer system. Each device is defined to the system with an **EQPDECK** entry. An optional parameter on the **EQPDECK** entry declares an **APRDECK** for the device. When you create an **APRDECK** entry for a device, modify the **EQPDECK** entry for the device to include the reference to the **APRDECK** entry. For more information, refer to **EQPDECK** in section 3, **Deadstart Decks**.

FORMAT Utility for 881/883 Disk Packs

When you use the **FLAW** utility or an **APRDECK** entry to flaw an area of an 881 or 883 disk pack, this flawing is only logical. That is, the flaw information is kept in the **TRT** and is discarded on subsequent deadstart initializations of the device. This information is also lost during online initializations if it was not possible to recover the equipment.

For 881 and 883 disk packs, which contain factory-recorded utility maps, you have the option of storing flaw information on the utility map on the disk. To modify or add information to the utility map, use the **FORMAT** utility. For more information, refer to appendix G.

Media Defect Example

Suppose a job encounters a media defect on a local file. The file is located on an 895 disk device with EST ordinal 15 and the bad track is 5010. Use the following procedure to record the media defect and flaw the defective area.

1. When the system displays the message SEE *A,OPERATOR* use the A,OPERATOR command to view the operator screen. This message is displayed:

```
13 SEE A,ERROR LOG.
```

2. To view the error log, enter this command:

```
A,ERROR LOG.
```

The following message is displayed on the error log screen:

```
EQ015,TK5010,MEDIA DEFECT ENCOUNTERED.
```

3. Record the information in the media defect logbook.
4. Enter this command to tell the system you are taking care of the problem:

```
LOG,13.
```

5. When the job using the local file releases the file, the following message is displayed on the error log screen:

```
EQ015,TK5010,MEDIA DEFECT FLAWED.
```

6. Record the information in the media defect logbook.
7. Suppose there is no APRDECK for the device with EST ordinal 15. Add the AP=nn parameter to the EQPDECK entry for device 15. Here is an example of what the EQPDECK entry should look like:

```
EQ015=type,ST=ON,CH=ch,AP=00.
```

This entry assigns the first APRDECK to device 15.

8. Create the APRDECK for device 15:

```
APRD00.  
SLF=5010.
```

9. Use the LIBEDIT command to place the APRDECK entry and the updated EQPDECK entry on the deadstart file. (Remember that the number of the APRDECK record is determined by its position in relation to the APRINST record and not by the number on the record name; it is a good idea to name the records according to their position.)

Releasing Special Files

Because a media defect cannot be flawed automatically by the system until the file assigned to the defective track is released, you may have to release certain file types:

File	Description
Direct access file	When a flaw is detected on a track assigned to a direct access file, purge the file so that the system can perform automatic track flawing for the bad track.
Dayfile	Job dayfiles are released at job termination; system dayfiles are released when they are initialized or if you perform a DFTERM and then purge the files created by DFTERM.
System file	The system file is released at deadstart. Thus, at deadstart a defective track assigned to the system file will be flawed automatically.

Reinitializing A Device

Some track chains are not released unless the device is reinitialized:

- Indirect access file chain.
- Disk deadstart file.
- Catalog track chain.

If a storage media defect is encountered in one of these files, perform the following steps:

1. Create or update an APRDECK for the device.
2. Modify the EQPDECK entry for the device to include the APRDECK reference.
3. Dump the files on the device to tape.
4. Reinitialize the device.
5. Reload the files to the device.

Display Disk File (DDF) Utility

J

Bringing Up DDF	J-1
DDF Display Left Screen	J-2
DDF Display Right Screen	J-3
Keyboard Input	J-4
Console Messages	J-5
DDF Commands	J-6
Print Sector Data Commands	J-11
Change Sector Data Commands	J-11
DDF Examples	J-12
Pack Recovery	J-16
Pack Recovery Display Changes	J-16
Recovery Table Format	J-16
Pack Recovery Commands	J-19
Pack Recovery Hints	J-21

Display Disk File (DDF) Utility

J

The Display Disk File (DDF) utility provides the capability to read, display, change, and print the contents of physical disk sectors. It is intended to be used by analysts as an aid when working on system problems that involve data stored on mass storage devices.

CAUTION

Changing disk table contents should be done carefully and only by analysts who understand the effects of the changes. Unless extreme care is exercised, use of DDF can lead to system hangs or loss of permanent files.

Bringing Up DDF

The DSD display for the DDF utility is brought up on the system console by entering the following command:

X.DDF.

NOTE

Since DDF allows access to all disk space without regard for security access levels and categories, use of DDF in a secured system requires that the system be in SECURITY-UNLOCK state (refer to Secured System Control Commands in section 5, DSD Commands).

DDF Display Left Screen

The left screen displays the following disk and control point information:

- Disk EST ordinal and equipment type.
- Current disk track and sector address.
- Central memory address and byte number of the TRT link byte for the current track. The link byte is intensified if it is not consistent with the sector linkage bytes.
- TRT reservation, interlock, and preserved file status for the current track.
- Family name, pack name, control point number, and JSN of the associated job.
- Track numbers for the first track of the permit chain and the first track of the indirect access file data chain for the specified family/pack.
- First and second control point messages, if issued. The second control point message is usually a disk error and is intensified.
- Contents of DDF scratch areas, labeled A through F (refer to the DDF STORE command).
- The recovery track pointer and track recovery word, when pack recovery is in progress.
- Various status and errors may be displayed at the bottom of the screen.

DDF Display Right Screen

The right screen displays the contents of the specified disk sector. The byte numbers displayed in parenthesis after the central memory address are the actual byte numbers for the physical disk sector. The byte numbers are used in the memory changing commands.

You can change the right screen format with the right blank key on the CC545 console or the tab key on the CC634B console.

The following right screen displays are available:

- **Five Bytes Per Line.** The contents of the sector are displayed in 5 bytes per line. The CC545 console displays one half sector per page and uses the / to toggle between the first and second half of the sector. The CC598B and CC634B consoles display one quarter of a sector per page and uses the / to increment the addresses.
- **Ten Bytes Per Line.** The contents of the sector are displayed in 10₈ bytes per line with byte numbers. The CC545 console displays the entire sector on one page. The CC598B and CC634B consoles display one half sector per page and uses the / to toggle between the first and second half of the sector.
- **Text Mode.** The contents of the sector are displayed in text mode. The CC545 console displays the entire sector on one page. The CC598B and CC634B consoles display one half sector per page and uses the / to toggle between the first and second half of the sector.
- **Interpreted mode.** The data is formatted in one of three ways depending on the linkage bytes:
 - **System Sector.** The file name table (FNT) word and file status table (FST) word of the file are displayed along with the date and time when the sector was last updated. If the file type is PMFT, the copy of the permanent file catalog entry (PFC) contained in the system sector is also displayed.
 - **PFC Entry or Data Sector.** The contents of one permanent file catalog entry (20₈ words) are displayed. Some fields such as dates are decoded. The / is used to advance to the next PFC entry in the sector.
 - **End of Information Sector.** This display shows a short format of the system sector display.
- **Legal Commands.** This display shows a list of the DDF commands.

Keyboard Input

You can use the following keys to interact with the DDF display.

Key	Action Initiated
+	Read the next sector. If positioned at EOI or the end of the track, the current sector is reread.
-	Read the previous sector. If positioned at the beginning of the track, the current sector is reread.
.	Advance to the track specified in the control bytes and set the sector to 0. If the control bytes do not contain a track link, the current sector is reread.
,	Advance to the next track in the TRT chain. The sector number is not changed. If currently at the last track in the chain, the current sector is reread.
=	Read the next sector, continuing past EOI. If positioned at the end of an EOI track, the current sector is reread.
CR or NEXT	Initiate processing of an entered command. This key also sets REPEAT ENTRY if a complete command has not yet been entered.
Space bar	Read the current sector into the buffer. The current sector is the sector to which the current EST ordinal, track, and sector values point.
*	Toggle between DSD and DDF.
8	Increment the track number by one and read the sector.
9	Decrement the track number by one and read the sector.
Right blank, →, or Tab	Change the right screen displays on the CC545 console. This function is done by the Tab key on the CC598B console and the → (tab) key on the CC634B console.
/	Toggle or increment the sector displays on the right screen.
BKSP, ←, or Back Space	Delete the previous character typed on the CC545 console. This function is done by the Back Space key on the CC598B console and by the back arrow ← (backspace) key on the CC634B console.
Left blank, ←, or Esc	Delete the current line being entered on the CC545 console. This function is done by the Esc key on the CC598B console and the ← (back tab) key on the CC634B console.

Console Messages

The following messages may appear at the control point.

Message	Description
WAITING FOR MEMORY	DDF is waiting for central memory to store the display and command processors.
READING EQxxx TPxxxx CTxxxx	This message shows the status of checkout during a pack recovery operation.

The following messages may appear above the command line at the bottom of the left screen.

Message	Description
FORMAT ERROR.	A format error has been detected during translation of the entry.
INCORRECT ENTRY.	The command is not valid.
INCORRECT EQUIPMENT.	EST is not mass storage, nor is it a null equipment.
INCORRECT PARAMETER.	The parameter in the entry is invalid or too long.
REPEAT ENTRY.	The entry will not be cleared after execution.
SYSTEM BUSY.	DDF is waiting for the system to process a request.
BYTE xxxx.	The data specified is in byte number xxxx.
LENGTH = nnnn.	The SKIPEI command has skipped nnnn octal sectors.

Some commands read several sectors. If a disk error is encountered while processing one of these commands, processing stops and an error code appears on the left screen.

DDF Commands

All DDF commands are displayed on the bottom of the left screen as they are entered. When you enter a command, it must be followed by a carriage return before any action is taken.

DDF commands are processed interpretively. After the period (.), any alphanumeric characters can be entered: however, the message **FORMAT ERROR** or **INCORRECT PARAMETER** will appear (after the carriage return is entered) if you make an error.

Most DDF commands do not actually cause a sector to be read and displayed. After entering a command that changes the EST ordinal, track, or sector, you must enter a SPACE key to read the sector. This allows data from one sector to be read and then written to another sector.

After entering a command that changes the equipment or track, the track and sector numbers are checked. If either the track or sector number is invalid, it will be set to zero.

The **FIND** commands scan the catalog track starting at the current position and search for the specified catalog entry. The EST ordinal and track number of the catalog track must first be set using the appropriate commands (**FAMILY**, **PACKNAM**, **UI**, etc.). The first sector of the catalog track must also be read, by entering the space bar, before the **FIND** command is used.

Command	Description
AUTOREAD.nnnn.	Read the sector every nnnn seconds. If nnnn is not entered, 1 is used. AUTOREAD is terminated by clearing the command entry with the left blank key or back tab key.
BLDEOI.	Creates an EOI sector in the buffer.
BOT.nnnn.	Back up one track. Search the TRT starting at track number nnnn for a track that points to the current track. If one is found, the current track is set to that value. If no track is found that points to the current track, the message INCORRECT PARAMETER is displayed. If nnnn is not entered, the search starts at track 0.
CTB.¹	Clear the track interlock bit for the current track and equipment.
DEP.	Disable error processing for calls to the mass storage driver to read a sector. (By default, error processing is disabled.)
DIS.	Drop the DDF display and call DIS to the control point.
DROP.	Drop the DDF display and PP.
DTK.¹	Drop tracks to the end of the chain starting with the current track.

1. The keyboard must be unlocked to use this command (refer to the **UNLOCK** command in section 5. DSD Commands).

Command	Description
DTK.ssss. ²	Drop tracks starting with the current track and set the EOI sector in the TRT to the value ssss.
E.EP.	Enable error processing for calls to the mass storage driver to read a sector. (By default, error processing is disabled.) If a disk error is encountered when reading or writing a sector with error processing enabled, the screens will not be refreshed until the driver has finished performing error processing. This may result in the screens being blank for many seconds.
EJT.ejt.	Enter the disk address from EJT ordinal ejt ($0 \leq \text{ejt} \leq \text{largest EJT ordinal}$). The equipment, track, and sector are set to the beginning of the file. The message INCORRECT PARAMETER is displayed if the EJT entry is not used or the file does not have any tracks assigned.
EST.est.	Enter the EST ordinal est. The equipment must be a mass storage device.
FAMILY.familyname.	Use permanent file family familyname. The family is set internally in DDF. PFCW in the control point area is not changed. The family is initially set to the family name specified in PFCW.
FIND.pfn.userindex.	Search for permanent file pfn with user index userindex.
FIND.pfn..	Search for permanent file pfn with a zero user index (i.e., a hole).
FIND.pfn.	Search for permanent file pfn with any nonzero user index.
FIND..userindex.	Search for any permanent file with user index userindex. The file name is ignored.
FIND...	Search for any permanent file with a zero user index. The file name is ignored (i.e., search for any hole).
FINDISS.userindex.	Starting at the current position, search for the next system sector with user index userindex on the chain. If userindex is not specified, the user index is not checked and the next system sector is displayed. This command is intended to be used on the indirect access permanent file chain. Set the EST ordinal, track, and sector to point to a position within the indirect chain before entering this command.

