



**NOS VERSION 2
ANALYSIS HANDBOOK**

**CDC® COMPUTER SYSTEMS:
CYBER 180
CYBER 170
CYBER 70
MODELS 71, 72, 73, 74
6000**



**NOS VERSION 2
ANALYSIS HANDBOOK**

**CDC® COMPUTER SYSTEMS:
CYBER 180
CYBER 170
CYBER 70
MODELS 71, 72, 73, 74
6000**

REVISION RECORD

REVISION	DESCRIPTION
<p>A (04-26-82)</p>	<p>Manual released at NOS 2.0, PSR level 562. NOS Version 2 is the successor product to NOS Version 1. The file name table reorganization and support of models 825, 835, and 855 is documented. In addition, parameter changes and enhancements are included for the permanent file utilities, queue file utilities, mass storage subsystem utilities, data gathering utilities, memory dump directives, and system accounting information. The ASTIM, NSTIM, and VALNET commands have been deleted and a glossary has been added.</p>
<p>B (01-27-83)</p>	<p>Manual released at NOS 2.1, PSR level 580. New features include support of the Remote Host Facility, Remote Diagnostic Facility, Multi Host, and models 815, 865, and 875. Section 4 has been revised to include the NAMI and COLLECT utilities. Also the LFG and NDA utility descriptions have been revised and the Network Products Stimulator and NETUVSN utility have been deleted. New CMR dump directives have been added to section 10. The Controlware Loading utility described in section 14 has been expanded and rewritten. Section 16 has been added to describe the new RHF utilities. New diagnostic messages and glossary terms have also been added. Due to extensive changes, change bars and dots are not used, and all pages reflect the latest revision level. This edition obsoletes all previous editions.</p>
<p>C (10-11-83)</p>	<p>Manual released at NOS 2.2, PSR level 596. New features include support of project prologues and epilogues, enhanced system security, service class assignment by users, expanded equipment status table, reformatted TRACER output, network enhancements, and support of model 845. Also the K displays have been revised for permanent file utilities, queue file utilities, MODVAL, and PROFILE. Due to extensive changes, change bars and dots are not used, and all pages reflect the latest revision level. This edition obsoletes all previous editions.</p>
<p>D (10-05-84)</p>	<p>Manual released at NOS 2.3, PSR level 617. New features include support of the 721 console, CYBER 180 Computer Systems, 834 Disk Storage Subsystem, 639 Magnetic Tape Units, and Mass Storage Archival Subsystem.</p> <p>The NOS 2 Analysis Handbook, formerly the System Maintenance Reference Manual, no longer contains user validation and system accounting information (that information is now included in the NOS 2 Administration Handbook). All error messages have been moved to the NOS 2 Operations Handbook. New to this manual is the analysis information from part II of the Operator/Analyst Handbook, such as the DSD commands, DIS operations, K-display and L-display utilities, and network operations. Also incorporated in this manual are the deadstart decks (CMRDECKs, EQPDECKs, APRDECKs, IPRDECKs, and LIBDECKs), which were in the Installation Handbook. Due to extensive changes, change bars and dots are not used, and all pages reflect the latest revision level. This edition obsoletes all previous editions.</p>
<p>Publication No. 60459300</p>	

REVISION LETTERS I, O, Q, S, X AND Z ARE NOT USED.

Address comments concerning this manual to:

Control Data Corporation
 Publications and Graphics Division
 4201 North Lexington Avenue
 St. Paul, Minnesota 55112

© 1982, 1983, 1984
 by Control Data Corporation
 All rights reserved
 Printed in the United States of America

or use Comment Sheet in the back of this manual.

LIST OF EFFECTIVE PAGES

New features, as well as changes, deletions, and additions to information in this manual, are indicated by bars in the margins or by a dot near the page number if the entire page is affected. A bar by the page number indicates pagination rather than content has changed.

PAGE	REV	PAGE	REV	PAGE	REV	PAGE	REV	PAGE	REV
Front Cover	-	3-2	D	3-61	D	4-5	D	5-44	D
Title Page	-	3-3	D	3-62	D	4-6	D	5-45	D
2	D	3-4	D	3-63	D	4-7	D	5-46	D
3	D	3-5	D	3-64	D	4-8	D	5-47	D
4	D	3-6	D	3-65	D	4-9	D	5-48	D
5/6	D	3-7	D	3-66	D	4-10	D	5-49	D
7	D	3-8	D	3-67	D	4-11	D	5-50	D
8	D	3-9	D	3-68	D	4-12	D	5-51	D
9	D	3-10	D	3-69	D	4-13	D	Tab	-
10	D	3-11	D	3-70	D	4-14	D	Divider	-
11/12	D	3-12	D	3-71	D	4-15	D	6-1	D
13	D	3-13	D	3-72	D	4-16	D	6-2	D
14	D	3-14	D	3-73	D	4-17	D	6-3	D
15	D	3-15	D	3-74	D	4-18	D	6-4	D
Tab	-	3-16	D	3-75	D	Tab	-	6-5	D
Divider	-	3-17	D	3-76	D	Divider	-	6-6	D
1-1	D	3-18	D	3-77	D	5-1	D	6-7	D
1-2	D	3-19	D	3-78	D	5-2	D	6-8	D
1-3	D	3-20	D	3-79	D	5-3	D	6-9	D
Tab	-	3-21	D	3-80	D	5-4	D	6-10	D
Divider	-	3-22	D	3-81	D	5-5	D	6-11	D
2-1	D	3-23	D	3-82	D	5-6	D	6-12	D
2-2	D	3-24	D	3-83	D	5-7	D	6-13	D
2-3	D	3-25	D	3-84	D	5-8	D	6-14	D
2-4	D	3-26	D	3-85	D	5-9	D	6-15	D
2-5	D	3-27	D	3-86	D	5-10	D	6-16	D
2-6	D	3-28	D	3-87	D	5-11	D	6-17	D
2-7	D	3-29	D	3-88	D	5-12	D	6-18	D
2-8	D	3-30	D	3-89	D	5-13	D	6-19	D
2-9	D	3-31	D	3-90	D	5-14	D	6-20	D
2-10	D	3-32	D	3-91	D	5-15	D	6-21	D
2-11	D	3-33	D	3-92	D	5-16	D	6-22	D
2-12	D	3-34	D	3-93	D	5-17	D	6-23	D
2-13	D	3-35	D	3-94	D	5-18	D	6-24	D
2-14	D	3-36	D	3-95	D	5-19	D	6-25	D
2-15	D	3-37	D	3-96	D	5-20	D	6-26	D
2-16	D	3-38	D	3-97	D	5-21	D	6-27	D
2-17	D	3-39	D	3-98	D	5-22	D	6-28	D
2-18	D	3-40	D	3-99	D	5-23	D	6-29	D
2-19	D	3-41	D	3-100	D	5-24	D	6-30	D
2-20	D	3-42	D	3-101	D	5-25	D	6-31	D
2-21	D	3-43	D	3-102	D	5-26	D	6-32	D
2-22	D	3-44	D	3-103	D	5-27	D	6-33	D
2-23	D	3-45	D	3-104	D	5-28	D	6-34	D
2-24	D	3-46	D	3-105	D	5-29	D	6-35	D
2-25	D	3-47	D	3-106	D	5-30	D	6-36	D
2-26	D	3-48	D	3-107	D	5-31	D	6-37	D
2-27	D	3-49	D	3-108	D	5-32	D	6-38	D
2-28	D	3-50	D	3-109	D	5-33	D	6-39	D
2-29	D	3-51	D	3-110	D	5-34	D	6-40	D
2-30	D	3-52	D	3-111	D	5-35	D	6-41	D
2-31	D	3-53	D	3-112	D	5-36	D	6-42	D
2-32	D	3-54	D	3-113	D	5-37	D	6-43	D
2-33	D	3-55	D	Tab	-	5-38	D	6-44	D
2-34	D	3-56	D	Divider	-	5-39	D	6-45	D
2-35	D	3-57	D	4-1	D	5-40	D	6-46	D
Tab	-	3-58	D	4-2	D	5-41	D	6-47	D
Divider	-	3-59	D	4-3	D	5-42	D	6-48	D
3-1	D	3-60	D	4-4	D	5-43	D	6-49	D

PAGE	REV	PAGE	REV	PAGE	REV	PAGE	REV	PAGE	REV
6-50	D	8-58	D	11-17	D	12-9	D	15-1	D
6-51	D	8-59	D	11-18	D	12-10	D	15-2	D
6-52	D	8-60	D	11-19	D	12-11	D	15-3	D
6-53	D	8-61	D	11-20	D	12-12	D	15-4	D
6-54	D	8-62	D	11-21	D	12-13	D	15-5	D
6-55	D	8-63	D	11-22	D	12-14	D	15-6	D
6-56	D	8-64	D	11-23	D	12-15	D	15-7	D
6-57	D	8-65	D	11-24	D	12-16	D	15-8	D
6-58	D	8-66	D	11-25	D	12-17	D	15-9	D
6-59	D	8-67	D	11-26	D	12-18	D	15-10	D
Tab	-	8-68	D	11-27	D	12-19	D	15-11	D
7-1	D	8-69	D	11-28	D	12-20	D	15-12	D
Tab	-	8-70	D	11-29	D	12-21	D	15-13	D
Divider	-	8-71	D	11-30	D	12-22	D	15-14	D
8-1	D	8-72	D	11-31	D	12-23	D	15-15	D
8-2	D	8-73	D	11-32	D	12-24	D	15-16	D
8-3	D	8-74	D	11-33	D	12-25	D	15-17	D
8-4	D	8-75	D	11-34	D	12-26	D	15-18	D
8-5	D	8-76	D	11-35	D	12-27	D	15-19	D
8-6	D	8-77	D	11-36	D	12-28	D	15-20	D
8-7	D	8-78	D	11-37	D	12-29	D	15-21	D
8-8	D	8-79	D	11-38	D	12-30	D	15-22	D
8-9	D	Tab	-	11-39	D	12-31	D	15-23	D
8-10	D	Divider	-	11-40	D	12-32	D	15-24	D
8-11	D	9-1	D	11-41	D	12-33	D	15-25	D
8-12	D	9-2	D	11-42	D	12-34	D	15-26	D
8-13	D	9-3	D	11-43	D	12-35	D	15-27	D
8-14	D	9-4	D	11-44	D	12-36	D	15-28	D
8-15	D	9-5	D	11-45	D	12-37	D	15-29	D
8-16	D	9-6	D	11-46	D	12-38	D	15-30	D
8-17	D	9-7	D	11-47	D	12-39	D	15-31	D
8-18	D	9-8	D	11-48	D	12-40	D	15-32	D
8-19	D	9-9	D	11-49	D	12-41	D	15-33	D
8-20	D	9-10	D	11-50	D	12-42	D	15-34	D
8-21	D	Tab	-	11-51	D	12-43	D	15-35	D
8-22	D	Divider	-	11-52	D	12-44	D	15-36	D
8-23	D	10-1	D	11-53	D	12-45	D	15-37	D
8-24	D	10-2	D	11-54	D	12-46	D	15-38	D
8-25	D	10-3	D	11-55	D	12-47	D	15-39	D
8-26	D	10-4	D	11-56	D	12-48	D	15-40	D
8-27	D	10-5	D	11-57	D	12-49	D	15-41	D
8-28	D	10-6	D	11-58	D	12-50	D	15-42	D
8-29	D	10-7	D	11-59	D	12-51	D	15-43	D
8-30	D	10-8	D	11-60	D	12-52	D	15-44	D
8-31	D	10-9	D	11-61	D	12-53	D	15-45	D
8-32	D	10-10	D	11-62	D	12-54	D	Tab	-
8-33	D	10-11	D	11-63	D	12-55	D	Divider	-
8-34	D	10-12	D	11-64	D	12-56	D	16-1	D
8-35	D	10-13	D	11-65	D	Tab	-	16-2	D
8-36	D	10-14	D	11-66	D	Divider	-	16-3	D
8-37	D	10-15	D	11-67	D	13-1	D	16-4	D
8-38	D	10-16	D	11-68	D	13-2	D	16-5	D
8-39	D	10-17	D	11-69	D	13-3	D	16-6	D
8-40	D	Tab	-	11-70	D	13-4	D	16-7	D
8-41	D	Divider	-	11-71	D	13-5	D	16-8	D
8-42	D	11-1	D	11-72	D	13-6	D	16-9	D
8-43	D	11-2	D	11-73	D	13-7	D	16-10	D
8-44	D	11-3	D	11-74	D	13-8	D	16-11	D
8-45	D	11-4	D	11-75	D	13-9	D	16-12	D
8-46	D	11-5	D	11-76	D	13-10	D	16-13	D
8-47	D	11-6	D	11-77	D	13-11	D	16-14	D
8-48	D	11-7	D	Tab	-	13-12	D	16-15	D
8-49	D	11-8	D	Divider	-	13-13	D	16-16	D
8-50	D	11-9	D	12-1	D	13-14	D	16-17	D
8-51	D	11-10	D	12-2	D	13-15	D	16-18	D
8-52	D	11-11	D	12-3	D	Tab	-	16-19	D
8-53	D	11-12	D	12-4	D	Divider	-	16-20	D
8-54	D	11-13	D	12-5	D	14-1	D	16-21	D
8-55	D	11-14	D	12-6	D	14-2	D	16-22	D
8-56	D	11-15	D	12-7	D	Tab	-	16-23	D
8-57	D	11-16	D	12-8	D	Divider	-	16-24	D

PAGE	REV	PAGE	REV	PAGE	REV	PAGE	REV	PAGE	REV
16-25	D	17-48	D	20-2	D	B-9	D		
16-26	D	17-49	D	20-3	D	B-10	D		
16-27	D	17-50	D	Tab	-	B-11	D		
16-28	D	17-51	D	Divider	-	B-12	D		
16-29	D	17-52	D	21-1	D	B-13	D		
16-30	D	17-53	D	21-2	D	B-14	D		
16-31	D	17-54	D	21-3	D	B-15	D		
16-32	D	17-55	D	21-4	D	B-16	D		
16-33	D	17-56	D	21-5	D	B-17	D		
16-34	D	17-57	D	21-6	D	B-18	D		
16-35	D	17-58	D	21-7	D	B-19	D		
16-36	D	Tab	-	21-8	D	B-20	D		
16-37	D	Divider	-	21-9	D	B-21	D		
16-38	D	18-1	D	21-10	D	B-22	D		
16-39	D	18-2	D	21-11	D	B-23	D		
16-40	D	18-3	D	21-12	D	B-24	D		
16-41	D	18-4	D	21-13	D	B-25	D		
16-42	D	18-5	D	21-14	D	B-26	D		
16-43	D	18-6	D	21-15	D	C-1	D		
16-44	D	18-7	D	21-16	D	C-2	D		
16-45	D	18-8	D	21-17	D	C-3	D		
16-46	D	18-9	D	21-18	D	D-1	D		
Tab	-	18-10	D	21-19	D	D-2	D		
Divider	-	18-11	D	21-20	D	D-3	D		
17-1	D	18-12	D	21-21	D	D-4	D		
17-2	D	18-13	D	21-22	D	D-5	D		
17-3	D	18-14	D	21-23	D	D-6	D		
17-4	D	18-15	D	21-24	D	D-7	D		
17-5	D	18-16	D	21-25	D	E-1	D		
17-6	D	18-17	D	21-26	D	E-2	D		
17-7	D	18-18	D	21-27	D	E-3	D		
17-8	D	18-19	D	21-28	D	E-4	D		
17-9	D	18-20	D	21-29	D	E-5	D		
17-10	D	18-21	D	21-30	D	E-6	D		
17-11	D	18-22	D	21-31	D	E-7	D		
17-12	D	18-23	D	21-32	D	F-1	D		
17-13	D	18-24	D	21-33	D	F-2	D		
17-14	D	18-25	D	21-34	D	F-3	D		
17-15	D	18-26	D	21-35	D	F-4	D		
17-16	D	18-27	D	21-36	D	F-5	D		
17-17	D	18-28	D	21-37	D	F-6	D		
17-18	D	18-29	D	21-38	D	F-7	D		
17-19	D	18-30	D	21-39	D	F-8	D		
17-20	D	18-31	D	21-40	D	F-9	D		
17-21	D	18-32	D	21-41	D	F-10	D		
17-22	D	18-33	D	21-42	D	Tab	-		
17-23	D	18-34	D	Tab	-	Index-1	D		
17-24	D	18-35	D	Divider	-	Index-2	D		
17-25	D	18-36	D	A-1	D	Index-3	D		
17-26	D	18-37	D	A-2	D	Index-4	D		
17-27	D	18-38	D	A-3	D	Index-5	D		
17-28	D	18-39	D	A-4	D	Index-6	D		
17-29	D	18-40	D	A-5	D	Index-7	D		
17-30	D	18-41	D	A-6	D	Index-8	D		
17-31	D	18-42	D	A-7	D	Index-9	D		
17-32	D	18-43	D	A-8	D	Index-10	D		
17-33	D	18-44	D	A-9	D	Index-11	D		
17-34	D	18-45	D	A-10	D	Index-12	D		
17-35	D	Tab	-	A-11	D	Index-13	D		
17-36	D	Divider	-	A-12	D	Index-14	D		
17-37	D	19-1	D	A-13	D	Index-15	D		
17-38	D	19-2	D	A-14	D	Comment			
17-39	D	19-3	D	A-15	D	Sheet	D		
17-40	D	19-4	D	B-1	D	Back Cover	-		
17-41	D	19-5	D	B-2	D				
17-42	D	19-6	D	B-3	D				
17-43	D	19-7	D	B-4	D				
17-44	D	19-8	D	B-5	D				
17-45	D	Tab	-	B-6	D				
17-46	D	Divider	-	B-7	D				
17-47	D	20-1	D	B-8	D				

PREFACE

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

- CDC® CYBER 180 Computer Systems Models 810, 830, 835, 845, and 855
- CDC CYBER 170 Computer Systems Models 171, 172, 173, 174, 175, 176, 720, 730, 740, 750, 760, 815, 825, 835, 845, 855, 865, and 875
- CDC CYBER 70 Computer Systems Models 71, 72, 73, and 74
- CDC 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.

Since the sections of the manual are self-contained in that they do not build on each other, the sections are ordered alphabetically by title. 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, and disk pack reformatting for 881/883 units.

CONVENTIONS

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 model 176 is large central memory extended (LCME). Extended memory for models 865 and 875 and CYBER 180-class machines is unified extended memory (UEM). Extended memory for models 865 and 875 may also include either extended core storage (ECS) or extended semiconductor memory (ESM). Extended memory for all other NOS computer systems is either ECS or ESM. ECS and ESM are the only forms of extended memory that can be shared in a linked shared device multiframe complex and can be accessed by a distributive data path (DDP).

In this manual, ECS refers to both ECS and ESM, 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 DDP access, UEM and LCME are excluded.

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

RELATED PUBLICATIONS

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

Control Data Corporation
Literature Distribution Services
308 North Dale Street
St. Paul, Minnesota 55103

The following is a list of NOS operating system manuals.

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

<u>Control Data Publication</u>	<u>Publication Number</u>
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 publications provide additional information about NOS and its product set that may be useful to the reader.

<u>Control Data Publication</u>	<u>Publication Number</u>
BML Message Formats	60459940
COBOL Version 5 Reference Manual	60497100
Communications Control Program Version 3 Diagnostic Handbook	60471500
COMPASS Version 3 Reference Manual	60492600
CYBER Cross System Version 1 Build Utilities Reference Manual	60471200
CYBER Initialization Package User's Handbook	60457180
CYBER Loader Version 1 Reference Manual	60429800

<u>Control Data Publication</u>	<u>Publication Number</u>
CYBER Supermini Operations Handbook	60459850
CYBER Record Manager Advanced Access Methods Version 2 Reference Manual	60499300
CYBER Record Manager Basic Access Methods Version 1.5 Reference Manual	60495700
CYBER 70 Model 71 Computer System Hardware Reference Manual	60453300
CYBER 70 Model 72 Computer System Hardware Reference Manual	60347000
CYBER 70 Model 73 Computer System Hardware Reference Manual	60347200
CYBER 70 Model 74 Computer System Hardware Reference Manual	60347400
CYBER 170 Computer Systems Models 171 through 175 and 176 (Level A) Hardware Reference Manual	60420000
CYBER 170 Computer Systems Models 720, 730, 750, 760, and 176 (Level B) Hardware Reference Manual	60456100
CYBER 170 Computer Systems Model 825 Hardware Reference Manual	60469350
CYBER 170/180 Computer Systems Models 835, 845, and 855 Hardware Operator's Guide	60458390
CYBER 170/180 Computer Systems Models 835, 845, 855, and 990 (CYBER 170 State) Hardware Reference Manual	60469290
CYBER 170 Computer Systems Models 865 and 875 Hardware Reference Manual	60458920
CYBER 180 Models 810 and 830 Hardware Operator's Guide	60469440
CYBER 180 Models 810 and 830 (CYBER 170 State) Hardware Reference Manual	60469420
FORTRAN Extended Version 4 Reference Manual	60497800
FORTRAN Version 5 Reference Manual	60481300
Message Control System Version 1 Reference Manual	60480300
Modify Version 1 Instant	60450200
Modify Version 1 Reference Manual	60450100
MSSI Version 3 Reference Manual	60458820
Network Products Network Access Method Version 1 Network Definition Language Reference Manual	60480000
Network Products Network Access Method Version 1/Communications Control Program Version 3 Host Application Programming Reference Manual	60499500

<u>Control Data Publication</u>	<u>Publication Number</u>
Network Products Network Access Method Version 1/Communications Control Program Version 3 Terminal Interfaces Reference Manual	60480600
Network Products Remote Batch Facility Version 1 Reference Manual	60499600
Network Terminal User's Instant	60459380
NOS Full Screen Editor User's Guide	60460420
NOS Online Maintenance Software Reference Manual	60454200
NOS Screen Formatting Reference Manual	60460430
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 Manual Abstracts†	60485500
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
Remote Host Facility Access Method Reference Manual	60459990
Remote Host Facility Usage Manual	60460620
SCOPE Version 2.1 Operator's Guide	60455090
Software Publications Release History††	60481000

†The NOS Manual Abstracts is a pocket-sized manual containing brief descriptions of the contents and intended audience of all NOS and NOS product manuals. The abstracts can be useful in determining which manuals are of greatest interest to a particular user.

††The Software Publications Release History lists all of the software manuals and revision packets Control Data has issued. The history specifies the revision level of a particular manual that corresponds to the level of software installed at the site.

<u>Control Data Publication</u>	<u>Publication Number</u>
TAF Version 1 Reference Manual	60459500
TAF Version 1 User's Guide	60459520
TAF/CRM Data Manager Version 1 Reference Manual	60459510
Text Editor Reference Manual	60436100
Update 1 Reference Manual	60449900

Programming information for the various forms of extended memory can be found in the COMPASS Reference Manual 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.

<u>Control Data Publication</u>	<u>Publication Number</u>
Extended Semiconductor Memory Hardware Reference Manual	60455990
Extended Core Storage Reference Manual	60347100
Extended Core Storage II and Distributive Data Path Reference Manual	60430000

DISCLAIMER

NOS and its product set are intended to be used only as described in this document. Control Data cannot be responsible for the proper functioning of undescribed features or parameters.

CONTENTS

<p>1. CONTROLWARE LOADING UTILITY 1-1</p> <p>Control Module Controlware 1-1</p> <p>Disk Controlware 1-1</p> <p>Network Access Device (NAD) Controlware 1-2</p> <p>2. DEADSTART 2-1</p> <p>Modifying the CMRDECK 2-2</p> <p>Modifying the EQPDECK 2-3</p> <p>Modifying the APRDECKs 2-15</p> <p>Modifying the IPRDECK 2-16</p> <p>Loading the System 2-17</p> <p>Initiating Job Processing 2-19</p> <p>Preparing for Recovery Deadstart 2-21</p> <p>Level 0 Initial Deadstart 2-28</p> <p>Error Processing 2-28</p> <p>3. DEADSTART DECKS 3-1</p> <p>CMRDECK 3-1</p> <p>EQPDECK 3-8</p> <p>APRDECK 3-69</p> <p>IPRDECK 3-73</p> <p>LIBDECK 3-112</p> <p>4. DIS OPERATIONS 4-1</p> <p>DIS Job Dayfile Display (A) 4-3</p> <p>DIS Exchange Package Display (B) 4-4</p> <p>DIS Memory Displays (C, D, F, G) 4-5</p> <p>DIS Job Status Display (J) 4-7</p> <p>DIS Directory Display (Z) 4-9</p> <p>Console Operation 4-10</p> <p>Display Selection Commands 4-12</p> <p>DIS Keyboard Entries 4-13</p> <p>Memory Entry Commands 4-16</p> <p>PP Call Commands 4-18</p> <p>5. DSD COMMANDS 5-1</p> <p>Dayfile Commands 5-2</p> <p>Queued File Utility Commands 5-3</p> <p>Job Processing Control Commands 5-3</p>	<p>1-1</p> <p>1-1</p> <p>1-1</p> <p>1-2</p> <p>2-1</p> <p>2-2</p> <p>2-3</p> <p>2-15</p> <p>2-16</p> <p>2-17</p> <p>2-19</p> <p>2-21</p> <p>2-28</p> <p>2-28</p> <p>3-1</p> <p>3-1</p> <p>3-8</p> <p>3-69</p> <p>3-73</p> <p>3-112</p> <p>4-1</p> <p>4-3</p> <p>4-4</p> <p>4-5</p> <p>4-7</p> <p>4-9</p> <p>4-10</p> <p>4-12</p> <p>4-13</p> <p>4-16</p> <p>4-18</p> <p>5-1</p> <p>5-2</p> <p>5-3</p> <p>5-3</p>	<p>Peripheral Equipment Control Commands 5-18</p> <p>Subsystem Control Commands 5-33</p> <p>System Control Commands 5-37</p> <p>Secured System Control Commands 5-46</p> <p>Memory Entry Commands 5-47</p> <p>Channel Control Commands 5-50</p> <p>Extended Memory Flag Register Commands 5-51</p> <p>6. EXPRESS DEADSTART DUMP INTERPRETER (DSDI) 6-1</p> <p>Calling the Express Deadstart Dump Interpreter 6-2</p> <p>Input Directives 6-5</p> <p>Interactive Use of DSDI 6-42</p> <p>Printer Output Listing Examples 6-49</p> <p>7. INSTALL COMMAND 7-1</p> <p>8. K-DISPLAY UTILITIES 8-1</p> <p>CYBERLOG K Display 8-2</p> <p>Responding to CYBERLOG Displays 8-2</p> <p>Flaw K Display 8-10</p> <p>Initialize K Display 8-14</p> <p>Machine Recovery (MREC) Utility K Display 8-20</p> <p>MREC Procedures 8-21</p> <p>MREC Unit and Controller Reservations 8-26</p> <p>Mass Storage Archival Subsystem (MSAS) K Display 8-27</p> <p>Mass Storage Subsystem (MSS) K Display 8-28</p> <p>Network Access Method (NAM) K Display 8-29</p> <p>Queue File Transfer Facility (QTF) K Display 8-42</p> <p>Redefine K Display 8-43</p> <p>Remote Batch Facility (RBF) K Display 8-54</p> <p>Remote Host Facility (RHF) K Display 8-57</p> <p>SCOPE 2 Station Facility (SSF) K Displays 8-65</p> <p>File Transfer Limit Commands 8-67</p> <p>Transaction Facility (TAF) K Displays 8-68</p>
---	--	--

9. L-DISPLAY UTILITIES	9-1	Network Control by DOP	15-30
FOTD L Display	9-2	Recent History Command	15-36
LIDOU L Display	9-3	Report Unsolicited Status Command	15-37
QDSPLAY L Display	9-6	Send Message Command	15-38
SCTD L Display	9-8	Status Network Element Command	15-39
SUBSYST L Display	9-10	Summary of Network Operation Commands	15-42
10. LID/RHF CONFIGURATION FILES	10-1	16. NETWORK UTILITIES	16-1
LID Configuration File	10-1	NAM Initialization (NAMI)	16-1
RHF Configuration Files	10-8	Network Dump Collector (COLLECT)	16-4
11. MASS STORAGE ARCHIVAL SUBSYSTEM (MSAS)	11-1	Load File Generator (LFG)	16-6
Introduction	11-1	NPU Dump Analyzer (NDA)	16-9
MSAS Utilities	11-11	PIP Dump Analyzer (LISTPPM)	16-25
MSAS Operational Procedures	11-68	AIP Trace With IAF	16-26
12. MASS STORAGE SUBSYSTEM (MSS)	12-1	Network Description Files	16-28
Introduction	12-1	Stimulator	16-29
MSS Utilities	12-10	17. PERMANENT FILE UTILITIES	17-1
MSS Operational Procedures	12-45	Introduction	17-1
13. MULTIMAINFRAME OPERATIONS	13-1	Calling the Utilities	17-23
Introduction	13-1	Parameters for the Permanent File Utilities	17-30
Linked SDM Operation	13-2	Permanent File Utility Routines	17-38
Independent SDM Operation	13-10	18. QUEUE/DAYFILE UTILITIES	18-1
14. NAD MAINTENANCE UTILITIES	14-1	Queue File Utilities	18-1
Maintenance Logging Transfer Facility (MLTF)	14-1	Dayfile Dumping Utilities	18-43
Dump NAD Memory (DMPNAD)	14-1	19. SYSEEDIT	19-1
15. NETWORK OPERATIONS	15-1	Directives	19-2
Network Organization	15-1	Record Types	19-8
NAM Startup	15-3	20. SYSTEM FILE INITIALIZATION	20-1
NAM Shutdown	15-5	ISF	20-1
Network Control By HOP	15-6	Deadstart Sequencing	20-3
Network Control By NOP	15-16	21. TRACER/PROBE UTILITIES	21-1
		TRACER Utility	21-1
		PROBE Utility	21-33

INDEX

APPENDIXES

A. CHARACTER SETS	A-1	E. 580 PROGRAMMABLE FORMAT CONTROL	E-1
Character Set Anomalies	A-2	Building Programmable Format Control	
Character Set Tables	A-2	Arrays	E-1
		Adding Programmable Format Control	
		Arrays	E-2
B. GLOSSARY	B-1		
C. SCOPE 2 STATION FACILITY	C-1	F. 881/883 PACK REFORMATTING UTILITY	F-1
Equipment Configuration	C-1	FORMAT Command	F-2
SSF File Transfers	C-2	Input Formats	F-3
		Output Formats	F-5
D. STATUS/CONTROL REGISTER SIMULATOR	D-1	Accessing Disk Devices	F-9
Using the Simulator	D-1		
Simulator Commands	D-5		

CONTROLWARE LOADING UTILITY

1

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

CONTROLWARE LOADING UTILITY

1

By using the LOADBC command, 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 command to load controlware into a control module for the 834 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.

<u>Parameter</u>	<u>Description</u>
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 2 Installation Handbook). If F=loadfile is omitted, controlware is read from the system library SYSTEM.

DISK CONTROLWARE

You initiate downloading of disk controlware only by entering a console command. Also, you can load the disk controlware from 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=cc,F=loadfile,D=dumpfile.

<u>Parameter</u>	<u>Description</u>
C=cc	cc is a 2-digit octal number of the channel from 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.

<u>Parameter</u>	<u>Description</u>
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 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 7155/7255 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 command 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.

<u>P_i</u>	<u>Description</u>
C=cc	The 2-digit octal number of the channel from 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 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.

<u>Parameter</u>	<u>Description</u>
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 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.

<u>value</u>	<u>Description</u>
170	CYBER 170 controlware (380-170)
200	CYBER 200 controlware (380-200)
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 FOF0 connected to TCU 0 of the local NAD on channel 5, enter the following commands.

```
X.DMPNAD(CH=05,ND=7F,AC=FOF0,LT=1000)
X.LOADBC(C=05,ND=7F,AC=FOF0,LT=1000)
```

Modifying the CMRDECK	2-2
Modifying the EQPDECK	2-3
Equipment Status Display	2-4
Mass Storage Status Display	2-7
Mass Storage Initialization Status Display	2-9
Controlware Status Display	2-11
Modifying the APRDECKs	2-15
Modifying the IPRDECK	2-16
Loading the System	2-17
Entering the Date and Time	2-18
Initiating Job Processing	2-19
Preparing for Recovery Deadstart	2-21
Level 3 Recovery Deadstart	2-23
Level 1 Recovery Deadstart	2-26
Level 2 Recovery Deadstart	2-27
Level 0 Initial Deadstart	2-28
Error Processing	2-28
Error Idle Recovery	2-35
Removable Device	2-35
Nonremovable Device	2-35

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.

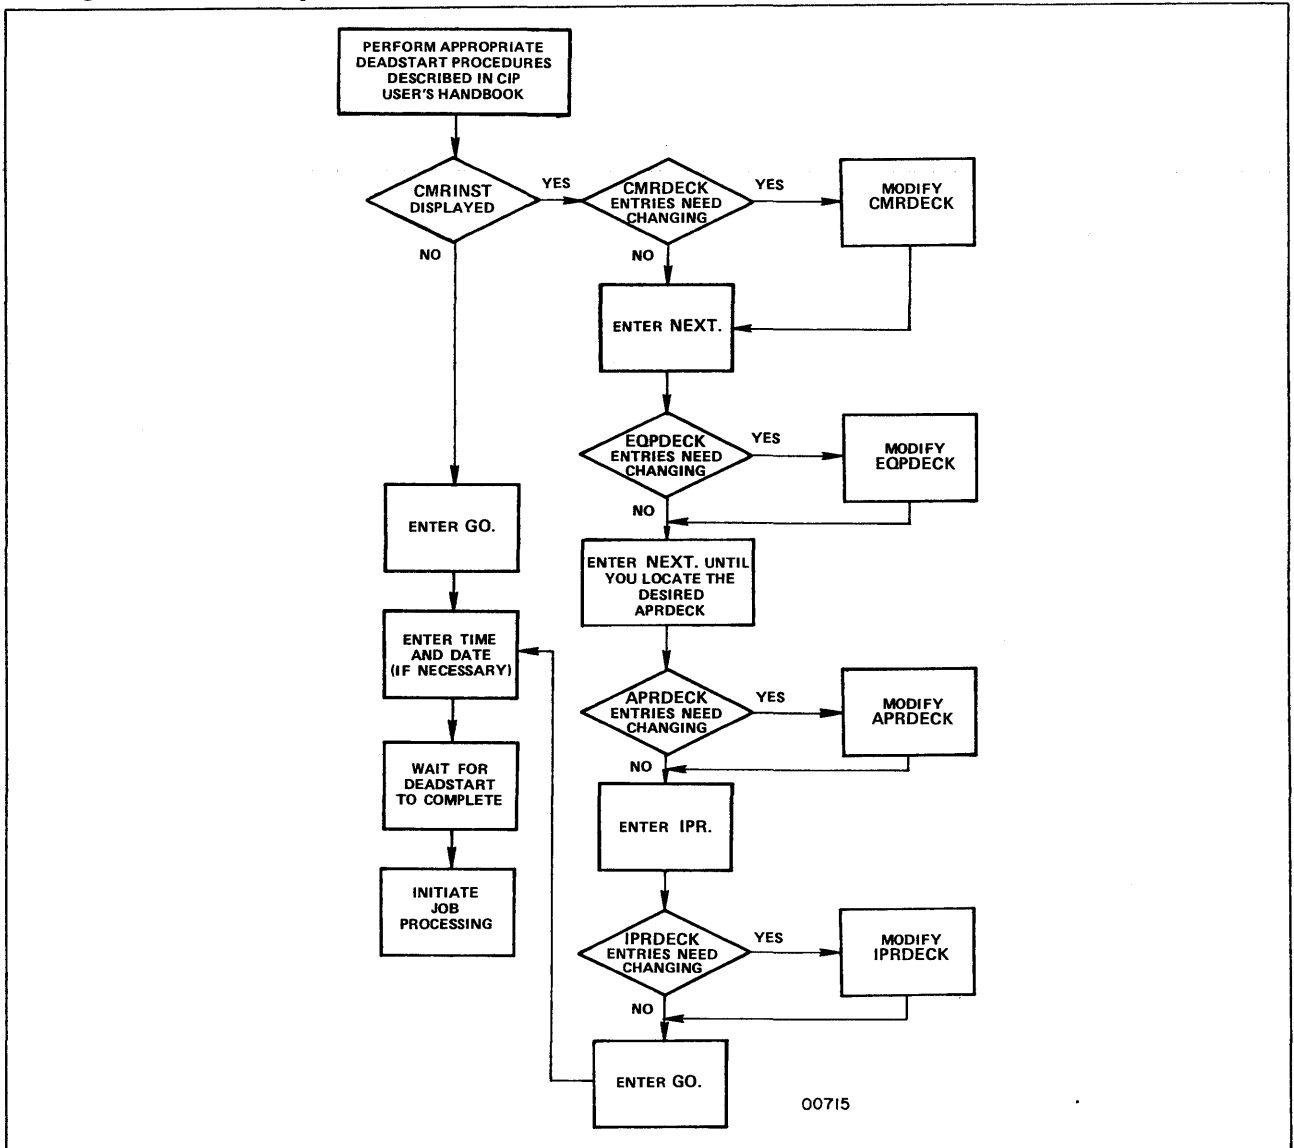


Figure 2-1. Deadstart Process

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.