². The keyboard must be unlocked to use this command (refer to the UNLOCK command in section 5. DSD Commands).

Command	Description
FINDO.octalnum.	Search from the current position for the octal number specified. The number can be 1 to 20 digits and is right justified in as many bytes as necessary to hold the number. The search starts on a byte boundary and the byte number is displayed if the search is successful. If the first part of the number is found at the end of the buffer, the search terminates even though the entire number is not found. If the carriage return is entered again, the search continues at the location of the first match.
FINDS.string.	Search from the current position for the specified string of characters. If the first part of the string matches the end of the buffer, the search terminates even though the entire string is not found. The byte number of the beginning of the string is displayed. If the carriage return is entered again, the search continues at that point.
FINDSS.userindex.	Starting from the current track and searching to the end of the TRT, find and display the next system sector with user index userindex. If userindex is not specified, the user index is not checked.
FNT.fnt.	Enter disk information from the system FNT ordinal fnt ($0 < \text{fnt} \leq \text{largest FNT ordinal}$). The EST ordinal, track, and sector are set to the beginning of the file. The message INCORRECT PARAMETER is displayed if the FNT entry is not used or the file does not have any tracks assigned.
FNTL.fnt.	Enter the disk address from the local FNT ordinal fnt ($0 \leq \text{fnt} \leq \text{largest FNT ordinal}$). The EST ordinal, track, and sector are set to the beginning of the file. The message INCORRECT PARAMETER is displayed if the FNT entry is not used or the file does not have any tracks assigned.
FNTLC.fnt.	Enter the disk address from the local FNT ordinal fnt ($0 \leq \text{fnt} \leq \text{largest FNT ordinal}$). The EST ordinal, track, and sector are set to the current position of the file. The message INCORRECT PARAMETER is displayed if the FNT entry is not used or the file does not have any tracks assigned.
GETTRT.nnnn.	This command is intended to be used on the label track. The EST ordinal and track must be set for the label track before the command is used. GETTRT reads the sector that contains the checkpoint TRT information for track number nnnn and displays the byte number within the sector for that track.
HOLD.	Release the display and wait for the operator to reassign the display.
LOAD.c.	Reload the EST ordinal, track, sector, and display selection from scratch area c. The scratch area is specified by a single alphabetic character. Valid characters are A through F. (Refer to the STORE command.)

Command	Description
PACKNAM.packname.	Use permanent file pack packname. The pack name is set internally in DDF. PKNW in the control point area is not changed. The pack name is initially set to the name specified in PKNW.
PREAD. ³	Read the current sector using the read protected sector function. This command is used instead of the space bar when reading protected sectors.
PTK.nnnn. ³	Enter the protected track number nnnn. This command works the same as the TK command except the track number entered is not checked.
PWRITE. ³	Write the current sector using the write protected sector function. This command should be used instead of the WRITE command when writing protected sectors.
QFT.qft.	Enter the disk address from QFT ordinal qft ($0 \leq \text{qft} \leq \text{largest QFT ordinal}$). The EST ordinal, track, and sector are set to the beginning of the file. The message INCORRECT PARAMETER is displayed if the QFT entry is not used or the file does not have any tracks assigned.
RANDOM.nnnn.addr.	Set the track and sector for random address addr using number nnnn as the first track. The current track is used as the first track if nnnn is not specified. The message INCORRECT PARAMETER is displayed if the random address is not on the chain.
RANDOM..addr.	Set the track and sector for random address addr using the current track as the first track. The message INCORRECT PARAMETER is displayed if the random address is not on the chain.
RANDOM.c.addr.	Set the track and sector for random address addr using the track specified by scratch area c as the first track. The scratch area is specified by a single alphabetic character. Valid characters are A through F. (Refer to the STORE command.) The message INCORRECT PARAMETER is displayed if the random address is not on the chain.
SC.nnnn.	Enter the sector number nnnn. The message INCORRECT PARAMETER is displayed if the sector number is too large.
SC.*.	Enter the sector number of the last sector on the track.
SCAN.	Scan from the current position until the EOI control bytes are encountered. Scanning stops if, at some point, the control bytes are incorrect or the track is not reserved.

3. The keyboard must be unlocked to use this command (refer to the UNLOCK command in section 5, DSD Commands).

Command	Description
SCAN.*.	Scan from the current position until the end of information indicated in the TRT is reached. This is intended to be used for the indirect access permanent file chain. Note that when PFM delinks a track in the middle of the chain, it does not update the track pointer in the preceding track; which causes SCAN to stop at this point and display an error.
SHOWPF.	Display the permanent file whose catalog entry is currently displayed. The PFC display must be on the right screen to use this command. The device, track, and sector from the PFC currently displayed are used to display the file. If the device is not present, the message INCORRECT PARAMETER is displayed.
SKIPEI.	Set the track and sector to EOI based on the current position and information in the TRT. The number of sectors skipped is displayed on the left screen.
SKIPF.	Read the file starting at the current position until an EOF is encountered.
SKIPR.	Read the file starting at the current position until an EOR or EOF is encountered.
STB. ⁴	Set the track interlock bit for the current track.
STORE.c.comment	Store the current EST ordinal, track, sector, and display selection into scratch area c. The scratch area is specified by a single alphabetic character (valid characters are A through F). The comment is copied to the scratch area and displayed on the left screen. The comment is for convenience only and is truncated after 10 characters. The EST ordinal, track, sector, and display selection can be reloaded from the scratch area using the LOAD command.
TK.nnnn.	Enter the track number nnnn. The message INCORRECT PARAMETER is displayed if the track number is too large.
UI.userindex.	Set the EST ordinal and track for the catalog entries for user index userindex (userindex ≤ 377777). The message INCORRECT PARAMETER is displayed if the catalog is not found (this may be the result of entering the wrong family name or pack name). If either the family name or pack name is changed after entering the UI command, the UI command must be reentered.
WRITE. ⁴	Write the contents of the buffer to the sector currently displayed.

⁴. The keyboard must be unlocked to use this command (refer to the UNLOCK command in section 5, DSD Commands).

Print Sector Data Commands

The following commands are used to print the contents of the disk sector and manipulate the listing file. These commands use a CPU resident helper program called DDFILE. If DDF is unable to load DDFILE (due to system activity or other constraints), the use of these commands will be disabled.

Command	Description
LISTING.filename.	Set the listing file name to filename. The default listing file name is OUTPUT.
OUT.	Release the listing file to the output queue.
PRINT.	Print the current equipment type, track, sector, TRT information, and the contents of the sector.
PRINT.num.	Print the next num (octal) sectors starting with the current sector. Printing stops if EOI is encountered before num sectors have been printed. One page of output is produced for each sector printed.
RETURN.	Return the listing file.
REWIND.	Rewind the listing file.
SETID.id.	Set the identifier for the listing file to id.
SKIPL.	Advance from the current position of the listing file to the end of the file.

Change Sector Data Commands

The following commands are used to change the data in the current sector. If the comma (,) is replaced by a plus (+), the byte number xxxx is incremented after the entry is processed. The actual data on the disk is not changed until you enter the WRITE command.

Command	Description
xxxx,yyyy.	Enter value yyyy into byte number xxxx.
xxxx,Dcc.	Enter display code characters cc into byte number xxxx.

DDF Examples

The following examples illustrate DDF usage.

Example 1:

Read the system sector for permanent file ABC under user index 1234 on family XYZ.

X.DDF.	Bring up the DDF display.
FAMILY.XYZ.	
UI.1234.	Set the EST ordinal, track, and sector of the first catalog entry on this catalog track.
Space bar	Read the first catalog sector.
FIND.ABC.1234.	Search the catalog track for the PFC entry.
SHOWPF.	Set the EST ordinal, track, and sector of the file.
Space bar	Read the system sector of the file.

Example 2:

Look at the PFC entries under user index 1234 on family XYZ.

X.DDF.	Bring up the DDF display.
FAMILY.XYZ.	
UI.1234.	Set the EST ordinal, track, and sector of the first catalog entry on this catalog track.
Space bar	Read the first catalog sector.
FIND..1234.	Search for the first file under the user index 1234.

Example 3:

While using DIS, look at the contents of a local file. Be sure to remember the FNT ordinal fnt of the file.

DDF.	Switch control from DIS to DDF.
FNTL.fnt.	Set the EST ordinal, track, and sector to the system sector of the file.
or	
FNTLC.fnt.	Set the EST ordinal, track, and sector to the current position of the file.
Space bar	Read the system sector of the file (if FNTLC was entered, read the sector to which the file is currently positioned).
Use +, -, and /	Look at different parts of the file.
DIS.	Switch control from DDF to DIS.

Example 4:

Display a system file (e.g., the VALIDUS file). Look at the DSD H display to determine the FNT ordinal fnt.

X.DDF.	Bring up the DDF display.
FNT.fnt.	Set the EST ordinal, track, and sector to the system sector of the file.
Space bar	Read the system sector of the file.
Use +, -, and /	Look at different parts of the file.

Example 5:

Determine the length of a file and display the EOI sector. Display the system sector of the file (see previous examples).

SKIPEI.	Set the track and sector to the file EOI and display the number of sectors skipped above the command.
Left blank	Clear the message and command.
Space bar	Read the EOI sector.

Example 6:

Verify that a file can be read and that the linkage bytes are correct. Display the first sector of the file (see previous examples).

SCAN.	Read every sector from the current position in the file to the EOI. If any errors are encountered, SCAN stops and displays a message. SCAN verifies that the track linkage bytes match the TRT.
-------	---

Example 7:

Print the contents of sectors. Display the first sector (see previous examples).

PRINT.num.	Print the next num sectors (num must be an octal number) to a listing file. The default for num is 1. The default for the listing file is OUTPUT unless the LISTING command has been used to specify another file name.
OUT.	Put the listing file in the output queue to be printed.
DIS.	File OUTPUT contains the listing file of sectors.

Example 8:

Often it is necessary to look at several different places on the disk. The DDF utility has a facility to remember disk locations.

X.DDF.	Bring up the DDF display. FAMILY.ABC.
UI.113.	Set the EST ordinal, track, and sector of the first sector of the catalog track.
Space bar	Read the first sector of the catalog track.
FIND.XYZ.113	Find the PFC entry for the file XYZ on user index 113.
STORE.A.XYZ PFC	Save the current EST ordinal, track and sector in scratch area A. Save the characters XYZ PFC as a comment. This information is displayed on the left screen.
SHOWPF.	Set the EST ordinal, track, and sector for the system sector of the file.
Space bar	Read the system sector of the file.
STORE.B.XYZ SS	Save the current EST ordinal, track, and sector in scratch area B. Save the characters XYZ SS as a comment.
SKIPFI.	Skip to the EOI sector of the file.
Left blank	
Space bar	Read the EOI sector of the file.
STORE.C.XYZ EOI	Save the current EST ordinal, track, and sector in scratch area C. Save the characters XYZ EOI as a comment.
LOAD.A.	Reset DDF to display the PFC entry.
Space bar	Read the PFC sector.
LOAD.B.	Reset DDF to display the system sector.
Space bar	Read the system sector.
LOAD.C.	Reset DDF to display the EOI sector.
Space bar	Read the EOI sector.

Example 9:

Change data on the disk. Display the sector to be changed. Determine from the display the byte number of the byte to be changed.

CAUTION

Be very careful and be sure you know what you are doing before changing data on the disk. Unless extreme care is exercised, use of DDF to change data on the disk can lead to system hangs or loss of permanent files.

*	Return control to DSD.
UNLOCK.	Allow privileged commands to be entered.
*	Return control to DDF.
XXXX,YYYY.	Set byte number xxxx to value yyyy. If a mistake is made, press the space bar to reread the sector and start over.
WRITE.	Write what is displayed to the disk.

CAUTION

Be sure another program does not change the sector between the time DDF reads it and writes it; otherwise, the changes may be lost. Depending on what you are changing, it may be best to make changes only when the system is idle. For changes to permanent files and catalog entries, the STB and CTB commands can be used to prevent other accesses to the sector; however, an understanding of permanent file interlocking is advisable before you use these commands.

Example 10:

Display the disk flaw map. First, determine (from the appropriate hardware manual) the logical track and sector of the flaw map. This is not currently operational for buffered devices.

X.DDF.	Bring up the DDF display.
EST.est.	Set the EST ordinal est.
*	Return control to DSD.
UNLOCK.	Allow privileged commands to be entered.
*	Return control to DDF.
PTK.nnnn.	Set the protected track number nnnn. DDF does not verify that the track number is valid.
SC.nnnn.	Set the sector number nnnn.
PREAD.	Read the sector using a protected read function. (Do not attempt to read the sector using the space bar; this would cause a disk error because this sector is protected.)

Pack Recovery

In the event that the label sector and TRT on a device have been overwritten and no suitable backup is available, a special set of DDF commands are available to allow the reconstruction of these tables. Note that this process requires extensive analyst intervention (on the order of several hours of dedicated time).

To effectively use these commands, an analyst must have a thorough and detailed understanding of the NOS mass storage and permanent file table structures, both in CMR and on disk. Recovery achieved by these commands is not automatic; the analyst will have to make manual changes to the label sector and TRT during the process of the reconstruction.

Refer to Pack Recovery Hints later in this section for an outline of the procedure to be used.


Pack Recovery Display Changes

When pack recovery is in progress, the recovery track pointer and the track recovery word are displayed on the left screen.

Recovery Table Format


The pack recovery process uses a recovery table in central memory to record linkage information. The table is initialized by the BEGINR command. The other commands use or modify the data in the table to effect device recovery. The table (which is stored within DDF's central memory field length) is formatted as follows:

- Image of MST for the device (MSTL words long).
- FNSS/FTSS/DTSS words from the system sector of the current chain.
- Recovery pointer word.




Field	Description
FT	First track of the current chain.
PT	Previous track in the chain.
SC	Current sector for the EOI search after an error.

- Indirect chain pointer word (not currently used).



Field	Description
FT	First track of the indirect chain.
PT	Previous track of the indirect chain.


- Build pointer word.



Field	Description
TT	Destination track for writing the TRT.
SC	Sector for the current piece of the TRT.
FT	First track for the current piece of the TRT.
NF	DAF file count.
TL	TRT length from the MST.

- Two reserved words.

- Track recovery table (one word per track).



Field	Description																														
FLAGS	<p>The following flags are valid:</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 10px;"> <thead> <tr> <th style="text-align: left;">Flag</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>TY</td> <td>59-54</td> <td>File type from system sector (for tracks beginning with a system sector).</td> </tr> <tr> <td>RE</td> <td>53</td> <td>Read error detected in track.</td> </tr> <tr> <td>SS</td> <td>52</td> <td>System sector in track at other than sector zero.</td> </tr> <tr> <td>FL</td> <td>51</td> <td>Flawed track indicator.</td> </tr> <tr> <td></td> <td>50</td> <td>Not used.</td> </tr> <tr> <td>EI</td> <td>49-48</td> <td>EOI status for track (0, 1, or 2).</td> </tr> <tr> <td></td> <td></td> <td style="padding-left: 20px;">0 No EOI.</td> </tr> <tr> <td></td> <td></td> <td style="padding-left: 20px;">1 Normal EOI in track.</td> </tr> <tr> <td></td> <td></td> <td style="padding-left: 20px;">2 EOI (written by IMS).</td> </tr> </tbody> </table>	Flag	Bits	Description	TY	59-54	File type from system sector (for tracks beginning with a system sector).	RE	53	Read error detected in track.	SS	52	System sector in track at other than sector zero.	FL	51	Flawed track indicator.		50	Not used.	EI	49-48	EOI status for track (0, 1, or 2).			0 No EOI.			1 Normal EOI in track.			2 EOI (written by IMS).
Flag	Bits	Description																													
TY	59-54	File type from system sector (for tracks beginning with a system sector).																													
RE	53	Read error detected in track.																													
SS	52	System sector in track at other than sector zero.																													
FL	51	Flawed track indicator.																													
	50	Not used.																													
EI	49-48	EOI status for track (0, 1, or 2).																													
		0 No EOI.																													
		1 Normal EOI in track.																													
		2 EOI (written by IMS).																													
FT	First track of the chain linked to this track.																														
PT	First track encountered that is linked to this track. (PT equals the current track if at the start of the file.)																														
NT	Next track or EOI sector number.																														
LC	Linkage count. (Number of tracks linked to this track.)																														

Pack Recovery Commands

The BEGINR command must be entered as the first command of a pack recovery sequence; until the BEGINR command is entered, none of the other pack recovery commands will be accepted. All of these commands require that the keyboard be unlocked.

Command	Description
BEGINR.	<p>Begin track recovery for the current selected equipment. The equipment to be recovered must be removable and unloaded. The MST in CMR for the selected equipment is used to initialize the MST for the label track built by the recovery commands. Be sure the equipment mnemonic in the E,M display matches that of the pack to be recovered. (The correct mnemonic can be set with the DSD INITIALIZE command.)</p> <p>The first use of this command initializes the recovery process by:</p> <ul style="list-style-type: none"> • Setting the recovery track pointer to zero. • Reserving and clearing the central memory (CM) recovery table. • Setting up the MST skeleton in the recovery table. • Enabling the use of the other recovery commands. <p>If the BEGINR command is used again, recovery restarts for the current equipment.</p>
BLDSL.	<p>Build a label sector image in the data buffer from data in the CM recovery table. Once the other recovery commands have completed, memory change commands can be used to fill in any missing data. The buffer can then be written to disk using the WRITE command.</p>
BLDSTRT.	<p>Build the next sector of the TRT data for the label track. The track and sector are set for the WRITE command. If any linkage errors are detected, the message ERROR IN TRACK STATUS is issued. After the last sector is built, the pointers are reset to the first sector again. You must manually generate the EOI sector for the label track using the BLDEOI command.</p>
BLDSTRT.nnnn.	<p>Build sector number nnnn of the TRT data for the label track (nnnn ≥ 1).</p>

Command	Description
NEXTAT.	Locate the next available track on the current equipment and read the first sector of that track. The search begins at the displayed recovery track pointer +1. By resetting the track pointer to track zero and using this command, unreserved tracks can be found and checked for being part of the catalog track chain or the indirect access file chain.
NEXTSS.	Locate the next track in the recovery table that begins with a system sector. The search begins at the displayed recovery track pointer +1.
RECOVER.RECOVER.c.	<p>Scan the entire device, saving linkage information in the track recovery table. The first sector of each track is read until a system sector is found. For non-SYFT system sectors and for SYFT files other than DATA, read each sector (following the sector linkage) until either an EOI sector, a multiple linkage, an embedded system sector, or a read error is encountered. For SYFT files LABEL, DATA, and PERMIT, set the ALGL track pointer in the MST skeleton.</p> <p>If the scan stops for any reason, it can be restarted by entering either the RECOVER or RECOVER.c command (where c is any character). Entering RECOVER.c resumes the linkage chain scanning from the sector after the one in which the scan stopped. Entering RECOVER causes DDF to go back to the last track on which a system sector was found, advance the track by one, and resume scanning the first sector of each track looking for another system sector. Once the RECOVER command has been used to scan the entire device (which may take several hours) the other commands can be used to resolve any linkage conflicts or other loose ends that remain.</p>
NOTE	
DDF releases the display to DSD while executing the RECOVER command to improve performance (minimize lost revolutions). The display is automatically rerequested when the scan stops for any reason.	
SETRW.b.nnnn.	Change byte b of the current track recovery word to number nnnn. This is the method used to resolve linkage conflicts before using the BLDSL and BLDSTRT commands.
SETTP.nnnn.	Set the recovery track pointer for the RECOVER, NEXTAT, NEXTSS, and SETRW commands to track number nnnn.

Pack Recovery Hints

The general process for recovering a pack is as follows:

Example	Explanation
BEGINR.	
RECOVER.	Repeat as many times as required.
SETRW.	Repeat as necessary to correct linkage conflicts and oversights.
BLDSL.	
WRITE.	
BLDSTRT.	Repeat until TRT is complete.
WRITE.	Repeat until TRT is complete.
SC.nnnn.	Set EOI sector address.
BLDEOI.	
WRITE.	

There are, however, several manual operations required in addition to the process outlined above.

- Several fields in the MST image in the label sector must be entered manually after the BLDSL command and before the WRITE command. These include:
 - Family/pack name.
 - Device number.
 - Device masks.
 - Number of catalog tracks.
- The catalog tracks are not automatically recovered. The NEXTAT command can be used to search for the catalog tracks after the device has been completely scanned by the RECOVER command. In general, the catalog tracks are allocated on a device immediately prior to the first track of the indirect chain. (The first track of the indirect chain can be found by checking byte 0 of ALGL in the MST, built by the BLDSL command.) Once the catalog tracks are found, the track linkage information for these tracks must be set in the track recovery words by using the SETRW command.

The label track must be linked to the first catalog track, which is linked to the second original catalog track, which is linked to the third original catalog track, and so forth. All overflow catalog tracks are linked after the last of the original catalog tracks.

There can be one or more overflow catalog tracks for each of the original catalog tracks. These overflow tracks are linked from their corresponding catalog track via the sector linkage in the last sector of that track. This means that the sector linkage does not correspond to the TRT linkage for these tracks.

- The indirect chain is not automatically recovered. Tracks in the indirect chain can be located using the NEXTAT command and the RECOVER.c or FINDISS commands. Tracks for which bit 52 is set in the recovery word (which indicates that a system sector was found in a sector other than sector zero on this track) are probably part of the indirect chain. Examining the sector linkage of the last sector of each track that has been identified as containing indirect access file data can be used to establish the ordering of these tracks. Note that if the last sector of an indirect chain track is a system sector, byte NSSS contains the pointer to the next track. Alternatively, the PFCs on the catalog tracks can be examined to determine the tracks that are on the indirect chain and the ordering of these tracks.
- Instances in which multiple tracks are linked to a single track in the track recovery table must be resolved in order to correctly build the TRT. Note that the BLDSTRT command will inform you of these conflicts. You must inspect the contents of the various tracks and make an empirical judgment as to the correct linkage.

The SETRW and SETTP commands must be used to eliminate the conflicting linkage. If a track is to be left unreserved, its track recovery word should be zeroed.

- Flawed tracks may be indicated by setting the flawed track indicator (FL flag) in the track recovery word for that track.

PACKER Utility

K

Operational Overview	K-1
Lost Space Processing	K-1
Hole Processing	K-2
Dayfile Statistics	K-3
Command Format	K-4
PACKER Examples	K-6

The PACKER utility enables you to manage holes within the indirect access permanent file chain (IAPF chain) on a permanent file device. This chain is the set of tracks that contains all the indirect access permanent files on a permanent file device. Space within the chain is allocated to individual files.

Holes are the spaces created within the IAPF chain when indirect access permanent files are deleted. These spaces may be reused later for other indirect access permanent files; however, over time these spaces tend to become fragmented into smaller and smaller pieces. Eventually, a large amount of space may be in small, unusable holes.

The PACKER utility reorganizes the IAPF chain so that more of the holes are available for storage. It combines adjacent holes, moves files to allow holes to be combined, and releases entire tracks when possible. These actions make more tracks available on the device (relieving a TRACK LIMIT condition, if one exists) as well as making the remaining holes larger and thus more usable.

Operational Overview

A single call to the PACKER utility only processes a single permanent file device, but you can process multiple devices at the same time by calling a different copy of PACKER for each device.

NOTE

PACKER sets the PF Utility Interlock while it is processing a device. This prevents any access to files on that device by PFDUMP, PFLOAD, or PFM.

PACKER begins processing a device by reading all the permanent file catalog entries (PFC entries) on the specified device and building a set of tables that allows it to operate on the IAPF chain in sequential order. PACKER then scans the IAPF chain from the beginning, searching for holes.

Lost Space Processing

If PACKER encounters any lost space during its scan, it attempts to reclaim the space for future use. Lost space is a condition where space exists on the IAPF chain that is not pointed to by a permanent file catalog (PFC) entry. Without a PFC entry, the permanent file manager (PFM) is unaware of the space. Lost space is usually the result of a previous PACKER run being interrupted by a system failure.

If the lost space is preceded by or followed by a hole, PACKER merges the lost space with one of the holes. If the lost space is preceded by a file and not followed by a hole, PACKER creates a new PFC entry for it, transforming it into a normal IAPF hole. If the lost space is smaller than the absolute minimum IAPF hole size of three PRUs, it cannot be reclaimed and is left as lost space.

Hole Processing

When PACKER encounters a hole during its scan, a sequence of events begins. Unless the collection hole is already open, the newly found hole is turned into the collection hole. The collection hole is the focal point of the major operations of combining holes and moving files. The PFC pointing to the collection hole is changed into a purged direct access permanent file (DAPF) PFC. This PFC is changed back into an IAPF PFC pointing to the hole when the collection hole is changed back into a normal hole; until that occurs, the collection hole is lost space.

PACKER uses this procedure to provide for system failures during PACKER operation. If such a failure occurs, the system is unaware of the changes PACKER is making, all files are intact, and the lost space is recovered the next time PACKER is run.