When processing the EQPDECK, four displays are available, which you can toggle through in a circular manner. You can make changes when the EQPDECK, EQPINST, or any one of the following four displays is being displayed at the console screen.

<u>Display</u>	<u>Description</u>
Equipment Status	Describes the hardware configurations (refer to figure 2-2).
Mass Storage Status	Describes how the system allocates files on a mass storage device (refer to figure 2-3).
Mass Storage Initialization Status	Describes initialization status of the devices (refer to figure 2-4).
Controlware Status	Describes the status of the controlware loaded (refer to figure 2-5).

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 equipment display, mass storage display, mass storage initialization display, and controlware display 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.

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 this 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.

- CC Satellite Coupler.
- CM Control module for an 834 Disk Storage Subsystem.
- CP 415 Card Punch.
- CR 405 Card Reader.
- CS MSS Cartridge Storage Unit.
- CT MSS Cartridge Transport.

HeaderDescription

DB-i	885-42 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).
DD-i	834 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).
DE	Extended memory.
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).
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).
LP	Any line printer.
LR	580-12 Line Printer.
LS	580-16 Line Printer.
LT	580-20 Line Printer.
MT	Magnetic Tape Drive (7 track).
NC	380-170 Network Access Device.
NP	255X Network Processing Unit.
NT	Magnetic Tape Drive (9 track).
RM	Two-port multiplexer (models 865, 875, and CYBER 180-class machines).
SS	MSAS Controller.
TT	Internal stimulation device.

Header

Description

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.

- DS Display console (EST ordinal 1g).
- NE Null equipment (EST ordinal 2g).
- RD Used for online reconfiguration of mass storage (EST ordinal 0).
- TE Tape equipment (EST ordinal 3g).
- TT Used for assignment of terminal files (EST ordinal 4g).

st Equipment status (ON, OFF, or DWN†).

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 Unit number (serves as ID code for unit record devices). The identifier code (un parameter) provides a method of grouping peripheral 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.

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.

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.

† DWN is DOWN status. An equipment cannot be logically turned ON when in DWN status.

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 this format:

est type status files

Header

Description

est EST ordinal.

type Device type:

DB-i 885-42 Disk Storage Subsystem (1 ≤ i ≤ 3; full track).

DD-i 834 Disk Storage Subsystem (1 ≤ i ≤ 8; full track).

DE Extended memory.

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).

DL-i 844-41/44 Disk Storage Subsystem (1 ≤ i ≤ 8; full track).

HeaderDescription

DM-i	885-11/12 Disk Storage Subsystem ($1 \leq i \leq 3$; half track).
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 Subsystem (single density).
DW	819 Disk Subsystem (double density).
status	Status conditions. Any combination of conditions can exist. The following codes are listed in the order in which they appear on the display.
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.

† If a device is shared by two or more mainframes (status M), the mainframe identification is shown on the screen as the mainframe accesses the shared 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.

- 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.

DB-i 885-42 Disk Storage Subsystem ($1 \leq i \leq 3$; full track).

DD-i 834 Disk Storage Subsystem ($1 \leq i \leq 8$; full track).

DE Extended memory.

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).

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).

options Initialize options. Maximum of eight options are displayed.

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.

t Type of unit device.

F Family device.

X Auxiliary device.

iam Indirect access mask.

<u>Header</u>	<u>Description</u>
dam	Direct access mask.
fm/pn	Family name/pack name.
	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

Controlware status display describes the status of the controlware loaded.

Figure 2-5 illustrates the controlware status display.

CONTROLWARE STATUS	
CHANNEL	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

<u>Header</u>	<u>Description</u>
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.

Table 2-1. EQPDECK Entries (Sheet 1 of 3)

Entry	Function
AUTOLOAD.	<p>Toggles the selection of buffer controller autoloading for all 7054/7154/7152/7155/7255 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 7054/7154/7152/7155/7255 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.</p>
GRENADE.	<p>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 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.</p>
INITIALIZE,op,est ₁ ,est ₂ ,...,est _n .	<p>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.</p> <p style="margin-left: 40px;">op Level of initialization:</p> <p style="margin-left: 100px;">AL Total initialization. PF Permanent files. QF Queued files. DF System dayfile. AF Account dayfile. EF Error log. FP Format pack (844). MF Binary maintenance log.</p>

Table 2-1. EQPDECK Entries (Sheet 2 of 3)

Entry	Function
	<p><code>est1</code> 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.</p> <p>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 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.</p> <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>

Table 2-1. EQPDECK Entries (Sheet 3 of 3)

Entry	Function
<p>PRESET,n.</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> <p>Initializes tables (MST, TRT, MRT, and DAT) on the link device that are required for the management of shared multiframe mass storage devices. The entry is valid only for level 0 deadstarts by the first mainframe in the multiframe 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 style="padding-left: 40px;">n Number of shared devices.</p> <p>If you do not specify n, the link device is preset, and the amount of table space reserved for the shared devices is determined by the number of shared device entries in the EQPDECK.</p>
<p>PRESET=est₁,est₂,...,est_n.</p>	<p>Presets independent shared devices in a multiframe complex. It is valid only on a level 0 deadstart by the first mainframe in the multiframe 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> <p style="padding-left: 40px;">est_i EST ordinal of independent shared device.</p>
<p>RESET=est₁,est₂,...,est_n.</p>	<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>

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 (APR000) 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 to 120 seconds) if 834 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.

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. 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 presence of a wall clock chip. If the clock is present, the date and time are automatically read from it. If the clock is not present, 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.

yy Year; 00 through 99.

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.

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 either:

AUTO.

or

MAINTENANCE.

Following entry of the AUTO or MAINTENANCE 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)

<u>Parameter</u>	<u>Description</u>
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.

<u>Control Point Number†</u>	<u>Job Sequence Name</u>	<u>Activity</u>
1	IAF	Interactive Facility.
2	NAM	Network Access Method.
3	RHF	Remote Host Facility.
4	SMF	Screen Management Facility.
5	NEXT	Available for automatic system assignment.
.	.	
.	.	
.	.	
n-3	NEXT	
n-2	MAG	Magnetic tape subsystem executive routine.
n-1	BIO	Central site batch input/output.
n	RBF	Remote Batch Facility.
n+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,0 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.

†n is the number of control points defined 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,0 display, and stops recovery of the device. To alter the EOI for the file and proceed with recovery, enter:

GO,SYS.

To terminate recovery of the device, enter:

PAUSE,SYS.

Table 2-2. Levels of System Deadstart

Deadstart Level	Information Recovered				
	Jobs	Queued Files	Permanent Files	System	System Dayfiles
0	Input and output queues recovered.†	No.	Yes.	No.	Yes.††
1	Recovered from last checkpoint.	Recovered from last checkpoint.	Yes.	Recovered from last checkpoint.	Yes.
2	Recovered from last checkpoint.	Recovered from last checkpoint.	Yes.	No.	Yes.
3	Recovered from CM copy of EJT.†††	Recovered from CM copy of QFT.	Yes.	Yes.	Yes.

†The input and output queues are recovered. Rollout queues are also recovered as input files. Input files associated with jobs currently executing are returned to the input queues, unless the device they reside on is initialized. Permanent files are recovered unless the device is initialized.

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

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

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

CAUTION

Before attempting any level of recovery deadstart (including level 0), examine the current status codes listed for each mass storage device in the disk status display (E,M). Delay deadstart if status code C (checkpoint requested) appears for any device. When the system has processed the request, status code C is cleared (within 30 seconds). Refer to the NOS 2 Operations Handbook for complete information concerning the disk status display (E,M). Failure to observe this caution may result in the loss of permanent file information.

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.

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=Y 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 2-6) and instructions on the right screen (figure 2-7).

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

Figure 2-6. 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 2-7. 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 multiframe mode.

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

<u>Command</u>	<u>Description</u>
E,M.	Displays the disk status display (E,M).
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 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 tapes† 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.

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

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

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

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†:

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.

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

Deadstart from disk:

- 1a. For an 844 or 885 disk, ensure that the disk is spinning, the READY light is on, and the SELECT light is on.
- 1b. For an 834 disk, ensure that the start button is pushed in and the power switch is on.
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.

The specific recovery function performed depends upon the level of deadstart selected. Table 2-3 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 2 Operations Handbook for information concerning all deadstart messages.

Table 2-3. Mass Storage Device Recovery (Sheet 1 of 5)

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 EQPDECK.	0	Device being recovered is not ready and, therefore, cannot be read.	Removable.	Error code NR and status U set in MST; device is left unavailable.	None.
			Nonremovable.	Wait and retry. Error message appears on the right screen display.	Check to ensure that all nonremovable devices are ready.
		Read error occurred when attempting to read label (parity error, surface error, and so on).	All.	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.	None.
			Nonsystem device.	Error code CE and status U set in MST; device is left unavailable.	Examine E,M display after deadstart is complete to determine corrective action.
		System device (nonremovable with system residence).		- If permanent files do not reside on the device, it is initialized automatically using parameters in existing label.	None.
			Bad label. Information read was not recognized as a label sector.	Nonsystem device.	- 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.
		System device (nonremovable with system residence).			Error code LE and status U set in MST; device is left unavailable.
				System device (nonremovable with system residence).	Device is initialized automatically. Any permanent files that reside on device are lost.

Table 2-3. Mass Storage Device Recovery (Sheet 2 of 5)

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 844-41/44 (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.

Table 2-3. Mass Storage Device Recovery (Sheet 3 of 5)

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
Read TRT into central memory.	0	TRT could not be read successfully.	All.	The following message appears on the right console display and deadstart processing halts. LENGTH OF DEVICES TRT 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	TRT could not be read successfully.	Removable (no active direct access files). Nonremovable or removable with active direct access files.	Error code TL and status U set in MST; device is left unavailable. The following message appears on the right console display and deadstart processing halts. LENGTH OF DEVICES TRT BAD.	Examine the E,M display after deadstart is complete to determine corrective action. Perform level 0 deadstart and initialize the device.
Edit TRT to remove nonpermanent file information.	0	Edit was not successful. This could be caused by specifying a new SPF/SLF entry (flawed track) APRDECK.	All.	The following message appears on the right console display and deadstart processing halts. PERM. FILE LINKAGE ERROR 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. If SPF/SLF entry was specified, redeadstart and omit SPF/SLF entry.
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. EQest ₁ EQest ₂ DN. CONFLICTING DM. PN. est ₁ and est ₂ are the EST ordinals of the conflicting devices.	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).
			Nonsystem device.	Error code DN, DM, or PN and status U set in MST; 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.

Table 2-3. Mass Storage Device Recovery (Sheet 4 of 5)

Recovery Function	Level of Deadstart	Description of Error	Type of Device	System Response to Error	Action
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.	<p>Increment count of direct access files in error. The message appears in the dayfile in the following format.</p> <p>EQest files PRESERVED FILE ERRORS.</p> <p>est EST ordinal of device. files Number of files in error.</p>	Examine dayfile after deadstart completes. Message described in preceding column indicates the number of direct access files in error.
		Interlock data specified file was in WRITE mode and last sector for file in TRT was not an EOI sector.	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.</p> <p>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. count Total number of length errors.</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/CE.
Recovering queued files.	0	Last sector of file specified by TRT is not an EOI sector.	All.	<p>The following message is issued to the dayfile and error log:</p> <p>QF LENGTH ERROR filename.</p> <p>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. count File count.</p>	<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 TKxxxx LENGTH ERROR.</p>

Table 2-3. Mass Storage Device Recovery (Sheet 5 of 5)

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 CMRDECK. 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 CMRDECK is in use.	Extended memory.	The following message appears on the right console display and deadstart processing halts. MID CURRENTLY ACTIVE	Either redeadstart 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 CMR 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).

DEADSTART DECKS

3

CMRDECK	3-1
Central Memory Descriptions	3-2
EQPDECK	3-8
Dayfile Descriptions	3-10
General Description of Equipment Assignment Entries	3-12
Clear EST Assignment Entry	3-13
Equipment Assignments: Nonmass Storage	3-14
Nonstandard Equipment EST Entry	3-14
Pseudoequipment	3-14
Unit Record Equipment EST Entry	3-15
Magnetic Tape Equipment EST Entry	3-18
Mass Storage Facility Equipment EST Entry	3-20
Mass Storage Archival Subsystem EST Entry	3-22
Stimulator Equipment EST Entry	3-23
Network Processing Unit EST Entry	3-24
Network Access Device EST Entry	3-26
Two-Port Multiplexer EST Entry	3-26
MAP III or MAP IV Equipment EST Entry	3-27
6683 Satellite Coupler EST Entry	3-27
CYBERPLUS Ring Port EST Entry	3-28
Equipment Assignments: Mass Storage	3-28
NOS Mass Storage Concepts	3-29
Specifying Ranges of EST Ordinals	3-42
EQ - Disk Equipment EST Entry	3-42
EQ - Control Module EST Entry	3-46
EQ - Extended Memory EST Entry	3-47
MSAL - Mass Storage Allocation Control Entry	3-50
PF - Permanent Files Device Assignment Entry	3-51
SYSTEM - System Library Device Assignment Entry	3-54
ASR - Alternate System Library Device Assignment Entry	3-55
SCKP - System Checkpoint File Device Assignment Entry	3-55
DOWN - Logical Device Removal Entry	3-56
FAMILY - Default Family Name Assignment Entry	3-57
REMOVE - Removable Device Assignment Entry	3-57
SHARE - Shared Device Entry	3-58
ISHARE - Independent Shared Device Entry	3-58
LBC - Load Buffer Controllers Entry	3-59
XM - Declare Extended Memory Space for Buffers or User Access	3-61
ACCESS	3-62
EQPDECK Entries Made Only During Deadstart	3-63
AUTOLOAD	3-63
GRENADE	3-63
INITIALIZE - Initialization Entry	3-63
PRESET - Preset the Link Device Entry	3-68
PRESET - Preset the Independent Shared Device Entry	3-68
RESET	3-69
APRDECK	3-69
APRDECK Format	3-69
CAF - Clear All Flaw Reservations	3-70
SLF - Set Logical Flaws on Any Mass Storage Device	3-71
CLF - Clear Logical Flaws on Any Mass Storage Device	3-71
SPF - Set Physical Extended Memory Track or Disk Area Flaws	3-71
CPF - Clear Physical Extended Memory Track or Disk Area Flaws	3-72

(Continued)

(Continued)

IPRDECK	3-73
Job Control Information	3-73
Job Scheduling	3-74
Job Control	3-75
Memory Control	3-77
Example of Job Control Parameters	3-77
IPRDECK Entries Made Only During Deadstart	3-81
COMLIB	3-81
CPM	3-82
CSM	3-84
DSD	3-85
EI	3-85
EXTENDED STACK PURGING	3-86
KEYPM	3-86
MEMORY CLEARING	3-87
MICRO	3-87
NAMIAF	3-88
PROBE	3-88
SCP	3-88
SCRSIM	3-89
SECCATS	3-89
SPC	3-90
SPD	3-90
SPL	3-91
SPW	3-91
SUBCP	3-91
TCVM	3-92
TDEN	3-92
TDTR	3-93
Deadstart Entries and DSD Commands	3-93
Subsystems	3-94
CLASS	3-96
DEBUG	3-97
DELAY	3-97
DDP ROLLOUT PATH	3-98
ENGR	3-98
FILE STAGING	3-99
LOCK	3-99
LOGGING	3-99
MASTER MAS	3-100
MASTER MSS	3-100
MS VALIDATION	3-101
OQSH	3-101
PCLASS	3-102
PF VALIDATION	3-102
PRIVILEGED RDF	3-103
QUEUE	3-103
REMOVABLE PACKS	3-105
RESIDENT RDF	3-105
SECONDARY USER COMMANDS	3-105
SECURES	3-106
SERVICE	3-107
SPINDOWN	3-110
SRST	3-111
UNLOCK	3-111
USER EXTENDED MEMORY	3-111
LIBDECK	3-112

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

<u>Deck</u>	<u>Description</u>
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 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 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 200g words per control point.

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

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
CLT=number.	0	<p>Specifies the octal number of common library table (CLT) entries allowed. The number can range from 0 to 100g.</p> <p>If the table exists, the first entry 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.</p> <p>Common library table entries are described in the IPRDECK entry COMLIB.</p>

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
EJT=number.	620g	<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 7777g; the minimum value is 3.</p>
EQP=number.	Number of CMRDECK	<p>Specifies the number of the EQPDECK to use at deadstart. The EQPDECK contains equipment assignments. Up to 100g EQPDECKs can exist on a deadstart file; the number can range from 00 to 77g.</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>
FNT=number.	23g	<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>

$$\text{number} = \text{SY} + \text{RS} + (\text{VL} + \text{PR}) * \text{FM} + \text{SFA}$$

<u>Variable</u>	<u>Description</u>
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.

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 201g.

The maximum value for number is 7777g; the minimum value is 3.

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
FOT=number.	10g	<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 100g; the minimum value is 3. The size of the FOT need not be the same for each mainframe in a multimainframe environment.</p>
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 100g IPRDECKs can exist on a deadstart file; the number can range from 00 to 77g.</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 1100g. The value to be specified can be determined by using this formula:</p> $(3+lid)*pid$ <p>lid is the maximum number of LIDs allowed per PID.</p> <p>pid is the total number of PIDs in all networks defined on this mainframe.</p> <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 61g, central memory words that should be allocated to the LID table.</p>
LIB=number.	0	<p>Specifies the number of the LIBDECK to use at deadstart; the number can range from 00 to 77g. LIBDECK is a directive record used by SYSEDIT.</p>
MID=id.	AA	<p>Specifies the 2-character machine identification (id) that is associated with the mainframe. The id characters must be alphanumeric.</p>

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
MINCM=size.	1400g	Reserves an amount of central memory for system operation (UEM is defined in the extended memory equipment EST entry). The minimum and default value for size is 49K.

<u>Parameter</u>	<u>Description</u>
size	The 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; this value represents the amount of central memory in hundreds (octal); the value cannot be 0 and cannot exceed the total number of words of memory present in the machine.

<u>Memory</u>	<u>Size</u>
49K	1400
65K	2000
98K	3000
131K	4000
198K	6000
262K	10000

NAME=date line.	CDC 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.
-----------------	------------------------------------	--

<u>Parameter</u>	<u>Description</u>
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.	12g	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.
-------------	-----	--

<u>Parameter</u>	<u>Description</u>
number	Number of control points available in central memory; number can be from 2 to 34g.

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
OPSECM=n.	0	Specifies the operating system security mode. The value for this entry should be supplied by the site security administrator. Refer to the NOS 2 Security Administrator's Handbook for information on this entry.

<u>n</u>	<u>Description</u>
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.

PPB=number.	1	Reserves space in central memory for accessing files on buffered I/O devices.
-------------	---	---

<u>Parameter</u>	<u>Description</u>
number	The number of blocks (each block has a length of 1028 words) reserved for PP I/O. It must be sufficient to handle the demand for PP access to the buffered device I/O handler. Ideally, there should be one block for each pool PP. More blocks mean better performance. The maximum value for the number parameter is 178.

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
PPU=pp ₁ ,pp ₂ ,...,pp _n .	All available PPs are active	Toggles the active status of the physically available PPs, † 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.

<u>Parameter</u>	<u>Description</u>
------------------	--------------------

pp _i	Number of the PP; from 0 to 11g and 20g to 31g. Specifying an asterisk (*) on the entry instead of pp _i toggles the status of PPs 20g to 31g.
-----------------	--

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).

PPU=3,4.

QFT=number.	620g	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 7777g.
VERSION=name.	NOS 2.3 617/587	Specifies the system version that is displayed on the system console display.

<u>Parameter</u>	<u>Description</u>
------------------	--------------------

name	Alphanumeric string; must be fewer than 19 characters. If name is not terminated with a period, one will be appended. If name is 18 characters, the last character must be a period.
------	--

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

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.
- Multiframe 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 100g 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 2 Reference Set, Volume 3 for more information concerning these commands). EQPDECKs are named EQPDnn, where nn is from 00 to 77g.

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 deadstart is performed, except for a level 3 deadstart. That is, changes to the EQPDECK 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 EQPDECK on the deadstart file using SYSGEN (refer to the NOS 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 777g. This range depends on the value of ESMX (refer to PPCOM parameters in the NOS 2 Installation Handbook). Ordinals 1 through 4 are reserved for psuedoequipment EST entries that are automatically defined by the system and cannot be changed. Ordinal 5 is reserved for extended memory equipment (refer to EQ - Extended Memory EST Entry later in this section). This entry need not be defined, but cannot be used for any other device 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, Multiframe 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.

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>						
ACCOUNT=est,length.	,400g	<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><u>Parameter</u></th> <th><u>Description</u></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 100g and less than or equal to 1000g. If 0 (zero) is specified, messages issued to the account dayfile are discarded.</td> </tr> </tbody> </table>	<u>Parameter</u>	<u>Description</u>	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 100g and less than or equal to 1000g. If 0 (zero) is specified, messages issued to the account dayfile are discarded.
<u>Parameter</u>	<u>Description</u>							
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 100g and less than or equal to 1000g. If 0 (zero) is specified, messages issued to the account dayfile are discarded.							

<u>Entry Format</u>	<u>Released Default Value</u>
DAYFILE=est,length.	,400g

Significance

Sets the residence of the system dayfile and the length of the system dayfile buffer.

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.

<u>Parameter</u>	<u>Description</u>
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 100g and less than or equal to 1000g. If 0 is specified, messages issued to the system dayfile are discarded.

ERRLOG=est,length.	,100g
--------------------	-------

Sets the residence of the error log and the length of the error log buffer.

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.

<u>Parameter</u>	<u>Description</u>
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.
length	The 3- or 4-digit octal length of the error log buffer in CMR; must be a multiple of 100g and less than or equal to 1000g. If 0 is specified, messages issued to the error log are discarded.

<u>Entry Format</u>	<u>Released Default Value</u>	<u>Significance</u>
MAINLOG=est,length.	,100g	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.

<u>Parameter</u>	<u>Description</u>
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 100g and less than or equal to 1000g. If 0 (zero) is specified, messages issued to the binary maintenance log are discarded.

GENERAL DESCRIPTION OF EQUIPMENT ASSIGNMENT ENTRIES

The general format of the EQ entry is:

EQest=type,keyword₁=value₁,...,keyword_n=value_n.

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 3-1.

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.

Table 3-1. Required And Optional Keywords

Equipment Type	Required Keywords	Optional Keywords	Equipment Type	Required Keywords	Optional Keywords
CC	EQ,CH	ST	DV	UN,CH,HT	ST,AP
CM	EQ,CH	CW	DW	UN,CH,HT	ST,AP
CP	EQ,CH	ST,ID,FC	LP	EQ,CH	ST,ID,FC,TN,PS
CR	EQ,CH	ST,ID	LR	EQ,CH	ST,ID,FC,TN,PS
CS	EQ,UN,CH,CI,MI,UI,PT	ST	LS	EQ,CH	ST,ID,FC,TN,PS
CT	EQ,UN,CH,CP,MI,UI,PT	ST	LT	EQ,CH	ST,ID,FC,TN,PS
DB	UN,CH,HT	ST,AP	MP	CH	ST
DD	UN	ST,AP	MT	EQ,UN,CH,TF	ST
DE	SZ	ST,MA,ET,AP,MC	NC	CH	ST
DI	UN,CH	ST,AP	NP	EQ,CH,PI,ND,SA	ST
DJ	UN,CH	ST,AP	NT	EQ,UN,CH,TF	ST
DK	UN,CH	ST,AP	RM	CH,PT	ST
DL	UN,CH	ST,AP	RP	CH	ST
DM	UN,CH	ST,AP	SS	EQ,UN,CH	ST
DP	SZ,CH	ST,MA,ET,AP,MC	TT	EQ,CH	ST,PT
DQ	UN,CH	ST,AP	**	WO	WI

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 EQ=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 7778.

EQUIPMENT ASSIGNMENTS: NONMASS STORAGE

The following EST entries are described in this subsection.

- Nonstandard equipment
- Pseudoequipment
- Unit record equipment
- Magnetic tape equipment
- Mass Storage Facility equipment
- Mass Storage Archival Subsystem equipment
- Stimulator equipment
- Network processing unit
- 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
```

<u>Parameter</u>	<u>Description</u>
est	EST ordinal of the equipment; est can be from 6 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.

Pseudoequipment

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.

Ordinal 0 (RD)

This ordinal reserves an EST entry and MST entry to be used for device reconfiguration.

Ordinal 1 (DS)

This ordinal reserves an entry for the system display console.

Ordinal 2 (NE)

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.

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:

```
X.DIS.  
ASSIGN,NE,TAPE.  
PFDUMP.
```

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 2 Operations Handbook.

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 form 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, stores the number of network terminals.

Unit Record Equipment EST Entry

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

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

or

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

Parameter

Description

est EST ordinal of the unit record equipment; from 6 to 777g.

type Unit record equipment type; NOS supports the following unit record equipments.

<u>type</u>	<u>Equipment</u>
	Card reader
CR	405-3447/3649
	Card punch
CP	415-3446/3644
CP	415-30
	Line printer
	Any line printer
LP	580-12
LR	580-16
LS	580-20
LT	

P Specifies that a 580 printer is equipped with a programmable format controller.

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

<u>tn</u>	<u>Print Train</u>	<u>Description</u>
1	595-1/596-1	CDC 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 used as an ASCII 63/64-character set (as on a 595-5/596-5).

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 you specify 4, the actual value is 5.

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

<u>status</u>	<u>Description</u>
ON	Unit record equipment is available.
OFF	Unit record equipment is ignored during system operation.

EQ=eq Controller number for equipment; from 0 to 7. To determine the number, look at the dial switch on the controller.

Parameter

Description

CH=ch Number of channel to which unit record equipment is connected; from 0 to 13g and from 20g to 33g.

NOTES

When performing a coldstart, a card reader must be available on channel 12g or 13g.

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 67g. 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.

PS=ps Paper size, L or S, for local batch line printer. Default is L.

L Long (11 inch) paper will be mounted.

S Short (8 1/2 inch) paper will be mounted.

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. The minimum number of magnetic tape units that NOS requires is two 667s, 669s, 677s, or 679s. The format of the entry is:

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

or

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

<u>Parameter</u>	<u>Description</u>
est	EST ordinal of the tape unit; from 6 to 7778. Refer to the MT-n or NT-n parameter.
MT-n	Equipment type; MT specifies 7-track tape units, and NT specifies 9-track tape units. n is the total number of magnetic tape units connected to the controller, from 1 to 208 for 677 and 679 units with a 7021-31/32 controller and from 1 to 108 for 667 and 669 units with a 7021-21/22 controller. 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.
or	
NT-n	

For 639 units, n should not be specified; if it is specified, it should be given a value of 0 (zero).

NOTE

To clear an MT-n or NT-n assignment, enter an EQest=0 entry for all n units. 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:

<u>status</u>	<u>Description</u>
ON	Magnetic tape unit is available for access.
OFF	Magnetic tape unit is ignored during system operation.

Parameter

Description

NOTE

If a magnetic tape unit is accessible from more than one mainframe, concurrent use must be disabled by one of these methods:

- Set access switches on the tape controller so that only one mainframe can access the unit at any time.
- Make sure that the EST entry for each shared unit is defined as ON in only one of the mainframes.

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.
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 17g for 677 and 679 units. The unit number is on the front of the tape unit, either as a switch or a label.
CH=ch _i	Number of the channel to which the tape unit is connected; from 0 to 13g and from 20g to 33g. A controller can be connected to from one to four channels, depending on the controller model. However, a maximum of four channels can be handled regardless of the number of controllers.
TF=tf	Hardware features available. This parameter specifies the following hardware characteristics of the tape unit(s) being defined; select one option.

677/679 units:

<u>tf</u>	<u>Description</u>
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, indicating the availability of the GE feature.
ATSGE	The unit(s) being defined can process 6250-cpi GE tapes. Specify this value 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.

Parameter

Description

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.

639 units:

tf

Description

639

For 639 units, specify TF=639.

Mass Storage Facility Equipment EST Entry

The Mass Storage Facility (MSF) hardware consists of a cartridge storage unit (CSU) and from one to four mass storage transports. The CSU and each transport is represented by a unique EST entry. The EST entries for the transports on a CSU must immediately follow the EST entry for that CSU. There must be no other EST entry between the entry describing a CSU and those describing its transports.

A mainframe that runs MSSEEXEC in slave mode must have a CSU equipment type entry, even though no MSF hardware is actually configured for this mainframe.

CSU Entry

The CSU entry format is:

EQ=est,ST=status,EQ=eq,UN=un,PT=pt,CH=ch,CI=ci,MI=mi,UI=ui.

Parameter

Description

est

EST ordinal of CSU; from 6 to 7778. Although the software allows an ordinal of 7778, the ordinal must be lower than that value because the entries for transports must follow the CS entry.

CS

Indicates a CSU.

ST=status

Specifies whether the CSU is available for use; enter one of these values:

status

Description

ON

CSU is available.

OFF

CSU is ignored during system operation.

EQ=eq

Number of the mass storage coupler (MSC); from 0 to 7.

<u>Parameter</u>	<u>Description</u>
UN=un	Mass storage adapter (MSA) unit number; from 0 to 3.
PT=pt	Number of the MSA port to which the CSU is connected; from 0 to 7.
CH=ch	Number of the channel to which the CSU is connected; from 0 to 13g and from 20g to 33g.
CI=ci	Identifier of the CSU; one of the letters A through M.

NOTE

The following two parameters are the MSA and CSU identifiers that are hardwired into the MSA and CSU equipment.

MI=mi	Number from 0 to 17g. Obtain this number from the customer engineer. Use 0 if the FCO that supports this parameter is not installed.
UI=ui	Number from 0 to 377g. Obtain this number from the customer engineer. Use 0 if the FCO that supports this parameter is not installed.

Mass Storage Transport Entry

The format of the mass transport entry is:

EQest=CT,ST=status,EQ=eq,UN=un,PT=pt,CH=ch,CP=cp,MI=mi,UI=ui.

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of the transport; from 6 to 777g.						
CT	Indicates the mass storage transport.						
ST=status	Specifies whether the transport is available for use; enter one of these values:						
	<table> <thead> <tr> <th><u>status</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>Transport is available.</td> </tr> <tr> <td>OFF</td> <td>Transport is ignored during system operation.</td> </tr> </tbody> </table>	<u>status</u>	<u>Description</u>	ON	Transport is available.	OFF	Transport is ignored during system operation.
<u>status</u>	<u>Description</u>						
ON	Transport is available.						
OFF	Transport is ignored during system operation.						

<u>Parameter</u>	<u>Description</u>
EQ=eq	Number of MSC; from 0 to 7.
UN=un	MSA unit number; from 0 to 3.
PT=pt	Number of the MSA port to which the transport is connected; from 0 to 7.
CH=ch	Number of the channel to which the transport is connected; from 0 to 13g and from 20g to 33g.
CP=cp	Physical position of the transport in the CSU; from 0 to 3. Obtain this number from the customer engineer.

NOTE

The following two parameters are the MSA and transport identifiers that are hardwired into the MSA and transport equipment.

MI=mi	Number from 0 to 17g. Obtain this number from the customer engineer. Use 0 if the FCO that supports this parameter is not installed.
UI=ui	Number from 0 to 377g. Obtain this number from the customer engineer. Use 0 if the FCO that supports this parameter is not installed.

Mass Storage Archival Subsystem EST Entry

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

A mainframe that runs SSEXEC in slave mode must have an MSAS EST entry, even though no MSAS hardware is actually configured for this mainframe.

The MSAS EST entry format is:

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

<u>Parameter</u>	<u>Description</u>
est	EST ordinal of the 7990 controller; from 6 to 77g.
SS	Indicates a 7990 controller.

<u>Parameter</u>	<u>Description</u>										
ST=status	Specifies whether the 7990 controller is available for use; enter one of these values:										
	<table border="0"> <thead> <tr> <th><u>status</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>7990 controller is available.</td> </tr> <tr> <td>OFF</td> <td>7990 is ignored during system operation.</td> </tr> </tbody> </table>	<u>status</u>	<u>Description</u>	ON	7990 controller is available.	OFF	7990 is ignored during system operation.				
<u>status</u>	<u>Description</u>										
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:										
	<table border="0"> <thead> <tr> <th><u>eq</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0</td> <td>First 7990 controller.</td> </tr> <tr> <td>2</td> <td>Second 7990 controller.</td> </tr> <tr> <td>4</td> <td>Third 7990 controller.</td> </tr> <tr> <td>6</td> <td>Fourth 7990 controller.</td> </tr> </tbody> </table>	<u>eq</u>	<u>Description</u>	0	First 7990 controller.	2	Second 7990 controller.	4	Third 7990 controller.	6	Fourth 7990 controller.
<u>eq</u>	<u>Description</u>										
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 13g and from 20g to 33g. You can specify up to two channels.										

Stimulator Equipment EST Entry

NOS 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.

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of the interactive stimulator; from 6 to 77g.						
TT	Indicates an interactive stimulator.						
ST=status	Specifies whether the interactive stimulator is available for use; enter one of these values:						
	<table border="0"> <thead> <tr> <th><u>status</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>Stimulator is available.</td> </tr> <tr> <td>OFF</td> <td>Stimulator is ignored during system operation.</td> </tr> </tbody> </table>	<u>status</u>	<u>Description</u>	ON	Stimulator is available.	OFF	Stimulator is ignored during system operation.
<u>status</u>	<u>Description</u>						
ON	Stimulator is available.						
OFF	Stimulator is ignored during system operation.						

<u>Parameter</u>	<u>Description</u>
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 13g and from 20g to 33g. The channel/controller combination must not have any equipment attached to it.
PT=lines	Number of lines to stimulate; from 1 to 1000g. If this parameter is omitted, 100g is used.

Network Processing Unit EST Entry

The format of the EST entry is:

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

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of the NPU; from 6 to 77g.						
NP	Indicates a 255X NPU.						
ST=status	Specifies whether the NPU is available for use; enter one of these values:						
	<table border="1"> <thead> <tr> <th><u>status</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>NPU is available.</td> </tr> <tr> <td>OFF</td> <td>NPU is ignored during system operation.</td> </tr> </tbody> </table>	<u>status</u>	<u>Description</u>	ON	NPU is available.	OFF	NPU is ignored during system operation.
<u>status</u>	<u>Description</u>						
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 four NPUs).						
CH=ch	Number of the channel to which the NPU is connected; from 0 to 13g and from 20g to 33g.						
ND=node	Node number of coupler associated with the NPU being defined; from 1 to 77g. 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.						

<u>Parameter</u>	<u>Description</u>
SA=sa	System autostart module (SAM) flag. This parameter is required.

<u>sa</u>	<u>Description</u>
ON	SAM is present.
OFF	SAM is absent.

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 or remotely connected to more than one host, the NPU should have a SAM. The NPU EST entries for channel-connected hosts should specify SA=ON.

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.

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: COUPLER NODE=2,HNAME=HOST1.

CPL2: COUPLER NODE=8,HNAME=HOST1.

CPL3: COUPLER NODE=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 (10g), and 11 (13g).

Network Access Device EST Entry

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

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

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of the NAD; from 6 to 777g.						
NC	Indicates a NAD.						
ST=status	Specifies whether the NAD is available for use; enter one of these values:						
	<table><thead><tr><th><u>status</u></th><th><u>Description</u></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>	<u>status</u>	<u>Description</u>	ON	NAD is available.	OFF	NAD is ignored during system operation.
<u>status</u>	<u>Description</u>						
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 13g and from 20g to 33g.						

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.

<u>Parameter</u>	<u>Description</u>						
est	The 1- to 3-digit octal EST ordinal of the TPM; 6 to 777g.						
RM	TPM equipment type.						
ST=status	Specifies whether the TPM is available for use; enter one of these values:						
	<table><thead><tr><th><u>status</u></th><th><u>Description</u></th></tr></thead><tbody><tr><td>ON</td><td>TPM is available.</td></tr><tr><td>OFF</td><td>TPM is ignored during system operation.</td></tr></tbody></table>	<u>status</u>	<u>Description</u>	ON	TPM is available.	OFF	TPM is ignored during system operation.
<u>status</u>	<u>Description</u>						
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.						

MAP III or MAP 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.

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of MAP; from 6 to 77g.						
MP	Indicates MAP equipment.						
ST=status	Specifies whether the MAP is available for use; enter one of these values:						
	<table><thead><tr><th><u>status</u></th><th><u>Description</u></th></tr></thead><tbody><tr><td>ON</td><td>MAP is available.</td></tr><tr><td>OFF</td><td>MAP is ignored during system operation.</td></tr></tbody></table>	<u>status</u>	<u>Description</u>	ON	MAP is available.	OFF	MAP is ignored during system operation.
<u>status</u>	<u>Description</u>						
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 13g and from 20g to 33g.						

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.

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of 6683 coupler; from 6 to 777g.						
CC	Indicates 6683 coupler.						
ST=status	Specifies whether the 6683 coupler is available for use; enter one of these values:						
	<table><thead><tr><th><u>status</u></th><th><u>Description</u></th></tr></thead><tbody><tr><td>ON</td><td>6683 coupler is available.</td></tr><tr><td>OFF</td><td>6683 coupler is ignored during system operation.</td></tr></tbody></table>	<u>status</u>	<u>Description</u>	ON	6683 coupler is available.	OFF	6683 coupler is ignored during system operation.
<u>status</u>	<u>Description</u>						
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 13g and from 20g to 33g.						

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.

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of the ring port; from 6 to 777g.						
RP	Indicates ring port.						
ST=status	Specifies whether the ring port is available for use; enter one of these values:						
	<table border="0"> <thead> <tr> <th><u>status</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>ON</td> <td>Ring port is available.</td> </tr> <tr> <td>OFF</td> <td>Ring port is ignored during system operation.</td> </tr> </tbody> </table>	<u>status</u>	<u>Description</u>	ON	Ring port is available.	OFF	Ring port is ignored during system operation.
<u>status</u>	<u>Description</u>						
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 13g and from 20g to 33g. This parameter is required.						

EQUIPMENT ASSIGNMENTS: MASS STORAGE

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

<u>Equipment</u>	<u>Entry</u>
Disk units	EQ
Control module	EQ
Extended memory	EQ
<u>Equipment Assignment</u>	<u>Entry</u>
Mass storage allocation control	MSAL
Permanent files device	PF
System library device	SYSTEM
Alternate system library device	ASR
System checkpoint file device	SCKP
Logical device removal	DOWN
Default family name	FAMILY
Removable device	REMOVE
Shared device	SHARE

<u>Equipment Assignment</u>	<u>Entry</u>
Independent shared device	ISHARE
Load buffer controllers	LBC
Extended memory allocation	XM

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 3-2 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 3-2), 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 3-2. Mass Storage Device Functions

Other Possible Functions	Device Type								
	System	Alternate System	Containing Temporary Files	Containing Direct Access Files	Containing Indirect Access Files	Auxiliary	Default Family	Shared	Link
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	-	-

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 3-3.

Table 8-3. 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:

<u>Entries</u>	<u>Comments</u>
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:

<u>Entries</u>	<u>Comments</u>
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 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 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 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 for the orderly management of 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 and indirect access files 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 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).

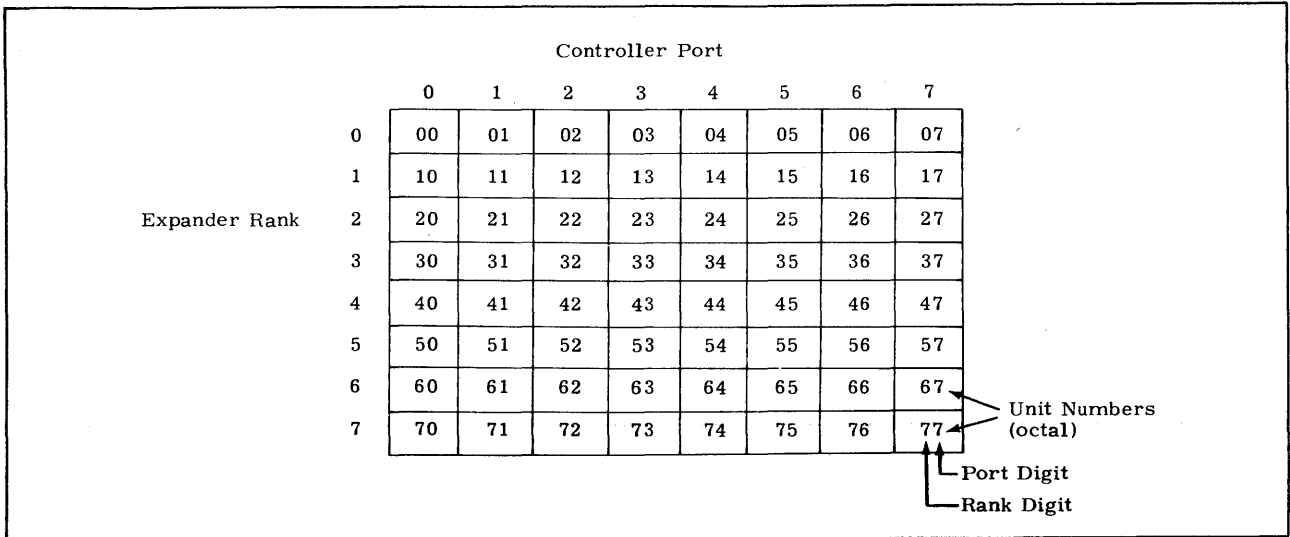


Figure 3-1. Expander Addressing Map

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, 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.

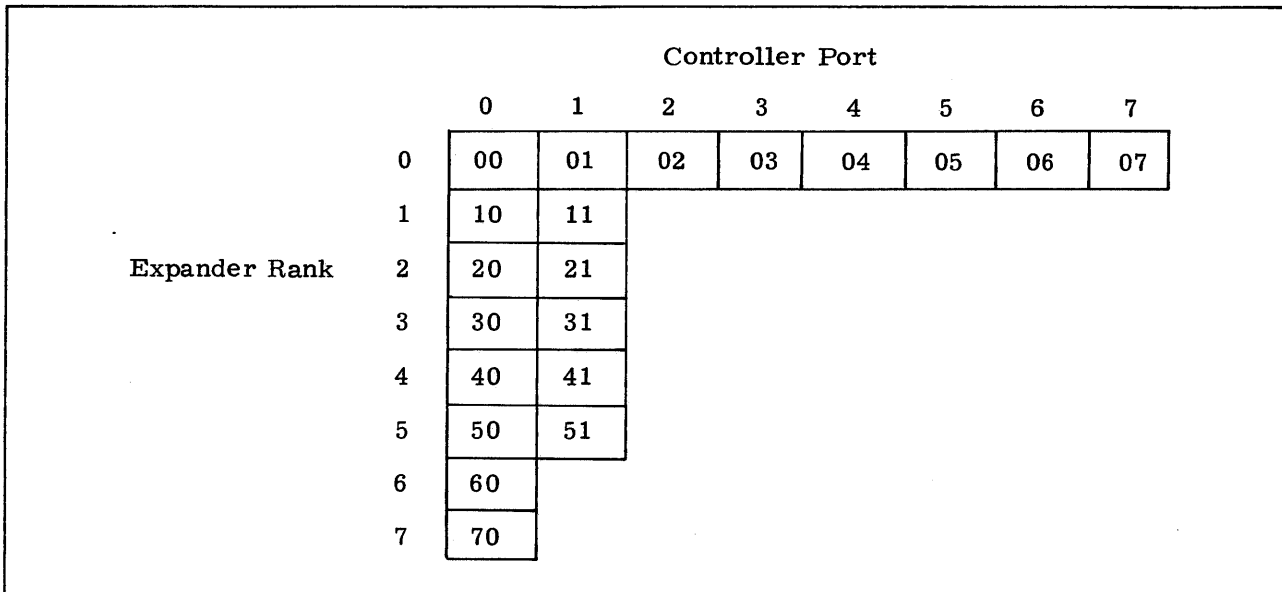


Figure 3-2. 844 Expander Configuration With 20 Drives

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.

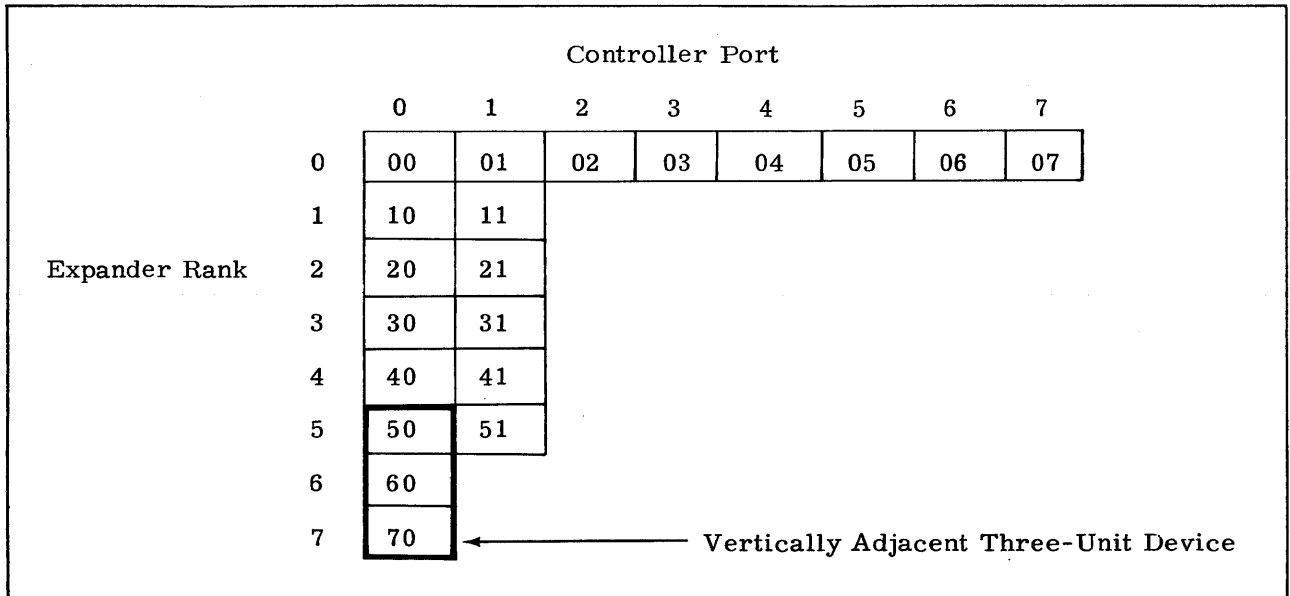


Figure 3-3. Vertically Adjacent Three-Unit Device

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.

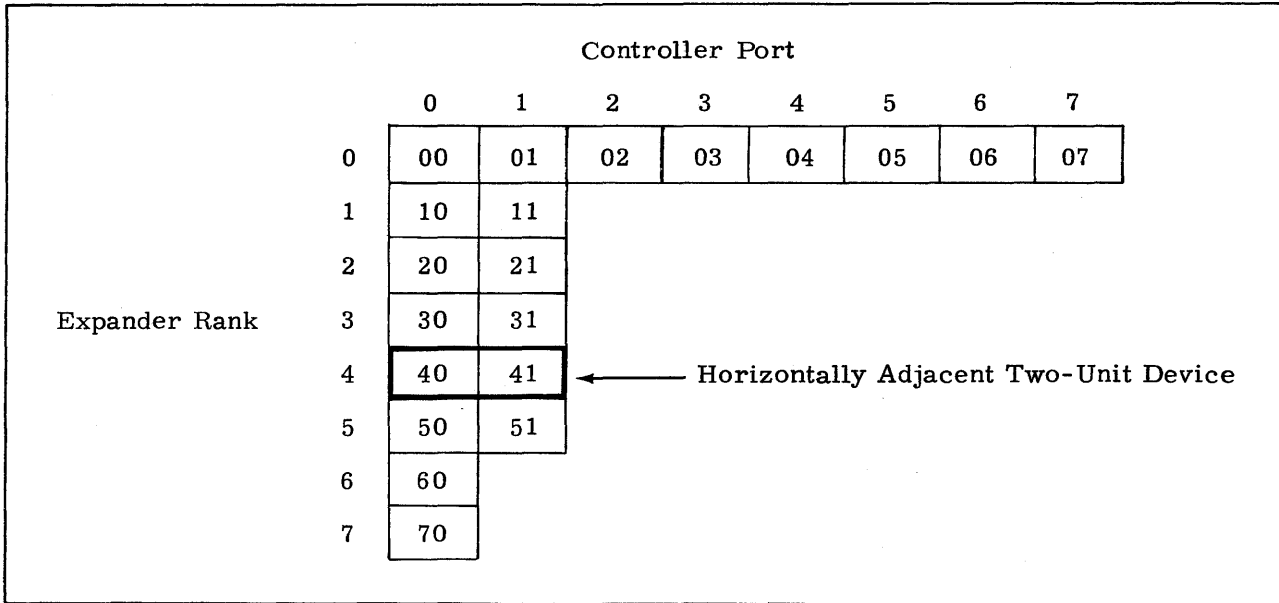


Figure 3-4. Horizontally Adjacent Two-Unit Device

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.

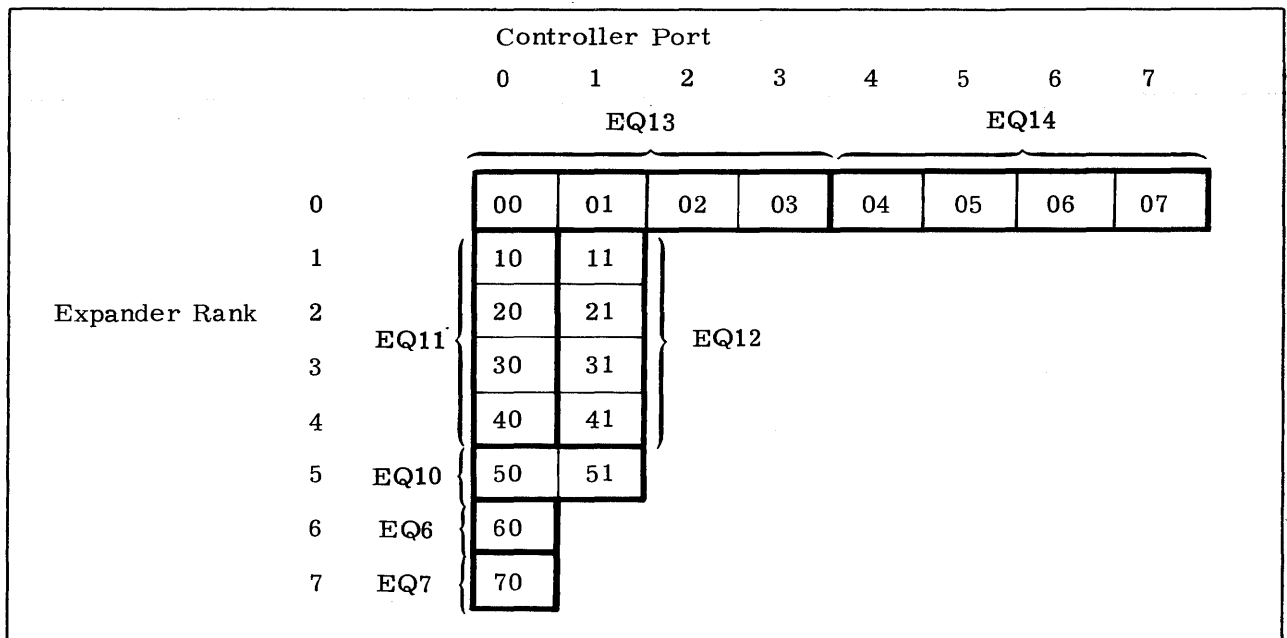


Figure 3-5. Hardware Configured into Seven Devices

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. It 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 Mass Storage Devices

Buffered mass storage devices are those devices for which the system provides an extended memory data buffer. For the model 176 only, an 819 mass storage device with the buffer in LCME is a buffered device. For all models except models 176 and CYBER 180-class machines, the buffered device is an 885-42 Mass Storage Subsystem with the buffer in ESM. 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 device I/O request processing. The system maintains the standard disk I/O interface to the user's programs.

The buffered device error reporting process logs read/write errors in the binary maintenance log but not in the system error log or dayfile. When an unrecovered read/write error occurs, the 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 following considerations are important to system performance.

- The amount of I/O buffer space in extended memory. Refer to XM - Declare Extended Memory Space for Buffers or User Access later in this section.
- The size of the hash buffer parameter. Refer to EQ - Mass Storage Equipment EST Entry later in this section.
- The amount of space reserved in central memory for providing PP access to buffered devices. Refer to the PPB=number entry under Central Memory Descriptions earlier in this section.

There may be situations when a particular user job runs slower with buffered devices. 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.

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 844 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, DOWN, and XM. 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=ch₁/ch₂,AP=ap,HT=ht.

<u>Parameter</u>	<u>Description</u>		
est	EST ordinal of the disk units; from 6 to 777 ₈ .		
type	Equipment type. The following disk units are supported by NOS.		
	<u>type</u>	<u>Equipment</u>	<u>Number of Units</u>
	DB†	885-42, 7155-401 (full track)	1 through 3
	DD††	834, 7255-1 (full track)	1 through 8
	DI	844-21, 7054/7154 (half track)	1 through 8
	DJ	844-41/44, 7054/7154 (half track), 7155-1	1 through 8
	DK	844-21, 7154 (full track)	1 through 8
	DL	844-41/44, 7154 (full track), 7155-1	1 through 8
	DM	885-11/12, 7155-1 (half track)	1 through 3
	DQ	885-11/12, 7155-1 (full track)	1 through 3
	DV†††	Single-density 819	1
	DW†††	Double-density 819	1

†Not applicable for CYBER 180-class machines.

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

†††Applicable to model 176 only.

Parameter

Description

ST=status Specifies whether or not the equipment is available for access; enter one of these values:

status

Description

ON	Equipment is available.
OFF	Equipment is not accessed during system operation. Specify OFF if the equipment is malfunctioning and access is not desirable.

If the equipment is removable and is not available at deadstart, the system determines that 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.

UN=units

Defines unit number(s).

For 834 units, unit numbers are generated as follows:

$$\text{unit number} = (\text{control module equipment number}) * 10g + (\text{834 disk equipment number})$$

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.

For 844 and 885 units, unit numbers are from 0 to 77g.

You can define the 834, 844, and 885 physical units with a separate EQ entry for each unit, or, if more continuous 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 is defining more than one unit of a multispindle device, unit numbers can be specified as UN=un₁-un_n, if units are contiguous, or as UN=un₁/un₂/.../un_n, if units are noncontiguous. For 834 units, contiguous means that the units are all on the same control module.

For example, three 844 units (0, 1, and 2) to be accessed as three units are defined as follows:

EQest=DI,ST=status,UN=0,CH=ch,AP=ap.

EQest=DI,ST=status,UN=1,CH=ch,AP=ap.

EQest=DI,CH=ch,UN=2,AP=ap,ST=status.

Parameter

Description

Three 844 units (0, 1, and 2) to be accessed as one logical unit are defined as follows:

EQest=DI,ST=status,UN=0-2,CH=ch,AP=ap.

or

EQest=DI,ST=status,UN=0/1/2,CH=ch,AP=ap.

An advantage to accessing the three units as one logical unit is that less space is used in CMR (662g words for the 844-21). A disadvantage is that if any of the units malfunctions or is destroyed, all units are affected.

CH=ch₁

For all devices except 819s and 834s, number of the channel or channels to which the controller is connected; from 0 to 13g and from 20g to 33g. For 819s, the channel pair for input/output (I/O) multiplexer (primary access); enter one of these values:

<u>ch₁</u>	<u>Description</u>
2	Channels 2 and 3.
4	Channels 4 and 5.
6	Channels 6 and 7.

For 834s, the channel number is specified on the control module EQ entry.

/ch₂

For all devices except 819s and 834s, indicates dual-access unit; cannot be 0; ch₁ and ch₂ should be connected to different controllers. Ch₂ is assigned first if it is free.

For 819s, the secondary access channel pair; one of the values described for ch₁.

For 834s, dual access is indicated on the control module EQ entry.

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.

HT=ht

Length of hash table segment for this unit; must be a power of 2 and from 10g to 400g. The hash table is maintained in extended memory. The hash table keeps track of the extended memory copies of disk segments for DV, DW, and DB disk units. The ht parameter should be as large as possible for maximum efficiency. If the hash buffer is too small, the system's ability to recognize disk data already in buffers is diminished, thus causing unnecessary disk reads.

EQ — Control Module EST Entry

The 834 Disk Storage Subsystem consists of from one to eight 834 disk units and from one to eight 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 disk unit is described under EQ - Disk Equipment EST Entry earlier in this section. The description of the EQ entry for the control module follows.

A control module for the 834 Disk Storage Subsystem is a controller that drives up to four 834 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 drives defined on the control module. Channel entries are not allowed on the 834 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.

<u>Parameter</u>	<u>Description</u>
est	EST ordinal of the control module; from 6 to 777g.
EQ=eq	Physical equipment number of the control module; from 0 to 7.
CH=ch ₁	Number of the channel to which the control module is connected; from 0 to 13g and from 20g to 33g.
/ch ₂	Indicates a dual-access control module; from 1 to 13g and from 20g to 33g.
CW=cw	Indicates if controlware is to be installed in the control module. 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:

<u>cw</u>	<u>Description</u>
CM	Install controlware (default).
NCM	Do not install controlware.

EQ — Extended Memory EST Entry

Three possible extended memory configurations exist. The first configuration is for ECS, LCME, and ESM. It allows a CPU (or two, if dual CPUs are available) to communicate with extended memory. The second configuration is the same as the first except additionally it uses the distributive data path (DDP) or low-speed port (LSP). This configuration allows PPs in the system, as well as the CPU, to communicate directly with ECS or ESM. It does not apply to LCME. NOS supports up to two DDPs or LSPs per mainframe. (Refer to both User Extended Memory and XM - Declare Extended Memory Space for Buffers or User Access later in this section for information regarding user-access extended memory.) For the third configuration, part of central memory is used as unified extended memory (UEM) on models 865 and 875 and CYBER 180-class machines.

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 TRACER utility 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. EST ordinal 5 cannot be used to define any other equipment.

The format of the entry is:

EQ5=type,ST=status,MA=mode,ET=xmem/ddp,SZ=size,CH=ch₁/ch₂,AP=ap,MC=mc.

Parameter

Description

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).

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.

Parameter

Description

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

Type of large memory and ddp. If you do not specify xmem, E1 is assumed. If you do not specify ddp, D1 is assumed if type is DP (if type is DE no value is assumed). If you specify ddp, you must also specify xmem.

xmem

Description

- E1 ECS I for all mainframes except models 176 and CYBER 180-class machines.
- E2 ECS II or ESM for all mainframes except models 176 and CYBER 180-class machines.
- 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 accessed through maintenance port. If ES is specified, you must also specify the MC=mc parameter.

If you specify MA=ON, the portion of ESM accessed for maintenance mode is the portion that is available to the system, not the portion that online diagnostics can reference.

ddp

Description

- D1 DC135 DDP.
- D2 DC145 (parity enhanced) DDP, or ESM low-speed port (LSP).

Parameter

Description

SZ=size Size of extended memory; one of the values in the following table for ECS I. For ECS II, ESM, and LCME, size ranges from a minimum value of 10g, which specifies 4K (K represents 1024 words) through a maximum value of 10000g, which specifies 2048K. For UEM, size ranges from a minimum value of 10g, which assigns 4K for UEM, through a maximum value of 10000g-MINCM, which assigns 2000K-MINCM (convert MINCM to decimal) for UEM. Size can be specified in increments of 1, if necessary (for example, 10g, 11g, 12g, 401g, or 2003g).

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

CH=ch₁,/ch₂ Numbers of the channels to which the DDP is connected; from 0 to 13g and from 20g to 33g.

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.

If the equipment type is DP, specify either one or two channels. The second channel cannot be 0 (zero). If a DDP is present, the loading of CPU programs residing in ECS or ESM still occurs via the CPU. A DDP must be connected to a channel by itself.

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 0 to 13 and from 20g to 33g. This channel is for ESM only. If MC=mc is specified, you must also specify MA=ON.

Examples:

- EQ5=DE,ST=OFF,SZ=1000,AP=21.
- EQ5=DP,ST=ON,SZ=2000,CH=2/3,AP=10.
- EQ5=DE,ST=ON,MA=ON,ET=E2,SZ=1000.
- EQ5=DP,ST=OFF,ET=E1/D2,SZ=2000,CH=2,AP=6.
- EQ5=DP,ST=ON,ET=ES/D2,SZ=10000,CH=5/6,MC=7.

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.

<u>Parameter</u>	<u>Description</u>																				
t	File type; one of these values: <table><thead><tr><th><u>t</u></th><th><u>Description</u></th></tr></thead><tbody><tr><td>B</td><td>LGO files.</td></tr><tr><td>D</td><td>Job dayfiles.</td></tr><tr><td>I</td><td>Input files.†</td></tr><tr><td>L</td><td>Local files.</td></tr><tr><td>O</td><td>Output files.†</td></tr><tr><td>P</td><td>Primary files.</td></tr><tr><td>R</td><td>Rollout files.</td></tr><tr><td>S</td><td>Secondary rollout files.</td></tr><tr><td>T</td><td>Temporary files.</td></tr></tbody></table>	<u>t</u>	<u>Description</u>	B	LGO files.	D	Job dayfiles.	I	Input files.†	L	Local files.	O	Output files.†	P	Primary files.	R	Rollout files.	S	Secondary rollout files.	T	Temporary files.
<u>t</u>	<u>Description</u>																				
B	LGO files.																				
D	Job dayfiles.																				
I	Input files.†																				
L	Local files.																				
O	Output files.†																				
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 57g. Ranges of ordinals can be specified.																				

NOTE

In most cases, using ESM as a secondary rollout device improves system performance. UEM, on the other hand, may not improve system performance since it takes more CPU overhead. Use the TRACER utility to determine the best use of system resources.

† 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 queue files to the specified devices.

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:

<u>Entries</u>	<u>Comments</u>
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.

PF — Permanent Files Device Assignment Entry

Before initializing a mass storage device (with the INITIALIZE entry in the EQPDECK), add a PF entry 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; 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.

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.

<u>Parameter</u>	<u>Description</u>						
est	EST ordinal of the device; from 5 to 777g.						
type	Type of device; one of these values:						
	<table border="0"> <thead> <tr> <th style="text-align: left;"><u>type</u></th> <th style="text-align: left;"><u>Description</u></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 377g. It can contain direct access files if the sm parameter is from 1 to 377g. 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>	<u>type</u>	<u>Description</u>	F	Family device. It can contain indirect access files if the dm parameter is from 1 to 377g. It can contain direct access files if the sm parameter is from 1 to 377g. 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.
<u>type</u>	<u>Description</u>						
F	Family device. It can contain indirect access files if the dm parameter is from 1 to 377g. It can contain direct access files if the sm parameter is from 1 to 377g. 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.						

Refer to table 3-4 for dependencies.

Table 3-4. PF Entry

Type of PF Device	Files Permitted on Device	PF Entry Parameter Settings				
		type	dm	name	device	sm
Auxiliary †	Indirect and/or direct	X	Omit	pack	Omit	Omit
Family	Direct only	F	0	family	1-778	1-3778
	Indirect only (master device)	F	1-3778	family	1-778	0
	Indirect and direct (master device)	F	1-3778	family	1-778	1-3778

† 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

dm Specifies the unit's device mask; from 0 to 3778. Set according to information 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 3778. Set according to information 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.

name 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.

The family name describes the permanent file devices available to a user.† A family may consist of from 1 to 63 logical devices; however, the master devices within the family must have device masks totaling 3778 if all possible user indexes are to be accommodated.

† 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 2 Reference Set, Volume 3 for a discussion of FA=A.

Parameter

Description

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.

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).

device

Number of the device in the family; from 1 to 778. Omit this parameter if the device is an auxiliary device.

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.

nc

Number of catalog tracks (optional) used only for master devices; from 1 to 200g. 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.

<u>Default</u>			<u>Default</u>		
<u>nc</u>	<u>Equipment</u>	<u>Type</u>	<u>nc</u>	<u>Equipment</u>	<u>Type</u>
40	834	DD	10	819	DV/DW
40	844-21	DI/DK	4	Extended memory	DE
40	844-41/44	DJ/DL	4	ECS or ESM with DDP	DP
10	885-11/12	DM/DQ			
10	885-42	DB	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 mass 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.

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. For example, if there are two system devices and the equipment EST ordinal for one of them specifies DI, the equipment EST ordinal for the other one must also specify DI.
- 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,...,est_n.

<u>Parameter</u>	<u>Description</u>
est _i	EST ordinal of the device to contain a copy of the system library on the deadstart file; from 5 to 7778. 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 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,...,est_n.

<u>Parameter</u>	<u>Description</u>
est _i	EST ordinal of mass storage device to be used as an alternative system device; from 5 to 777 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,...,est_n.

<u>Parameter</u>	<u>Description</u>
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.

DOWN — Logical Device Removal Entry

The DOWN entry specifies which EST ordinals will be logically removed (or down) at deadstart time. You can use this entry for all types of equipment that are defined as nonshared and removable. It is particularly useful either in allowing online diagnostics of nonshared extended memory or in removing mass storage devices that cannot be removed online (perhaps due to hardware malfunctioning).

The DOWN entry format is:

DOWN=est₁,est₂,...,est_n.

<u>Parameter</u>	<u>Description</u>
est _i	EST ordinal of the device to be logically removed; from 5 to 7778. Ranges of ordinals can be specified.

The DOWN entry has the following characteristics.

- It can be entered anywhere after the EQ entries for nonmass storage devices.
- For removable, nonshared mass storage devices, it can be entered anywhere after the EQ and REMOVE entries.
- Shared devices (including the link device) cannot be defined as DOWN in an MMF environment.
- It forces down and off status for the EST ordinal, regardless of what was specified in the EQ entry.
- The deadstart device cannot be defined as DOWN at deadstart time.
- If extended memory is not shared, it can be specified as DOWN at deadstart time, and it will remain DOWN until another level 0 deadstart is performed. Extended memory cannot be brought up online.
- DOWN cannot be specified for a device that has initialization requested.

If a control module is to be downed, all of the above characteristics must be true for each 834 disk unit defined on that control module.

DOWN entry examples follow.

```
EQ6=DI,ST=ON,UN=3,CH=32/26.  
REMOVE=6.  
DOWN=6.
```

```
EQ25=CR,ST=OFF,EQ=4,CH=12.  
EQ26=LP,ST=ON,EQ=6,CH=12.  
DOWN=25,26.
```

```
EQ5=DE,ST=OFF,SZ=400.  
DOWN=5.
```


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 (200g 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 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 family device.

The FAMILY entry format is:

FAMILY=est.

<u>Parameter</u>	<u>Description</u>
est	EST ordinal number of the mass storage device that the system automatically uses to determine your family when the you do not specify a family name at login or job initiation; from 5 to 777g.

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,T=, 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,...,est_n.

<u>Parameter</u>	<u>Description</u>
est _i	EST ordinal of mass storage device that is to be removable; from 6 to 777g. 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 or ESM) 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 - Present the Link Device Entry later in this section.

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

<u>Equipment</u>	<u>Type</u>
Extended memory	DE
ECS or ESM with DDP	DP
844-21	DI/DK
844-41/44	DJ/DL
885-11/12	DM/DQ

The format of SHARE is:

SHARE=est₁,est₂,...,est_n.

<u>Parameter</u>	<u>Description</u>
est _i	EST ordinal of the mass storage device being shared; from 5 to 777 ₈ . Ranges of ordinals can be specified.

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

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.

<u>Equipment</u>	<u>Type</u>
834	DD
844-21	DI/DK
844-41/44	DJ/DL
885-11/12	DM/DQ

The ISHARE entry format is:

ISHARE=est₁,est₂,...,est_n

<u>Parameter</u>	<u>Description</u>
est ₁	EST ordinal of the mass storage device shared; from 6 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:

<u>Device Type</u>	<u>Controlware Version Number</u>
DB †	MA722
DD ††	MA462
DI	MA710
DJ	MA710
DK	MA401
DL	MA401
DM	MA721
DQ	MA721

† Not applicable for CYBER 180-class machines.

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

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

LBC,type,c₁,c₂,...,c_n.

Parameter

Description

type	Controlware to be installed; one of these values.																												
	<table border="0"> <thead> <tr> <th style="text-align: left;"><u>type</u></th> <th style="text-align: left;"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>FM</td> <td>Install controller with full-track (MA721) controlware.</td> </tr> <tr> <td>FT</td> <td>Install controller with full-track (MA401) controlware.</td> </tr> <tr> <td>HT</td> <td>Install controller with half-track (MA710) controlware.</td> </tr> <tr> <td>ID</td> <td>Install 7255-1 adapter with full-track (MA462) controlware.</td> </tr> <tr> <td>M8</td> <td>Install 7990 controller with MB466 controlware.</td> </tr> <tr> <td>NF</td> <td>Identify channel(s) as having full-track controller, but do not install the controlware.</td> </tr> <tr> <td>NH</td> <td>Identify channel(s) as having half-track controller, but do not install the controlware.</td> </tr> <tr> <td>NI</td> <td>Identify channel(s) as having full-track 7255-1 adapter, but do not install controlware.</td> </tr> <tr> <td>NM</td> <td>Identify channel(s) as having full-track 7155-1 controller but do not install the controlware.</td> </tr> <tr> <td>NN</td> <td>Identify channel(s) as having NADs, but do not install the controlware.</td> </tr> <tr> <td>NP</td> <td>Identify channel as having 7155-401 controller, but do not install the controlware.</td> </tr> <tr> <td>N8</td> <td>Identify channel(s) as having 7990 controller, but do not install the controlware.</td> </tr> <tr> <td>PH</td> <td>Install 7155-401 controller with MA722 controlware.</td> </tr> </tbody> </table>	<u>type</u>	<u>Description</u>	FM	Install controller with full-track (MA721) controlware.	FT	Install controller with full-track (MA401) controlware.	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 as having 7155-401 controller, 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.
<u>type</u>	<u>Description</u>																												
FM	Install controller with full-track (MA721) controlware.																												
FT	Install controller with full-track (MA401) controlware.																												
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 as having 7155-401 controller, 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:

<u>Adapter Controller</u>	<u>Controlware Version Number</u>
7054	MA710
7152	MA710, MA401
7154	MA710, MA401
7155-1	MA721
7155-401	MA722
7255-1	MA462
380-170	MG101
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 UEM or buffered devices. If UEM is the only use for extended memory, you do not need to specify EST ordinal 5 in the EQPDECK.

The XM entry format is:

XM=id,iob,uec,EM.

<u>Parameter</u>	<u>Description</u>
id	Identifier of the mainframe that will access the reserved extended memory space; id is required. In a multimainframe complex, you must specify XM entries for each mainframe in the complex in all EQPDECKs, enabling the first mainframe to reserve ECS or ESM space for all mainframes in the complex.
iob	Number of words/1000 ₈ reserved for buffers for DB, DW, or DV devices; from 40 ₈ to 7777 ₈ . The optimum number of words to reserve is the amount left after the alternate system library and user ECS 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 7774 ₈ .
EM	For models 865 and 875 and CYBER 180-class machines, user access extended memory can be allocated in UEM without allocating UEM as an equipment. If you want only user-access extended memory in UEM, clear or do not specify EST ordinal 5 in the EQPDECK and enter XM with the uec size and the EM keyword. For models 865 and 875, if you specify EM, user access extended memory is allocated in UEM regardless of the presence of ECS or ESM; if you omit EM, user access extended memory is allocated in the device defined in the extended memory EST entry.

ACCESS

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 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,list.

<u>Parameter</u>	<u>Description</u>
lower	Access level name specified in deck COMSMLS (refer to the NOS 2 Installation Handbook) corresponding to the desired lower limit.
upper	Access level name specified in deck COMSMLS corresponding to the desired upper limit.
list	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

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

AUTOLOAD.

GRENADE

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 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.

The format of INITIALIZE is:

INITIALIZE,op,est₁,est₂,...,est_n.

<u>Parameter</u>	<u>Description</u>																		
op	Level of initialization; one of these values: <table><thead><tr><th><u>op</u></th><th><u>Description</u></th></tr></thead><tbody><tr><td>AF</td><td>Initialize account dayfile.</td></tr><tr><td>AL</td><td>Total initialization.</td></tr><tr><td>DF</td><td>Initialize system dayfile.</td></tr><tr><td>EF</td><td>Initialize error log.</td></tr><tr><td>FP</td><td>Initialize format pack (an automatic selection of AL also occurs).</td></tr><tr><td>MF</td><td>Initialize binary maintenance log.</td></tr><tr><td>PF</td><td>Initialize permanent files.</td></tr><tr><td>QF</td><td>Initialize queued files.</td></tr></tbody></table>	<u>op</u>	<u>Description</u>	AF	Initialize account dayfile.	AL	Total initialization.	DF	Initialize system dayfile.	EF	Initialize error log.	FP	Initialize format pack (an automatic selection of AL also occurs).	MF	Initialize binary maintenance log.	PF	Initialize permanent files.	QF	Initialize queued files.
<u>op</u>	<u>Description</u>																		
AF	Initialize account dayfile.																		
AL	Total initialization.																		
DF	Initialize system dayfile.																		
EF	Initialize error log.																		
FP	Initialize format pack (an automatic selection of AL also occurs).																		
MF	Initialize binary maintenance log.																		
PF	Initialize permanent files.																		
QF	Initialize queued files.																		
est _i	EST ordinal of mass storage device to be initialized; from 5 to 777 ₈ . 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 do not specify a PF entry when initializing a device, the default family name is SYSTid (where id is the machine identifier), and the default device numbers begin at 1 and increase by 1 for each device that you initialize without a PF entry. If you initialize equipment 5 without a PF entry, the device mask and secondary mask are set to 377₈. For all other equipment, the default masks are set to 0 (zero).

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.

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 300 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, and DAT) 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 two formats:

PRESET.

or

PRESET,n.

<u>Parameter</u>	<u>Description</u>
n	Number of shared devices; from 1 to 778.

If you specify n, space is allocated for the specified number of shared devices. Use this entry when the total number of shared devices is greater than the number of shared devices defined in the EQPDECK of the first mainframe in the multimainframe complex to do a level 0 deadstart.

If you do not specify n, the link device is preset, and the amount of table space reserved for the shared devices is determined by the number of shared entries in the EQPDECK.

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,...,est_n.

<u>Parameter</u>	<u>Description</u>
est ₁	EST ordinal of the ISHARE device; from 6 to 778.

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

RESET

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. If the DOWN entry was specified, EQest will be OFF, regardless of what was specified on the ST= parameter on the EQ entry. 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 100g 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
```

nn is the number identifying the APRDECK; from 00 to 77g.

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 FLAW 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
STK=4173.
STK=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.

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 F 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 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.

<u>Parameter</u>	<u>Description</u>
track	Logical track number. Refer to table 3-5 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.

<u>Parameter</u>	<u>Description</u>
track	Logical track number. Refer to table 3-5 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.

<u>Parameter</u>	<u>Description</u>
Address	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.
Address ₁ - Address 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 3-5 for the number of cylinders, tracks, and sectors for each device.

<u>Parameter</u>	<u>Description</u>
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 3-5. Information for Setting or Clearing Flaw Areas

Device	Device Type	Logical Track Range†	Cylinders/ Device †	Tracks/ Cylinder †	Sectors/ Track †
844-21, half track	DI	4000-7137	630	22	30
844-41/44, half track	DJ	4000-7147	1464	23	30
844-21, full track	DK	4000-7137	630	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
834	DD	4000-7133	1456	12	40
819, single density	DV	4000-5465	633	12	21
819, double density	DW	4000-7153	1466	12	24
885-42, full track	DB	4000-7221	1511	12	40
2048K extended memory	DE/DP	4000-7620			

† Numbers are in octal.

IPRDECK

The IPRDECK contains the system installation parameters that determine the system's operation mode. From 1 to 1008 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

The modified IPRDECK remains in effect only until the next deadstart is performed; that is, changes to the IPRDECK are not recovered across deadstart unless you create a new deadstart file to reflect those changes.

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 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 not retained after any level of recovery 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.

$$p = \frac{(ct-et)}{WF} + LP$$

<u>Variable</u>	<u>Description</u>
p	Priority; $LP \leq p \leq UP$. UP (a parameter on the QUEUE entry) is the highest priority.
WF	Weighting factor; a parameter on the QUEUE entry.
ct	Current time in seconds.
et	Time in seconds at which the job entered the queue.
LP	Lowest priority; a 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 10₈ 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.

Queue control for job execution determines how much continuous execution time a job gets. The amount of execution time depends both on the parameters specified on the QUEUE and SERVICE entries in the IPRDECK and whether other jobs are 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. When the job is selected for execution, it executes either until completion or until it is preempted by a higher priority job. When the job is preempted by a higher priority job, it goes to the rollout queue with the queue priority it had at the control point. While the job is in the rollout queue, its queue priority increases like any other job in the rollout queue. When a job exceeds its first time slice, its priority becomes the initial slice priority (IL; a parameter on the QUEUE entry). On expiration of subsequent time slices the job reenters the execution queue with the lowest priority (LP; a parameter on the QUEUE entry). The IL and LP parameters provide two opportunities for you to change the job's scheduling priority.

For the interactive service class, there is an additional execution queue priority, TP (a parameter on the QUEUE 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 the interactive service class, 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 to each newly initiated interactive job step. 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 table 3-6 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 table 3-6.

Job Control

The parameters on the SERVICE entry in the IPRDECK direct control of the jobs. The parameters set the following conditions:

- Initial CPU priority at job initiation.
- CPU time slice in milliseconds divided by 64.
- Central memory time slice in seconds.
- Number of jobs per service class.
- Time-out delay for interactive and detached jobs.

A job leaves a central memory control point when the following conditions are present:

- A job completes, aborts, or is suspended.
- A system request causes a job to be rolled out.

A system request includes a job request for a tape or disk pack, a ROLLOUT command from the generator, and execution of the ROLLOUT macro. Whenever a job rolls out, it is assigned whatever priority it has at that point. 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 becomes 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 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 gives 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.

- Terminal input/output is required.

A job leaves central memory when the following conditions occur:

- The system requests terminal input.
- 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 central processor or central memory time each type of job can use when it is at a control point. This ensures that one job does not monopolize system resources.

If a job exceeds either the central processor or central memory time slice, and it is not a subsystem, the scheduling priority is set to the lower bound priority (the IL 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.

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. Each control point needs 200g 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, then 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. 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.

Memory Control

You can control the maximum memory allowed for job types and for service classes with the parameters you specify on the SERVICE entry in the IPRDECK.

These parameters specify these lengths:

- Maximum field length divided by 100g for a job in a service class.
- Maximum field length divided by 100g for all jobs of the specified service class.
- Maximum extended memory length in words divided by 1000g for a job in a service class.
- Maximum extended memory length in words divided by 1000g for all jobs of the specified service class.

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 of these memory constraints, it attempts to schedule a job a second time. During this second attempt, any job that requires other jobs to be rolled out is not scheduled; otherwise, the constraints are ignored, and the job is scheduled at the lower bound priority, IL (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 because batch jobs are at a maximum and no other jobs are available, the scheduler attempts to schedule the batch jobs.

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 table 3-6. These entries fall within the example of ranges of service class priorities shown in figure 3-6. 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. Most system jobs and all subsystems are coded to adjust their priorities correctly, once execution begins. 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 batch, remote batch, and detached jobs are similar. The time slice for detached jobs is shorter than for the batch and remote batch jobs. The assumption is that detached jobs need less CPU time than either 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. Table 3-6 shows almost no overlaps of values except that 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.

Table 3-6. Example Set of Job Control Parameters†

Service Class	Input Queue QUEUE Parameters			Execution Queue QUEUE Parameters					Interactive Job Initial Scheduling TP QUEUE Parameter
	LP	UP	WF	LP	UP	WF	IP	IL	
SY (system)	7770	7776	1	2000	7000	1	7000	4000	-
BC (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)††	10	4000	1	1000	4004	1	2000	2000	-
Service Class	Output Queue QUEUE Parameters			Time Slice SERVICE Parameters		CPU Priority PR SERVICE Parameter	Number of Jobs NJ SERVICE Parameter	Time-Out Delay TD SERVICE Parameter	
	LP	UP	WF	CP	CM				
SY (system)	7000	7776	1	100	20	30	7777	-	
BC (batch)	1	7000	1	400	200	30	7777	-	
RB (remote batch)	1	7000	1	400	200	30	7777	-	
CT (communication task)	7000	7776	1	400	200	30	7777	-	
TS (interactive)	1	7000	1	40	10	30	7777	113†††	
NS (network supervision)	1	7000	1	400	200	31	7777	-	
DI (detached)	1	7000	1	100	20	30	7777	341††††	
SS (subsystem)	7400	7776	1	100	20	30	7777	-	
MA (maintenance)	7000	7776	1	100	20	2	7777	-	
I _n (installation)††	1	7000	1	400	200	30	7777	-	

††All values are octal; DELAY parameters are JS=1, CR=30, AR=1750, MN=10, MX=20, and JQ=2.
†††Installation classes I₁, I₂, I₃, and I₄ all have release values the same as class BC.
††††A TS suspended job times out after 10 minutes when the time-out delay is 113.
†††††A DI suspended job times out after 30 minutes when the time-out delay is 341.




Figure 3-6. Example of Ranges of Service Class Priorities

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.

IPRDECK ENTRIES MADE ONLY DURING DEADSTART

You can enter the following IPRDECK entries on the IPRDECK 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.

<u>Parameter</u>	<u>Description</u>
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.
D	If D is specified, the associated username and family are deleted from the common library table.

CPM

Format :

CPM, s₁=n₁, s₂=n₂.

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 2 Administration Handbook for a discussion of multiplier use.) The values of n_i range from 1 to 45₈ and are used as indexes to values defined in COMSSRU in order to determine the multiplier value. The default values are listed in table 3-7.

Table 3-7. CPM Default Values for n_i (Sheet 1 of 2)

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

Table 3-7. CPM Default Values for n_i (Sheet 2 of 2)

n_i (Octal)	COMSSRU Name	Mainframe Model	COMSSRU Default Multiplier Value
20	C720	720	1.0
21	C730	730	1.0
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	C845	845	1.0
33	C855	855	1.0
34	C865	865	1.0
35	C875	875	1.0
36 through 40	Reserved for Control Data		
41	ICM1	Model on which you are installing.	1.0
42	ICM2	Model on which you are installing.	2.0
43	ICM3	Model on which you are installing.	3.0
44	ICM4	Model on which you are installing.	4.0
45	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.

<u>csm</u>	<u>Description</u>
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.

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. 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.

<u>Parameter</u>	<u>Description</u>
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.

Example:

DSD,0,MAI%X.QREC(PO=N)

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 2 Reference Set, Volume 4 for a description of instruction-stack purging.

KEYPM

Format:

KEYPM=cc.

Default:

26

Significance:

This entry specifies the keypunch mode to be assumed during system operation.

<u>cc</u>	<u>Description</u>
26	026 keypunch mode.
29	029 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:

200g

Significance:

This entry specifies the number of network terminals that can be connected to IAF at one time.

<u>Parameter</u>	<u>Description</u>
maxt	Total number of network terminals; maxt can range from 1 to 1440g.

PROBE

Formats:

ENABLE,PROBE.

DISABLE,PROBE.

Default:

Disabled.

Significance:

These entries enable and disable the data gathering facility of CPUMTR.

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, MAS, MCS, MSS, NAM, RBF, RHF, 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 2 Security Administrator's Handbook.

Parameter

Description

cat ₁	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.

<u>Parameter</u>	<u>Description</u>
d	Density in lines per inch; d can be 6 or 8 lines per inch.
lines	Lines per page; lines can range from 16 to 255.

SPD

Format:

SPD=d.

Default:

6

Significance:

This entry specifies the assumed density for printed output.

<u>d</u>	<u>Description</u>
6	6 lines per inch.
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.

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.

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 conversion mode to be assumed during system operation.

<u>mode</u>	<u>Description</u>
AS	ASCII 9-track conversion.
US	ANSI (previously known as USASI) 9-track conversion (same as AS).
EB	EBCDIC 9-track conversion.

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.

<u>density</u>	<u>Description</u>
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).

TDTR

Format:

TDTR=tracktype.

Default:

NT

Significance:

This entry sets the default track type.

<u>tracktype</u>	<u>Description</u>
MT	7 track.
NT	9 track.

DEADSTART ENTRIES AND DSD COMMANDS

You can enter the following IPRDECK entries on the IPRDECK 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

Before you can enable or disable a subsystem that uses a system control point, 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 must enable a subsystem before you can initiate it. 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 2 Operations Handbook.)

The IPRDECK entries that enable and disable subsystems also assign a default 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 default control point for the subsystem.

Once you have enabled a subsystem, you can initiate the subsystem using either of the following methods.

- Call the subsystem procedure file by entering the name of the procedure file.
- Name the subsystem procedure file the same name as the subsystem and store the file as an indirect access permanent file under the system user name (SYSTEMX) with user index 377777g. Then, the subsystem will be automatically initiated when you enter the AUTO or MAINTENANCE commands.

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

BIO	Central site batch I/O for line printers, card readers, and card punches.
CDC	CYBER Database Control System.
IAF	Interactive Facility. Do not specify the cp parameter on the ENABLE entry for IAF.

TDTR**Format:**

TDTR=tracktype.

Default:

NT

Significance:

This entry sets the default track type.

<u>tracktype</u>	<u>Description</u>
MT	7 track.
NT	9 track.

DEADSTART ENTRIES AND DSD COMMANDS

You can enter the following IPRDECK entries on the IPRDECK 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

Before you can enable or disable a subsystem that uses a system control point, 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 must enable a subsystem before you can initiate it. 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 2 Operations Handbook.)

The IPRDECK entries that enable and disable subsystems also assign a default 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 default control point for the subsystem.

Once you have enabled a subsystem, you can initiate the subsystem using either of the following methods.

- Call the subsystem procedure file by entering the name of the procedure file.
- Name the subsystem procedure file the same name as the subsystem and store the file as an indirect access permanent file under the system user name (SYSTEMX) with user index 377778. Then, the subsystem will be automatically initiated when you enter the AUTO or MAINTENANCE commands.

Use these IPRDECK entries to enable or disable a subsystem:

ENABLE,subsystem,cp.

DISABLE,subsystem,cp.

<u>Parameter</u>	<u>Description</u>
subsystem	Three characters that select the desired subsystem; one of these values:

<u>subsystem</u>	<u>Description</u>
BIO	Central site batch I/O for line printers, card readers, and card punches.
CDC	CYBER Database Control System.
IAF	Interactive Facility. Do not specify the cp parameter on the ENABLE entry for IAF.

Parameter

Description

subsystem

Description

MAG	Magnetic Tape Subsystem. Enable MAG if removable auxiliary packs are used. Disabling MAG frees a control point for other use.
MAP	MAP III or MAP IV.
MAS	Mass Storage Archival Subsystem.
MCS	Message Control System.
MSS	Mass Storage Subsystem.
NAM	Network Access Method.
RBF	Remote Batch Facility.
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	STIMULA. Do not specify the cp parameter on the ENABLE entry for STM.
TAF	Transaction Facility.

NOTES

When MCS and RBF are started by the NAMI startup master file, the ENABLE and DISABLE commands are ignored.

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 control point.

Omit cp for IAF, RDF, and STM.

CLASS

Format:

CLASS,ot,sc1,sc2,...,scn.

Default:

<u>Origin Type</u>	<u>Service Class</u>
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 2 Systems Programmer's Instant).

<u>Parameter</u>	<u>Description</u>
ot	Origin type; must be BC (batch), RB (remote batch), or IA (interactive). This parameter is required.
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:

<u>sc</u>	<u>Description</u>
BC	Batch.
CT	Communication task.
DI	Detached interactive.
I0	Installation class 0.
I1	Installation class 1.
I2	Installation class 2.
I3	Installation class 3.
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SS	Subsystem.

<u>Parameter</u>	<u>Description</u>
<u>sc</u>	<u>Description</u>
SY	System.
TS	Interactive.
ALL	Causes validation of all service classes except subsystem.
NUL	Clears validation for all service classes.

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,CRcr,ARar,MXmx,MNmn,JQjq,JSjs.

Default:

None.

Significance:

This entry specifies the system delay parameters. Refer to table 3-6 for an example set of parameter entries. Refer to figure 3-6 for an example of ranges of service class priorities.

<u>Parameter</u>	<u>Description</u>
CRcr	CPU recall delay in milliseconds; cr is from 1 to 7777g.
ARar	PP/auto recall delay in milliseconds; ar is from 1 to 7777g.
MXmx	Maximum job switch delay in milliseconds; mx is from 1 to 7777g.
MNmn	Minimum job switch delay in milliseconds; mn is from 1 to 7777g.
JQjq	Exponent used to determine the input job (QFT to EJT) scheduling delay; jq is from 0 to 14g seconds. The delay in seconds between the scheduling of input jobs is calculated as follows: delay=2jg
JSjs	Job scheduler delay in seconds; js is from 1 to 7777g.

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.

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.

FILE STAGING

Formats:

ENABLE,FILE STAGING.

DISABLE,FILE STAGING.

Default:

Disabled.

Significance:

These entries specify whether permanent files that reside on the alternate storage facility (MSS or MSAS) are staged to disk. If disabled, jobs attempting to access MSS/MSAS resident files are aborted.

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 2 Operations Handbook.

MASTER MAS

Formats:

ENABLE,MASTER MAS

DISABLE,MASTER MAS

Default:

Disabled.

Significance:

These entries specify whether the MSAS executive program (SSEXEC), when initialized, is to run in master (enabled) or slave (disabled) mode.

MASTER MSS

Formats:

ENABLE,MASTER MSS.

DISABLE,MASTER MSS.

Default:

Disabled.

Significance:

These entries specify whether the MSS executive program (MSSEXEC), when initialized, is to run in master (enabled) or slave (disabled) mode.

FILE STAGING

Formats:

ENABLE,FILE STAGING.

DISABLE,FILE STAGING.

Default:

Disabled.

Significance:

These entries specify whether permanent files that reside on the alternate storage facility (MSS or MSAS) are staged to disk. If disabled, jobs attempting to access MSS/MSAS resident files are aborted.

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 2 Operations Handbook.

MASTER MAS

Formats:

ENABLE,MASTER MAS

DISABLE,MASTER MAS

Default:

Disabled.

Significance:

These entries specify whether the MSAS executive program (SSEXEC), when initialized, is to run in master (enabled) or slave (disabled) mode.

MASTER MSS

Formats:

ENABLE,MASTER MSS.

DISABLE,MASTER MSS.

Default:

Disabled.

Significance:

These entries specify whether the MSS executive program (MSSEXEC), 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 60g 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 2 Installation Handbook) that corresponds to the desired output queue special handling level.

PCLASS

Format:

PCLASS,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 2 Reference Set, Volume 3).

<u>Parameter</u>	<u>Description</u>
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 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,LP1p,UPup,IPip,WFwf,IL1l,TPtp.

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 3-6 for an example set of parameter entries and to figure 3-6 for an example of ranges of service class priorities.

<u>Parameter</u>	<u>Description</u>
sc	Service class.
<u>sc</u>	<u>Description</u>
BC	Batch.
CT	Communication task.
DI	Detached interactive.
I0	Installation class 0.
I1	Installation class 1.
I2	Installation class 2.
I3	Installation class 3.
MA	Maintenance.

Parameter

Description

<u>sc</u>	<u>Description</u>
NS	Network supervisor.
RB	Remote batch.
SS	Subsystem.
SY	System.
TS	Interactive.

qt Job queue type.

<u>qt</u>	<u>Description</u>
EX	Executing jobs.
IN	Input queued files.
OT	Output queued files.

LPip Lowest priority at which a file or job can enter the specified queue; lp ranges from 0 to 7777₈.

UPup Highest priority a job or queued file can reach in the specified queue; up ranges from 0 to 7777₈.

IPip Initial priority only for executing jobs (EX). ip ranges from 0 to 7777₈.

WFwf Weighting factor for queue priority calculation; wf must be a power of 2, from 0 to 4000₈. The smaller the weighting factor, the faster the queue entry reaches its highest priority.

ILil Lower bound priority for the initial time slice; il ranges from 0 to 7777₈. This parameter is valid only for executing jobs (EX).

TPtp Initial scheduling priority for online interactive jobs; tp ranges from 0 to 7777₈.

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=la,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 la 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.

<u>Parameter</u>	<u>Description</u>
ot	Origin type.
	<u>ot</u> <u>Description</u>
	SY System.
	BC Batch.
	IA Interactive.
	RB Remote Batch.
la	Access level name specified in deck COMSMLS (refer to the NOS 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,PRpr,CPcp,CMcm,NJnj,FLfl,AMam,ECec,EMem,FCfc,CScs,FSfs,DSds,TDtd,DTsc.

Default:

None.

Significance:

This entry specifies the service limits associated with each service class. Refer to table 3-6 for an example set of parameter entries and to figure 3-6 for an example of ranges of service class priorities.

<u>Parameter</u>	<u>Description</u>																												
sc	Service class; one of these values: <table><thead><tr><th><u>sc</u></th><th><u>Description</u></th></tr></thead><tbody><tr><td>BC</td><td>Batch.</td></tr><tr><td>CT</td><td>Communication task.</td></tr><tr><td>DI</td><td>Detached interactive.</td></tr><tr><td>I0</td><td>Installation class 0.</td></tr><tr><td>I1</td><td>Installation class 1.</td></tr><tr><td>I2</td><td>Installation class 2.</td></tr><tr><td>I3</td><td>Installation class 3.</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>	<u>sc</u>	<u>Description</u>	BC	Batch.	CT	Communication task.	DI	Detached interactive.	I0	Installation class 0.	I1	Installation class 1.	I2	Installation class 2.	I3	Installation class 3.	MA	Maintenance.	NS	Network supervisor.	RB	Remote batch.	SS	Subsystem.	SY	System.	TS	Interactive.
<u>sc</u>	<u>Description</u>																												
BC	Batch.																												
CT	Communication task.																												
DI	Detached interactive.																												
I0	Installation class 0.																												
I1	Installation class 1.																												
I2	Installation class 2.																												
I3	Installation class 3.																												
MA	Maintenance.																												
NS	Network supervisor.																												
RB	Remote batch.																												
SS	Subsystem.																												
SY	System.																												
TS	Interactive.																												
PRpr	CPU priority; pr ranges from 2 to 778. All service classes except network supervisor (NS) and maintenance (MA) jobs are normally set to the same CPU priority. Since jobs with lowest priority access the CPU last, MA jobs are usually set to the lowest priority to prevent them from interfering with other system activity.																												
CPcp	CPU time slice/1008 in milliseconds. This parameter specifies the maximum amount of time a job of the specified service class can use the CPU before its scheduling priority is set to its lower bound priority. cp must be an octal number.																												

<u>Parameter</u>	<u>Description</u>
CMcm	Central memory time slice in seconds; cm ranges from 0 to 7777g. This parameter specifies the maximum amount of time a job of the specified service class can remain at a control point before it becomes eligible to be rolled out.
NJnj	Maximum number of jobs; nj ranges from 0 to 7777g. This parameter specifies the number of jobs that can be active at one time in the specified service class.
FLfl	Maximum field length/100g for any job of the specified service class; fl ranges from 0 to 7777g. Jobs with field length requirements that exceed this value are not considered for scheduling, unless no other jobs are to be run and sufficient unused memory is available to run the job without rolling out a running job. However, an interactive service class job that exceeds its maximum field length aborts. 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.
AMam	Maximum field length/100g for all jobs of the specified service class; am ranges from 0 to 7777g. This parameter partitions central memory by limiting the field length available to each service class. For example, if a job whose field length exceeds that specified for its service class is scheduled to a control point, 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 normally exceed the field length for its service class can be scheduled to a control point if not enough jobs exist to fill the field length specified for another service class. The system attempts to use central memory to its greatest capacity.
ECec	Maximum extended memory length in words divided by 1000g for any job of the specified service class; ec ranges from 0 to 7777g.
EMem	Maximum extended memory length in words divided by 1000g for all jobs of the specified service class; em ranges from 0 to 7777g.
FCfc	Number of permanent files allowed. fc indicates a limit value that is the maximum number of permanent files allowed.

<u>fc</u>	<u>Limit Value</u>
0	Unlimited
1	10g
2	20g
3	30g
4	40g
5	50g
6	100g
7	Unlimited

Parameter

Description

CScs

Cumulative size in PRUs allowed for all indirect access permanent files. cs indicates a limit value for the cumulative size.

<u>cs</u>	<u>Limit Value</u>
0	Unlimited
1	1000g
2	2000g
3	5000g
4	10000g
5	50000g
6	100000g
7	Unlimited

FSfs

Size in PRUs allowed for individual indirect access permanent files. fs indicates a limit value for the size of the files.

<u>fs</u>	<u>Limit Value</u>
0	Unlimited
1	10g
2	30g
3	50g
4	100g
5	150g
6	300g
7	Unlimited

Parameter

Description

DSds Size in PRUs allowed for individual direct access permanent files.
ds indicates the limit value for the size of the files.

<u>ds</u>	<u>Limit Value</u>
0	Unlimited
1	1000g
2	2000g
3	5000g
4	10000g
5	50000g
6	100000g
7	Unlimited

TDtd Suspension time-out delay; td ranges from 0 to 7777g. A suspended job will not be timed out for td*10g seconds. The maximum delay is approximately 9 hours.

DTsc Service class to which a job of the service class specified by sc will be assigned if the job is detached. The default is detached interactive (DI).

SPINDOWN

Formats:

ENABLE, SPINDOWN.

DISABLE, SPINDOWN.

Default:

Enabled.

Significance:

These entries enable or disable the spindown of 834 disk units. When spindown is enabled, entering a CHECK POINT SYSTEM command will cause all 834 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 disk units.

SRST

Format:

SRST=n.

Default:

0

Significance:

This entry specifies the secondary rollout sector threshold. Any rollout file smaller than n sectors ($0 \leq n \leq 77778$) 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.

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 1008 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 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.

<u>Directive Format</u>	<u>Significance</u>
*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 376000g central memory words. You cannot use more than approximately 326000g central memory words (376000g-50000g 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 100g).
*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.
*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 2 Reference Set, Volume 3).
*/ comments	Defines comment lines.

The following record types may be specified in SYSEDIT directives. Some directives do not allow all types.

<u>Type (ty₁)</u>	<u>Description</u>
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.
REL	Relocatable central processor program.
TEXT	Unrecognizable as one of the other types.
ULIB	User library program.

DIS Job Dayfile Display (A)	4-3
DIS Exchange Package Display (B)	4-4
DIS Memory Displays (C, D, F, G)	4-5
DIS Job Status Display (J)	4-7
DIS Directory Display (Z)	4-9
Console Operation	4-10
Display Selection Commands	4-12
DIS Keyboard Entries	4-13
Memory Entry Commands	4-16
PP Call Commands	4-18

DIS displays information about a single job. Under DIS, the B 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,f1 (f1=field length desired) or X.DIS (field length of 60000g assumed by default). This brings DIS to an empty control point to initiate utility programs.
- Enter a command in the form DIS (the job must be system origin or have system origin privileges).
- On a secured or unsecured system, before you can call DIS, the console must be unlocked (refer to the DSD UNLOCK command in section 5). For a secured system the console must be unlocked by the security administrator. When the system display console shows UNLOCK status or SECURITY UNLOCK status, you can call DIS by entering the command:

DIS,jsn.

jsn is the job sequence name to which you want DIS assigned.

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 J 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,† the job is stopped after each command is processed. That is, only one command can be processed at a time. Under DIS, the J 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 B 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.

†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:

<u>Display Identifier</u>	<u>Description</u>
A	Job dayfile. Messages and files (local FNT entries) attached to the job.
B	Exchange package. Breakpoint address and the exchange package.
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 equivalents.
H	File status. All files assigned to the job as well as equipment assigned to files.
J	Job status. Individual job status, equipment assigned, current messages, and command buffer.
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.
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.
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 C, D, J, 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 J 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).




Figure 4-1. DIS Job Dayfile Display (A)

The level field, shown in the figure, is displayed only on a secured system.

DIS EXCHANGE PACKAGE DISPLAY (B)

The DIS B 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 4-2 shows the DIS exchange package display (B).

DIS B. 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 4-2. DIS Exchange Package Display (B)

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.							
Address	Memory contents				COMPASS mnemonic translation		
00000000	0000	0000	0000	0000	PS		PS
00000001	05160	42000	00000	00000	NE	BB	PS
00000002	23312	32405	15000	00001	AX	XB DXX+X	BXX*-X PS
00000003	00000	00000	00000	00000	PS		PS
00000004	00000	00000	00000	00000	PS		PS
00000005	00000	00000	00000	00000	PS		PS
00000006	00000	00000	00000	00000	PS		PS
00000007	00000	00000	00000	00000	PS		PS
00000010	00000	00000	00000	00000	PS		PS
00000011	00000	00000	00000	00000	PS		PS
00000012	00000	00000	00000	00000	PS		PS
00000013	00000	00000	00000	00000	PS		PS
00000014	00000	00000	00000	00000	PS		PS
00000015	00000	00000	00000	00000	PS		PS
00000016	00000	00000	00000	00000	PS		PS
00000017	00000	00000	00000	00000	PS		PS

Figure 4-4. DIS Program Storage Display (G)

DIS JOB STATUS DISPLAY (J)