The following sequence is repeated until the collection hole is closed.

1. Any holes following the collection hole have their PFCs changed into purged DAPF PFCs and their space is added to the collection hole. This combining of contiguous holes is continued until PACKER encounters a file.
2. When a file is encountered, PACKER checks to see if the collection hole is big enough for tracks to be delinked from it. If so, it delinks and returns the tracks to the system. This premove delinking limits the amount of data (files) that are moved by limiting the size of the collection hole.
3. The utility looks for the largest file that fits into the collection hole by starting from the end of the IAPF chain and scanning toward the current position. If it finds such a file, PACKER copies the file to the beginning of the collection hole, updates its PFC to point to the new position, and updates the collection hole's starting address and length. (This is called the fill-move technique.)

If the collection hole is completely filled, PACKER considers it closed and begins scanning again. If the move is successful, PACKER repeats it until no further files can be moved using the fill-move technique.

4. The file adjacent to the collection hole (which was excluded from the fill-move search) is then examined to see if it can be slid across the collection hole. This file must completely fit within the collection hole, for the file would be destroyed if the file was copied on top of itself and the system failed during the copy.

PACKER examines the sizes of all the files until the next hole to ensure that they fit in the collection hole (if any do not fit, the effect of the moves would only be to move the position of the hole, which in itself is of no benefit).

If all the files fit, PACKER copies them, one by one, across the collection hole, and updates the address of the collection hole. (This is called the slide-move technique.) PACKER combines the newly adjacent hole with the collection hole and resumes scanning. If all the files between the collection hole and the next hole do not fit into the collection hole, PACKER closes the collection hole and scans for the next hole.

5. PACKER terminates the scan when it encounters the end of the IAPF chain. If the collection hole is open, PACKER drops it off the end of the IAPF chain by adjusting the EOI.

Dayfile Statistics

At termination, PACKER issues the following set of statistical dayfile messages. These messages document the original state of the device, changes made by PACKER, final state of the device, and resources used by PACKER.

ONSET FILES	NNNNNN,	SSSSSS	PRUS.
ONSET HOLES	NNNNNN,	SSSSSS	PRUS.
FILL MOVES	NNNNNN,	SSSSSS	PRUS.
SLIDE MOVES	NNNNNN,	SSSSSS	PRUS.
FILES MOVED	NNNNNN,	SSSSSS	PRUS.
LOST SPACE	NNNNNN,	SSSSSS	PRUS.
HOLES LEFT	NNNNNN,	SSSSSS	PRUS.
HOLES FREED	NNNNNN,	SSSSSS	PRUS.
TRACKS FREED	NNNNNN,	SSSSSS	PRUS.
FLAW SPACE	NNNNNN,	SSSSSS	PRUS.
SRUS		23456.890	UNTS.
CPU SECONDS		23456.890	SECS.
MAXIMUM CM FL		234567890	SCM.
MANAGED TABLE MOVES		234567890	UNTS.

Command Format

The PACKER command calls the utility that reorganizes the indirect access permanent files on a specified permanent file device. The command has the following format:

PACKER, P₁=f₁, P₂=f₂, . . . , P_n=f_n.

P_i	Description
DN=devicenumber	Device number. This parameter is required if a family device is to be processed. It cannot be used with the PN parameter.
FM=familyname	Family name. This parameter cannot be used with the PN parameter. If you omit both the FM and PN parameters, PACKER processes the system default family.
NM	No moves. If you specify this parameter, PACKER does not move any files; it limits its work to the combining of contiguous holes and the reclamation of lost space. Since this parameter greatly decreases the time required to run PACKER, you can use it when attempting to relieve a TRACK LIMIT condition during production hours.
PN=packname	Pack name. This parameter is required if the device to be processed is an auxiliary pack. This parameter cannot be used with the FM or DN parameter.
TL=timelimit	Time limit. If you specify this parameter, PACKER stops running after the specified number of seconds (wall clock time). If a value is specified, the system assumes it is octal. If the TL parameter is specified but not equivalenced (that is, TL is specified without a value), the system uses a limit of 300D.

The following command parameters are not needed for normal operations, but may be useful in special situations.

P_i	Description
EF	Exact fit. If you specify this parameter, an exact fit is required when moving a file with the fill-move technique. Use of this parameter might increase the number of files moved using the slide-move technique, which can result in more space being released to the operating system at the expense of moving an increased number of files.
IX=index	Index. This parameter specifies the starting index into PACKER's sorted PFC table. You can use it to skip the specified number of PFC entries before beginning processing. In some situations, this parameter provides a useful method of shortening the execution time of PACKER. If a value is specified, the system assumes it is octal. (The use of this parameter has an impact on the end-of-run statistics, since they do not reflect the portion skipped.)

P_i	Description
MR=moveratio	<p>Move ratio. This parameter specifies the maximum ratio between the total length of the files to be moved in a slide move and the size of the hole into which the files are to be moved. A value of 0 prevents all files from being moved using the slide-move technique; however, it is recommended that you use the NS parameter instead of MR=0. If a value is specified, the system assumes it is octal. If you omit the MR parameter, the system uses a value of 100B for the move ratio; if the MR parameter is specified, but not equivalenced, the system uses a value of 377777B.</p>
NC	<p>No changes. If you specify this parameter, PACKER determines what operations ought to be performed in response to the other specified parameters, but does not actually make any changes to the device.</p> <p>Since PACKER still issues dayfile messages at termination detailing the changes that would have been made, you may use this option to preview the potential effect of a parameter setting. You can compare alternative parameter combinations to determine the best way to reach a desired result.</p>
NF	<p>No fill moves. If you specify this parameter, PACKER does not move any files using the fill-move technique. Use of this parameter forces all files to be moved using the slide-move technique. This can result in more space being released to the operating system at the expense of moving an increased number of files.</p>
NP	<p>No premove delinks. If you specify this parameter, PACKER does not delink tracks before attempting to move files. Delinking tracks before moving files limits the size to which the collection hole is allowed to grow, which in turn limits the amount of data moved by PACKER in its attempts to fill the collection hole. These delinks might shorten the time PACKER executes at the expense of possibly leaving small holes where tracks were delinked.</p> <p>This parameter does not affect the delinking of tracks from holes that remain after PACKER has completed its work.</p>
NS	<p>No slide moves. If specified, PACKER does not move any files using the slide-move technique. This might increase the speed of PACKER execution by leaving more holes unfilled. You may also use the MR parameter to control slide moves in a more general fashion; however, to totally disable slide moves, you should use the NS parameter rather than specifying MR=0.</p>

PACKER Examples

The following examples illustrate how you can use the PACKER utility to perform periodic maintenance on permanent file devices.

Example 1:

For normal periodic maintenance, the PACKER utility can be executed when the production load is light and there is little interactive usage. For example, a site might want to use the following command every day in the early morning:

```
PACKER (FM=ffffff, DN=nn)
```

Example 2:

To get some space on a device that is at or near track limit, PACKER can be run quickly even during active production hours by specifying the NM parameter. Specifying this parameter reduces the amount of space reclaimed, but also greatly reduces the amount of time required to run PACKER. The following command may be entered:

```
PACKER (FM=ffffff, DN=nn, NM)
```

Example 3:

If a site determines that daily PACKER runs have too great an impact on production, the site can specify the NM parameter (as in example 2) when running PACKER on weekdays, and only do full PACKER runs (as in example 1) once a week.

Index

- * DIS command 4-17
- *. DSDI directive 6-9
- 533/536, 537, and 585 EVFU F-1
- 5380 Disk Array Subsystem 3-52
- 580 programmable format control E-1
- 5830 dual-channel access 3-53
- 6683 satellite coupler EST entry 3-37
- 7990 block header and trailer 11-69
- 7990 catalog
 - Analysis 11-44
 - Backup 11-83
 - Chain problems 11-88
 - Description 11-5
 - PFC mismatch 11-88
 - Problems 11-87
 - SM map mismatch 11-87
- 7990 control unit 11-1
- 7990 equipment configuration constraints 11-73
- 7990 files 11-11
- 7990 hardware components 11-1
- 7990 space management 11-79
- 7990 stripe header and trailer 11-68
- 7991 storage module 11-1
- 819 disk, buffer in LCME 3-52
- 834 disk, control module 3-65
- 836 disk, control module 3-65
- 844 disk drives 3-46
- 844 expander 3-46, 3-48
- 881/883 pack reformatting G-1
- 885-42 disk, buffer in ESM 3-52
- 887 disk, buffer in UEM 3-52
- 895 disk, buffer in UEM 3-52
- 9853 disk, buffer in UEM 3-52
- 9853 dual-channel access 3-53
- 99 DSD command 5-69

- A**
- ACCESS EQPDECK entry 3-94
- Access level limits 3-94
- Account dayfile 3-15
- ACCOUNT DSDI directive 6-24
- ACCOUNT EQPDECK entry 3-15
- ACN DSD command 5-74

- ACS
 - Components 22-6
 - Configurations 22-9
 - Constraints 22-16
 - Cross-coupled controllers 22-13
 - Description 22-6
 - Example 22-15
 - Maintenance 22-17
 - Performance considerations 22-17
 - Requirements 22-16
- ACS TAPE PF STAGING IPRDECK entry 3-137
- AIP trace with IAF 16-29
- ALLMEM DSDI directive 6-11
- Allocation information 3-45
- Alter memory file 15-6
- Alternate permanent file family 3-43
- Alternate storage 12-1, 17-2
- Alternate storage address 12-1
- Alternate system device 3-41
- Alternate system library entry 3-85
- AP DSDI directive 6-17
- APRDECK
 - Description 3-103
 - Display 2-21
 - Format 3-103
 - Modifying 2-21
- APRDECK entries
 - CAF 3-104
 - CLF 3-105
 - CPF 3-106
 - SLF 3-105
 - SPF 3-105
- APRINST display 2-21
- Archive file
 - Description 17-2
 - Multifile 17-2
 - Preassignment 17-61
- Arrays
 - Adding E-3
 - Building E-2
- ASR EQPDECK entry 3-85
- ASSIGN DSD command 5-24

ATF IPRDECK entry 3-136
 AUTO DSD command 5-56
 AUTOLOAD EQPDECK entry 2-4
 Automated Cartridge Subsystem (see ACS) 3-138
 Automated cartridge tape subsystem (see ACS) 22-6
 Automated Tape Facility (see ATF) 3-136
 Automated tape facility subsystem (ATF) 22-13
 Automatic label checking 3-151
 AUTORESTART IPRDECK entry 3-138
 Auxiliary device 3-44
 Auxiliary packs 3-136

B

B DSDI directive 6-24
 BATCHIO DSDI directive 6-39
 BCDUMP DSDI directive 6-51
 BEGIN DIS command 4-14
 Binary maintenance log 3-17
 BIO DSD command 5-50
 BIO DSDI directive 6-39
 BIO IPRDECK entry 3-136
 BKP DIS command 4-14
 BKP DSD command 5-76
 BKPA DIS command 4-14
 BKSP DSD command 5-24
 BKSPF DSD command 5-25
 BKSPRU DSD command 5-25
 Breakpoint package
 CPUMTR commands 5-80
 PP commands 5-75
 Breakpoint processing 5-75
 BSP CMRDECK entry 3-3
 Buffer access table 13-7
 Buffer request table 13-7
 Buffered disks 3-52

C

C DSDI directive 6-14
 CAF APRDECK entry 3-104
 Cancel alternate NPU load file command 15-11
 Cartridge for 7990
 Description 11-3
 Removal 11-86

Restoration 11-86
 CARTRIDGE PF STAGING IPRDECK entry 3-138
 Cartridge tape subsystem (see CTS) . 22-1
 Catalog archive file (PFATC) 17-67
 Catalog image record 17-78
 Catalog information 3-45
 Catalog permanent file device (PFCAT) 17-67
 Catalog track 17-9
 CBT DSDI directive 6-25
 CCT DSDI directive 6-25
 CDC IPRDECK entry 3-136
 CDCfff DSD command 5-50
 CDCNET
 Device interface EST entry 3-34
 MDI 15-2
 Network 15-1
 TDI 15-2
 CEF DIS command 4-14
 Central memory
 Descriptions 3-3
 Resident 3-3
 Table of sizes 3-12
 CFO DSD command 5-9
 CFR DSD command 5-75
 Change NPU load file command 15-12
 Channel control commands 5-74
 CHECK POINT SYSTEM DSD command 5-56
 CKP DSD command 5-4
 CLASS DSD command 5-10
 CLASS IPRDECK entry 3-139
 Clear all flaw reservations 3-104
 Clear disk area flaws 3-106
 Clear logical flaws 3-105
 Clear physical extended memory track 3-106
 Clear unit reservations 3-96
 Clearing permanent file error flags 11-85
 CLF APRDECK entry 3-105
 CLT CMRDECK entry 3-3
 CM DSDI directive 6-11
 CMRDECK
 Description 3-2
 Display 2-2
 Listing 3-2