The DIS J 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 4-5 shows the DIS job status display (J).

```
DIS J. JOB STATUS.

JSN = AAAC          EJTO = 5          P = 2046
UI = 2755           SRUA = 1          RA = 1741
FM = MLSTEST        SRUL = 777777     FL = 26
PN = 721C           CS = BATCH        RAE = 160
                   CONN = 0          FLE = 0
                   LEVEL = LVLO      CPA = 600
                                     } Job status

EST = 1             } Equipment assigned

MS1 = DEVICE UNAVAILABLE. } Current messages
MS2 =

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 4-5. DIS Job Status Display (J)

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.

<u>Item</u>	<u>Description</u>
JSN	Job sequence name of the DIS job.
UI	User index.
FM	Current family name.
PN	Current pack name.
EJTO	Executing job table ordinal.
SRUA	System resource units accumulator.
SRUL	Account block limit for system resource units.
CS	Connection status (interactive jobs only).
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 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.
- B EXCHANGE PACKAGE.
- 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.
- Y MONITOR FUNCTIONS.
- Z DIRECTORY.

LEFT AND RIGHT SCREEN DISPLAYS

- C CENTRAL MEMORY. FIVE GROUPS OF FOUR.
- D CENTRAL MEMORY. FIVE GROUPS OF FOUR.
- J JOB STATUS.
- 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 both the 721 and CC545 console types are listed along with a description of the use of each key.

<u>721</u>	<u>CC545</u>	<u>Description</u>
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.
+	+	Pages the left screen forward for the A, C, D, F, G, H, M, T, and U displays. Increments the breakpoint address for the B display.
-	-	Pages the left screen backward for the C, D, F, G, M, T, and U displays. Decrements the breakpoint address for the B display. Resets the A and H displays to the beginning.
((Pages the right screen forward for the C and D displays.
))	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.
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	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)	Advances the left screen display sequence established by the SET,screen. command (refer to Display Selection Commands later in this section).
NEXT	CR (carriage return)	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 or the ← key.

The following keys are interpreted as control characters by DIS.

<u>721</u>	<u>CC545</u>	<u>Description</u>
← or ERASE	Left blank (erase)	Clears current DIS keyboard entry and any resultant error message; AUTO MODE (automatic command processing) is also cleared.
←	BKSP (backspace)	Deletes last character displayed and clears error message (if one exists).
NEXT	CR (carriage return)	Initiates processing of an entered command.

The following keyboard messages may appear above your entry.

<u>Message</u>	<u>Description</u>
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.
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 10000g and less than 377700g 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 71g.
INSUFFICIENT FIELD LENGTH.	Your field length is not long enough to support either extended memory or 026.
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.

<u>Message</u>	<u>Description</u>
REPEAT ENTRY.	The command in the command buffer is repeated each time you press CR or NEXT. This can be cleared by pressing the left blank or ←key.
REQUEST EXCEEDS MAXIMUM FILE.	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 026 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 026 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.

<u>Command</u>	<u>Description</u>																
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 C, D, J, 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="487 1323 1347 1869"> <thead> <tr> <th><u>z</u></th> <th><u>Description</u></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:</td> </tr> <tr> <td>0-3</td> <td>Changes the specified group to display the eight words beginning at location loc.</td> </tr> <tr> <td>4</td> <td>Changes the entire display so that all 8-word groups are contiguous locations beginning at location loc.</td> </tr> <tr> <td>5</td> <td>Advances the display by loc locations.</td> </tr> <tr> <td>6</td> <td>Decrements 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>	<u>z</u>	<u>Description</u>	x	Display identifier (C, D, F, G, or M).	z	Type of display modifications:	0-3	Changes the specified group to display the eight words beginning at location loc.	4	Changes the entire display so that all 8-word groups are contiguous locations beginning at location loc.	5	Advances the display by loc locations.	6	Decrements the display by loc locations.	loc	Location parameter (maximum of eight digits for central memory address or seven digits for extended memory address).
<u>z</u>	<u>Description</u>																
x	Display identifier (C, D, F, G, or M).																
z	Type of display modifications:																
0-3	Changes the specified group to display the eight words beginning at location loc.																
4	Changes the entire display so that all 8-word groups are contiguous locations beginning at location loc.																
5	Advances the display by loc locations.																
6	Decrements the display by loc locations.																
loc	Location parameter (maximum of eight digits for central memory address or seven digits for extended memory address).																

<u>Command</u>	<u>Description</u>
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 (display) or → key 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.

<u>Command</u>	<u>Description</u>
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.
DCP.	Drops the central processor and displays the exchange package area on the B display.
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.
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,f1.	Sets central memory field length FL=f1 in the exchange package area ($0 \leq f1 \leq 777777_8$). f1 must be at least 10000g if user extended memory is assigned.

<u>Command</u>	<u>Description</u>
ENFLE,file.	Sets extended memory field length FLE to file000 in the exchange package area ($1 \leq \text{file} \leq 7777_8$). If user extended memory is assigned (file=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 job priority to pr ($2 \leq \text{pr} \leq 70_8$).
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 (7777_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 (erase) or ← key twice clears DIRECT CPU INPUT mode.
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. You will need a CC545 console to use the O26 file editor. It does not work with a 721 console. Refer to the NOS 2 Systems Programmer's Instant for a description of the editing commands.

<u>Command</u>	<u>Description</u>
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 777g is specified for spr, 777g 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 (377700g). 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.
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.

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 77775g.

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.

<u>Command</u>	<u>Description</u>
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.†
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.†

†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.

<u>Command</u>	<u>Description</u>
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.†
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.†
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.†

†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

Dayfile Commands	5-2
Queued File Utility Commands	5-3
Job Processing Control Commands	5-3
General Job Control Commands	5-4
Interactive Job Control Commands	5-6
Job Communication Commands	5-7
Job Flow Commands	5-8
Peripheral Equipment Control Commands	5-18
Subsystem Control Commands	5-33
System Control Commands	5-37
Secured System Control Commands	5-46
Memory Entry Commands	5-47
Channel Control Commands	5-50
Extended Memory Flag Register Commands	5-51

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.

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.

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 purposes. 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 commands 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 in maintaining the hardware, the system enters the first characters (up to 25) of the following commands into the error log after each execution.

DOWN,param.

FORM,param.

INITIALIZE,param.

LOG,param.

OFF,eq.

ON,eq.

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 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.

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.

<u>Command</u>	<u>Description</u>
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 2 Operations Handbook for descriptions of dayfile displays.

QUEUED FILE UTILITY COMMANDS

The following commands are used to control queued input and output files.

<u>Command</u>	<u>Description</u>
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.

<u>Command</u>	<u>Description</u>
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 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.

DROP,jsn,qt,ujn.	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.
------------------	--

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.

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).

The queue type is one of the following.

<u>qt</u>	<u>Queue Type</u>
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.

<u>Command</u>	<u>Description</u>
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, the 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.
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>

NOTE

Before initiating the command, ensure that the correct job sequence name has been specified.

In some cases, a KILL command will be intercepted by a job's relieve processing. If the job does not terminate after finishing its relieve processing, issue another KILL command to terminate the job.

OVERRIDE,jsn.	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.
---------------	--

NOTE

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.

<u>Command</u>	<u>Description</u>
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.	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. 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.

INTERACTIVE JOB CONTROL COMMANDS

The following commands apply only to interactive online jobs. The interactive facility subsystem must be active at control point 1.

<u>Command</u>	<u>Description</u>
DIAL,jsn, messagetext.	Sends message messagetext (48 characters maximum) to the terminal currently assigned to the job with job sequence name jsn. 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.
WARN,messagetext.	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,0 display. When sent to an interactive terminal, the message messagetext is always preceded by the statement hh.mm.ss. WARNING specifying the time (hours.minutes.seconds) when you entered the WARN command.

<u>Command</u>	<u>Description</u>
	For example, if you enter WARN,SYSTEM SHUTDOWN AT 1500. the following information would be transmitted to all interactive terminals. hh.mm.ss. WARNING SYSTEM SHUTDOWN AT 1500. This command is typically used to notify interactive users of an interruption in service or system shutdown.
WARN.	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.

JOB COMMUNICATION COMMANDS

The following job communication commands are used to communicate with a job currently in the executing job table.

<u>Command</u>	<u>Description</u>
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+70g through RA+74g 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,s1, s2,...,s6.	Turns off one or more sense switch(es) s_i ($1 \leq s_i \leq 6$) 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, MAS, and TAF Subsystems.
ONSW,jsn,s1, s2,...,s6.	Turns on sense switch(es) s_i ($1 \leq s_i \leq 6$) 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, MAS, 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 2 Operations Handbook for more information on the S display).

<u>Command</u>	<u>Description</u>
CLASS,ot,sc ₁ , sc ₂ ,...,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.

<u>Parameter</u>	<u>Description</u>																								
ot	Origin type (must be batch, remote batch, or interactive). This parameter is order dependent.																								
	<table border="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>ot</u></th> <th style="text-align: left;"><u>Origin Type</u></th> </tr> </thead> <tbody> <tr> <td>BC</td> <td>Batch.</td> </tr> <tr> <td>RB</td> <td>Remote batch.</td> </tr> <tr> <td>IA</td> <td>Interactive.</td> </tr> </tbody> </table>	<u>ot</u>	<u>Origin Type</u>	BC	Batch.	RB	Remote batch.	IA	Interactive.																
<u>ot</u>	<u>Origin Type</u>																								
BC	Batch.																								
RB	Remote batch.																								
IA	Interactive.																								
sc _i	Authorizes service classes for the specified origin type.																								
	<table border="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>sc</u></th> <th style="text-align: left;"><u>Service Class</u></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 ≤ n ≤ 3).</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>SY</td> <td>System.</td> </tr> <tr> <td>TS</td> <td>Interactive.</td> </tr> <tr> <td>ALL</td> <td>Sets all the service classes except subsystem (SS) and deadstart (DS).</td> </tr> <tr> <td>NUL</td> <td>Clears all service classes.</td> </tr> </tbody> </table>	<u>sc</u>	<u>Service Class</u>	BC	Local batch.	CT	Communication task.	DI	Detached interactive.	In	Installation class n (0 ≤ n ≤ 3).	MA	Maintenance.	NS	Network supervisor.	RB	Remote batch.	SY	System.	TS	Interactive.	ALL	Sets all the service classes except subsystem (SS) and deadstart (DS).	NUL	Clears all service classes.
<u>sc</u>	<u>Service Class</u>																								
BC	Local batch.																								
CT	Communication task.																								
DI	Detached interactive.																								
In	Installation class n (0 ≤ n ≤ 3).																								
MA	Maintenance.																								
NS	Network supervisor.																								
RB	Remote batch.																								
SY	System.																								
TS	Interactive.																								
ALL	Sets all the service classes except subsystem (SS) and deadstart (DS).																								
NUL	Clears all service classes.																								

<u>Command</u>	<u>Description</u>
DELAY,p ₁ ,p ₂ , ...,p _n .	Alters current system delay parameters. Examine the S display to determine the current delay parameter values.

<u>P_i</u>	<u>Delay</u>
ARint	PP auto recall interval in milliseconds. This parameter specifies the time interval at which peripheral processor units (PPs) in auto recall are recalled. One PP routine is activated each time period. The acceptable range for int is from 1 to 7777 ₈ .
CRrp	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. The acceptable range for rp is from 1 to 7777 ₈ .
JQdel	Exponent used to determine the input job scheduler delay in seconds. The acceptable range for del is from 0 to 14 ₈ . The delay in seconds between scheduling of input jobs is calculated as follows: $\text{delay} = 2^{**}\text{del}$
JSint	Job scheduler interval in seconds. This parameter specifies the interval at which the job scheduler and priority evaluation routines are called. The scheduler may also be called at other times. The acceptable range for int is from 1 to 7777 ₈ .
MNsi	Minimum CPU job switch interval in milliseconds. The acceptable range for si is from 1 to 7777 ₈ .
MXsi	Maximum CPU job switch interval in milliseconds. The acceptable range for si is from 1 to 7777 ₈ .

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 S 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	
MN	
MX	

Figure 5-1. Record of Original Values in S Display (Delay Values)

Command

Description

PCLASS,sc0,sc1,
...,sc7.

Selects the service class associated with each priority level (p0 through P7) 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.

Parameter

Description

sc_i The 2-character service class parameter.

<u>sc</u>	<u>Service Class</u>
BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class n (0 ≤ n ≤ 3).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SY	System.
TS	Interactive.

QUEUE,sc,qt,qp1,
qp2,....,qp_n.

Alters the queue priorities qp_i associated with the input, executing, and output queues qt for each service class sc. Examine the S display to determine the priority values currently associated with each service class. These changes are valid only until the next deadstart.

<u>sc</u>	<u>Service Class</u>
BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class n (0 ≤ n ≤ 3).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SS	Subsystem.
SY	System.
TS	Interactive.

CommandDescription

qt Job Queue Type

EX Executing.

IN Input.

OT Output.

qp_i

Queue Priority

ILpr Lowest priority pr an executing job can have for an initial time slice. The acceptable range for pr is from 0 to 7777₈.

IPpr Initial scheduling priority pr for an executing job. The acceptable range for pr is from 0 to 7777₈.

LPpr Lowest priority pr at which a job can enter the queue. The acceptable range for pr is from 0 to 7777₈.

UPpr Highest priority pr a job can reach in the specified queue; aging stops when this priority is reached. The job is also given this priority when initially assigned to a control point. The value of pr, which can range from 0 to 7777₈, is normally greater than LPpr but cannot exceed the maximum of the job's service class.

WFnn Weighting factor nn for queue priority calculations. nn must be a power of 2, from 0 to 4000₈.

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 S display. 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	IP	IL	LP	UP	WF	IP	LP	UP	WF	IP
SY													
BC													
RB													
TS													
DI													
NS													
SS													
MA													
CT													
I0													
I1													
I2													
I3													

Figure 5-2. Record of Original Values in S Display (Queue Priorities)

<u>Command</u>	<u>Description</u>
SERVICE,sc,p ₁ ,p ₂ , ...,p _n .	Alters the service limits p _i associated with each service class sc. These changes are valid only until the next deadstart.

<u>sc</u>	<u>Service Class</u>
BC	Local batch.
CT	Communication task.
DI	Detached interactive.
In	Installation class n (0 ≤ n ≤ 3).
MA	Maintenance.
NS	Network supervisor.
RB	Remote batch.
SS	Subsystem.
SY	System.
TS	Interactive.

<u>p_i</u>	<u>Service Limits</u>
AMf1	Maximum field length divided by 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 exceeds the field length specified 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 normally exceed the field length for its service class can be scheduled to a control point if there are not enough jobs to fill the field length specified for another service class. The system always attempts to use central memory to its greatest capacity. The acceptable range for f1 is from 0 to 7777 ₈ .
CMts	Central memory time slice in seconds. This parameter specifies the maximum amount of time a job of the specified service class can remain at a control point before the job becomes eligible to be rolled out. A job is not necessarily rolled out when its time slice is exceeded. The job is set to a lower priority and then may be replaced by a job of higher priority. The value of ts can range from 0 to 7777 ₈ .
CPts	CPU time slice (milliseconds*64). This parameter specifies the maximum amount of time a job of the specified service class can use the CPU before its queue priority is set to the lower boundary. The value of ts can range from 1 to 7777 ₈ .

Command

Description

<u>P_i</u>	<u>Service Limits</u>																		
CSv	Cumulative size in PRUs allowed for all indirect access permanent files; v is used to specify the corresponding values shown next: <table><thead><tr><th><u>v</u></th><th><u>Limit Value in Octal</u></th></tr></thead><tbody><tr><td>0</td><td>Unlimited</td></tr><tr><td>1</td><td>1000</td></tr><tr><td>2</td><td>2000</td></tr><tr><td>3</td><td>5000</td></tr><tr><td>4</td><td>10000</td></tr><tr><td>5</td><td>50000</td></tr><tr><td>6</td><td>100000</td></tr><tr><td>7</td><td>Unlimited</td></tr></tbody></table>	<u>v</u>	<u>Limit Value in Octal</u>	0	Unlimited	1	1000	2	2000	3	5000	4	10000	5	50000	6	100000	7	Unlimited
<u>v</u>	<u>Limit Value in Octal</u>																		
0	Unlimited																		
1	1000																		
2	2000																		
3	5000																		
4	10000																		
5	50000																		
6	100000																		
7	Unlimited																		
DSv	Size in PRUs allowed for individual direct access permanent files; v is used to specify the corresponding values shown next: <table><thead><tr><th><u>v</u></th><th><u>Limit Value in Octal</u></th></tr></thead><tbody><tr><td>0</td><td>Unlimited</td></tr><tr><td>1</td><td>1000</td></tr><tr><td>2</td><td>2000</td></tr><tr><td>3</td><td>5000</td></tr><tr><td>4</td><td>10000</td></tr><tr><td>5</td><td>50000</td></tr><tr><td>6</td><td>100000</td></tr><tr><td>7</td><td>Unlimited</td></tr></tbody></table>	<u>v</u>	<u>Limit Value in Octal</u>	0	Unlimited	1	1000	2	2000	3	5000	4	10000	5	50000	6	100000	7	Unlimited
<u>v</u>	<u>Limit Value in Octal</u>																		
0	Unlimited																		
1	1000																		
2	2000																		
3	5000																		
4	10000																		
5	50000																		
6	100000																		
7	Unlimited																		
DTsc	Detached service class for interactive jobs. The default value for sc is DI (detached interactive).																		
ECf1	Maximum extended memory divided by 1000g for any job of the specified service class. The acceptable range for f1 is from 0 to 7777 ₈ .																		
EMf1	Maximum extended memory divided by 1000g for all jobs of the specified service class. The acceptable range for f1 is from 0 to 7777 ₈ .																		

Command

Description

<u>P_i</u>	<u>Service Limits</u>																		
FCv	Number of permanent files allowed; v is used to specify the corresponding values shown next: <table><thead><tr><th><u>v</u></th><th><u>Limit Value in Octal</u></th></tr></thead><tbody><tr><td>0</td><td>Unlimited</td></tr><tr><td>1</td><td>10</td></tr><tr><td>2</td><td>20</td></tr><tr><td>3</td><td>30</td></tr><tr><td>4</td><td>40</td></tr><tr><td>5</td><td>50</td></tr><tr><td>6</td><td>100</td></tr><tr><td>7</td><td>Unlimited</td></tr></tbody></table>	<u>v</u>	<u>Limit Value in Octal</u>	0	Unlimited	1	10	2	20	3	30	4	40	5	50	6	100	7	Unlimited
<u>v</u>	<u>Limit Value in Octal</u>																		
0	Unlimited																		
1	10																		
2	20																		
3	30																		
4	40																		
5	50																		
6	100																		
7	Unlimited																		
FLf1	Maximum field length divided by 100g for any job of the specified service class. Jobs with field lengths that exceed this value are not considered for scheduling unless no other jobs are waiting to be scheduled and sufficient field length is available without rolling out any jobs. The acceptable range for f1 is from 0 to 7777 ₈ .																		
FSv	Size in PRUs allowed for individual indirect access permanent files; v is used to specify the corresponding values shown next: <table><thead><tr><th><u>v</u></th><th><u>Limit Value in Octal</u></th></tr></thead><tbody><tr><td>0</td><td>Unlimited</td></tr><tr><td>1</td><td>10</td></tr><tr><td>2</td><td>30</td></tr><tr><td>3</td><td>50</td></tr><tr><td>4</td><td>100</td></tr><tr><td>5</td><td>150</td></tr><tr><td>6</td><td>300</td></tr><tr><td>7</td><td>Unlimited</td></tr></tbody></table>	<u>v</u>	<u>Limit Value in Octal</u>	0	Unlimited	1	10	2	30	3	50	4	100	5	150	6	300	7	Unlimited
<u>v</u>	<u>Limit Value in Octal</u>																		
0	Unlimited																		
1	10																		
2	30																		
3	50																		
4	100																		
5	150																		
6	300																		
7	Unlimited																		
NJn	Maximum number of jobs. For each service class, this parameter specifies the number of jobs that can be executing in the system. The value for n can range from 0 to 7777 ₈ .																		

Command

Description

<u>P_i</u>	<u>Service Limits</u>
PRpr	CPU priority ($2 \leq pr \leq 778$). Jobs with the highest priority get the CPU first. All service classes except maintenance (MA) and network supervisor (NS) are normally set to the same CPU priority. MA jobs are run at the lowest CPU priority. NS jobs are run at a higher priority than any of the other job service classes to provide good response time to network users. Subsystems control their own priorities.
TDs	Suspension time-out delay. A suspended job is timed out for s times 10g seconds. s can range from 0 to 7778. The maximum time-out delay is approximately 9.1 hours.
TPpr	Initial scheduling priority. This is the priority assigned to the start of a new command or after terminal I/O for an interactive terminal. The acceptable range for pr is from 2 to 7778.

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 S display. For additional information concerning the SERVICE command, refer to section 3, Deadstart Decks.

SRST,t. Changes the secondary rollout sector threshold to the value t ($0 \leq t \leq 7778$). Any rollout file smaller than t sectors 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.

JOB SERVICE CLASS	SERVICE LIMITS														
	PR	CP	CM	NJ	TD	FL	AM	TP	DT	EC	EM	DS	FC	CS	FS
SY															
BC															
RB															
TS															
DI															
NS															
SS															
MA															
CT															
I0															
I1															
I2															
I3															

Figure 5-3. Record of Original Values in S Display (Service Limits)

PERIPHERAL EQUIPMENT CONTROL COMMANDS

The commands described in this category provide control of the peripheral equipment available to the system.† 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).
- Family status display (E,F).
- 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 2 Operations Handbook.

<u>Command</u>	<u>Description</u>
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,0 display), and the ASSIGN command must be entered to assign a tape unit to the job.
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.
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.

†Operation of peripheral equipment is described in the NOS 2 Operations Handbook.

<u>Command</u>	<u>Description</u>
CP,est,id.	Assigns a numeric identifier id to the BIO card punch defined by EST ordinal est. The value of id can range from 00 to 67g. Only those files in the punch queue with an identifier equal to id are directed to card punch est. Refer to the descriptions of the LOAD and CR commands to assign an identifier to a job or queue type file.
CR,est,id.	Assigns a numeric identifier id to the card reader defined by EST ordinal est. The value of id can range from 00 to 67g. All subsequent jobs loaded from card reader est are assigned the identifier id.
DOWN,CHcc,EQest. or	The first form of the command discontinues use of channel cc for the equipment specified by EST ordinal est. If the EQest parameter is omitted, channel cc is discontinued for all equipments in the EST.
DOWN,MCHcc,EQest. or	The second form of the command discontinues use of MUX channel cc for the equipment specified by EST ordinal est (this form only applies to model 176). If the EQest parameter is omitted, MUX channel cc is discontinued for all equipments in the EST.
DOWN,EQest.	The third form of the command discontinues the use of equipment with EST ordinal est for all channels. If channel cc is the only channel available to a mass storage equipment, its use is not discontinued. If channel cc is <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) the command is ignored and the following message appears on the left screen. INCORRECT ENTRY Correct the entry and retry the command.

Command

Description

If you attempt to DOWN an equipment

- already with a DOWN status
- assigned to a job
- not listed as removable in the EST†
- listed in the EST as removable but not globally unloaded†
- that is an independent shared device (multimainframe)†

the following message appears.

INCORRECT ENTRY

Correct the entry and retry the command.

NOTE

This command should be used with caution since it directly affects the operation of system peripheral equipment.

END,est,rc.

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.

If est defines a card reader that is stopped because of a compare error when END,est. is entered, perform the following steps:

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.

The job terminates and if another card deck follows the end-of-information card, it is processed normally.

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.

†Applies only to mass storage equipment.

<u>Command</u>	<u>Description</u>
FORM,est,fc.	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 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).
IDLEFAMILY,est.	This command performs one of the two following functions, depending on the status of the family. <ul style="list-style-type: none"> • If the family is active, the IDLEFAMILY command causes all new jobs and USER commands for the family on the equipment specified by EST ordinal est to be aborted. Jobs in progress are allowed to complete. • If the family is inactive, the IDLEFAMILY command allows jobs to access the family on the equipment specified by EST ordinal est. <p>The famc column of the E,F display shows the number of jobs in progress on the equipment.</p>
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.

<u>op</u>	<u>Level of Initialization</u>
AF	Initialize inactive account dayfile.
AL	Total initialization.
DF	Initialize inactive system dayfile.
EF	Initialize inactive error log.
FP	Format pending.†
FT	Total initialization as full-track device.††
HT	Total initialization as half-track device.††
MF	Initialize binary maintenance log.
PF	Initialize permanent files.
QF	Initialize inactive queue files.

† Sets format pending bit in the MST for 881/883 pack reformatting. Refer to the description of the FORMAT program in appendix F.
†† Device must be removable.

Command

Description

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.

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,0). Refer to the description of this message in the NOS 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.

est is the EST ordinal.

id is the machine identifier of the first machine found without unload status set.

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 (B,0):

INITIALIZE PENDING ON THIS DEVICE.

<u>Command</u>	<u>Description</u>
	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.
LOAD,est,id.	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 67g. 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.
LP,est,id. or LR,est,id. or LS,est,id. or LT,est,id.	Assigns a numeric identifier id to the line printer defined by EST ordinal est. The value of the identifier can range from 00 to 67g. Only those files in the print queue with an identifier equal to id are directed to line printer est. Refer to the descriptions of the LOAD and CR commands in this section to assign an identifier to a job or queue type file. The user usually assigns the identifier to the file.
MSAL=est ₁ ,est ₂ , ...,est _n .	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: <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.

Command

Description

MSAL,t=est₁,est₂,
...,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 57 octal est_is may be specified on a single entry. Examine the disk status display (E,M) to determine if the device is nonremovable.

<u>t</u>	<u>File Type</u>	<u>t</u>	<u>File Type</u>
B	LGO	P	Primary
D	Job dayfile	R	Rollout
I	Input	S	Secondary rollout
L	Local	T	Temporary
O	Output		

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 7777g) 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.

```

EQPDECK
.
.
.
EQ005=DP,ST=ON,SZ=4000,CH=30.    Equipment 5 is extended memory
                                  via DDP.
.
.
.
MSAL,S=5.                          Assigns secondary rollout files
                                  to extended memory.
.
.
.
IPRDECK
.
.
.
SRST=20.                            Sets the threshold count so
.                                  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.

<u>Command</u>	<u>Description</u>
MOUNT,est,P.	<p>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 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>When you specify P in the MOUNT command for an independent shared device in a multiframe 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>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 style="text-align: center;">INCORRECT EQUIPMENT.</p> <p>If the device is shared in a multiframe 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,0):</p> <p style="text-align: center;">INITIALIZE PENDING ON THIS DEVICE.</p>
OFF,est.	<p>Logically turns off the device defined by EST ordinal est. This command allows you to logically remove a device from the operating environment. Examine the E,A display to determine the EST ordinal and current status (ON or OFF) of the device. If est specifies a mass 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 mass storage devices have system residency or allow system allocation of temporary files. In addition, if an MSAL entry is currently specified for a mass 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>
ON,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 drive, it will be automatically spun up. Examine the E,A display to determine the EST ordinal and current status (OFF or ON) of the device.</p>
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,est,rc.†	<p>Repeats the current operation on the BIO equipment defined by EST ordinal est the number of times specified by rc. The maximum value that can be entered for rc is 77g.</p>

†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.

<u>Command</u>	<u>Description</u>
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 < pr00 < service class maximum; refer to the S display in the NOS 2 Operations Handbook; pr value is multiplied by 100g 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 < pr00 < service class maximum; refer to the S display in the NOS 2 Operations Handbook; pr value is multiplied by 100g internally). If pr is not specified, the service class default priority is assigned.
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.	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.

NOTE

On a secured system this command is accepted only if a security administrator has placed the system in security unlock status.

<u>Command</u>	<u>Description</u>
SKIP,est,rr.	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.
SKIPF,est,ff.	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.
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 10g PRUs (the current buffer size) plus the number of PRUs remaining in the buffer. If the buffer was empty, ss would be limited to 20g 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.
SPINUP,est.	Spins up the 834 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).
SPINDOWN,est.	Spins down the 834 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).

NOTE

Spinning down an 834 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.
SUPPRESS,est.	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.

Command

Description

TRAIN,est,t.

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, and LT is a 580-20 line printer. LP is applicable to all line printer types. Print trains supported for the 580 printers are 596-1, 596-5, and 596-6. The t field specifies the print train.

<u>t</u>	<u>Print Train</u>
0	596-1 (CDC graphic 63/64-character set); default.
1	596-1 (CDC graphic 63/64-character set).
2	Reserved for future use. †
3	Reserved for future use. †
4	Reserved for future use. †
5	596-5 (ASCII graphic 63/64-character set).
6	596-6 (ASCII graphic 95-character set).
7	596-6 (ASCII graphic 63/64-character set or ASCII graphic 95-character set).

UNLOAD,est.

Physically unloads a tape or logically removes a removable mass 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 mass storage device is specified, you can dismount the disk pack on that device.

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).

† These values are allowed but will default to 596-1.

Command

Description

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.

- 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.

NOTE

If a multispindle family is mounted on a single spindle device, only the first device shows the global unload status.

- 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.

NOTE

A device should be physically dismantled only if global unload status (N) is displayed on all machines accessing the device.

Command

Description

If a removable pack is dismounted before the N status is displayed, the following may occur.

- Mass storage device status errors.
- Permanent file errors when the pack is remounted at some later date.
- 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.

NOTE

If the Mass Storage Subsystem (MSS) or Mass Storage Archival Subsystem (MSAS) is active, it must be idled before unloading a removable family pack which has MSS/MSAS files. After dismounting the family pack, MSS/MSAS can be initialized again.

UP,CHcc,EQest.

or

UP,MCHcc,EQest.†

or

UP,EQest.

Resumes normal use of channel cc or the equipment with EST ordinal est by tape units and/or mass storage devices, reversing the DOWN command. If EQest is specified, use of channel cc resumes only for the mass storage equipment with EST ordinal est. The equipment specification is valid only for mass storage equipments. If you attempt to UP a channel or equipment that

- is already UP,
- has a DOWN status and is assigned to a maintenance job,
- was given a DOWN status by another mainframe that is sharing the equipment,

this error message appears:

INCORRECT ENTRY

If the channel or equipment is already up, no further action is necessary. If the equipment is assigned to a maintenance job, you cannot bring up the equipment until the maintenance job releases it. If the DOWN,EQest command was entered from another mainframe in a multimainframe environment, that mainframe must enter the UP,EQest command.

†Model 176 only.

Command

Description

Putting an equipment in an UP status does not completely reverse the DOWN status. The equipment must also be turned on using the ON,est command.

NOTE

This command should be used with caution since it directly affects the operation of system peripheral equipment.

VALIDATE,est.	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.
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,.
VSN,est,vsn.†	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) or 9 track (HD, PE, or GE), on which the tape is to be mounted, the required VSN, 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.

†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

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

INCORRECT ENTRY

appears on the left console screen. To change a VSN previously assigned by this command, clear the first VSN by entering:

VSN,est.

Parameter

Description

est EST ordinal of the tape unit.

This command can then be entered:

VSN,est,vsn.

Parameter

Description

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 multireel 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 reel of the file is on a 669 tape unit on channels 13 and 33, all subsequent reels 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.

If two or more unassigned tapes having identical VSNs are mounted on units of the same track type, the message

REQUEST,dt,vsn

appears on the B,0 display. The dt field is either the device type MT or the 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. However, if the job cannot be rolled out, the command used to call the subsystem would not be valid. In this case, either terminate the job (if the subsystem required 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).

<u>Command</u>	<u>Description</u>						
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 2 Operations Handbook).</p> <p>Response to line printer errors is controlled by setting sense switch 1 or 2 as follows:</p> <table><thead><tr><th><u>Entry</u></th><th><u>Response</u></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>	<u>Entry</u>	<u>Response</u>	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.
<u>Entry</u>	<u>Response</u>						
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.†	<p>Calls the CYBER Database Control System (CDCS). CDCffff calls a procedure file named CDCffff under the system user index to initiate the CDCS Subsystem.</p>						
IAFffff.†	<p>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.</p> <p>Ensure that NAM is active at a control point (refer to the B,0 or B,A display before you enter the IAFffff command.</p>						

† The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

Command

Description

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.)

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. |

IDLE,sub. Sets idledown status for subsystem sub, except for NAM and RHF. NAM is idled by using the procedure shown in NAM Shutdown (refer to section 15, Network Operations). RHF is idled by using the RHF K display (refer to section 8, K-Display Utilities).

The subsystem terminates when idledown conditions are met. MAG terminates when no tapes are assigned. MSS/MSAS terminates when no requests are outstanding and no MSS/MSAS 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.

Command

Description

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.

MAGffff.†

Calls the Magnetic Tape Subsystem (MAG) to the control point specified in the ENABLE command. 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. MAGffff executes a procedure file named MAGffff under the system user index to initiate the MAG Subsystem.

NOTE

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.

MAPffff.†

Calls the Matrix Algorithm Processor (MAP) Subsystem to the control point specified in the ENABLE command. MAPffff calls a procedure file named MAPffff under the system user index to initiate the MAP Subsystem.

MASffff.†

Calls the Mass Storage Archival Subsystem (MSAS) processing program to the control point specified in the ENABLE command. MASffff calls a procedure file named MASffff under the system user index to initiate the MSAS.

You can turn the trace mode on or off by using the following sense switch entries after the subsystem is initiated.

<u>Entry</u>	<u>Response</u>
ONSW,MAS,1.	Turn the trace mode on. The trace mode issues all MSAS messages to the job dayfile.
OFFSW,MAS,1.	Turn the trace mode off. Default is trace mode off.

MCSffff.†

Calls the Message Control System (MCS) to the control point specified in the ENABLE command. MCSffff calls a procedure file named MCSffff under the system user index to initiate the MCS Subsystem. NAM must be active at a control point to use MCS.

MSSffff.†

Calls the Mass Storage Subsystem (MSS) processing program to the control point specified in the ENABLE command. MSSffff calls a procedure file named MSSffff under the system user index to initiate the MSS.

† The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

<u>Command</u>	<u>Description</u>
NAMffff. †	<p>Calls the network access method (NAM) to the control point specified in the ENABLE command. NAMffff calls a procedure file named NAMffff 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>
RDFffff. †	<p>Calls the Remote Diagnostic Facility (RDF) Subsystem to control point 1. RDFffff calls a procedure file named RDFffff or IAFffff (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,0 or B,A display) before you enter the RDFffff command.</p>
RHFffff. †	<p>Calls the Remote Host Facility (RHF) Subsystem to the control point specified in the ENABLE command. RHFffff calls a procedure file named RHFffff under the system user index to initiate the RHF Subsystem.</p>
SMFffff. †	<p>Calls the Screen Management Facility (SMF) Subsystem to the control point specified in the ENABLE command. SMFffff calls a procedure file name SMFffff under the system user index to initiate the SMF Subsystem. IAF must be active to use SMF.</p>
SSFffff. †	<p>Calls the SCOPE 2 Station Facility (SSF) Subsystem to the control point specified in the ENABLE command. SSFffff calls a procedure file named SSFffff under the system user index to initiate the SSF Subsystem.</p>
STMffff. †	<p>Calls the stimulator to the last control point. STMffff calls a procedure file named STMffff under the system user index to initiate the stimulator.</p>
STOP,sub.†	<p>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.</p>

† The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

<u>Command</u>	<u>Description</u>
TAFffff. †	Calls the Transaction Facility (TAF) Subsystem to the control point specified in the ENABLE command. TAF must be active for transaction processing to take place.

TAFffff calls a procedure file named TAFffff 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.

Response to an operator drop or fatal error is controlled by setting sense switches as follows:

<u>Entry</u>	<u>Response</u>
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 STOP,TAF command.
ONSW,TAF,5.	Dump the entire field length and release the output file after the transaction is dropped or aborted.
ONSW,TAF,6.	Print the job dayfile upon termination.

Activity at the transaction control point can be monitored on the 0 displays (refer to the NOS 2 Operations Handbook).

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.

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.

<u>Command</u>	<u>Description</u>
AUTO.	Calls enabled subsystems to control points, initiates automatic job processing, and spins up any 834 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 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.

† The characters ffff are optional; if required, installation personnel must supply the necessary 1 to 4 alphanumeric characters.

<u>Command</u>	<u>Description</u>
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"> 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. 2. All job scheduling is inhibited. (This has the same effect as if the IDLE command was entered.) 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. 4. The system moves the system dayfile buffers maintained in CMR to disk. 5. All subsystems except the magnetic tape executive (MAG) are aborted. 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. 7. The system is left in an idle state. If SPINDOWN is enabled, all 834 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.

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.

<u>Command</u>	<u>Description</u>
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).

<u>Parameter</u>	<u>Description</u>
yy	Year; 00 through 99.
mm	Month; 01 through 12.
dd	Day; 01 through 31.

DEBUG.	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.
--------	--

NOTE

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).

DIS,jsn.	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. Refer to section 4, DIS Operations, for complete information concerning the DIS package.
----------	---

NOTE

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).

DISABLE,op.	Disables or enables option op.
or	
ENABLE,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.

<u>op</u>	<u>Description</u>
DDP ROLLOUT PATH	Enables or disables the use of the DDP for job rollout and rollin. If DDP equipment is not present, this entry has no effect.

Command

Description

<u>op</u>	<u>Description</u>
ENGR	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).

NOTE

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).

FILE STAGING	Enables or disables the staging of MSF/7990-resident permanent files to disk. Disabling FILE STAGING causes job attempts to access MSF/7990-resident files to be aborted. If the MSSEEXEC/SSEEXEC is running, enabling FILE STAGING allows MSF/7990-resident files to be staged to disk when accessed.
LOGGING	Enables or disables logging of dayfile messages concerned with program efficiency.
MASTER MAS	Enables or disables master mainframe mode for MAS processing. When MAS is brought to a control point, the SSEEXEC program runs if master mainframe mode is enabled. The SSSLV program runs if master mainframe mode is disabled. This entry has no effect unless MAS processing is activated.
MASTER MSS	Enables or disables master mainframe mode for MSS processing. When MSS is brought to a control point, the MSSEEXEC program runs if master mainframe mode is enabled. The MSSSLV program runs if master mainframe mode is disabled. This entry has no effect unless MSS 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.

Command

Description

<u>op</u>	<u>Description</u>
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 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 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 with no exit processing. An interactive origin type job that issues a USER command is logged off. 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 the 834 disk storage devices during a system checkpoint.
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.

<u>Command</u>	<u>Description</u>
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 or MAINTENANCE command is entered. When the BIO Subsystem is idle, it is terminated.
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.
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.
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.
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.† 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>

† Monitor functions are described in the NOS 2 Systems Programmer's Instant.

<u>Command</u>	<u>Description</u>
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).

<u>Parameter</u>	<u>Description</u>
hh	Hour; 00 through 23.
mm	Minute; 00 through 59.
ss	Second; 00 through 59

<u>Command</u>	<u>Description</u>
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.</p> <p>DEBUG.†</p> <p>DIS,jsn.†</p> <p>DISABLE,SECONDARY USER COMMANDS.</p> <p>ENABLE,ENGR.†</p> <p>ENABLE,SECONDARY USER COMMANDS.</p> <p>OVERRIDE,jsn.</p> <p>QDSPLAY,jsn.†</p> <p>SPINDOWN,est.</p> <p>SPINUP,est.</p> <p>STEP.</p> <p>STEP,jsn,ff,b,v.</p> <p>STOP,sub.</p> <p>TIME.hh.mm.ss.</p> <p>UNLOAD,est. (est specifies a nonremovable shared mass storage device)</p> <p>UNSTEP.</p> <p>All memory entry commands†</p> <p>All channel control commands</p> <p>All extended memory flag register commands</p>

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) or to enter the UNSTEP command after a power or environmental interrupt.

†Refer to Secured System Control Commands later in this section for information on unlocking the console keyboard on a secured system.

<u>Command</u>	<u>Description</u>
UNSTEP.	Clears step mode. This command clears the effect of any format of STEP command. Unlock the console before entering this command (refer to the UNLOCK command later in this section). Depending on current syntax overlay residency, you may have to enter the 99 command before entering the UNSTEP command (refer to the 99 command description later in this section). 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,d)	Initiates the monitor display driver (MDD). MDD is a PP program and is independent of the operating system. Refer to the NOS Online Maintenance Software Reference Manual for more information.

<u>Parameter</u>	<u>Description</u>
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.
d	Time-out parameter. If you enter 0 (zero), MDD waits 15 minutes for a terminal to be connected before timing out. If you enter a nonzero, MDD waits indefinitely for the terminal connection. The default is 0.
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 where (parameters) specifies the parameters. 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 38 characters following X. are used.
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.

<u>Command</u>	<u>Description</u>
QQSH=level.	Specifies the output queue special handling (QQSH) level on a secured system. The (QQSH) level is set initially during deadstart by the QQSH IPRDECK entry. The current QQSH level appears on the DSD S display. The QQSH 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 QQSH 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.
RELEASE,jsn.	Allows the operator to release a file from the output queue whose access level is equal to or above the QQSH 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.
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>

<u>Command</u>	<u>Description</u>
UNLOCK,username, password.	Unlocks the console keyboard. When this command is active, the message SECURITY UNLOCK appears in the header on the left screen display. The following commands are restricted to entry only when the console is in security unlock status.
	DEBUG.
	DIS,jsn.
	ENABLE,ENGR.
	QDSPLAY,jsn.
	SECURES,ot,LA=lowerlevel,UA=upperlevel.
	SECUREQ,est,LA=lowerlevel,UA=upperlevel.
	All memory entry commands.
	All memory display commands.

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). 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 7777775g.

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 the absolute memory you can see only the system tables.

Formats and descriptions for these commands are:

<u>Command</u>	<u>Description</u>
addr,cont. or addr+cont.†	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.†	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.

†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.

<u>Command</u>	<u>Description</u>
addr,Dcont. or addr+Dcont.†	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.†	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.†	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.†	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.

† 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 13g in a 10 PP system and 0 to 13g, 20g to 33g in a 20 PP system.

NOTE

Extreme caution must be exercised when using the following channel control commands.

<u>Command</u>	<u>Description</u>
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 1700g 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).

NOTE

Extreme caution must be exercised when using these commands.

<u>Command</u>	<u>Description</u>
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 ($0 \leq bb \leq 17$).