Modifying	2-2	CPF APRDECK entry	3-106
CMRDECK entries		CPM default values	3-121
BSP	3-3	CPM IPRDECK entry	3-120
CLT	3-3	CPO DSDI directive	6-28
DFT	3-4	CPTT DSD command	5-11
EJT	3-4	CPTT IPRDECK entry	3-140
EMB	3-4	CPU	
EQP	3-4	Scheduling	3-112
FNT	3-5	CPUMTR breakpoint	
FOT	3-5	C display	5-81
INB	3-6	Commands	5-80
IPD	3-6	Precautions	5-83
LDT	3-6	CPUPFM transfer threshold	3-140
LIB	3-6	CS information command	15-46
MID	3-6	CSM IPRDECK entry	3-123
MINCM	3-7	CT DSDI directive	6-28
NAME	3-8	CTS	
NCP	3-8	Components	22-1
OPSECM	3-9	Constraints	22-16
PCP	3-9	Cross-coupled controllers	22-13
PPU	3-10	Description	22-1
QFT	3-10	Dual-state configuration	22-4
VE	3-11	Example	22-14
VERSION	3-12	Maintenance	22-17
CMRINST display	2-2	Minimum configuration	22-2
COMLIB IPRDECK entry	3-120	Performance considerations	22-17
COMMENT DSD command	5-9	Requirements	22-16
Console operation under DIS	4-10	CYBERPLUS ring port EST entry	3-38
CONTINUE DSD command	5-25	CYBfff DSD command	5-50
Control module		D	
Controlware	1-1	D DSDI directive	6-14
Control of jobs	3-110	Data storage display (F)	4-6
Control points		Database control system, CYBER	3-136
Number	3-111	DATE DSD command	5-57
Usage	3-110	Date/time entry	2-24
Controlware		Dayfile	
Disk	1-2	Account	3-15
Dumping utility	1-4	Commands	5-3
Loading utility	1-1	Description	3-15
Network Access Device	1-2	Dumping utilities	18-55
Status display	2-18	Dumps	5-3
Copy archive file (PFCOPY)	17-72	Error log	3-16
Coupler		Logging messages	3-145
Description	15-2	System	3-16
Disable or enable	15-9	DAYFILE DSDI directive	6-28
CP DSDI directive	6-26		
CPB DSD command	5-80		

Dayfile dumping commands		
AFD	18-55	
DFD	18-55	
ELD	18-55	
GETLOG	18-55	
MAINLOG	18-55	
DAYFILE EQPDECK entry	3-16	
DB DSDI directive	6-28	
DBW DSDI directive	6-29	
DCH DSD command	5-74	
DCN DSD command	5-74	
DCP DIS command	4-14	
DDB DSDI directive	6-29	
DDF		
Bringing up	J-1	
Change sector data commands	J-11	
Commands	J-6	
Console messages	J-5	
Description	J-1	
Examples	J-12	
Keyboard input	J-4	
Left screen display	J-2	
Pack recovery	J-16	
Print sector data commands	J-11	
Right screen display	J-3	
DDF DIS command	4-14	
DDF pack recovery		
Commands	J-19	
Display changes	J-16	
Hints	J-21	
Recovery table format	J-16	
DDP rollout path	3-142	
DDP ROLLOUT PATH IPRDECK entry		
	3-142	
Deadstart		
Displays	2-8	
Error processing	2-35	
Level 0 recovery	2-35	
Level 1 recovery	2-33	
Level 2 recovery	2-34	
Level 3 recovery	2-30	
Preparation	2-27	
Process	2-1	
Deadstart decks	3-1	
Deadstart dump interpreter (see DSDI)	6-1	
Deadstart file	3-1	
Deadstart sequencing	20-4	
DEBUG DSD command	5-58	
DEBUG IPRDECK entry	3-140	
Debug mode	3-140	
Debugging commands	5-84	
Default family name assignment entry	3-86	
DELAY DSD command	5-11	
DELAY IPRDECK entry	3-141	
Destage dump optical disk volume		
Characteristics	12-6	
Data format	12-6	
Description	12-1	
Destaging files	12-7	
Destage dump tape		
Characteristics	12-3	
Data format	12-3	
Description	12-1	
Destaging files	12-3	
Modifying	12-10	
Recycling	12-10	
Restoring	12-9	
TMS environment	12-5	
Destage/release processing		
Algorithms	11-33	
Destaging	12-1	
Excluding	11-32	
Destaging files	12-3, 12-7, 17-3	
Device		
Auxiliary	3-44	
Family	3-45	
Link	3-45	
Master	3-45	
Multispindle	3-46	
Nonremovable	3-51	
Private	3-44	
Public	3-44	
Removable	3-51	
Shared	3-51	
System	3-51	
Temporary file	3-51	
Device access table	13-6	
Device index tables	13-15	
Device mask	17-5	
DFLIST routine	18-27	
DFPT IPRDECK entry	3-142	
DFT CMRDECK entry	3-4	
DFTERM routine	18-27	

Diagnostic operator (see DOP)	15-45	RCS	4-16
DIAL DSD command	5-8	RNS	4-16
Direct access files	3-46	ROLLOUT	4-16
Directory display (Z)	4-9	RSS	4-16
DIS		SCS	4-16
Command	4-1	SET	4-13
Commands	4-14	SUI	4-16
Console operation	4-10	T	4-16
Display identifiers	4-2	U	4-17
Display selection commands	4-13	UCC	4-17
Keyboard entries	4-14	V	4-17
Memory displays	4-6	X	4-17
Memory entry commands	4-18	DIS DIS command	4-14
Messages	4-11	DIS displays	
Operations	4-1	Data storage	4-6
PP call commands	4-20	Directory	4-9
DIS commands		Display (A)	4-3
*	4-17	Display (B)	4-4
BEGIN	4-14	Display (F)	4-6
BKP	4-14	Display (G)	4-7
BKPA	4-14	Display (X)	4-8
CEF	4-14	Display (Z)	4-9
DCP	4-14	Displays (C, D, F, G)	4-6
DDF	4-14	Exchange package	4-8
DIS	4-14	Job dayfile	4-3
DROP	4-14	Job status	4-4
ELS	4-14	Memory	4-6
ENAi	4-14	Program storage	4-7
ENBi	4-14	DIS DSD command	5-58
ENEM	4-14	DISABLE DSD command	5-59
ENFL	4-14	Disable host element command	15-15
ENFLE	4-15	Disable network element command	15-29
ENP	4-15	Disable subsystems	3-136
ENPR	4-15	DISABLE,subsystem IPRDECK entry	3-136
ENS	4-15	Disk Array Subsystem (see DAS)	3-52
ENTER	4-15	Disk controlware	1-2
ENTL	4-15	Disk equipment EST entry	3-58
ENXi	4-15	Disk idle mode	I-2
ERR	4-15	Disk space	
GO	4-15	Detecting failing devices	I-1
HOLD	4-15	Known media defects	I-3
M	4-15	Management	11-78
N	4-15	Release	12-8
O26	4-16	Disk thresholds display	2-19
OFFSW	4-16	DISK VALIDATION IPRDECK entry	3-123
ONSW	4-16	Display disk file (see DDF)	J-1
RCP	4-16		

Display selection commands (DIS)	4-13	DCH	5-74
Display selection commands (DSD)	5-2	DCN	5-74
DISPOSE DSDI directive	6-9	DEBUG	5-58
DMPCCC utility	1-4	DELAY	5-11
DMPNAD command	14-1	Description	5-1
Documentation		DIAL	5-8
Audience	13	DIS	5-58
Conventions	14	DISABLE	5-59
Organization	13	DOWN	5-25
Related publications	15	DROP	5-5
Software support hotline	20	END	5-27
Submitting comments	20	ENPR	5-5
DOP		ENQP	5-5
Commands	15-45	FCN	5-74
Network control	15-41	FORM	5-27
Responsibilities	15-41	GO	5-9
Status message formats	15-43	IAFfff	5-50
Unsolicited status reports	15-45	IAN	5-74
DOWN channel entry	3-79	ID	5-27
DOWN DSD command	5-25	IDLE	5-28, 5-52, 5-64
DOWN EQPDECK entry	3-79	IDLEFAMILY	5-29
DP DSDI directive	6-29	INITIALIZE	5-30
DROP DIS command	4-14	K	5-64
DROP DSD command	5-5	KILL	5-6
DSD commands		L	5-64
99	5-69	LDC	5-74
ACN	5-74	LOAD	5-32
ASSIGN	5-24	LOCK	5-64
AUTO	5-56	MAGfff	5-52
BIO	5-50	MAINTENANCE	5-65
BKP	5-76	MAPfff	5-52
BKSP	5-24	MCH	5-74
BKSPF	5-25	MCSfff	5-52
BKSPRU	5-25	MOUNT	5-33
Categories	5-1	MSAL	5-34
CDCfff	5-50	MSEfff	5-53
CFO	5-9	NAMfff	5-53
CFR	5-75	NEXTVSN	5-35
CHECK POINT SYSTEM	5-56	OAN	5-74
CKP	5-4	OFF	5-36
CLASS	5-10	OFFSW	5-9
COMMENT	5-9	ON	5-37
CONTINUE	5-25	ONSW	5-9
CPB	5-80	OQSH	5-70
CPTT	5-11	OVERRIDE	5-6
CYBfff	5-50	PAUSE	5-9
DATE	5-57		

PCLASS	5-13	X.AFD	5-3
PRSIZE	5-37	X.DFD	5-3
QUEUE	5-14	X.ELD	5-3
RDFfff	5-53	X.MDD	5-68
RELEASE	5-70	X.QMOVE	5-4
REPEAT	5-37	X.QREC	5-4
REPRINT	5-38	DSD IPRDECK entry	3-124
REPUNCH	5-38	DSD V display	5-77
RERUN	5-7	DSDI	
RETRY	5-38	Buffer controller directive	6-51
RHFfff	5-54	CMR dump directives	6-24
ROLLIN	5-7	Command	6-3
ROLLOUT	5-7	Description	6-1
SCHEDULE	5-65	Directive format	6-7
SCRATCH	5-39	Dump control directives	6-11
SECUREQ	5-39	Examples	6-60
SECURES	5-70	Features	6-1
SERVICE	5-17	File manipulation and control	
SFR	5-75	directives	6-9
SKIP	5-39	Hardware register dump directives	6-43
SKIPF	5-39	Input directives	6-6
SKIPRU	5-40	Interactive use	6-52
SMFfff	5-54	List control directives	6-8
SPINDOWN	5-40	Memory dump directives	6-14
SPINUP	5-40	PP dump directives	6-17
SRST	5-22	Subsystem dump/analysis directives	6-39
STEP	5-65	Terminal output directives	6-54
STMfff	5-54	DSDI directives	
STOP	5-40, 5-54	*	6-9
SUPPRESS	5-41	ACCOUNT	6-24
TAFfff	5-55	ALLMEM	6-11
TEMP	5-41	AP	6-17
TERMINATE	5-41	B	6-24
THRESHOLD	5-42	BATCHIO	6-39
TIME	5-66	BCDUMP	6-51
TRACE	5-84	BIO	6-39
TRAIN	5-43	C	6-14
TRAP	5-87	CBT	6-25
UNLOAD	5-45	CCT	6-25
UNLOCK	5-67, 5-71	CM	6-11
UNSTEP	5-68	CP	6-26
UP	5-47	CPO	6-28
V	5-77	CT	6-28
VALIDATE	5-47	D	6-14
VSN	5-47	DAYFILE	6-28
WARN	5-8	DB	6-28
X	5-68		

DBW	6-29	PROCA	6-45
DDB	6-29	PROCW	6-45
DISPOSE	6-9	PROMR	6-45
DP	6-29	PROPM	6-46
E	6-15	PRORF	6-46
EC	6-12	PROSM	6-46
EICB	6-29	PROXP	6-47
EJ	6-8	PST	6-35
EJOFF	6-8	PUT	6-36
EJON	6-8	PX	6-20
EJT	6-30	Q	6-21
EPB	6-30	QF	6-21
ERRLOG	6-30	QFT	6-36
EST	6-30	QOA	6-22
FMFREG	6-43	QOD	6-22
FNT	6-30	QXA	6-23
FOT	6-31	QXD	6-23
HAT	6-31	RA	6-12
I	6-15	RAC	6-13
IAF	6-40	RCL	6-36
IOUCR	6-43	READ	6-10
IOUMR	6-43	REWIND	6-10
JC	6-31	RHF	6-42
LC	6-31	RPL	6-37
LDIS	6-32	SAB	6-37
LIDT	6-32	SC	6-47
LPVA	6-44	SECEDED	6-37
MAG	6-41	SETCPU	6-47
MAGNET	6-41	SETIOU	6-48
MAINLOG	6-32	SETJPS	6-48
MCT	6-32	SETRMA	6-49
MEMMR	6-44	SETVEP	6-49
MPP	6-18	SST	6-38
MST	6-33	TBDUMP	6-38
MTR	6-33	TRACEBK	6-50
MTRQUEUE	6-34	UEC	6-13
ODIS	6-34	W	6-16
OUTPUT	6-9	XP	6-50
P	6-18	Dual-channel access	3-53
PCP	6-34	Dual-state,NOS-NOS/VE	10-7
PD	6-8	Dump	
PF	6-19	Full	11-81, 17-74
PLD	6-35	Incremental	11-82, 17-73
PMS	6-19	Partial	17-74
PO	6-20	Secured system	17-73
PP	6-35	Tape management	11-82
PROBE	6-35		

Dump NAD memory	
Description	14-1
Local NAD dumping	14-2
Remote NAD dumping	14-2
Dump permanent file (PFDUMP)	17-72

E

E DSDI directive	6-15
EC DSDI directive	6-12
EI IPRDECK entry	3-125
EICB DSDI directive	6-29
EJ DSDI directive	6-8
EJOFF DSDI directive	6-8
EJON DSDI directive	6-8
EJT CMRDECK entry	3-4
EJT DSDI directive	6-30
Electronic vertical format unit (see EVFU)	F-1
ELS DIS command	4-14
EMB CMRDECK entry	3-4
Enable application command	15-16
Enable network element command	15-32
Enable subsystems	3-136
ENABLE,subsystem IPRDECK entry	3-136
ENAi DIS command	4-14
ENBi DIS command	4-14
END DSD command	5-27
ENEM DIS command	4-14
ENFL DIS command	4-14
ENFLE DIS command	4-15
Engineering mode	3-143
ENGR IPRDECK entry	3-143
ENP DIS command	4-15
ENPR DIS command	4-15
ENPR DSD command	5-5
ENQP DSD command	5-5
ENS DIS command	4-15
ENTER DIS command	4-15
ENTL DIS command	4-15
ENXi DIS command	4-15
EPB DSDI directive	6-30
EQP CMRDECK entry	3-4
EQPDECK	
Description	3-13
Display	2-3
Entries during deadstart	3-96
Listing	3-13

Modifying	2-3
Restrictions	3-14
EQPDECK entries	
ACCESS	3-94
ACCOUNT	3-15
ASR	3-85
AUTOLOAD	2-4
DAYFILE	3-16
DOWN	3-79
ERRLOG	3-16
FAMILY	3-86
GRENADE	2-4
INITIALIZE	2-4
ISHARE	3-89
LBC	3-90
MAINLOG	3-17
MAXSHARE	3-88
MSAL	3-77
PF	3-80
PRESET	3-101, 2-7
REMOVE	3-86
RESET	3-102, 2-7
SCKP	3-85
SHARE	3-87
SYSTEM	3-84
THRESHOLD	3-95
UEMIN	3-93
UP	3-79
XM	3-92
EQPINST	
Description	3-13
EQPINST display	2-8
Equipment assignments	
Clearing	3-18
Entries	3-18
Mass storage	3-40
Nonmass storage	3-21
Equipment status display	2-8
ERR DIS command	4-15
ERRLOG DSDI directive	6-30
ERRLOG EQPDECK entry	3-16
Error idle recovery	2-43
Error log	3-16
Error processing	
Deadstart from disk	2-36
Deadstart from tape	2-35

Error idle recovery	2-43
Recovery functions	2-36
EST DSDI directive	6-30
EST entries	
6683 satellite coupler	3-37
CDCNET device interface	3-34
Control module	3-65
CYBERPLUS ring port	3-38
Disk equipment	3-58
Extended memory	3-66
Magnetic tape equipment	3-27
MAP equipment	3-37
MSE	3-30
Network access device	3-35
Network processing unit	3-32
Nonstandard equipment	3-21
Optical disk	3-39
Pseudoequipment	3-22
Stimulator equipment	3-31
Two-port multiplexer	3-36
Unit record equipment	3-23
EST ordinal ranges	3-58
EVFU	
Description	F-1
Directives for 533/536 printers	F-1
Format control characters	F-2
Format control directives	F-1
Load file example	F-3
Load image	F-5
Exchange package display (X)	4-8
Extended core storage	3-55
Extended memory	
Clearing	3-126
EST entry	3-66
Examples	3-70
Field length	3-155
Flag registers	5-75
Logical configuration	3-56
Overview	3-55
Physical configuration	3-55
Reserve for buffers	3-92
Simultaneous usage	3-57
UEM size	3-57
User	3-160
Extended semiconductor memory	3-55
EXTENDED STACK PURGING	
IPRDECK entry	3-125

F

Family device	3-45
FAMILY EQPDECK entry	3-86
Family name assignment entry	3-86
FCN DSD command	5-74
File dumping	11-81
File reloading	11-84
Flaw entries	
CLF	8-5
CPF	8-5
SLF	8-5
SPF	8-5
FLAW K display	8-3
FLEXIBLE PARTITIONS IPRDECK	
entry	3-143
FMFREG DSDI directive	6-43
FNT CMRDECK entry	3-5
FNT DSDI directive	6-30
FORCED FORMATTING IPRDECK	
entry	3-144
FORM DSD command	5-27
FORMAT command	G-2
FOT CMRDECK entry	3-5
FOT DSDI directive	6-31
FOTD L display	9-2

G

GENPFD	
Command	12-12
File selection processing	12-16
Input directives	12-13
Limitations	12-16
GO DIS command	4-15
GO DSD command	5-9
GRENADE EQPDECK entry	2-4
Group	11-12

H

HARDWARE FAULT INJECTION	
IPRDECK entry	3-125
HAT DSDI directive	6-31
HOLD DIS command	4-15
HOP	
Commands	15-11
Control commands	15-56

Network control	15-9
Status message formats	15-9
Unsolicited status reports	15-10
Host operator (see HOP)	15-11

I

I DSDI directive	6-15
----------------------------	------

IAF

Control point assignment	3-136
IAF DSDI directive	6-40
IAF IPRDECK entry	3-136
IAF stimulator	
Description	16-33
Dynamic login	16-39
Multiple sessions	16-46
Parameter output	16-46
Postprocessing	16-48
Session file	16-34, 16-42
STIMULA command	16-34
Task definitions	16-41
Think time	16-40
IAFfff DSD command	5-50
IAFTM procedure file	16-29
IAFTR procedure file	16-30
IAN DSD command	5-74
ID DSD command	5-27
IDLE DSD command	5-28, 5-52, 5-64
Idle host element command	15-16
IDLEFAMILY DSD command	5-29
INB CMRDECK entry	3-6
Independent SDM operation	
Checkpoint	13-17
CMRDECK entry	13-11
Deadstart	13-11
Description	13-11
EQPDECK entry	13-11
MSE	13-17
Recovery	13-16
Unload	13-16
Independent shared device entry	3-89
Indirect access files	3-46
Initialization entry	3-96
INITIALIZE command	8-8
INITIALIZE DSD command	5-30
INITIALIZE EQPDECK entry	2-4
INITIALIZE K display	8-8
Initiating job processing	2-25

INSTALL command	7-1
Interactive facility (see IAF)	3-136
Interactive job control commands	5-8
Interactive Stimulator	3-137
IOUCR DSDI directive	6-43
IOUMR DSDI directive	6-43
IPD CMRDECK entry	3-6
IPRDECK	
Deadstart entries	3-135
Deadstart only entries	3-120
Description	3-107
Display	2-22
Entries	3-107
Listing	3-107
Modifying	2-22
Subsystem entries	3-135
IPRDECK entries	
ACS TAPE PF STAGING	3-137
ATF	3-136
AUTORESTART	3-138
BIO	3-136
CARTRIDGE PF STAGING	3-138
CDC	3-136
CLASS	3-139
COMLIB	3-120
CPM	3-120
CPTT	3-140
CSM	3-123
DDP ROLLOUT PATH	3-142
DEBUG	3-140
DELAY	3-141
DFPT	3-142
DISABLE,subsystem	3-136
DISK VALIDATION	3-123
DSD	3-124
EI	3-125
ENABLE,subsystem	3-136
ENGR	3-143
EXTENDED STACK PURGING	3-125
FLEXIBLE PARTITIONS	3-143
FORCED FORMATTING	3-144
HARDWARE FAULT INJECTION	3-125
IAF	3-136
KEYPM	3-126
LOCK	3-145
LOGGING	3-145

MAG	3-136	TRACE	3-134
MAP	3-136	UNLOCK	3-160
MASTER MSE	3-146	USER EXTENDED MEMORY	3-160
MCS	3-136	IPRINST display	2-22
MEMORY CLEARING	3-126	ISF command	20-1
MICRO	3-127	ISHARE EQPDECK entry	3-89
MS VALIDATION	3-146		
MSE	3-136	J	
NAM	3-136	JC DSDI directive	6-31
NAMIAF	3-127	Job communication commands	5-9
NVE	3-136	Job control	
OQSH	3-147	Commands	5-4
PCLASS	3-147	Information	3-108
PF VALIDATION	3-148	Parameter example	3-114
PRIVILEGED ANALYST MODE	3-148	Job dayfile display (A)	4-3
PRIVILEGED RDF	3-149	Job flow commands	5-10
PROBE	3-127	Job processing	
QUEUE	3-149	Initiating	2-25
RBF	3-136	Job scheduling	3-108
RDF	3-137	Job status display (B)	4-4
REMOVABLE PACKS	3-151		
RESIDENT RDF	3-151	K	
RHF	3-137	K DSD command	5-64
SCP	3-128	K-Display utilities	
SCRSIM	3-128	FLAW	8-3
SECCATS	3-129	INITIALIZE	8-8
SECONDARY USER COMMANDS	3-151	MREC	8-16
SECURES	3-152	MSE	8-24
SERVICE	3-153	NAM	8-26
SMF	3-137	QTF	8-40
SPC	3-129	RBF	8-70
SPD	3-130	REDEFINE	8-59
SPINDOWN	3-158	RHF	8-73
SPL	3-130	SSF	8-82
SPW	3-130	TAF	8-86
SRST	3-158	K-display utilities	
SSF	3-137	Overview	8-1
STM	3-137	KEYPM IPRDECK entry	3-126
SUBCP	3-131	Keypunch mode	3-126
SYSTEM DEBUG	3-159	KILL DSD command	5-6
TAF	3-137		
TAPE PF STAGING	3-160	L	
TCVM	3-131	L DSD command	5-64
TDEN	3-132	L-display utilities	
TDTY	3-132	FOTD	9-2
TMS	3-133		
TMSTO	3-133		

LIDOU	9-3	Listing NAD dumps	14-3
QDSPLAY	9-6	LISTLID command	10-7
SCTD	9-9	Load buffer controllers entry	3-90
SDSPLAY	9-11	LOAD DSD command	5-32
SUBSYST	9-19	Load file generator	
Label checking	3-151	Description	16-8
Large central memory extended	3-55	Input directives	16-9
LBC EQPDECK entry	3-90	Load partition	16-10
LC DSDI directive	6-31	Network load file	16-10
LCFFILE file	16-31	Summary listing	16-10
LCN examples	10-18	Load permanent file (PFLOAD)	17-76
LDC DSD command	5-74	LOADBC utility	
LDIS DSDI directive	6-32	Command	1-1
LDLIST routine	18-30	Loading microcode	3-127
LDT CMRDECK entry	3-6	Loading the system	2-23
Level 0 recovery deadstart	2-35	Local configuration file (LCF)	16-31
Level 1 recovery deadstart	2-33	LOCK DSD command	5-64
Level 2 recovery deadstart	2-34	LOCK IPRDECK entry	3-145
Level 3 recovery deadstart	2-30	LOGGING IPRDECK entry	3-145
LIB CMRDECK entry	3-6	LPVA DSDI directive	6-44
LIBDECK		M	
*/ entry	3-162	M DIS command	4-15
*AD entry	3-161	Machine recovery utility (see MREC)	8-16
*CM entry	3-161	MAG DSDI directive	6-41
*FL entry	3-161	MAG IPRDECK entry	3-136
*PROC entry	3-162	MAGffff DSD command	5-52
*SC entry	3-162	MAGNET command	12-11
LID configuration file	10-1	MAGNET DSDI directive	6-41
LID configuration file creation	10-1	Magnetic tape equipment EST entry	3-27
LID table		Magnetic tape subsystem (see MAG)	3-136
Building or Rebuilding	10-6	MAINLOG DSDI directive	6-32
Listing	10-7	MAINLOG EQPDECK entry	3-17
Specifying size	10-6	MAINTENANCE DSD command	5-65
LIDOU L display	9-3	Maintenance host facility (MHF)	14-2
LIDT DSDI directive	6-32	Management of storage media defects	I-1
Link device	3-45	MAP	
Linked SDM operation		EST entry	3-37
Checkpoint	13-10	IPRDECK entry	3-136
CMRDECK entry	13-3	MAP III or IV	3-136
Configuration	13-3	MAP III or IV equipment EST entry	3-37
Deadstart	13-3	MAPffff DSD command	5-52
Description	13-2	Mass Storage	
Device access tables	13-6	Allocation control entry	3-77
EQPDECK entry	13-3	Mass storage	
MSE	13-10	Clear logical flaws	3-105
Recovery	13-9		
Unload	13-9		