Calling the Express Deadstart Dump Interpreter	6-2
Input Directives	6-5
Directive Format	6-6
List Control Directives	6-7
EJ - Force Page Eject	6-7
EJOFF - Turn Off Auto Page Eject	6-7
EJON - Turn On Auto Page Eject	6-8
PD - Reset Print Line Density	6-8
*. - Comment in Subtitle Line	6-8
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
Memory Dump Directives	6-14
PP Dump Directives	6-16
AP - Dump Analysis of PP and PP Memory in Octal Line Format	6-16
MPP - Move PP	6-16
P - Dump PP Memory in Octal Block Format	6-17
PF - Dump FLPP Memory in Octal Block Format	6-17
PMS - Read PP Select Switches	6-17
PO - Dump 16-Bit PP Memory in Octal Block Format	6-18
PX - Dump 16-Bit PP Memory in Hexadecimal Block Format	6-18
Q - Dump PP Memory in Octal Line Format	6-18
QF - Dump FLPP Memory in Octal Line Format	6-19
QQA - Dump 16-Bit PP Memory in Octal/ASCII Line Format	6-19
QOD - Dump 16-Bit PP Memory in Octal/Display Line Format	6-20
QXA - Dump 16-Bit PP Memory in Hexadecimal/ASCII Line Format	6-20
QXD - Dump 16-Bit PP Memory in Hexadecimal/Display Line Format	6-20
CMR Dump Directives	6-21
ACCOUNT - Dump Account Dayfile Buffer	6-21
CBT - Dump Control Buffer Table	6-22
CCT - Dump Channel Control Table	6-22
CP - Dump Active Control Point Areas	6-23
CPO - Reset Default List Options	6-25
CT - Dump Channel Tables	6-25
DAYFILE - Dump System Dayfile Buffer	6-25
DB - Dump Disk Buffers	6-25
DBW - Dump Buffered Device/Buffer Statistics/PP-I/O Buffer Link Tables	6-26
DDB - Dump Dayfile Dump Buffer	6-26
DP - Dump Dayfile Buffer Pointers	6-26
EJT - Dump Executing Job Table	6-26
EPB - Dump Extended Memory/PP Buffer	6-26
ERRLOG - Dump Error Log Buffer	6-27
EST - Dump Equipment Status Table	6-27
FNT - Dump System File Name/File Status Table	6-27
FOT - Dump Family Ordinal Table	6-27

(Continued)

(Continued)

HAT - Dump Hash Table	6-28
JC - Dump Job Control Area for Each Service Class	6-28
LC - Dump Low Central Memory	6-28
LDIS - Dump L-Display Buffer	6-28
LIDT - Dump Logical Identifier Table	6-29
MAINLOG - Dump Binary Maintenance Log Buffer	6-29
MST - Dump Mass Storage/Track Reservation Tables	6-29
MTR - Dump CPU Monitor	6-29
MTRQUEUE - Dump CPUMTR Request and Recall Queues	6-30
ODIS - Dump Operator Display Buffer	6-30
PLD - Dump Peripheral Library Directory	6-30
PP - Dump PP Communication Areas	6-31
PROBE - Dump PROBE Data Tables	6-31
PST - Dump Program Status Table and Entry Point Directory	6-31
PUT - Dump Physical Unit Table	6-31
QFT - Dump Queued File Table	6-32
RCL - Dump Resident Central Library	6-32
RPL - Dump Resident Peripheral Library	6-32
SDA - Dump Extended Statistical Data Area	6-32
SECEDED - Dump SECEDED Identifier Table	6-33
SST - Dump Subsystem Control Point/Subsystem Assignment Tables	6-33
Subsystem Dump/Analysis Directives	6-34
BATCHIO (BIO) - Dump Associated Memory for Analysis	6-34
IAF - Dump Associated Memory for Analysis	6-35
MAGNET (MAG) - Dump Associated Memory for Analysis	6-36
RHF - Dump Associated Memory for Analysis	6-37
Hardware Register Dump Directives	6-37
FMFREG - Dump PP Registers	6-37
IOUMR - Dump IOU Maintenance Registers	6-38
MEMMR - Dump Memory Maintenance Registers	6-38
PROCW - Dump Processor Controlware Part Number and Revision Level	6-38
PROMR - Dump Processor Maintenance Registers	6-39
PRORF - Dump Processor Register File	6-39
PROXP - Dump Processor Exchange Package	6-40
SC - Dump S/C Register	6-40
XP - Dump Deadstart Exchange Package	6-40
Buffer Controller Directive	6-41
Interactive Use of DSDI	6-42
Terminal Output Directives	6-43
C - Dump Memory in Instruction Parcel Format	6-43
CP - Dump Active Control Point Areas	6-44
D - Dump Memory in Byte Format	6-44
PP - Dump PP Communication Areas	6-45
Q - Dump PP Memory in Line Format	6-45
QOA, QOD, QXA, QXD - Dump 16-Bit PP Memory in Line Format	6-46
Example of DSDI Terminal Use	6-47
Printer Output Listing Examples	6-49

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, where applicable; and all PPs except for PPO, 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. This permits 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.

CALLING THE EXPRESS DEADSTART DUMP INTERPRETER

Processing of the EDD file is initiated with the DSDI command. The format of the command is:

DSDI, p_1, p_2, \dots, p_n .

Each parameter p_i is either a keyword or a keyword equated to a value. All parameters are optional and order independent.

<u>p_i</u>	<u>Description</u>
B=bmlfile	The 1- to 7-character name of the message file on which the maintenance register messages are to be written. This parameter is valid only if maintenance registers have been dumped to the EDD file.
B omitted	No message file is to be created.
D	This parameter creates a random dump file with the same name as the EDD file and returns the original EDD file. The created 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.
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 2 Reference Set, Volume 3). Use the following directives to dump selected portions of memory to be analyzed.

<u>Directive</u>	<u>Description</u>
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,0/EC/C,fwa,lwa	Dump extended memory in C format.
RA,0/EC/D,fwa,lwa	Dump extended memory in D format.

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:

BCDUMP/ops.

Refer to the BCDUMP directive described later in this section for a description of the ops parameter.

<u>P_i</u>	<u>Description</u>
F=dumpfile	The 1- to 7-character name of the EDD file that DSDI will interpret.
F omitted	Same as F=DUMP.
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 (for example, a CPUMTR error exit leaves an exchange package in memory locations 0 through 20g). 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).
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.

<u>P_i</u>	<u>Description</u>
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 20000g locations of central
              memory.
EC.           Set memory type to extended memory.
D,0,10000.    Dump the first 10000g locations of extended
              memory.
--EOI--
```

A request for the EDD tape must precede the DSDI command. Since EDD writes information on an unlabeled, 7- or 9-track tape at a density of 800 characters per inch (cpi) for 7 track and 1600 cpi for 9 track, the request should appear as follows:

```
LABEL,DUMP,D=density,tape,F=S,LB=KU,VSN=DUMP.
```

<u>Parameter</u>	<u>Description</u>
density	800 cpi for 7-track tape. 1600 cpi for 9-track tape.
tape	MT for 7-track tape. NT for 9-track tape.

It is recommended that the vsn parameter be specified in the request. If this is done, 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.

- List Control Directives
Enable 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 Directives
Enable you to specify alternate files for DSDI input directives and listing output.
- Central Memory/Extended Memory Dump Directives
Provide 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 Directives
Provide 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 Directives
Provide specially formatted dumps of selected areas in central memory resident. These areas are specified by name rather than by address.
- Subsystem Dump/Analysis Directives
Provide specially formatted dumps of subsystem control points and associated tables and buffers.
- Hardware Register Dump Directives
Provide 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,...,p_n.comments

<u>Field</u>	<u>Description</u>
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:

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 DSDI 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.

<u>Parameter</u>	<u>Description</u>
------------------	--------------------

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.

<u>Parameter</u>	<u>Description</u>
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

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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,*.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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. The second format is for models 865 and 875 and CYBER 180-class machines.

Format:

EC. Subsequent C, D, and E directives dump extended memory.

or

EC,fwa. Subsequent C, D, and E directives dump UEM.

<u>Parameter</u>	<u>Description</u>
fwa	First word address (fwa) divided by 1000g. fwa is added to the fwa and last word address (lwa) parameters in subsequent C, D, and E directives. 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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
fwa	First word address divided by 1000g 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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

Format:

I,fba,lba,ei.

<u>Parameter</u>	<u>Description</u>
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.

The output listing is read from top to bottom by column rather than across the page.

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.

Format:

W,fba,lba,ei.

<u>Parameter</u>	<u>Description</u>
fba	First byte address, hexadecimal, to be printed (mandatory).
lba	Last byte address, 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.

PP DUMP DIRECTIVES

PP dump directives provide the capability to obtain a memory dump of all or selected PPs. 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 directive cells. 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:

AP,n₁,n₂,...,n_m.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
n ₁	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) 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, n₁, n₂, ..., n_m.

<u>Parameter</u>	<u>Description</u>
n ₁	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.

<u>Parameter</u>	<u>Description</u>
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.

† 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, n_1, n_2, \dots, n_m .

<u>Parameter</u>	<u>Description</u>
n_i	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, n_1, n_2, \dots, n_m .

<u>Parameter</u>	<u>Description</u>
n_i	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 .

<u>Parameter</u>	<u>Description</u>
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 FLPP 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,n₁,n₂,...,n_m.

<u>Parameter</u>	<u>Description</u>
n _i	Number of FLPP to be dumped. If omitted, all FLPPs are dumped.

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,n₁,n₂,...,n_m.

<u>Parameter</u>	<u>Description</u>
n _i	Number of PP to be printed. If omitted, all PPs are printed.

†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.

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 .

<u>Parameter</u>	<u>Description</u>
n_i	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, n_1, n_2, \dots, n_m .

<u>Parameter</u>	<u>Description</u>
n_i	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, n_1, n_2, \dots, n_m .

<u>Parameter</u>	<u>Description</u>
n_i	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 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.

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, n_1, n_2, \dots, n_m .

<u>Parameter</u>	<u>Description</u>
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, n_1, n_2, \dots, n_m .

<u>Parameter</u>	<u>Description</u>
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. Refer to Interactive Use of DSDI later in this section for additional information concerning use of this directive from an interactive terminal.

Format:

CP,n₁/ops₁,n₂/ops₂,...,n_m/ops_m.

<u>Parameter</u>	<u>Description</u>
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:

<u>ops_i</u>	<u>Description</u>
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†	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†	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.
C†	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.

†This option does not dump information for the system control point.

Parameter

Description

<u>ops_i</u>	<u>Description</u>
D†	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†	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†	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†	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.

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.

†This option does not dump information for the system control point.

CPO — Reset Default List Options

The CPO directive selects a new string of default list options for the CP directive.

Format:

CPO,ops.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

EJT — Dump Executing Job Table

The EJT directive causes the executing job 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:

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 causes the equipment status 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:

EST.

FNT — Dump System File Name/File Status Table

The FNT directive causes the system file name/file status table (FNT/FST) 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:

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.

MST — Dump Mass Storage/Track Reservation Tables

The MST directive causes all mass storage and track reservation tables to be dumped unless EST ordinals are specified, in which case only the specified MSTs are 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. 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,...,est_n.

<u>Parameter</u>	<u>Description</u>
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 while 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.

<u>Parameter</u>	<u>Description</u>
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:

<u>idents</u>	<u>Queue Name</u>
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.

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, n_1, n_2, \dots, n_m .

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 causes the queued file 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:

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.

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.

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, MAG, and IAF). 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 LCD, LIO, 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,LCD,LIO,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, lTA, lTN, and lTO. This option also provides an analysis and dump of all PPs having resident copies of lRO and lRI 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,lTA,lTN,lTO,lRO,lRI.

The exception is that the AP directive also dumps all PPs having copies of lRO and lRI 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 three options (UQP) 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.
- P Provides analysis and full memory dump of all PPs having resident copies of lMT. The output listing generated is the same (in format and content) as that produced by entering the AP directive in this format:

AP,lMT.

Refer to the description of the AP directive earlier in 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.

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.

<u>Parameter</u>	<u>Description</u>
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.

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.

<u>Parameter</u>	<u>Description</u>
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.

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.

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.

<u>Parameter</u>	<u>Description</u>
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.

PRORF — Dump Processor Register File

The PRORF directive is valid on CYBER 180-class machines only. 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.

<u>Parameter</u>	<u>Description</u>
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.

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.

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.

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 directive provides the capability to dump specified buffer controllers and list verification information for controlware dump files produced by the LOADBC utility. 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.

ops_i

Interpretation

A	7-bit ASCII code interpretation.
D	6-bit display code interpretation (default).

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:

<u>Directive</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

Example of terminal output:

```
? 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 EEI19R
```

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.

Example of terminal output:

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	BIO		SYS	←Job Sequence Name at Control Point			
6200	6400	6600	7000	←Control Point Area Address			

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.

<u>Parameter</u>	<u>Description</u>
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.

Example of terminal output:

```
? 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 EEI19R
```

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.

Example of terminal output:

Program in Execution				PP Number	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	----	110-32	----	----	----	1MA-01	----	
7320	7330	7340	7350	7360	7370	7400	7410	7420	7430	
PP00	Input Register Address									

7440										

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.

<u>Parameter</u>	<u>Description</u>
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 parameters will be interpreted as PP numbers.

Example of terminal output:

```
? Q,5,0,100.
0000 0003 2020 3340 ---- 0614 ---- 4334 0117 CPP05 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 / ( )
```

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 (----).

<u>Directive Format</u>	<u>Numeric Representation</u>	<u>Character Representation</u>
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

<u>Parameter</u>	<u>Description</u>
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.

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.

Example of terminal output:

Program in Execution				PP Number	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	----	110-32	----	----	----	1MA-01	----	
7320	7330	7340	7350	7360	7370	7400	7410	7420	7430	
PP00	Input Register Address									

7440										

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.

<u>Parameter</u>	<u>Description</u>
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 parameters will be interpreted as PP numbers.

Example of terminal output:

```
? Q,5,0,100.
0000 0003 2020 3340 ---- 0614 ---- 4334 0117 CPP05 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 / ( )
```

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 (----).

<u>Directive Format</u>	<u>Numeric Representation</u>	<u>Character Representation</u>
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

<u>Parameter</u>	<u>Description</u>
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=ku,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 6-1).
/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 xxxxxxB)	DSDI is terminated by entering a carriage return in response to the ? prompt.
/	

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 6-1 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.

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

†All records in an alternate directives file must be in text format; the first word in each record contains only the record name.

Figure 6-1. Sample Alternate Directives File

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
```

<u>Header</u>	<u>Description</u>
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.
hh.mm.ss. yy/mm/dd. CDC NETWORK OPERATING SYSTEM. NOS 2	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:

The following sample illustrates the printer output listing produced by the PP directive (dump PP communication areas).

RA =	PP.	DUMP 01				DSDI	82/01/20. 16.22.05.				PAGE	1			
0000000	CM	16.52.11. 82/01/06. CDC NETWORK OPERATING SYSTEM.				NOS 2									
PP00	MTR	1524	2234	0000	0000	0000	MTR1	PP01	DSD	0423	0434	0000	0000	0000	DSD1
		0000	0001	0000	0003	5570	A C			0000	0000	0000	0000	0000	
	CP34	0101	1501	0000	0000	0007	AAMA G		CP34	3404	1434	0010	0044	0142	1DL1 H 9A7
		0000	0000	0000	0004	0565	DE		CH10	0000	0000	0000	0000	0000	
		0000	0000	0000	0000	0000	Program in Execution			0000	0000	1401	0613	5400	LAFK=
		0000	0000	0000	0000	0000	Control Point Assignment			5534	3457	3737	5733	3357	11.44.00.
		0000	0000	0000	0000	0000	Channel Assignment			5543	3550	3334	5033	4157	82/01/06.
		0000	0000	0000	0000	0000				0040	0000	0000	0000	0000	5
PP02	QAC	2101	0356	0000	0000	0123	QAC, AS	PP03		0000	0000	0000	0000	7430	X
	UTEM	1077	0000	0000	0003	3734	H C41			0000	0000	0000	0000	0000	
	CP16	0006	0001	4600	0001	0000	F A- A		CP00	0000	0000	0000	0000	0000	
		0000	0002	0300	0000	0000	BC			0000	0002	0235	0022	0001	BB2 R A
		0000	0000	0000	0000	0000				0000	0000	0000	0000	0000	
		0013	0000	1765	0000	2265	K O R			0004	0000	3314	0000	7325	D OL U
		0301	6400	0000	0000	0000	CA			0724	0000	2127	1031	1200	GT QWHYJ
		0330	0000	0000	0000	0125	CX AU			0000	0000	0002	1030	1020	BHXHP
PP04		0000	0000	0000	0000	7350	/	PP05	1R0	3422	1705	0000	0000	0000	1ROE
		0000	0000	0000	0000	0000	Monitor Function		PIOM	2043	0400	0000	0001	5400	P8D A=
	CP00	0000	0000	0000	0000	0000			CP05	0154	0011	4440	0000	0400	A= 195 D
		0000	0000	0235	0022	0001	B2 R A			4000	0000	0504	2600	0000	5 EDV
		0000	0000	0000	0000	0000				0000	0000	0000	0000	0000	
		4002	0000	5502	0002	1500	5B B BM			0003	0000	1776	0000	2777	C O W
		0301	6400	0000	0000	0000	CA			0301	6400	0000	0000	0000	CA
		0330	0000	0000	0000	0333	CX CO			0000	0000	1401	0620	3735	LAFP42
PP06	1SJ	3423	1234	0003	0000	0000	1SJ1 C	PP07	1MT	3415	2431	3170	0000	0000	1MTYY
	LDAM	1055	0001	4472	0305	1073	H A9 CEH		LDAM	1055	0001	4472	0230	1073	H A9 BXH
	CP34	0000	0001	4472	0305	0000	A9 CE		CP31	0000	0001	4472	0230	0000	A9 BX
		0000	0002	0237	0005	0007	BB4 E G			0020	0000	0000	0000	0000	P
		0000	0000	0000	0000	0000				0000	0000	0000	0000	0000	
		0214	4400	0000	0000	0000	BL9			3441	4033	3635	5655	1116	165032, IN
		0301	6400	0000	0000	0000	CA			5700	0000	2100	3645	0100	. Q 3+A
		0330	0000	0000	0000	0333	CX CO			3547	3400	2000	3345	5415	2*1 P 0+=M
PP10		0000	0000	0000	0000	7370		PP11	QAC	2101	0354	0000	0000	0216	QAC= BN
		0000	0000	0000	0000	0000			UTEM	1077	0000	0003	0003	3734	H C C41
	CP00	0000	0000	0000	0000	0000			CP14	0004	0002	4476	0152	0000	D B9 A)
		0000	0002	0235	0004	0014	BB2 D L			0000	0000	0237	0004	0010	B4 D H
		0000	0000	0000	0000	0000				0000	0000	0000	0000	0000	
		4011	0005	5031	0005	5432	SI E/Y E=Z			0013	0000	4126	0000	4426	K 6V 9V
		0724	0000	2127	1031	1200	GT QWHYJ			0724	0000	2127	1031	1200	GT QWHYJ
		0000	0000	0002	1030	1020	BHXHP			0000	0000	0002	1030	1020	BHXHP
PP20	PIP	2011	2003	0000	0000	0136	PIPC A3	PP21	CPD	0320	0434	0000	0000	0225	CPD1 BU
		0000	0000	0000	0000	0000			LDAM	1055	0000	0000	0000	0000	H
	CP03	0311	1703	0000	0002	0331	CI0C BCY		CP34	0000	0004	6123	0131	0000	D SAY
	CH05	0000	0000	0000	0000	0000				0000	0005	0424	0021	0022	EDT Q R
		1130	6700	6017	1701	0574	IX 00AE			0000	0000	0000	0000	0000	
		0100	6202	0200	6155	0363	A BB C			0777	1777	3777	0100	7100	G O 4 A
		0100	6022	3004	1201	3101	A RXDJAYA			2000	0045	6001	3003	1013	P + AXCHK
		3406	1411	3204	0502	1412	1FLIZDEBLJ			0770	1401	3411	3412	3077	G LA11JX

Example 2:

The following sample illustrates the printer output listing produced by the AP directive (dump analysis of PP and PP memory in line format).

RA = AP,3. DUMP 01 DSDI 82/01/20. 16.22.09. PAGE 3
 000000 PP03 16.52.11. 82/01/06. CDC NETWORK OPERATING SYSTEM. NOS 2

ANALYSIS OF PP03

```

PP03      0000 0000 0000 0000 7430      X      LAST MAIN PROGRAM LOADED -
          0000 0000 0000 0000 0000      LAST OVERLAY LOADED - 3ME
          CPOO 0000 0000 0000 0000 0000      LAST MASS STORAGE DRIVER - 6DI
          0000 0002 0235 0022 0001      BB2 R A
          0000 0000 0000 0000 0000
          0004 0000 3314 0000 7325      D OL U
          0724 0000 2127 1031 1200      GT QWHYJ
          0000 0000 0002 1030 1020      BHXHP
  
```

RESIDENT ENTRY POINTS			LOW MEMORY CONSTANTS			
NAME	LOC	CALLER	NAME	LOC	ACTUAL	EXPECTED
FTN	0163	2205	RA	55	4004	----
DFM	0432	2341	FL	56	0021	----
EXR	0471	1477	ON	70	0001	0001
SMS	0506	0355	HN	71	0100	0100
RDS	0553	0370	TH	72	1000	1000
WDS	0556	1740	TR	73	0003	0003
EMS	0561	0415	CP	74	6200	----
			IA	75	7230	----
			OA	76	7231	----
			MA	77	7232	----
			R-REG		00010000	00000000

*** WARNING ***

PPU MEMORY

```

0000 0216 2210 2175 2616 0006 0001 4473 0315  BNRHQ VN F A9 CM 0010 ---- ---- ---- ---- ---- 2120 ---- ----      QP
0020 0065 0030 0027 ---- 4000 ---- ---- ----      X W S 0030 ---- 4000 ---- ---- ---- ---- 0301 ----      5      CA
0040 ---- 0002 5504 ---- 0203 0053 3101 0006      B D BC $YA F 0050 ---- ---- ---- ---- ---- 7430 4004 0021 0312      X5D QGJ
0060 ---- 3314 ---- ---- ---- ---- ----      OL 0070 0001 0100 1000 0003 6200 7230 7231 7232      AA H C X Y Z
0100 ---- ---- ---- ---- 0002 ---- 0343 0411      B C8DI 0110 ---- 0004 2711 ---- ---- 1006 0621 1002      DWI HFFQHB
0120 0704 2000 1701 0576 3010 2200 0177 0336  GDP OAE XHR A C3 0130 1457 6010 3010 3374 0513 1425 0327 2000      L. HXHO EKLUCWP
0140 1720 0676 3076 6010 3010 1301 0546 3074  OPF X HXHKAE-X 0150 1625 6010 2411 3013 3455 3014 3456 3076      NU HTIXK1 XL1,X
0160 6010 1400 0100 2205 0443 3410 3076 6210  HL A RED81HX H 0170 1427 3210 0644 0312 0711 3011 1277 5100      LWZHF9CJGIXIJ (
0200 0510 6010 3010 1015 0731 2005 1455 6010  EH HXHHMGYPEL H 0210 3010 3111 0572 2005 1447 6370 0111 2005      XHYIE PEL* AIPE
0220 1447 2610 6010 3014 0445 1720 0676 1502  L*VH HXLD+OPF MB 0230 3500 0665 1476 6210 0356 2001 2551 6173      2 F L HC,PAU(
0240 7762 ---- ---- ---- ---- ---- ---- ----      0250 2000 1704 0576 3075 6050 3050 0471 3077      P ODE X /X/D X
0260 6173 0020 3051 1340 1006 3474 3574 3374      PX(K5HF1 2 0 0270 3350 1006 0200 0321 0115 0005 2000 0036      O/HFB CQAM EP 3
0300 3412 2000 1505 3413 1473 0200 0163 0326  1JP ME1KL B A CV 0310 3311 1014 3112 6113 2120 1400 0200 0163      OIHLYJ KQPL B A
0320 0100 0474 5400 0302 1063 5400 0277 3076  A D = CBH = B X 0330 6010 3710 0543 3076 6210 3014 0502 3015      H4HE8X HXLEBXM
0340 3415 5400 0314 3011 1377 0542 3015 1702  1M= CLXIK E7XM0B 0350 3403 3076 6004 0200 0506 1440 5500 0110      1CX DB EFL5 AH
0360 4003 3416 5003 0001 3417 3003 0200 0553  5C1N/C A10XCB E$ 0370 0727 3607 5300 0106 0504 3407 4003 3406      GW3G$ AFED1G5C1F
0400 3016 4403 3017 5403 0001 2000 0500 3503  XN9CXO=C AP E 2C 0410 5715 0003 0545 0200 0561 0100 0315 2000      .M CE+B E A CMP
0420 2725 0302 3071 2342 0420 3401 1453 6173  WUCBX S7DP1AL$ 0430 7762 0100 2341 5400 0457 1701 3402 1063      A S6= D.OA1BH
0440 3412 1217 1111 0410 3602 4002 0575 3602  1JJ0I0H3B5BE 3B 0450 5200 0457 3411 1406 3401 3077 6301 2572      ) D.1ILF1AX AU
0460 1452 0200 0163 3014 0444 2001 2225 6173  L)B A XLD9PARU 0470 7764 1477 0200 0321 5000 0471 5415 0006      L B CQ/ D =M F
0500 0115 0007 1466 6170 0100 0100 0355 2000  AM GL A A C P 0510 7464 3105 6010 3014 1003 1606 6170 0102      YE HXLHCF AB
0520 5000 0110 1237 0556 5000 0104 5200 0107  / AHJ4E, / AD) AG 0530 0420 5500 0107 3413 1441 3412 1473 0200      DP AG1KL61JL B
0540 0163 3011 1014 3112 6113 0551 0340 ----  A XIHLYJ KE(C5 0550 0100 1056 0100 0370 0314 0100 1740 0334      A H,A C CLA 05C1
0560 0100 0415 1400 3413 4471 0200 0773 0370  A DML 1K9 B G C 0570 5400 0606 4771 0705 3107 5300 0624 0403      = FF* GEYGS FTDC
0600 0200 0663 1404 0200 0746 7106 2116 0532  B F LDB G- FQNEZ 0610 5000 0553 0335 5400 0640 1063 5400 0633      / ESC2= F5H = F0
0620 4771 0705 3107 2300 0342 0403 0200 0663  * GEYGS C7DCB F 0630 4071 0403 2000 ---- 1135 0200 0746 7306      5 DCP I2B G- F
0640 1131 0515 5000 0556 6606 0644 7546 3410  IYEM/ E, FF9 -1H 0650 1412 0200 0746 1401 7106 0013 0525 3013      LJB G-LA F KEUXK
0660 0524 0110 ---- 0602 3076 6010 3010 0574  ETAH FBX HXHE 0670 3077 6204 1455 0200 0163 3714 4471 3107      X DL B A 4L9 YG
  
```

Example 3:

The following sample 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 CM 22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM. NOS 2

0400 - CONTROL POINT 02

CP02 EXCHANGE PACKAGE

P	33357	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	46600	A2	26250	B2	31053	(A2)=0000	0000	0000	0000	0024	T	(B2)=0400	0320	2500	0000	0000	D CPU		
EM	70070000	A3	14424	B3	0	(A3)=0000	0000	0000	0000	0004	3521	D2Q	(B3)=0000	0000	0000	0000	0000		
RAE	0	A4	33070	B4	43	(A4)=0000	0000	0000	0003	3244	CZ9	(B4)=0000	0000	0000	0000	0000			
FLE	0	A5	33072	B5	31367	(A5)=0000	0000	0000	0003	3273	CZ	(B5)=0400	0320	5100	0000	0000	D CP C		
MA	400	A6	1	B6	33646	(A6)=0000	0000	0000	0000	0000		(B6)=5110	0336	3001	0003	3274	(HC3XA CZ		
.	0	A7	33070	B7	43515	(A7)=0000	0000	0000	0003	3244	CZ9	(B7)=0000	0004	0435	0204	3715	DD2BD4M		

X0	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	D2Q	(X3)=1000	0150	0000	0000	0000	H A/							
X4	0000	0000	0000	0003	3243	CZ8	(X4)=0400	0321	6630	2205	0301	D C4 XRECA							
X5	0000	0000	0000	0003	3273	CZ	(X5)=0000	0000	0000	0000	0000								
X6	2203	1400	0000	0000	0000	RCL	(X6)=0000	0000	0000	0000	0000								
X7	0000	0000	0000	0003	3244	CZ9	(X7)=0400	0320	1130	0530	0503	D CPIXEXC							

PARAMETER SUMMARY

USER NAME	NETOPS	SENSE SWITCHES	00
PRIMARY FILE		CCL R1	000000
ORIGIN TYPE	SY0T	CCL R2	000000
TIME USED	0000000022	CCL R3	000000
TIME REMAINING	1014223327	CCL EF	00
SUBSYSTEM ID	7770	CPU PRIORITY	0077
CPU STATUS	X	SERVICE CLASS xxxx	SSSC

MESSAGE 1 - NP/SN102 22.05.09.LLINK L003103,EN,RL=3,H-N,S
MESSAGE 2 - \$REQUEST *K* DISPLAY

CURRENT COMMAND - NIP(NIN=199,ISTP=YES,FASTP=YES,MC=500)

LAST DAYFILE MESSAGE - 22.05.09.NP/SN102 22.05.09.LLINK L003103,EN,RL=3,H-N,SN102 /103,SN102 /003

SPECIAL ENTRY POINTS -

X List Option

EJT ENTRY

025000	1601	1555	0137	7777	0002	NAM A4	B
025001	0000	0000	4000	0002	7776	5	B
025002	0000	0000	0000	0000	0000		
025003	0700	0000	0000	0000	0036	G	3

(Continued)

(Continued)

```

RA = CP,2. DUMP 00 DSDI 83/06/02. 22.36.17. PAGE 4
000000 CM 22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM NOS 2

LIMIT FOR SIZE OF IAFP 7
USER NAME UIDW 000470 16052417202300 NETOPS 4
CHARGE FLAG, USER INDEX 377772
NO EXIT FLAG EECW 000471 2 P B5H
RESERVED 000
REPRIEVE DATA 0077
TERMINAL INPUT POINTER TINW 000000
REPRIEVE DATA 024010
EJT ORDINAL OF JOB TFSW 000472 0017 0
PRIMARY FILE FNT OFFSET 0000
RESERVED 00
ROLLOUT TIME TERW 000
EVENT DESCRIPTOR 0000000
RESERVED B IA1
LOGOUT EPILOG, EPILOG REQ 0
PR CHARGE/USER/INHIBIT DEC 2
EOR FLAG/CS COUNT 00000011
NEXT STATEMENT INDEX 0134
LIMIT INDEX 0065
INPUT/SKIP FLAGS CSSW 000474 0 F7A7A B
EST ORDINAL 006
FIRST TRACK 4201
CURRENT TRACK 4201
CURRENT SECTOR 0002
OVERLAP WORD COUNT 0000

```

```

RECALL REQUESTS IR-S 000553 00000000000000000000
RECALL REQUESTS IR-S 000554 00000000000000000000
RECALL REQUESTS IR-S 000555 00000000000000000000
RECALL REQUESTS IR-S 000556 00000000000000000000
RECALL REQUESTS IR-S 000557 00000000000000000000
RECALL REQUESTS IR-S 000560 00000000000000000000
RECALL REQUESTS IR-S 000561 00000000000000000000
RECALL REQUESTS MB-S REPW 000562 00000000000000000000
RECALL REQUESTS MB-S 000563 00000000000000000000
RECALL REQUESTS MB-S 000564 00000000000000000000
RECALL REQUESTS MB-S 000565 00000000000000000000
RECALL REQUESTS MB-S 000566 00000000000000000000
RECALL REQUESTS MB-S 000567 00000000000000000000
RECALL REQUESTS MB-S 000570 00000000000000000000
RECALL REQUESTS MB-S 000571 00000000000000000000
RECALL REQUESTS MB-S 000572 00000000000000000000
RECALL REQUESTS MB-S 000573 00000000000000000000
RECALL REQUESTS MB-S 000574 00000000000000000000
RECALL REQUESTS MB-S 000575 00000000000000000000
RECALL REQUESTS MB-S 000576 00000000000000000000
RECALL REQUESTS MB-S 000577 00000000000000000000

```

T List Option resumes at top of right column on sheet 2 of 6.

DAYFILE POINTERS AND BUFFER ← A List Option

```

00120275 0011 7777 0000 0000 0046 I -
00120276 0000 0000 0114 0000 0104 AL AD
00120277 0007 4200 0010 0003 G7 7 H C

00117777 55353557334057334157 55232520052226112311 17165507011116050400 00000000000000000000 22.05.06. SUPERVISION GAINED
00120003 55353557334057334157 16205023163433355555 55353557334057334157 55030320552605222311 22.05.06.NP/SN102 22.05.06. CCP VERSI
00120007 17160055553636565514 05260514005540443756 55260122110116240055 4044370000000000000000 ON 33, LEVEL 594, VARIANT 594
00120013 55353557334057334157 16205023163433355555 55353557334057334157 55202205261117252355 22.05.06.NP/SN102 22.05.06. PREVIOUS
00120017 03235516170405005533 3335655202205261117 25235516235516170405 00553333300000000000 CS NODE 000, PREVIOUS NS NODE 000
00120023 55353557334057334457 16205023163433355555 55353557334057334457 14141116130055143333 22.05.09.NP/SN102 22.05.09.LLINK L00
00120027 36343336560516562214 54365610461656231634 333555550333365623 1634333555550343336 3103,EN,RL=3,H-N,SN102 /003,SN102 /103
00120033 00000000000000000000 55353557334057334457 16205023163433355555 55353557334057334457 22.05.09.NP/SN102 22.05.09.
00120037 14141116130055143333 36343336560516562214 54365610461656231634 3335555503433365623 LLINK L003103,EN,RL=3,H-N,SN102 /103,S
00120043 163433355555033336 00000000000000000000 162350555555555555 55353557333757373657 N102 /003 NS/ 22.04.43.
00120047 2316343335555565504 2515205516203333444 44550317152014052405 04570000000000000000 SN102, DUMP NP00199 COMPLETED.
00120053 55353557333757373657 162350555555555555 55353557333757373657 23163433355555555555 22.04.43.NS/ 22.04.43.SN102 , L
00120057 17010411160755511511 03555555257550000000 55353557333757374157 16235055555555555555 OADING (MIC ). 22.04.46.NS/
00120063 55353557333757374157 231634333555555565514 17010411160755510211 07141555525755000000 22.04.46.SN102 , LOADING (BIGLM ).
00120067 55353557333757404157 162350555555555555 55353557333757404157 231634333555555565514 22.04.56.NS/ 22.04.56.SN102 , L
00120073 17010411160755511603 02555555257550000000 55353557333757404257 16235055555555555555 OADING (NCB ). 22.04.57.NS/
00120077 55353557333757404257 231634333555555565514 1701045031715201405 24050457000000000000 22.04.57.SN102 , LOAD COMPLETED.
00120103 55353557334057334157 16205023163433355555 55353557334057334157 16202500552316343335 22.05.06.NP/SN102 22.05.06.NPU SN102
00120107 55555601035634333600 00000000000000000000 55353557334057334157 16205023163433355555 ,AC,103 22.05.06.NP/SN102

```

DAYFILE LINES IN BUFFER

```

NS/ 22.04.43.SN102 , DUMP NP00199 COMPLETED.
22.04.43.NS/ 22.04.43.SN102 , LOADING (MIC ).
22.04.46.NS/ 22.04.46.SN102 , LOADING (BIGLM ).

```

(Continued)

(Continued)

(Continued)

RA = CP,2. DUMP 00 DSDI 83/06/02. 22.36.18. PAGE 5
0000000 CM 22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM NOS 2

22.04.56.NS/ 22.04.56.SN102 , LOADING (NCB).
22.04.57.NS/ 22.04.57.SN102 , LOAD COMPLETED.
22.05.06.NP/SN102 22.05.06.NPU 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 NODE 000, PREVIOUS NS NODE 000
22.05.09.NP/SN102 22.05.09.LLINK L003103,EN,RL=3,H-N,SN102 /003,SN102 /103
22.05.09.NP/SN102 22.05.09.LLINK L003103,EN,RL=3,H-N,SN102 /103,SN102 /003

ATTACHED FILES ← F List Option

0304 INPUT IN FNT - 1116 2025 2400 0001 1700 INPUT AO FUT 0000 0000 0000 0000 0000
FST - 0006 4175 4175 0002 0005 F6 6 B E EST 7700 6210 0013 1377 0421 1277 H KK DQJ
TRACK CHAIN -
4175 0002

0312 NAM LO FNT - 1601 1500 0000 0000 1500 NAM M FUT 0000 0000 0000 0000 0000
FST - 0006 4176 4176 0004 0005 F6 6 D E EST 7700 6210 0013 1377 0421 1277 H KK DQJ
TRACK CHAIN -
4176 0005

0315 OUTFIL LO FNT - 1725 2406 1114 0000 1500 OUTFIL M FUT 0000 0000 0000 0000 0000
FST - 0007 4205 4205 0001 0105 G7E7E AAE EST 7702 6210 0013 1377 0421 1302 H KK DQKB
TRACK CHAIN -
4205 0003

0320 ZZZZC0 LI FNT - 3232 3232 3203 3301 1007 ZZZZC0AHG FUT 0000 0000 0000 0000 0000
FST - 0006 4175 4175 0001 0005 F6 6 A E EST 7700 6210 0013 1377 0421 1277 H KK DQJ
TRACK CHAIN -
4175 0002

0323 ZZZZC2 LO FNT - 3232 3232 3203 3500 1507 ZZZZC2 MG FUT 0000 0000 0000 0000 0000
FST - 0006 4177 4177 0002 0307 F6 6 BCG EST 7700 6210 0013 1377 0421 1277 H KK DQJ
TRACK CHAIN -
4177 0002

0326 OUTPUT PM FNT - 1725 2420 2524 0000 1200 OUTPUT J FUT 0000 0000 0700 0000 0000 G
FST - 0010 6742 6742 0001 0005 H 7 7 A E EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE
TRACK CHAIN -
6742 0001

0331 NRF1 LO FNT - 1622 0634 0000 0000 1500 NRF1 M FUT 0000 0000 0000 0000 0000
FST - 0007 4210 4210 0002 0705 G7H7H BGE EST 7702 6210 0013 1377 0421 1302 H KK DQKB
TRACK CHAIN -
4210 0003

0334 NRF2 LO FNT - 1622 0635 0000 0000 1500 NRF2 M FUT 0000 0000 0000 0000 0000
FST - 0006 4200 4200 0001 0705 F7 7 AGE EST 7700 6210 0013 1377 0421 1277 H KK DQJ
TRACK CHAIN -
4200 0003

0337 LIST PM FNT - 1411 2324 0000 0000 1200 LIST J FUT 0000 0000 0700 0000 0000 G
FST - 0010 7052 7052 0001 0005 H)) A E EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE
TRACK CHAIN -
7052 0001

0342 ZZZZPP PM FNT - 3232 3232 3220 2000 1200 ZZZZPP J FUT 0000 0000 0700 0000 0000 G

RA = CP,2. DUMP 00 DSDI 83/06/02. 22.36.18. PAGE 6
0000000 CM 22.15.04. 83/06/02. CDC NETWORK OPERATING SYSTEM NOS 2

FST - 0010 7053 7053 0001 0005 H \$ \$ A E EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE
TRACK CHAIN -
7053 0001

0345 ZZZZDMB PM FNT - 3232 3232 0415 0200 1200 ZZZZDMB J FUT 0000 0000 0700 0000 0000 G
FST - 0010 7054 7054 0001 0005 H = = A E EST 7704 4210 3226 1377 0412 1305 7HZVK DJKE
TRACK CHAIN -
7054 0001

0350 ZZZZDN LO FNT - 3232 3232 3204 1600 1500 ZZZZDN M FUT 0000 0000 0000 0000 0000
FST - 0007 4216 4216 0022 0703 G7N7N RGC EST 7702 6210 0013 1377 0421 1302 H KK DQKB
TRACK CHAIN -
4216 0022

Example 5:

The following sample 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 CM 16.52.11. 82/01/06. CDC NETWORK OPERATING SYSTEM. NOS 2

0000050 0000 0000 0000 0000 0000 0000061 7770 0000 0000 0000 0000 0000072 7467 56475 12000 00000 J
----- 0000062 0000 0000 01065 34137 AF$64 0000073 00033 36401 00007 30005 CO A E
0000052 0000 0000 0000 07240 5 0000063 0000 0000 0000 0000 0000 0000074 00033 56404 00035 00025 C2 D C/ U
0000053 30766 01030 10057 43001 X HXHE XA 0000064 0000 0000 0000 51426 ELV 0000075 00035 56406 00057 50002 C F E B
0000054 34131 44234 12147 30200 1KL71JL B 0000065 00042 71100 0000 0000 DWI 0000076 00000 00000 00000 00000
0000055 01633 01401 00034 17761 A XLA C6 0000066 00000 00000 00000 00000 00000077 00007 75400 00400 07200 = 5
0000056 00000 00000 00000 07164 ----- 0000100 00000 00000 00000 00000
0000057 77775 12230 00320 00000 (RX Z 0000070 00000 03300 24040 40000 0 TDD -----
0000060 00751 10000 00000 00000 I 0000071 74500 00122 25000 07421 / ARU Q 0000107 00054 13700 01065 34117 E64 AF$60

CM D,50,110. D - FORMAT DUMP.

0000050 0000 0000 0000 0000 0000 0000061 7777 0000 0000 0000 0000 0000072 7464 7564 7512 0000 0000 J
----- 0000062 0000 0000 0001 0653 4137 AF$64 0000073 0003 3364 0100 0073 0005 CO A E
0000052 0000 0000 0000 0000 7240 5 0000063 0000 0000 0000 0000 0000 0000074 0003 3564 0400 0350 0025 C2 D C/ U
0000053 3076 6010 3010 0574 3001 X HXHE XA 0000064 0000 0000 0000 0005 1426 ELV 0000075 0003 5564 0600 0575 0002 C F E B
0000054 3413 1442 3412 1473 0200 1KL71JL B 0000065 0004 2711 0000 0000 0000 DWI 0000076 0000 0000 0000 0000 0000
0000055 0163 3014 0100 0341 7761 A XLA C6 0000066 0000 0000 0000 0000 0000 0000077 0000 7754 0000 4000 7200 = 5
0000056 0000 0000 0000 0000 7164 ----- 0000100 0000 0000 0000 0000 0000
0000057 7777 5122 3000 3200 0000 (RX Z 0000070 0000 0033 0024 0404 0000 0 TDD -----
0000060 0075 1100 0000 0000 0000 I 0000071 7450 0001 2225 0000 7421 / ARU Q 0000107 0005 4137 0001 0653 4117 E64 AF$60

CM E,50,110. E - FORMAT DUMP.

0000050 000000000000000000 000000000000000000 0000000000000007240 30766010301005743001 5X HXHE XA
0000054 34131442341214730200 01633014010003417761 0000000000000007164 77775122300032000000 1KL71JL B A XLA C6 (RX Z
0000060 00751100000000000000 77770000000000000000 00000000000106534137 00000000000000000000 I AF$64
0000064 0000000000000051426 00042711000000000000 00000000000000000000 00000000000000000000 ELV DWI
0000070 000003300240400000 74500001222500007421 74647564751200000000 00033364010000730005 O TDD / ARU Q J CO A E
0000074 00033564040003500025 00035564060005750002 00000000000000000000 00007754000040007200 C2 D C/ U C F E B = 5
0000100 000000000000000000 000000000000000000 000000000000000000 000000000000000000 000000000000000000
0000104 000000000000000000 000000000000000000 000000000000000000 00054137000106534117 E64 AF$60

```

Example 6:

The following sample illustrates the printer output listing of the system file name table produced by the FNT CMR 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      DSDI      83/06/12. 12.33.11.    PAGE    1
0000000  CM      12.31.31. 83/06/12. CDC NETWORK OPERATING SYSTEM  NOS 2

      SYSTEM FILENAME TABLE

00024706 2331 2324 0515 0000 1000 SYSTEM H 00024723 0000 0000 0000 0000 0000      00024740 0000 0000 0000 0000 0016      N
00024707 0006 4005 0000 0000 0000 F5E      00024724 0000 0000 0000 0000 0010      H 00024741 0000 0000 0000 0000 0000
00024710 2223 3004 4137 0000 1300 RSXD64 K 00024725 0000 0000 0000 0000 0000      00024742 0000 0000 0000 0000 0017      O
00024711 0010 4463 0000 0000 0000 H9      00024726 0000 0000 0000 0000 0011      I 00024743 0000 0000 0000 0000 0000
00024712 2223 3026 4137 0000 1300 RSXV64 K 00024727 0000 0000 0000 0000 0000      00024744 0000 0000 0000 0000 0020      P
00024713 0010 4201 0000 0000 0000 H7A      00024730 0000 0000 0000 0000 0012      J 00024745 0000 0000 0000 0000 0000
00024714 2022 1706 1114 0300 1300 PROFILC K 00024731 0000 0000 0000 0000 0000      00024746 0000 0000 0000 0000 0021      Q
00024715 0010 5410 0000 0000 0000 H=H      00024732 0000 0000 0000 0000 0013      K 00024747 0000 0000 0000 0000 0000
00024716 2601 1411 0425 2300 1300 VALIDUS K 00024733 0000 0000 0000 0000 0000      00024750 0000 0000 0000 0000 0022      R
00024717 0010 6504 0000 0000 0000 H D      00024734 0000 0000 0000 0000 0014      L 00024751 0000 0000 0000 0000 0000
00024720 2601 1411 0425 2300 1300 VALIDUS K 00024735 0000 0000 0000 0000 0000
00024721 0005 4215 0000 0000 0000 E7M      00024736 0000 0000 0000 0000 0015      M
00024722 0000 0000 0000 0000 0007      G 00024737 0000 0000 0000 0000 0000

```


NOS allows multiple copies of the running system to reside on several disk devices of the same type. The system deadstart file (SDF) is defined to be a copy of the deadstart tape that resides on a deadstart disk. When you deadstart the system from this disk, the SDF is read to generate copies of the running system.

The INSTALL command installs a copy of the running system or a deadstart file you specify to a deadstart disk. The disk selected must be on, available, a single-unit device such as an 844 or 885-11/12,[†] and must have the Common Test and Initialization (CTI) software. (For more information about CTI, refer to the CIP 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.

The format of the command is:

```
INSTALL,filename,EQest.
```

<u>Parameter</u>	<u>Description</u>
filename	<p>Disk file (assigned to the job) to be installed as an SDF. Default is SYSTEM, which must be assigned to the job. SDF is a reserved file name and cannot be specified for filename.</p> <p>If the deadstart file to be installed is on tape, it must first be copied to disk before INSTALL can be used.</p>
est	EST ordinal of the deadstart disk on which file filename is to be installed.

[†]SDF cannot be installed to an 819 or 885-42 device.

CYBERLOG K Display	8-2
Responding to CYBERLOG Displays	8-2
Event Type	8-3
Reason	8-3
Component	8-4
Elapsed Time	8-7
Lost Time	8-7
Impact	8-8
Edit	8-9
Flaw K Display	8-10
Initialize K Display	8-14
Machine Recovery (MREC) Utility K Display	8-20
MREC Procedures	8-21
Using MREC With the K Display	8-21
Using MREC Without the K Display	8-24
MREC Unit and Controller Reservations	8-26
Mass Storage Archival Subsystem (MSAS) K Display	8-27
Mass Storage Subsystem (MSS) K Display	8-28
Network Access Method (NAM) K Display	8-29
Display Control Characters	8-31
NAM K-Display Operation	8-32
NAM Mode Maintenance Commands	8-34
DB Command	8-34
DE Command	8-35
DU Command	8-35
FL Command	8-36
LB Command	8-36
LE Command	8-37
LR Command	8-37
RS Command	8-38
ST Command	8-39
Queue File Transfer Facility (QTF) K Display	8-42
Redefine K Display	8-43
Remote Batch Facility (RBF) K Display	8-54
Remote Host Facility (RHF) K Display	8-57
RHF Initiation	8-58
Operator Interface	8-58
RHF Commands Under K Display	8-60
Application Table Display	8-61
Network Path Table Display	8-62
RHF Commands Available Under Application or Network Path Display	8-64
RHF Termination	8-64

(Continued)

(Continued)

SCOPE 2 Station Facility (SSF) K Displays	8-65
Operator Interface	8-65
Station Login	8-65
Enabling and Disabling File Transfers	8-65
Station Recovery	8-66
Station Disconnection and Logout	8-66
File Transfer Limit Commands	8-67
Staged File Transfer Commands	8-67
Spooled File Transfer Commands	8-67
Transaction Facility (TAF) K Displays	8-68
Initiating TAF K Display	8-68
Restarting TAF K Display	8-71
TAF K Display	8-71
TAF K-Display Commands	8-72
TAF/CRM Status K Displays and Commands	8-76

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, queue file utilities, and the status/control register simulator.

<u>K Display</u>	<u>Description</u>
CYBERLOG	Gathers information on the system's performance.
Flaw	Disk flaw mapping.
Initialize	Disk initialization.
Machine Recovery (MREC)	Multimainframe device recovery.
Mass Storage Archival Subsystem (MSAS)	K display of the MSAS.
Mass Storage Subsystem (MSS)	K display of the MSS.
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 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.

CYBERLOG K DISPLAY

The CYBERLOG utility gathers information on system performance. CYBERLOG gathers the information from operator responses after deadstart is complete. The information gathered contains the number and type of service interruptions, the elapsed and lost time from the interruption, and the level of severity or impact of each interruption.

Initiate CYBERLOG by entering:

X.CYBRLOG.

When CYBERLOG appears at a control point, a request message on the B,0 display asks you to assign the K display to the CYBERLOG job. Enter:

K,jsn.

jsn Job sequence name of CYBRLOG.

RESPONDING TO CYBERLOG DISPLAYS

The CYBERLOG displays request data about the event that causes the system to be down or degraded. Your selection of the event type, scheduled or unscheduled, determines the path through the displays. The data requested for the two event types are:

<u>Event Type</u>	<u>Data Requested</u>
Scheduled	Reason Elapsed time
Unscheduled	Reason Component Elapsed time Lost time Impact

EVENT TYPE

The initial display (figure 8-1) requests the type of event being recorded. Enter 0 for a scheduled event and 1 for an unscheduled event.

CYBERLOG

ENTER THE CODE FOR THE TYPE OF EVENT

CODE	DESCRIPTION
0	SCHEDULED
1	UNSCHEDULED

Figure 8-1. CYBERLOG Display for Type of Event

In the resulting reason display, the item you select appears on the line labeled EVENT TYPE, and the next list of options appears on the bottom.

REASON

After you identify the type of event that occurred, a display appears that asks you to identify the reason for the event. The specific display depends on the type of event (figure 8-2 or 8-3).

For a scheduled event, enter the code that corresponds to the reason. On the resulting elapsed-time request display, the item you select appears on the line labeled REASON.

For an unscheduled event, select the most likely reason. In the resulting display, the item you select appears on the line labeled REASON. Selecting 0, 1, 2, or 3 leads to a display that requests the component causing the failure. Selecting 4 (UNKNOWN) causes a branch to the elapsed time request display.

NOTE

Exact information is required. If you do not know the reason for the event, enter the code for unknown.

```

CYBERLOG

EVENT TYPE:  SCHEDULED

ENTER THE CODE FOR THE REASON FOR THE EVENT

CODE      DESCRIPTION
0         BEGINNING OF NORMAL DAY
1         HARDWARE RECONFIGURATION
2         SOFTWARE RECONFIGURATION
3         MAINTENANCE
4         OTHER

```

Figure 8-2. CYBERLOG Reason for Scheduled Event Display

```

CYBERLOG

EVENT TYPE:  UNSCHEDULED

ENTER THE CODE FOR REASON FOR EVENT

CODE      DESCRIPTION
0         OTHER
1         HARDWARE
2         SOFTWARE
3         COMMUNICATIONS
4         UNKNOWN

```

Figure 8-3. CYBERLOG Reason for Unscheduled Event Display

COMPONENT

After you enter the reason for an unscheduled event, a display appears that asks you to identify the component that caused the event. The specific display that appears depends on the reason you've entered (figure 8-4, 8-5, 8-6, or 8-7). Enter the code of the component that caused the failure. The component you select appears on the line labeled COMPONENT on the resulting elapsed-time request display.

NOTE

Exact information is required. If you do not know the component causing the failure, enter the code for UNKNOWN.

CYBERLOG

EVENT TYPE: UNSCHEDULED

REASON: OTHER

ENTER THE CODE FOR THE COMPONENT CAUSING FAILURE

CODE	DESCRIPTION
0	OTHER
1	TEMP - HUMIDITY
2	ELECTRICAL
3	POWER SUPPLY
4	PROCEDURAL
5	UNKNOWN

Figure 8-4. CYBERLOG Component Request Display (Reason OTHER)

CYBERLOG

EVENT TYPE: UNSCHEDULED

REASON: HARDWARE

ENTER THE CODE FOR THE COMPONENT CAUSING FAILURE

CODE	DESCRIPTION
0	OTHER
1	CPU
2	PPU/CHANNEL
3	MEMORY
4	EXTENDED MEMORY
5	TAPE SUBSYSTEM
6	DISK SUBSYSTEM
7	MASS STORAGE SUBSYSTEM
8	UNKNOWN

Figure 8-5. CYBERLOG Component Request Display (Reason HARDWARE)

CYBERLOG

EVENT TYPE: UNSCHEDULED

REASON: SOFTWARE

ENTER THE CODE FOR THE COMPONENT CAUSING FAILURE

CODE	DESCRIPTION
0	OTHER
1	OPERATING SYSTEM
2	COMPILER/ASSEMBLER
3	APPLICATION PROGRAMS
4	DATA MANAGEMENT
5	UNKNOWN

Figure 8-6. CYBERLOG Component Request Display (Reason SOFTWARE)

CYBERLOG

EVENT TYPE: UNSCHEDULED

REASON: COMMUNICATIONS

ENTER THE CODE FOR THE COMPONENT CAUSING FAILURE

CODE	DESCRIPTION
0	OTHER
1	SUBSYSTEM HARDWARE
2	HOST SOFTWARE
3	LINE ADAPTER
4	MODEM
5	SUBSYSTEM SOFTWARE
6	UNKNOWN

Figure 8-7. CYBERLOG Component Request Display (Reason COMMUNICATIONS)

ELAPSED TIME

After you either enter the reason for a scheduled event or enter the component causing the failure for an unscheduled event, a display appears that asks you to record the amount of time that has elapsed since the event occurred (figure 8-8).

Enter the elapsed time in hours and minutes in the format shown. The range of values for hours is from 0 to 99 and the range of values for minutes is from 00 to 59. A leading zero is not required for hours, but is required for minutes. A period is required to delimit the hours from the minutes. The time you enter appears on the line labeled ELAPSED TIME in the resulting display.

```

                                CYBERLOG

EVENT TYPE:  UNSCHEDULED

REASON:      HARDWARE

COMPONENT:   DISK SUBSYSTEM

ENTER THE ELAPSED TIME SINCE THE SYSTEM WAS FULLY OPERATIONAL
IN THE FORM

          HH.MM
```

Figure 8-8. CYBERLOG Elapsed-Time Request Display

LOST TIME

After you enter the elapsed time for an unscheduled event, a display appears that asks you to enter the length of time during which the system or part of the system was down or degraded (figure 8-9). For scheduled events the field is set to 00.00.

Enter the lost time in hours and minutes in the format shown. The same restrictions in format apply as in the elapsed-time entry.

```

                                CYBERLOG

EVENT TYPE:  UNSCHEDULED

REASON:      HARDWARE

COMPONENT:   DISK SUBSYSTEM

ELAPSED TIME: 03.15

ENTER THE TIME LOST IN THE FORMAT:

          HH.MM
```

Figure 8-9. CYBERLOG Lost-Time Request Display

IMPACT

After you enter the lost time for an unscheduled event, a display appears that asks you to identify the impact of the event (figure 8-10). Enter the code for degraded if the system resources were reduced; enter the code for down if the system resources were unavailable. The item you select appears on the line labeled IMPACT on the resulting edit display.

CYBERLOG	
EVENT TYPE:	UNSCHEDULED
REASON:	HARDWARE
COMPONENT:	DISK SUBSYSTEM
ELAPSED TIME:	03.15
LOST TIME:	03.15
ENTER THE CODE FOR THE IMPACT OF THE EVENT	
CODE	DESCRIPTION
0	DEGRADED
1	DOWN

Figure 8-10. CYBERLOG Request for Impact Display

EDIT

After you make all the entries, you have an opportunity to change any of your responses (figure 8-11). By entering the number that now precedes each line with an entry, that line is blanked; subsequent lines that depend on the blanked line are removed from the display. Those choices are then presented again.

To add a comment to an entry, enter:

CYB, comment.

The maximum length of the comment (excluding CYB, and the period) is 35 characters. The system inserts the word COMMENT on the display before the comment (figure 8-11).

After completing edit entries, enter NEXT or END to transfer the data to the error log. If you enter NEXT, the CYBERLOG program transfers the data to the error log and restarts from the beginning to allow a different event to be recorded. If you enter END, the CYBERLOG program terminates after the message is written to the error log.

```
                                CYBERLOG

0  EVENT TYPE:  UNSCHEDULED
1  REASON:      HARDWARE
2  COMPONENT:   DISK SUBSYSTEM
3  ELAPSED TIME: 03.15
4  LOST TIME:   03.15
5  IMPACT:      DEGRADED
   COMMENT CYB, THIS IS A SAMPLE COMMENT LINE.

TO CHANGE ANY LINE, ENTER THE LINE NUMBER

TO WRITE CURRENT MESSAGE AND REPEAT FOR NEW
MESSAGE, ENTER "NEXT"

TO EXIT, ENTER END

TO WRITE A COMMENT, ENTER *CYB,* FOLLOWED BY NO
MORE THAN 35 CHARACTERS, TERMINATED BY A PERIOD.
```

Figure 8-11. Sample Completed CYBERLOG Entry

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 8-12 and 8-13).

```

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 8-12. 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.)
= A000000 (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 8-13. 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 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 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.

xtk Specifies one of the following types of flaw entries.

<u>Entry</u>	<u>Description</u>
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. ta can be one of the following.

<u>ta</u>	<u>Description</u>
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. Refer to table 8-1.
Aaddr	Physical block (track) address for extended memory. (Use with SPF or CPF.) addr is the extended memory address divided by 10g. Refer to table 8-1.
Aaddr- Aaddr	Range of physical block addresses for extended memory. (Use with SPF or CPF.) Refer to table 8-1.
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. Refer to table 8-1.

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 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.

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 8-14) 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	
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			

SLF=4001.

Figure 8-14. FLAW Utility K Display (Left Screen)

Table 8-1. Flawing Information for Disks and Extended Memory

Device Type	MNE	Logical Track Range †	Cylinders/ Device †	Tracks/ Cylinder †	Sectors/ Track †
844-2X, half track	DI	4000-7137	630	22	30
844-4X, half track	DJ	4000-7147	1464	23	30
844-2X, full track	DK	4000-7137	630	23	30
844-4X, full track	DL	4000-7147	1464	23	30
885, half track	DM	4000-7221	1511	50	40
885, full track	DQ	4000-7221	1511	50	40
834	DD	4000-7133	1456	12	40
819, single density	DV	4000-5465	633	12	24
819, double density	DW	4000-7153	1466	12	24
885-42	DB	4000-7221	1511	12	40
2048K, extended memory	DE/DP	4000-7620			

† 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) flow 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 8-15) appears on the left console screen. All parameters required to initialize and (if necessary) flow the specified device are entered through the K display.

```
      K. INITIALIZE.      CMS

      MASS STORAGE DEVICE INITIALIZATION

      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 8-15. INITIALIZE Command K Display

The K display lists the options used to initialize a device (under the OPTION column as shown in figure 8-15). 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 8-15).

Table 8-2. Device Initialization Options (Sheet 1 of 2) †

Option	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 778) 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.

† Device initialization options may be changed only if the total initialization level (AL) is specified in the INITIALIZE command.

Table 8-2. Device Initialization Options (Sheet 2 of 2) †

Option	Description
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).
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).
<p>† Device initialization options may be changed only if the total initialization level (AL) is specified in the INITIALIZE command.</p>	

Table 8-3. Track Flawing Options †

Option ††	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 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.
† Flawing of 881 and 883 disk packs is automatic; only flaws additional to the current flaw information in the utility sector should be entered. †† 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.

<u>Message</u>	<u>Description</u>
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.

<u>Message</u>	<u>Description</u>
ENTER OLD FAMILY NAME	<p>This message appears only if more than one family of permanent file devices is currently active in the system. Enter the following response.</p> <p>K.FM=familyname.</p>
ENTER OLD PACK NAME	<p>This message appears only if the device to be initialized is an auxiliary device. Enter the following response.</p> <p>K.PN=packname.</p>
ENTER OLD USER NAME	<p>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.</p> <p>K.UN=username.</p> <p>The user name is written to the account dayfile when the device is mounted.</p> <p>If you discover that the wrong device was specified in the INITIALIZE command, you can clear the initialize status for that device by entering:</p> <p>K.CLEAR.</p> <p>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.</p>

- When the following message appears on the K display, enter the parameters 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 parameters 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.

6. After all the necessary parameters 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.

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 8-16.




Figure 8-16. 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:†
 - 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.

†Under DIS control, the command MREC calls the MREC utility.

3. Activate the K display for that control point by entering:

K,jsn.

jsn MREC job sequence name.

The K display for MREC (figure 8-17) appears on the left console screen.

```
*** MACHINE RECOVERY UTILITY ***  
  
SHARED DEVICES MID = 42.  
  
EQ TYPE UN DN FM/PN MID(S) SHARING DEVICE  
5 DE 0 5 LINK  
12 DQ 40 40 NOSCLSH  
14 DQ 42 42 NOSCLSH  
16 DQ 44 44 NOSCLSH  
20 DQ 46 46 NOSCLSH  
21 DQ 47 47 NOSCLSH  
22 DJ 4 0 DEVCTR  
23 DJ 3 0  
24 DJ 2 0 THETA  
25 DJ 1 0 THEZOO  
26 DJ 0 0 KILLER  
27 DJ 5 0 SYSDUMP  
30 DJ 5 0 MSE  
31 DJ 3 0 SYSPACK  
32 DJ 1 0 TASPAC  
33 DJ 0 1 FILES  
34 DI 7 0  
35 DI 17 0  
36 DI 27 0  
37 DI 36 0  
  
ID = MID OF DOWNED MACHINE  
EQ = ALL EQ(S) TO RECOVER
```

Figure 8-17. MREC K Display (Left 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 mid(s)
```

<u>Header</u>	<u>Description</u>
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.
mid(s)	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.

- 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 (figure 8-18), enter:

```
KK.
```

- Enter the MREC parameters in the following format.

```
K.param1=value1,param2=value2,...,  
  paramn=valuen.
```

The parameters entered (and error messages, if any) are displayed on the lower portion of the MREC left screen K display.

- 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.

- 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.

- After all MREC operations are complete, end the MREC utility by entering:

```
K.STOP.
```

*** MACHINE RECOVERY UTILITY ***

DESCRIPTION OF TABLE ENTRIES.

EQ - EST ORDINAL OF EQUIPMENT.
UN - UNIT NUMBER.
DN - DEVICE NUMBER.
FM/PN - FAMILY OR PACK NAME (IF ANY)
MID(S) SHARING DEVICE - MACHINE ID(S) OF OTHER
MACHINES CURRENTLY SHARING DEVICE.
(* BY MID INDICATES MACHINE HAS
BEEN DETERMINED TO BE DOWN.)

DESCRIPTION OF PARAMETERS.

ID - MID OF DOWNED MACHINE TO PROCESS.
EQ - EST ORDINALS OF DEVICES SHARED BETWEEN THIS
MACHINE AND DOWNED MACHINE TO PROCESS. I.E.
EQ=XXX,XXX,...,XXX. DEFAULT = ALL. I.E. EQ=ALL.
OP=I AND OP=R ARE PARAMETERS WHICH MAY BE ENTERED IF
844 RESERVE SITUATIONS OCCUR. OP=I WILL CAUSE THE
DEVICE TO BE IGNORED. OP=R WILL CAUSE ALL UNIT
RESERVES ON THAT CONTROLLER TO BE CLEARED. THESE
PARAMETERS ARE ONLY VALID FOR THE CURRENT RESERVED
844 UNIT AND MUST BE REENTERED UPON REOCCURENCE OF
ANOTHER RESERVE SITUATION.

DESCRIPTION OF K DISPLAY COMMANDS.

GO - INITIATE PROCESSING OF DEVICES SPECIFIED.
RERUN - REINITIALIZE K DISPLAY AND RERUN PROGRAM.
STOP - TERMINATE PROGRAM.
+ - PAGE RIGHT DISPLAY IF .GT. 24D SHARED DEVICES.

Figure 8-18. MREC K Display (Right Screen)

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)
```

param₁=value₁ Parameters 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.

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.						
	<table border="0"> <thead> <tr> <th style="text-align: center;"><u>x</u></th> <th style="text-align: center;"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">R</td> <td>Directs the MREC utility to release all unit reservations (using the GRENADE function; refer to section 2, Deadstart) for 834, 844-41/44, or 885-11/12 equipment.</td> </tr> <tr> <td style="text-align: center;">I</td> <td>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.</td> </tr> </tbody> </table>	<u>x</u>	<u>Description</u>	R	Directs the MREC utility to release all unit reservations (using the GRENADE function; refer to section 2, Deadstart) for 834, 844-41/44, or 885-11/12 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.
<u>x</u>	<u>Description</u>						
R	Directs the MREC utility to release all unit reservations (using the GRENADE function; refer to section 2, Deadstart) for 834, 844-41/44, or 885-11/12 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.
+	Toggles the right screen K display between the K-display instructions and the second page of device descriptions if there is more than one page of shared devices.

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.

est EST ordinal of the device.

cc Channel number.

uu Physical unit number (from 0 to 77₈).

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, activate the deadstart switch 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, activate the deadstart switch 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 834, 844, or 885 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 multiframe 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 multiframe 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 ARCHIVAL SUBSYSTEM (MSAS) K DISPLAY

MSAS uses the K display to present messages that require your action. You can use the MSAS K display to reply to these messages and to open and close staging/destaging operations.

When MSAS enters messages into the K display, a request for the K display appears on the B,0 display. Enter K,MAS to activate the display for MSAS. After you have responded to all K-display messages, the request for the K display terminates.

The MSAS K display provides space for four messages with up to three lines per message. Figure 8-19 shows a sample MSAS K display.

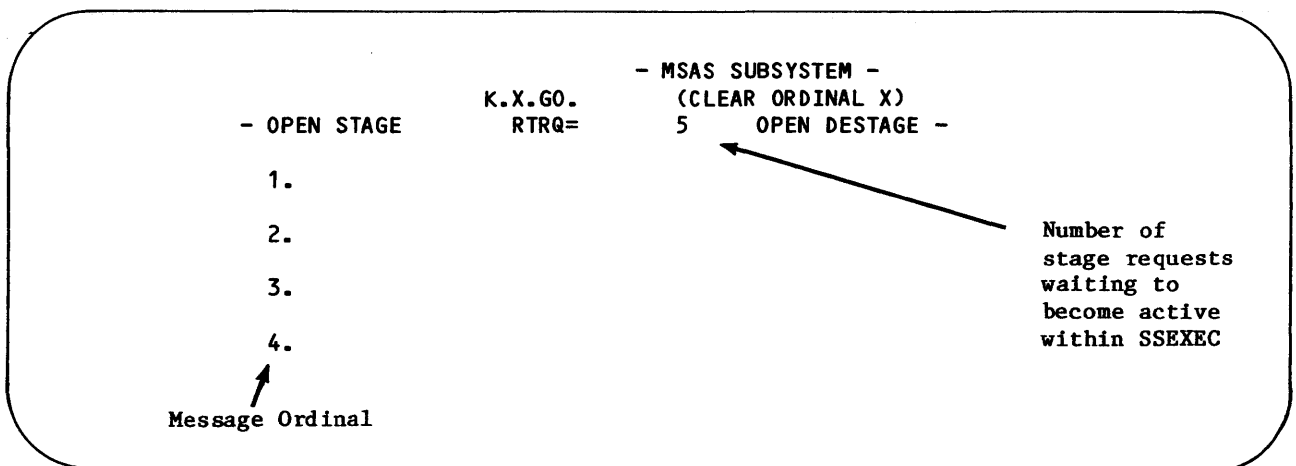


Figure 8-19. MSAS K Display

Valid MSAS commands are:

<u>Command</u>	<u>Description</u>
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).

MSAS 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 2 Operations Handbook for the appropriate action for each message.

MASS STORAGE SUBSYSTEM (MSS) K DISPLAY

MSS uses the K display to present messages that require your action. You can use the MSS K display to reply to these messages and to control the rate that files are staged/destaged between the MSF hardware and a disk.

When MSS enters messages into the K display, a request for the K display appears on the B,0 display. Enter K,MSS to activate the K display for MSS. After you have responded to all K-display messages, the request for the K display terminates.

The MSS K display provides space for four messages with up to three lines per message. If you enter an incorrect command, the incorrect command appears under a line containing *** REJECT ***. Figure 8-20 shows a sample MSS K display.

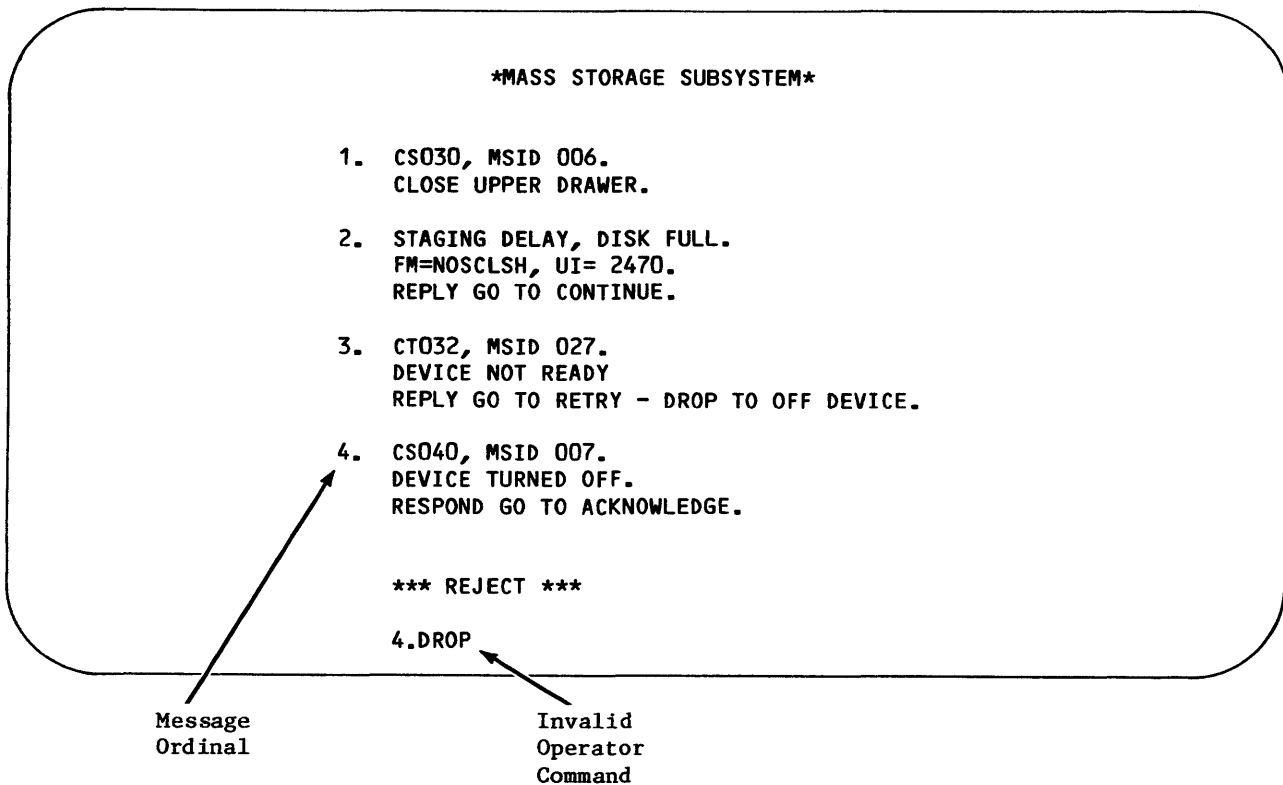


Figure 8-20. MSS K Display

Valid MSS commands are:

<u>Command</u>	<u>Description</u>
K.FILES,x.	Controls the rate that files are staged/destaged by changing the number, x, of concurrent files to be staged/destaged between the MSF and a disk. Increasing or decreasing the number of concurrent files by one results in a corresponding increase or decrease in field length of approximately 6500g central memory words. The number x must be a single digit.
K.m.GO.	Enters a GO response to the message at message ordinal m.
K.m.DROP.	Enters a DROP response to the message at message ordinal m.
K.	Clears the previously entered incorrect command and the *** REJECT *** line.