Concepts	3-41	MS VALIDATION IPRDECK entry	3-146
Configuration example	3-43	MSAL DSD command	5-34
Device functions	3-42	MSAL EQPDECK entry	3-77
Device recovery table	2-37	MSE	
Equipment assignments	3-40	EST entry	3-30
Hardware verification	3-123	Introduction	11-1
Initialization status display	2-16	IPRDECK entry	3-136
Set logical flaws	3-105	Operational procedures	11-77
Status display	2-12	Utilities	11-2
Thresholds	I-2	Utility terminology	11-3
Validation	3-146	MSE commands	8-24
Mass storage extended subsystem (see MSE)	3-30	MSE K display	8-24
Mass storage table	17-9	MSE processing	
Master device		7990 space management	11-79
Concepts	3-45	Backup and recovery	11-81
Master mode, MSE executive	3-146	Corrective actions	11-85
MASTER MSE IPRDECK entry	3-146	Disk space management	11-78
Matrix algorithm processor (see MAP)	3-37	Error conditions	11-85
Maximum shared device entry	3-88	Initialization	11-77
MAXSHARE EQPDECK entry	3-88	MSE utilities	
MCH DSD command	5-74	SSALTER	11-74
MCS		SSBLD	11-70
IPRDECK entry	3-136	SSDEBUG	11-59
MCSffff DSD command	5-52	SSDEF	11-13, 11-14
MCT DSDI directive	6-32	SSMOVE	11-24
MEMMR DSDI directive	6-44	SSUSE	11-51
Memory clearing	3-126	SSVAL	11-40
MEMORY CLEARING IPRDECK entry	3-126	MSEffff DSD command	5-53
Memory control	3-113	MST DSDI directive	6-33
Memory displays (DIS)	4-6	MTR DSDI directive	6-33
Memory entry commands (DIS)	4-18	MTRQUEUE DSDI directive	6-34
Memory entry commands (DSD)	5-72	Multimainframe operation	
Message control system	3-136	Deadstart	13-3
Messages, DIS	4-11	Description	13-1
MICRO IPRDECK entry	3-127	Multispindle device	3-46
MID CMRDECK entry	3-6		
MINCM CMRDECK entry	3-7	N	
MOUNT DSD command	5-33	N DIS command	4-15
MPP DSDI directive	6-18	NAD	
MREC		Controlware	1-2
Commands	8-20	EST entry	3-35
K display	8-16	NAD maintenance utilities	14-1
Parameters	8-20	NAM	
Procedures	8-17	Configuration file directives	10-21
Unit and controller reservations	8-22	Initialization	15-1, 16-1
		IPRDECK entry	3-136

Procedure file	16-1	CDCNET	15-1
Shutdown	15-8	Command summary	15-56
Startup	15-5	Description files	16-31
NAM K display		Operations	15-1
Control characters	8-28	Organization	15-1
Fields	8-26	Processing units (NPU)	15-1
Operation	8-29	Supervisor (NS)	15-1
NAM mode commands	8-29	Network Access Device (see NAD)	1-2
NAM mode maintenance commands		Network access device (see NAD)	14-1
DB	8-31	Network access method (see NAM)	3-136
DE	8-32	Network commands	
DU	8-32	COLLECT	16-5
FL	8-33	DEMUX	16-49
LB	8-33	IAFEX	16-29
LE	8-34	LFG	16-8
LR	8-34	LISTPPM	16-28
RS	8-35	NAMI	16-2
ST	8-36	NDA	16-11
NAME CMRDECK entry	3-8	STIMULA	16-34
NAMfff DSD command	5-53	Network configuration file (see NCF)	10-10
NAMIAF IPRDECK entry	3-127	Network control	
NAMNOGO		DOP	15-41
Procedure file	16-2	HOP	15-9
NCF		NOP	15-21
APPL statement	10-11	Network dump collector	16-5
CHARGE statement	10-17	Network operation commands	
DEBUG statement	10-16	ALERT	15-26
Examples	10-18	CFO	15-7
File	16-31	CONTROL	15-27
LNAD statement	10-14	DISABLE	15-15, 15-29
LNDR statement	10-11	DUMP	15-39
NPID statement	10-13	ENABLE	15-16, 15-32
PATH statement	10-13	FILE	15-12
RNAD statement	10-16	GO	15-36
Statements	10-11	HISTORY	15-49
NCFFILE file	16-31	IDLE	15-58
NCP CMRDECK entry	3-8	INFO	15-46
NCTFi file	16-31	LOAD	15-37
NDA		NAM	15-5
Description	16-11	NAMNOGO	15-6
Dump interpretation	16-28	NOFILE	15-11
Input directives	16-12	REPORT	15-50
NPU dump files	16-28	SEND	15-40, 15-51
Output file content	16-27	STATUS	15-12, 15-52
Parameters	16-11	Summary	15-56
Sample output	16-16	TST	15-47
Network			

Network operator utility (NETOU)	15-1	OPSECM CMRDECK entry	3-9
Network processing unit EST entry	3-32	Optical disk	12-1
NEXTVSN DSD command	5-35	Optical disk EST entry	3-39
NLID statement	10-3	OQSH DSD command	5-70
Node/path linkage statements	11-70	OQSH IPRDECK entry	3-147
Nonremovable mass storage device	3-51	Origin type	3-139
Nonstandard equipment EST entry	3-21	OUTPUT DSDI directive	6-9
NOP		Output queue special handling level	3-147
ALERT command	15-26	OVERRIDE DSD command	5-6
Commands	15-26		
CONTROL command	15-27	P	
DISABLE command	15-29	P DSDI directive	6-18
ENABLE command	15-32	Pack reformatting, 844	
Network control	15-21	Access for read and write	G-12
Responsibilities	15-21	Access for read only	G-11
SEND command	15-40	Pack reformatting, 881/883	
Status message formats	15-23	Control directives	G-4
Unsolicited status reports	15-25	Data statements	G-5
NOS/VE address formats	H-1	Description	G-1
NOS/VE Dual-state interface	3-136	Examples	G-7
NPID statement	10-2	FORMAT command	G-2
NPU		Input formats	G-4
Alternate load file command	15-11	Output formats	G-6
Diagnostic test command	15-47	PACKER command	K-4
DUMP command	15-39	PACKER utility	
GO command	15-36	Command format	K-4
LOAD command	15-37	Dayfile statistics	K-3
Load file command	15-12	Description	K-1
Load status command	15-12	Examples	K-6
NPU dump analyzer (see NDA)	16-11	Hole processing	K-2
NS control commands	15-11	Lost space processing	K-1
NS recent HISTORY command	15-14	Operational overview	K-1
NVE IPRDECK entry	3-136	PAUSE DSD command	5-9
NVF control commands	15-15	PCLASS DSD command	5-13
NVF recent HISTORY command	15-18	PCLASS IPRDECK entry	3-147
O		PCP CMRDECK entry	3-9
O26 DIS command	4-16	PCP DSDI directive	6-34
OAN DSD command	5-74	PD DSDI directive	6-8
ODIS DSDI directive	6-34	Peripheral equipment control commands	5-24
OFF DSD command	5-36		
OFFSW DIS command	4-16	Permanent file catalog	
OFFSW DSD command	5-9	Analysis	11-45
ON DSD command	5-37	Entry	17-11
ONSW DIS command	4-16	Format	17-11
ONSW DSD command	5-9	Problems	11-87
Operation under DIS	4-1	Permanent file commands	

FAMILY	17-26	PO DSDI directive	6-20
PFATC	17-67	PP breakpoint	
PFCAT	17-67	Commands	5-75
PFCOPY	17-72	Precautions	5-79
PFDUMP	17-72	V display	5-77
PFLOAD	17-76	PP call commands (DIS)	4-20
PFREL	17-89	PP DSDI directive	6-35
Permanent file recovery	11-85	PPU CMRDECK entry	3-10
Permanent files		Preserved file validation	3-148
Device assignment	3-80	PRESET EQPDECK entry	3-101, 2-7
Permit information	3-46	Preset independent shared device entry	
PF DSDI directive	6-19	3-102
PF EQPDECK entry	3-80	Preset link device entry	3-101
PF utilities		Primary VSN	12-2
Batch input	17-43	PRIVILEGED ANALYST MODE	
Calling	17-36	IPRDECK entry	3-148
Console input	17-36	PRIVILEGED RDF IPRDECK entry	3-149
DIS input	17-43	PROBE command	21-36
List	17-2	PROBE DSDI directive	6-35
Overview	17-1	PROBE IPRDECK entry	3-127
Parameters	17-45	PROBE utility	
Procedure files	17-44	Conditions monitored	21-36
Routines	17-62	Description	21-36
Summary file	17-29	Example	21-38
Terminal input	17-44	PROCA DSDI directive	6-45
Terminology	17-2	PROCW DSDI directive	6-45
PF VALIDATION IPRDECK entry	3-148	Program storage display (G)	4-7
PFDUMP		Programmable format control arrays	
Alternate storage environment	17-74	Adding	E-3
Description	17-72	Building	E-2
Example	17-82	PROMR DSDI directive	6-45
File interlocking	17-74	PROPM DSDI directive	6-46
Full	17-74	PRORF DSDI directive	6-46
Incremental	17-73	PROSM DSDI directive	6-46
Partial	17-74	PROXP DSDI directive	6-47
Protected files	17-76	PRSIZE DSD command	5-37
Secured system	17-73	Pseudo-control points	
PFLOAD		Number	3-111
Alternate storage environment	17-82	Usage	3-111
Catalog image record	17-78	Pseudoequipment EST entry	3-22
Description	17-76	PST DSDI directive	6-35
Example	17-82	PUT DSDI directive	6-36
Family reconfiguration	17-81	PX DSDI directive	6-20
Selection process	17-79		
PIP dump analyzer	16-28		
PLD DSDI directive	6-35		
PMS DSDI directive	6-19		
		Q	
		Q DSDI directive	6-21

QALTER routine	18-30	QMOVE	18-44
QDISPLAY commands	9-7	QREC	18-49
QDISPLAY L display	9-6	Queued file utilities	
QDUMP routine	18-36	Batch input	18-11
QF DSDI directive	6-21	Calling	18-2
QFT CMRDECK entry	3-10	Console input	18-2
QFT DSDI directive	6-36	Overview	18-1
QFTLIST routine	18-40	Parameters	18-11
QLIST routine	18-40	Preassignment	18-26
QLOAD routine	18-41	Routines	18-27
QMOVE routine	18-44	Terminal input	18-11
QOA DSDI directive	6-22	Queued file utility	
QOD DSDI directive	6-22	Commands	5-4
QREC routine	18-49	QXA DSDI directive	6-23
QTF		QXD DSDI directive	6-23
Configuration requirements	10-20		
Procedure file	10-22	R	
QTF K display		RA DSDI directive	6-12
Commands	8-46	RAC DSDI directive	6-13
File selection criteria	8-43	RBF	
File selection process	8-40	IPRDECK entry	3-136
File size index values	8-43	K display	8-70
NAM operation	8-45	RCFGEN utility	10-8
RHF operation	8-44	RCL DSDI directive	6-36
Selection classes	8-40	RCP DIS command	4-16
Usage	8-44	RCS DIS command	4-16
QTF K display commands		RDF	
CLASS	8-47	IPRDECK entry	3-137
DISABLE	8-48	Port number	3-36
ENABLE	8-48	Privileged mode	3-149
HELP	8-49	Resident mode	3-151
IDLE	8-49	RDFfff DSD command	5-53
INCLUDE	8-50	READ DSDI directive	6-10
SCHED	8-51	Recent HISTORY command	
STATUS	8-52	Description	15-49
QUEUE DSD command	5-14	NS command	15-14
Queue file transfer facility (see QTF)	8-40	NVF command	15-18
QUEUE IPRDECK entry	3-149	Reconfiguration	
Queued file commands		Commands	8-64
DFLIST	18-27	Equipment requirements	8-59
DFTERM	18-27	Parameters	8-63
LDLIST	18-30	Sequence	8-60
QALTER	18-30	REDEFINE K display	8-59
QDUMP	18-36	Release disk space (PFREL)	17-89
QFTLIST	18-40	RELEASE DSD command	5-70
QLIST	18-40	Releasing disk space	12-8
QLOAD	18-41		

Remote batch facility (see RBF) . . . 3-136
 Remote diagnostic facility (see RDF) 3-36,
 3-137
 Remote host facility (see RHF) . . . 3-137
 Removable device
 Assignment 3-51
 Assignment entry 3-86
 Removable mass storage device . . . 3-51
 Removable mass storage devices . . 3-151
 REMOVABLE PACKS IPRDECK entry
 3-151
 REMOVE EQPDECK entry 3-86
 REPEAT DSD command 5-37
 Report unsolicited status command 15-50
 REPRINT DSD command 5-38
 REPUNCH DSD command 5-38
 RERUN DSD command 5-7
 Reset device attributes 3-102
 RESET EQPDECK entry . . . 3-102, 2-7
 RESIDENT RDF IPRDECK entry 3-151
 RETRY DSD command 5-38
 REWIND DSDI directive 6-10
 RHF
 Applications 8-73
 Commands 8-76
 Configuration file directives . . . 10-20
 Configuration files 10-8
 Initiation 8-74
 IPRDECK entry 3-137
 K display 8-73
 Operator interface 8-74
 Termination 8-81
 RHF DSDI directive 6-42
 RHFfff DSD command 5-54
 RNS DIS command 4-16
 ROLLIN DSD command 5-7
 ROLLOUT DIS command 4-16
 ROLLOUT DSD command 5-7
 RPL DSDI directive 6-37
 RSS DIS command 4-16

S

SAB DSDI directive 6-37
 SC DSDI directive 6-47
 SCHEDULE DSD command 5-65
 Scheduling jobs 3-108
 SCKP EQPDECK entry 3-85

SCOPE 2 station facility (see SSF) 3-37,
 3-137, C-1
 SCP IPRDECK entry 3-128
 SCRATCH DSD command 5-39
 Screen management facility (see SMF)
 3-137
 SCRSIM
 Batch input D-4
 Calling format D-2
 Commands D-5
 Console input D-2
 Description D-1
 IPRDECK entry 3-128
 SCS DIS command 4-16
 SCTD commands 9-9
 SCTD L display 9-9
 SDSPLAY L display 9-11
 Commands 9-16
 System attributes 9-17
 SECCATS IPRDECK entry . . . 3-129
 SECEDED DSDI directive 6-37
 Secondary mask 17-7
 Secondary rollout sector threshold 3-158
 SECONDARY USER COMMANDS
 IPRDECK entry 3-151
 Secondary VSN 12-2
 Secured system control commands . 5-70
 SECUREQ DSD command 5-39
 SECURES DSD command 5-70
 SECURES IPRDECK entry . . . 3-152
 Security access categories 3-129
 Send message command . . . 15-40, 15-51
 Service class
 List 5-10
 Origin type 3-139
 Queue priorities 3-149
 Service limits 3-153
 Table 3-139
 SERVICE DSD command 5-17
 SERVICE IPRDECK entry 3-153
 Set access level limits 3-94
 SET DIS command 4-13
 Set disk area flaws 3-105
 Set disk thresholds 3-95
 Set logical flaws 3-105
 Set physical extended memory track 3-105
 SETCPU DSDI directive 6-47

SETIOU DSDI directive	6-48	Configuration constraints	11-73
SETJPS DSDI directive	6-48	Description	11-70
SETRMA DSDI directive	6-49	Example	11-72
SETVEP DSDI directive	6-49	Statements	11-70
SFR DSD command	5-75	SSDEBUG	
SHARE EQPDECK entry	3-87	Command	11-59
Shared device	3-51	Description	11-59
Shared device entry	3-87	Directive parameter	11-61
Shared mass storage		Directives	11-60
Buffered device	13-7	Restrictions	11-65
Configuring devices	13-6, 13-14	Usage	11-64
Initialization	13-8, 13-15	SSDEF command	11-13
Multimainframe	13-6, 13-13	SSEXEC	
Recovery tables	13-7	Description	11-12
Unloading	13-9, 13-16	Master mode	3-30
SKIP DSD command	5-39	SSF	
SKIPF DSD command	5-39	Coupler	3-37
SKIPRU DSD command	5-40	Equipment configuration	C-1
Slave mode, MSE executive	3-146	Error logout	C-3
SLF APRDECK entry	3-105	File transfers	C-1
SM Map		IPRDECK entry	3-137
Problems	11-87	SSF K display	8-82
SM map		File transfer commands	8-84
Analysis	11-43	File transfers	8-82
Backup	11-83	Operator interface	8-82
Description	11-3	Spooled file commands	8-85
SM/CU linkage statements	11-72	Staged file commands	8-84
SMF IPRDECK entry	3-137	Station disconnection	8-83
SMFfff DSD command	5-54	Station login	8-82
SPC IPRDECK entry	3-129	Station logout	8-83
SPD IPRDECK entry	3-130	Station recovery	8-83
SPF APRDECK entry	3-105	SSLABEL	
Spindown 834, 836, 887, and 9853 disk units	3-158	Command	11-14, 11-15
SPINDOWN DSD command	5-40	Description	11-14
SPINDOWN IPRDECK entry	3-158	Directive parameters	11-19
SPINUP DSD command	5-40	Directives	11-15
SPL IPRDECK entry	3-130	Example	11-23
Spooled file transfers	C-3	Restrictions	11-23
SPW IPRDECK entry	3-130	Update sequence	11-22
SRST DSD command	5-22	SSMOVE	
SRST IPRDECK entry	3-158	Command	11-25
SSALTER		Description	11-24
Console input	11-74	Directives	11-27
Description	11-74	Example	11-36
SSBLD		Restrictions	11-35
Command	11-70	Selection algorithms	11-32

Specific file directives	11-27	FORMAT utility	I-5
Value specifier directives	11-28	Keeping records	I-3
SST DSDI directive	6-38	Known defects	I-3
SSUSE		Management	I-1
Basic usage report	11-52	Reinitializing	I-7
Command	11-51	Releasing files	I-7
Description	11-51	Tools	I-4
Optional report A	11-54	STORNET	
Optional report B	11-56	EST entry	3-66
Optional report C	11-57	External memory	3-55
Optional report D	11-58	Subcontrol point processing	3-131
SSVAL		SUBCP IPRDECK entry	3-131
Command	11-40	Subfamily	11-12
Description	11-40	SUBSYST L display	9-19
Error classification	11-43	Subsystem control commands	5-50
Error detection	11-43	SUI DIS command	4-16
Examples	11-49	Summary file	17-29
Problem fixing	11-46	SUPPRESS DSD command	5-41
Problem reporting	11-43	SYSEdit	
Release processing	11-42, 11-46	Description	19-1
Restrictions	11-50	Directives	19-3
Typical runs	11-49	Record types	19-10
Validation report	11-47	SYSEdit command	19-1
Staged file transfers	C-2	SYSEdit directives	
Staging	12-2	*/	19-8
Staging files	17-3	*AD	19-5
Status host element command	15-18	*CM	19-3
Status message formats		*DELETE	19-8
DOP	15-43	*FILE	19-9
HOP	15-9	*FL	19-7
NOP	15-23	*IGNORE	19-9
Status network element command	15-52	*MS	19-4
Status/control register simulation	3-128	*PPSYN	19-9
Status/control register simulator (see		*SC	19-7
SCRSIM)	D-1	System	
STEP DSD command	5-65	Access level limits	3-152
STIMULA	3-31	Checkpoint entry	3-85
Stimulator equipment EST entry	3-31	Console lock	3-145
STM IPRDECK entry	3-137	Console unlock	3-160
STMffff DSD command	5-54	Control commands	5-56
STOP DSD command	5-40, 5-54	Control point facility	3-128
Storage media defects		Debug mode	3-159
APRDECK entries	I-5	Delay parameters	3-141
EQPDECK entries	I-5	Device	3-51
Example	I-6	File initialization	20-1
FLAW utility	I-4	Library	3-84
Flawed disk space	I-3		

Library device entry 3-84
 Tape density 3-132
 System dayfile 3-16
 SYSTEM DEBUG IPRDECK entry 3-159
 SYSTEM EQPDECK entry 3-84

T

T DIS command 4-16
 TAF IPRDECK entry 3-137
 TAF K display
 Commands 8-92
 CRMTASK commands 8-96
 Initialization 8-86
 Restarting 8-90
 Status 8-90
 TAFffff DSD command 5-55
 Tape
 Conversion mode 3-131
 Density 3-132
 Device type 3-132
 Track type 3-132
 Tape alternate storage 12-1
 Tape management system (TMS) . 3-133
 TAPE PF STAGING IPRDECK entry 3-160
 TBDUMP DSDI directive 6-38
 TCVM IPRDECK entry 3-131
 TDEN IPRDECK entry 3-132
 TDTY IPRDECK entry 3-132
 TEMP DSD command 5-41
 Temporary file device 3-51
 TERMINATE DSD command 5-41
 THRESHOLD DSD command 5-42
 THRESHOLD EQPDECK entry . . . 3-95
 TIME DSD command 5-66
 TMS IPRDECK entry 3-133
 TMSTO IPRDECK entry 3-133
 Toggle autoloading 3-96
 TRACE DSD command 5-84
 TRACE IPRDECK entry 3-134
 TRACEBK DSDI directive 6-50
 TRACER commands
 ACPD 21-4
 CPD 21-3
 ENDCPD 21-5
 ICPD 21-6
 TRACER utility
 Commands 21-4

Conditions monitored 21-1
 Data items reported 21-18
 Description 21-1
 Example 21-8
 Fast loop items 21-18, 21-21
 Medium loop items 21-19, 21-27
 Output file format 21-7
 Programs 21-3
 Slow loop items 21-20, 21-31
 Statistical summary 21-33
 Summary file format 21-17
 Track reservation table (TRT) . . 17-23
 TRAIN DSD command 5-43
 Transaction facility (see TAF) . . 3-137
 TRAP DSD command 5-87
 TST command 15-47
 Two-port multiplexer EST entry . . 3-36

U

U DIS command 4-17
 UCC DIS command 4-17
 UEC DSDI directive 6-13
 UEM equipment initialization . . . 3-93
 UEMIN EQPDECK entry 3-93
 Unified extended memory 3-56
 Unit record equipment EST entry . . 3-23
 UNLOAD DSD command 5-45
 UNLOCK DSD command 5-67, 5-71
 UNLOCK IPRDECK entry 3-160
 Unsolicited status reports
 DOP 15-45
 HOP 15-10
 NOP 15-25
 UNSTEP DSD command 5-68
 UP channel entry 3-79
 UP DSD command 5-47
 UP EQPDECK entry 3-79
 User extended memory 3-160
 USER EXTENDED MEMORY IPRDECK
 entry 3-160
 User index 17-4
 Utilities
 Controlware loading 1-1
 Dayfile dumping 18-55
 Deadstart dump interpreter 6-1
 Display disk file J-1

GENPFD	12-11	Secondary	12-2
K-display	8-1	VSN DSD command	5-47
L display	9-1		
MHF	14-2	W	
MSE	11-1	W DSDI directive	6-16
NAD maintenance	14-1	WARN DSD command	5-8
NETOU	15-1	Workstation commands	22-13
Network	16-1		
Permanent file	17-1	X	
PROBE	21-36	X DIS command	4-17
Queued file	18-1	X DSD command	5-68
RCFGEN	10-8	X.AFD DSD command	5-3
Tracer	21-1	X.DFD DSD command	5-3
		X.ELD DSD command	5-3
V		X.FLAW command	8-4
V DIS command	4-17	X.MDD DSD command	5-68
V DSD command	5-77	X.MREC command	8-17
VALIDATE DSD command	5-47	X.QMOVE DSD command	5-4
VE CMRDECK entry	3-11	X.QREC DSD command	5-4
VERSION CMRDECK entry	3-12	XM EQPDECK entry	3-92
Volume	11-12	XP DSDI directive	6-50
VSN			
Primary	12-1		

Control Data Systems, Inc.
ARH215
4201 Lexington Avenue N.
St. Paul, Minnesota 55126-6198

Title: NOS Version 2 Analysis Handbook

Publication No.: 60459300

Revision: U

Date: August 1994

Reason for Change: Manual updated to reflect NOS 2.8.3 at PSR level 840. This edition obsoletes all previous editions.

Note

If you need a binder for this manual, you can order
binder no. 60086240 from:

Control Data Systems, Inc.
Literature Distribution Services
ARHLDS
4201 Lexington Avenue N.
St. Paul, Minnesota 55126-6198

NOS Version 2 Analysis Handbook
60459300 U