MSS clears a message when an acceptable action is taken. Usually this action is entering the K.m.GO or K.m.DROP command. However, for actions such as closing or emptying the input or output drawers on a cartridge storage unit, the message is cleared automatically when the hardware status indicates the action has occurred.

If you enter an invalid command, the *** REJECT *** line and the command are both displayed. Valid commands are those described previously. However, the K.m.GO and K.m.DROP commands are invalid if there is no message displayed at message ordinal m or if the GO or DROP response is not appropriate. Clear the *** REJECT *** line by entering a valid command. Refer to appendix B of the NOS 2 Operations Handbook for the appropriate action for each message.

NETWORK ACCESS METHOD (NAM) K DISPLAY

The NAM K display provides a common, centralized interface by which network applications (CS, NS, NVF) can communicate with the host operator (HOP).

Figure 8-21 shows the format of the NAM K display. Table 8-6 explains the fields that appear on the NAM K display.

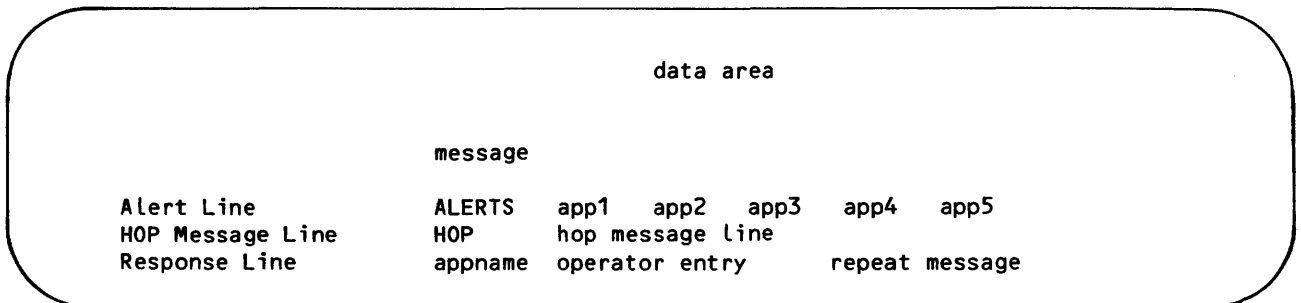


Figure 8-21. NAM K-Display Format

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, or NVF) 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 8-22 shows the NAM K display as it appears when you initially assign it.




Figure 8-22. NAM K Display

The NAM K display has two operating modes: NAM mode and 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.

<u>Character</u>	<u>Description</u>
*	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 8-21). 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, or NVF, 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 (Sheet 1 of 2)

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 8-21). appname is one of the supervisory programs CS, NS, or NVF. 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.

Table 8-7. NAM Mode Commands (Sheet 2 of 2)

Command	Description												
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=appname 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 8-23). Otherwise, if the K display is dropped, NAM automatically ends after 60 to 90 seconds.												
K.SEND,rrr,string.	<p>Causes NAM to send the message text if a logical link exists between the specified remote host and this host. Use the NAM ST command to determine which hosts are connected.</p> <table border="0" data-bbox="617 882 1494 1029"> <thead> <tr> <th data-bbox="617 882 747 913"><u>Parameter</u></th> <th data-bbox="1071 882 1234 913"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td data-bbox="617 934 673 966">rrr</td> <td data-bbox="779 934 1494 966">Host node number of the remote host (from 1 to 255).</td> </tr> <tr> <td data-bbox="617 987 714 1018">string</td> <td data-bbox="779 987 1364 1018">Message text of no more than 40 characters.</td> </tr> </tbody> </table> <p>The message</p> <p>hh.mm.ss. FROM sss string</p> <p>is displayed on the specified remote host's NAM K display following the alert line.</p> <table border="0" data-bbox="617 1228 1494 1375"> <tbody> <tr> <td data-bbox="617 1228 722 1260">hh.mm.ss</td> <td data-bbox="779 1228 1169 1260">Time the message is received.</td> </tr> <tr> <td data-bbox="617 1281 657 1312">sss</td> <td data-bbox="779 1281 1494 1312">Host node number of the sending host (from 1 to 255).</td> </tr> <tr> <td data-bbox="617 1333 706 1365">string</td> <td data-bbox="779 1333 1364 1365">Message text of no more than 40 characters.</td> </tr> </tbody> </table>	<u>Parameter</u>	<u>Description</u>	rrr	Host node number of the remote host (from 1 to 255).	string	Message text of no more than 40 characters.	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.
<u>Parameter</u>	<u>Description</u>												
rrr	Host node number of the remote host (from 1 to 255).												
string	Message text of no more than 40 characters.												
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 supervisory application or NAM and all applications using the following commands.

<u>Command</u>	<u>Description</u>
DB	Activates the online debugging code.
DE	Deactivates the online debugging code.
DU	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	Begins the logging of network traffic in the debug log file.
LE	Ends the logging of network traffic in the debug log file.
LR	Releases the debug log file.
RS	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 8-23).

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.

<u>mode</u>	<u>Description</u>
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.

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.

<u>mode</u>	<u>Description</u>
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.

<u>mode</u>	<u>Description</u>
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 MAXNIP build parameter described in the NOS 2 Installation Handbook.

The command format is:

FL=f1.

<u>Parameter</u>	<u>Description</u>
f1	Specified maximum field length in octal.

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/CCP Host Application Programming Reference Manual for information on creating the application's log file).

The command format is:

LB=mode.

<u>mode</u>	<u>Description</u>
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 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.

<u>mode</u>	<u>Description</u>
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.

<u>mode</u>	<u>Description</u>
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/CCP reference manuals 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.

<u>mode</u>	<u>Description</u>
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 8-23) 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	AAAC	000000		30		5	2	07.54.35												
RBF	AABC	000000		3		8	4	08.10.00												
TAF	AABY	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	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 8-23. 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
<u>Header</u>		<u>Description</u>						
app	Name of the application.							
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.							
	<u>Bit Number</u>	<u>Description</u>						
	0	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 following applications are also available as part of NAM.

<u>Application</u>	<u>Description</u>
PTF, PTFS	Permanent File Transfer Facility (PTF) and Permanent File Transfer Facility Sevicer (PTFS). PTF and PTFS provide access to remote permanent files. A local user activates PTF with the MFLINK command (refer to the NOS 2 Reference Set, Volume 3). When a remote user enters the MFLINK command, NAM activates a PTFS application on the local host to service the remote request.
QTF, QTFS	Queue File Transfer Facility (QTF) and Queue File Transfer Facility Servicer (QTFS). QTF and QTFS transfer input and output files to a remote system. When NAM is initiated, the system automatically activates QTF. When a remote host QTF application has a file to transfer, NAM initiates QTFS on the local host to service the remote request.

After initiation, QTF periodically rolls in to check for files to transfer. You can change the time interval between periodic executions of QTF. For information on changing this time interval, refer to the NOS 2 Installation Handbook.

The NAM status display shows the coupler status entries in the following format.

<u>Header</u>	<u>Description</u>
est	Equipment status table ordinal of the front end NPU.
hn	Host node number of the coupler.
nsm	Number of asynchronous supervisory messages currently queued for the coupler.
nhm	Reserved.
nln	Reserved.
ivtstat	Number of characters transferred downline on interactive connections in the last 30 seconds.
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.

<u>Header</u>	<u>Description</u>
hn	Host node number of the coupler.
tn	Terminal node number of the NPU on NPU-to-host 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	NPU regulation level as reported by the NPU. n can have the following values:

<u>n</u>	<u>Description</u>
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.

<u>Header</u>	<u>Description</u>
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 following procedure describes K-display usage for QTF operations under DSD control. QTF is used by both the RHF and NAM Subsystems.

1. When RHF or NAM is initiated, QTF is reinitiated automatically. If QTF is no longer running and must be restarted, either bring up the RHF application table display (described later in this section) and enter

ENABLE,ord.

ord QTF ordinal in the application table display.

or use the NAMI utility (described in section 16, Network Utilities).

2. Examine the DSD system status display (B,0). When QTF is scheduled to a control point, it is indicated on the B,0 display. The message

REQUEST *K* DISPLAY

appears on the B,0 display.

3. Activate the K display for that control point by entering:

K,jsn

jsn QTF job sequence name.

The QTF K display (figure 8-24) appears on the left screen.

* QUEUE FILE TRANSFER FACILITY *			
NUMBER OF FILES TRANSFERRED XXXX			
ACTIVE TRANSFERS			
FILE NAME	LID	PID	DC
ABCD	LD9	M90	LP
ABCG	LD1	M90	IN
ACCC	LD1	M90	PU
ADBA	LD9	M90	SB

Figure 8-24. QTF K Display

Each entry in the QTF K display has the following form.

filename lid pid dc

Header Description

filename Name of the file.

lid Logical identifier.

pid Physical identifier.

dc Disposition code.

PIDs must be defined in the LID configuration file. Also LIDs should be defined in the LID configuration file, but can be added or deleted using the LIDOU utility (refer to section 9, L-Display Utilities).

REDEFINE K DISPLAY

Use the REDEFINE utility to reconfigure 844 and 885-11/12† disk drives online and thereby logically eliminate a unit that is malfunctioning without performing a level 0 initial deadstart. A single unit or multiunit device †† 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.

† When reconfiguring an 885-11/12 disk drive, a customer engineer must be present.

†† When reconfiguring a multiunit device, all units of the device must be on the same controller(s).

Table 8-8. Equipment Requirements for Reconfiguration

Operation To Be Performed	Current Equipment	Replacement Equipment		Special Considerations
Add or Return a Unit	Must be unloaded, removable Yes	Must be in EST No	Must be unloaded, removable 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.

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.

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 REDEFINE K display (see figure 8-25) to the left console screen by entering:

K,jsn.

jsn Job sequence name of the job requesting the K display.

As shown in figure 8-25, 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 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 K display, and second, by the system responses and error messages.

**** MASS STORAGE FAILURE RECOVERY ****

K DISPLAY COMMANDS

CLEAR CLEARS IDLE AND SUSPEND ON CURRENT EQUIPMENT.
END CLEAR ALL DEVICE IDLES, AND END *CONFIG*.
GO INITIATES PROCESSING OF ENTERED CONFIGURATION.
RERUN RESTARTS *CONFIG* UTILITY PROCESSING.
RESET RESETS CURRENT EQUIPMENT TO DEFAULT PARAMETERS.
SUSPEND SUSPEND ALL SYSTEM OPERATION ON CURRENT EQUIPMENT.

PARAMETER DEFINITIONS

CH=C1,C2 C1 AND OPTIONALLY C2 ARE NEW CHANNELS FOR ACCESS.
EQ=EEE SET EST ORDINAL EEE TO BE PROCESSED.
UL=U1,U2 ..UN SET UNIT LIST AS SPECIFIED.
UR=UU UNIT UU IS TO BE RECABLED WITH A NEW DRIVE.

CURRENT EQUIPMENT CONFIGURATION.

034. DJ-1 CH05,07 UL=05 -- -- -- -- -- -- -- ST = I - - - -

REQUESTED EQUIPMENT CONFIGURATION.

034. DJ-1 CH05,07 UL=05 -- -- -- -- -- -- -- ST = I - - - -
 CHANGED UNITS

IDLED EQUIPMENTS

027. DJ-1 CH26,32 UL=07 -- -- -- -- -- -- -- ST = I - - - -

Figure 8-25. REDEFINE K Display (Left 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>
	<p style="text-align: center;">NOTE</p> <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 8-25).

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

REDEFINE,est.

est EST ordinal of the shared device.

at the console of each mainframe for which reconfiguration of the shared device is not desired.

2. Enter

K,jsn.

jsn Job sequence name of the job requesting the 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 8-26 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 8-26. Reconfiguration Run Output

Following are examples of the reconfiguration of disk devices and how their status changes in the REDEFINE K display.

Example 1, 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:

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
7	DJ-0	00	--	I---

Enter the following commands:

<u>Command</u>	<u>Description</u>
REDEFINE,7.	7 is the EST ordinal of the shared device.
K,jsn.	jsn is the job sequence name of the job requesting the 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 PROCESSING COMPLETE.

To end the reconfiguration, enter:

K.END

Disk configuration after the REDEFINE:

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
7	DJ-1	32, 26	7	----

Example 2, 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:

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
10	DJ-1	32	1	I---
11	DJ-1	13	2	----

To perform the reconfiguration, enter the following commands:

<u>Command</u>	<u>Description</u>
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 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 PROCESSING COMPLETE.

Enter the reconfiguration commands for the failing device as follows:

<u>Command</u>	<u>Description</u>
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.

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 PROCESSING COMPLETE.

To end the reconfiguration, enter:

K.END

Disk configuration after the REDEFINE:

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
10	DJ-0	00	--	----
11	DJ-1	32	1	----

Example 3, 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:

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
12	DI-1	26	3	----
13	DI-1	32	11	----

Enter the following commands:

<u>Command</u>	<u>Description</u>
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 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 PROCESSING COMPLETE.

Disk configuration at this point in example 3.

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
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:

<u>Command</u>	<u>Description</u>
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 PROCESSING COMPLETE.

To end the reconfiguration, enter:

K.END

Disk configuration after the REDEFINE:

<u>EST Ordinal</u>	<u>Equipment Type</u>	<u>Channel Number(s)</u>	<u>Unit Number</u>	<u>Status</u>
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 8-27).

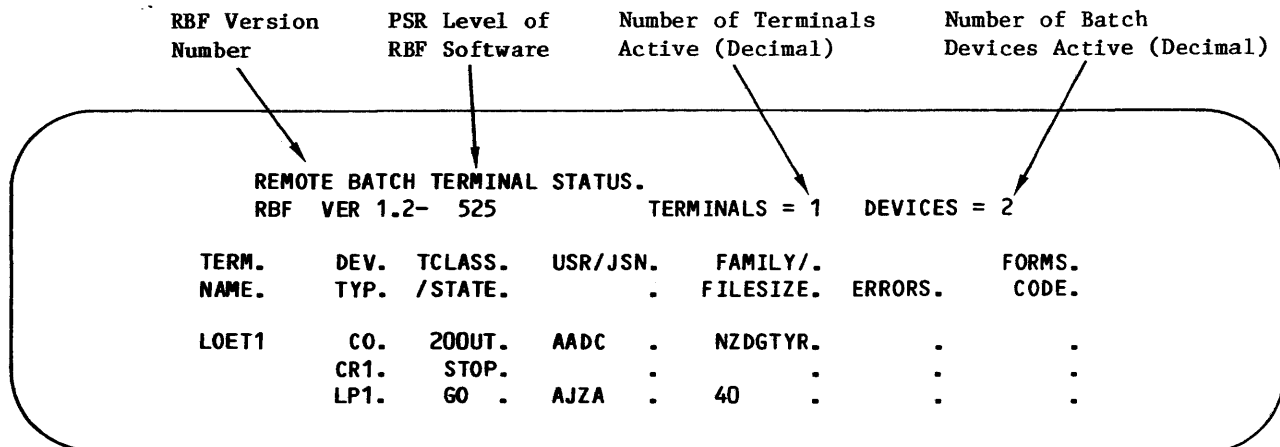


Figure 8-27. RBF K Display

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 of the terminal. A unique name, `termname`, assigned by the installation
name during network startup. This field is used for console devices only.

dev Device type. The device type code is one of the following:
typ

<u>Code</u>	<u>Description</u>
CO	Console device.
CPi	Card punch.
CRi	Card reader.
LPi	Line printer.
PLi	Plotter.

The ordinal number, `i`, of the device can range from 1 to 7.

tclass If this field is in the same line as the `termname`, the contents of the
/state 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:

<u>Code</u>	<u>Description</u>
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.

Header

Description

Code

Description

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.

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:

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:

<u>Application</u>	<u>Description</u>
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 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.
MLTF	Maintenance Logging Transfer Facility (MLTF). MLTF provides maintenance logging capabilities for local and remote mainframe NADs. All local NAD errors are logged to the binary maintenance log (BML). All remote mainframes that have error logging enabled (refer to the network path table display described later in this section) log errors. This application starts automatically when RHF is initiated.
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.

After initiation, both QTF and MLTF periodically roll in to check for files to transfer or to log NAD errors. You can change the time interval between periodic executions of either QTF or MLTF. For information on changing this time interval, refer to the NOS 2 Installation Handbook.

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, RHF generates jobs for loading NAD controlware for all local NADs that are defined in RHF's configuration and have an EST status of ON.

When RHF is initiated, it starts all enabled applications defined by the network configuration as autostart applications (refer to RHF - Configuration File Generation in the NOS 2 Installation Handbook). These applications are typically QTF and MLTF.

OPERATOR INTERFACE

The operator interface for RHF consists of using the following displays.

<u>Display</u>	<u>Description</u>
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 8-28.

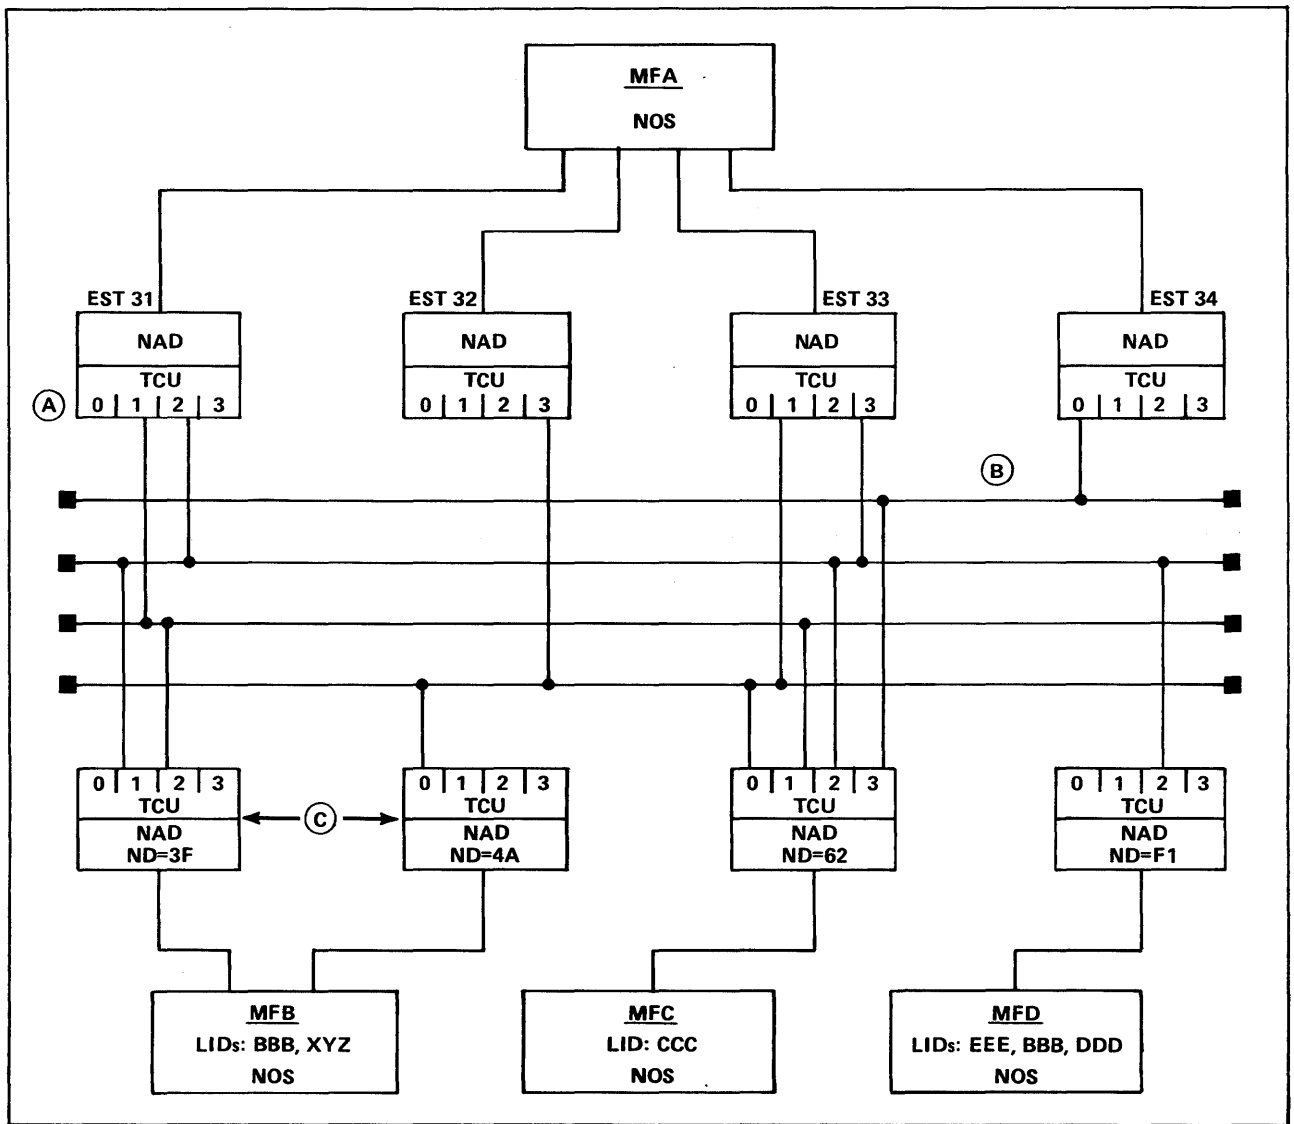


Figure 8-28. Sample LCN

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 8-28).

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 8-28).

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 8-28).

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 8-29) and the network path table display (see figure 8-30). 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:

<u>Command</u>	<u>Function</u>
APPL	Displays the application table (see figure 8-29).
IDLE	Begins the idle-down process of RHF and all its associated applications.
PATH	Displays the network path table (see figure 8-30).
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 8-29) lists all active applications.

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	MLTF	YES	1	0	0

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			

† USRAP represents an application written by the site.

Figure 8-29. Application Table Display

Each entry in the display appears in the following format.

```
ord  appl  enabled  maxcopies  activecopies  maxconnects
jobname  jobord  connects
```

<u>Header</u>	<u>Description</u>
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 (see figure 8-30) shows how the remote hosts are connected to the local host.

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

† 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 8-28.

Figure 8-30. 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 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.

<u>Header</u>	<u>Description</u>
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 MLTF 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.

<u>Parameter</u>	<u>Description</u>
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 MLTF periodically copy the trunk errors occurring on the specified remote NAD to the binary maintenance log (BML). MLTF 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 determines a NAD or a NAD trunk control unit (TCU) on a path is bad, it turns off the NAD or path and notifies you with a message on the B,0 display. It displays either

RHF,NAD ON EST est HAS BEEN TURNED OFF.

or

RHF,BAD TCU ON PATH xxx, PATH TURNED OFF.

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 MLTF 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.† 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,est command (refer to section 5, DSD Commands, for a description of the ON,est 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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
x	Specifies the SCOPE 2 mainframe PID.

†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 acknowledgements) 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.

<u>Parameter</u>	<u>Description</u>
pid	PID of the mainframe you want to disconnect (required only in a multi-mainframe 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 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:

<u>Command</u>	<u>Description</u>
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.

<u>Command</u>	<u>Description</u>
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 TAfffff command before issuing these commands.

INITIATING TAF 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.

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.

<u>Command</u>	<u>Description</u>
K.BFL=n.	Changes the starting and minimum field length allocated by TAF to CRM ($20000_8 \leq n \leq 100000_8$). Default is 70000_8 .
K.CMB=n.	Changes the maximum number of communication blocks allowed to the TAF Subsystem ($7 \leq n \leq 40$). Default is 40.
K.ECS=n.	Sets the extended memory field length to be used by the transaction executive; n is octal thousands of words. Default is 0.
K.EFL=n.	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 K display appears.

<u>Command</u>	<u>Description</u>										
K.ERO=CRF,op.	Specifies whether to override certain I/O and logic errors when processing the communication recovery files (CRF).										
	<table border="0" style="width: 100%;"> <thead> <tr> <th style="text-align: left;"><u>op</u></th> <th style="text-align: left;"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>NO</td> <td>Aborts if I/O or logic errors are encountered while processing the communication recovery files. This is the default setting.</td> </tr> <tr> <td>YES</td> <td>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.</td> </tr> <tr> <td></td> <td>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.</td> </tr> <tr> <td></td> <td>I/O or logic errors encountered on the CRF header record result in an unconditional abort of the TAF Subsystem.</td> </tr> </tbody> </table>	<u>op</u>	<u>Description</u>	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.
<u>op</u>	<u>Description</u>										
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.										
K.GO.	Ends input of the transaction executive initialization parameters. Initialization is complete when the TAF K display appears.										
K.INT=typ,fileid.	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.										
	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.										
	<table border="0" style="width: 100%;"> <thead> <tr> <th style="text-align: left;"><u>typ</u></th> <th style="text-align: left;"><u>Description</u></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>	<u>typ</u>	<u>Description</u>	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.				
<u>typ</u>	<u>Description</u>										
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.										

Command

Description

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 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.

NOTE

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.

- 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.

- K.MFL=n. Sets the maximum field length to be used by the transaction executive ($40000 \leq n \leq 376600$). Default is 376600.
- 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 ($2 \leq n \leq 31$). 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 ($10000g \leq n \leq 100000g$). Default is 30000g. 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 automatically restarts TAF by transferring control to the TAF automatic recovery program. The TAF automatic recovery program recovers the central memory pointers and variables defined during TAF initiation.

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 initial K display, except that values specified in the TAF initiation replace any default values that were in the initial display.

TAF K DISPLAY

When the TAF Subsystem is executing, the TAF K display 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 TAF K display appears on the system console as shown in figure 8-31.

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.




Figure 8-31. TAF K Display

The default parameter values for the CMDUMP and DSDUMP commands are given on the TAF K display shown in figure 8-31.

<u>Parameter</u>	<u>Description</u>
DB	This parameter is not used. It is retained for compatibility with previous releases of NOS.
EP	Exchange package: 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	Output queue: BC Local batch. RB Remote batch. PF Permanent file.
QD	Queue destination: User name (if OQ=BC). Equipment identifier (if OQ=RB). Permanent file name (if OQ=PF).

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 Reference Manual for additional information on the KPOINT request and the system task library.)

<u>Command</u>	<u>Description</u>
K.ASSIGN,est.	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.
or	
K.ASSIGN,est,db,n.	

<u>Command</u>	<u>Description</u>
	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*.
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 K Display, earlier in this section, for explanations of the parameters.
K.DUMP,fwa,lwa.	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.

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 Reference Manual).

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.

- 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.
- 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.
- 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.

<u>Command</u>	<u>Description</u>
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 K display shown in figure 8-31. 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 376600g 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>
K.NODEBUG.	<p>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.</p>

<u>Command</u>	<u>Description</u>
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,nnnnn.	Changes the amount of time that TAF will retain its field length between communication input messages; nnnnn is specified in units of milliseconds. Refer to the installation parameter ITRTL in the NOS 2 Installation Handbook.
K.SWITCH.	Causes the console K display to change to a display listing all allowable console commands. When K.SWITCH is entered a second time, the normal 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, U=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 CRM status K display enter:

K.DIS,CRMTASK.

After the K display is assigned to the task, the display in figure 8-32 is brought to the left screen.

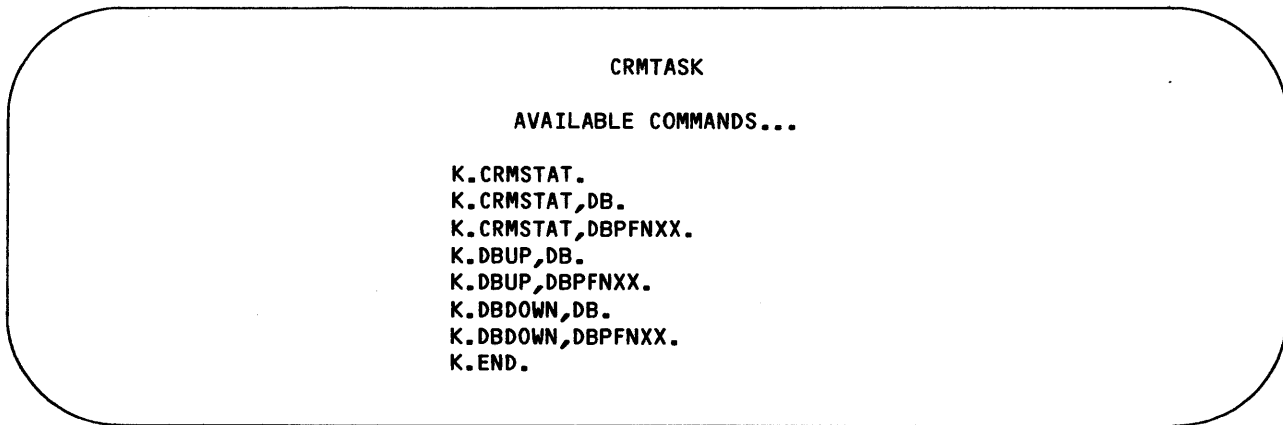


Figure 8-32. CRMTASK K Display

The following commands are available under CRMTASK.

<u>Command</u>	<u>Description</u>
K.CRMSTAT.	Selects the CRM status K display (see figure 8-33). This display shows the status of all CRM data bases.
K.CRMSTAT,db.	Selects the CRM data base status K display (see figure 8-34). 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 (see figure 8-35). 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.

<u>Command</u>	<u>Description</u>
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.

```

** CRM STATUS **

nnn TRANSACTIONS IN INPUT QUEUE.
nn ACTIVE TRANSACTIONS.
nnn TRANSACTIONS IN OUTPUT QUEUE.

CRM DATA BASES

DB STATUS      DB STATUS      DB STATUS
AA  UP         BB  IDEL         CC  DOWN
DD  DOWN       EE  DOWN         FF  UP
GG  UP

VALID COMMANDS ARE -

CRMSTAT. CRMSTAT,DB. CRMSTAT,DBPFN. END. +. -.
DBUP,DB.  DBUP,DBPFN. DBDOWN,DB.  DBDOWN,DBPFN.

```

Figure 8-33. CRM Status K Display

**** CRM DATA BASE STATUS ****

DATA BASE = db DATA BASE STATUS = UP
AFTER IMAGE FILE = arfname PRU-S REMAINING = nnn

BEFORE IMAGE FILES AND STATUSES

ZZAAB01-U ZZAAB02-U ZZAAB03-U ZZAAB04-D ZZAAB05-U
ZZAAB06-U ZZAAB07-D

FILE	STATUS	FILE	STATUS	FILE	STATUS
AAPFN01	UP	AAPFN02	DOWN	AAPFN03	DOWN
AAPFN04	IDLE	AAPFN05	UP	AAPFN06	UP
AAPFN07	UP				

VALID COMMANDS ARE -

CRMSTAT. CRMSTAT,DB. CRMSTAT,DBPFN. END. +. -.
DBUP,DB. DBUP,DBPFN. DBDOWN,DB. DBDOWN,DBPFN.

Figure 8-34. CRM Data Base Status K Display

**** CRM FILE STATUS ****

FILE NAME = AAPFN01
FILE STATUS = UP
RECOVERABLE = YES

PACK NAME = PACKNAM
DEVICE TYPE = DJ3
ATTACH MODE = RM

SIZE OF PRIMARY KEY = 80
NUMBER OF ALTERNATE KEYS = 3

ACTIVE USERS = 4
ACTIVE LOCKS = 8

VALID COMMANDS ARE -

CRMSTAT. CRMSTAT,DB. CRMSTAT,DBPFN. END.
DBUP,DB. DBUP,DBPFN. DBDOWN,DB. DBDOWN,DBPFN.

Figure 8-35. CRM File Status K Display

For the CRM status and CRM data base status displays, if all the information does not fit on one screen, you can bring up additional pages by entering the following command.

K.+.

To return to the first page of the display, enter the following command.

K.-.

At any time you can return to the initial CRMTASK K display (figure 8-32) by entering the following command.

K.MENU.

To end CRM status K-display processing, enter the following command.

K.END.

FOTD L Display	9-2
LIDOU L Display	9-3
QDSPLAY L Display	9-6
SCTD L Display	9-8
SUBSYST L Display	9-10

L-DISPLAY UTILITIES

9

This section documents the following L displays and the utilities used to present them.

<u>Utility</u>	<u>Description</u>
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.
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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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,L0=option.

or

FOTD,outfile,option.

<u>Parameter</u>	<u>Description</u>
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. 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.

FAMILY ORDINAL TABLE DISPLAY			
DEFAULT FAMILY = FAMILY1		AVAILABLE ENTRIES = 638.	
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 first line of the FOTD L display contains the default family name and the octal number of FOT entries that remain available for assignment. The rest of the display 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

LID,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). From this display you can select the HELP display shown in figure 9-3. You can toggle between these two L displays as described at the bottom of each 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.

<u>Parameter</u>	<u>Description</u>
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.

LID DISPLAY					
<u>PID</u>	<u>LID</u>	<u>ENABLED</u>	<u>ATTRIBUTE</u>	<u>MFTYPE</u>	<u>NETWORK</u>
MFB	---	YES	HV---	NOS V2	RHF - IN SERVICE
	MFB	YES	-----		NAM - OUT OF SERVICE
	BBB	YES	--S--		SSF - DISABLED
	XYZ	YES	--S--		

MFC	---	YES	L----	NOS V2	RHF - IN SERVICE
	MFC	YES	-----		NAM - IN SERVICE
	CCC	YES	--S-N		

MFD	---	YES	LV--N	NOS V2	NAM - DISABLED
	MFD	YES	-----		SSF - IN SERVICE
	DDD	NO	--S--		
	EEE	NO	--S--		
	BBB	YES	--S--		

(ENTER L.HELP TO TOGGLE TO THE HELP DISPLAY.)

Figure 9-2. LIDOU L Display (LID)

The information shown in figure 9-2 corresponds to the sample LCN shown in figure 8-28.

HELP DISPLAY

COMMAND	DESCRIPTION
L.END	TERMINATE LIDOU
L.BACK	FORMAT LID DISPLAY
L.HELP	FORMAT HELP DISPLAY
L.OUT	ROUTE A LISTING OF THE LID TABLE
L.+	PAGE *L* DISPLAY FORWARD
L.-	PAGE *L* DISPLAY BACKWARD
L.PID,PA=XXZ	SET (ALL APPLICABLE) PID ATTRIBUTES
L.PID,NT=YYY,NA=Z	ENABLE/DISABLE NETWORK STATUS
L.PID,MF=AAAAAAA	SET MAINFRAME DESCRIPTION
L.PID,LD=BBB,LA=CCCCZ	ADD OR SET ATTRIBUTES FOR LID
L.PID,LD=BBB,DELETE	DELETE LID

PID IS THE PHYSICAL IDENTIFIER, BBB IS AN ASSOCIATED LID.
X CAN BE V - VALIDATE, N - NOLIST.
Z CAN BE E - ENABLE, D - DISABLE.
YYY CAN BE RHF, NAM, SSF, OR ALL.
AAAAAAA IS A (1 TO 7 CHARACTER) MAINFRAME DESCRIPTOR.
C CAN BE S-STORE AND FORWARD, B-LOOPBACK, V-VALIDATE, N-NOLIST.

ATTRIBUTES

S - STORE AND FORWARD (LIDS FOR LINKED PIDS ONLY)
B - LOOPBACK (LIDS FOR HOST PID ONLY)
V - VALIDATE
D - DISABLE
N - NOLIST
E - ENABLE

(ENTER L.BACK TO TOGGLE TO THE LID DISPLAY).

Figure 9-3. LIDOU L Display (HELP)

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=aaaaaaa,NT=yyy,NA=z.
```

QDSPLAY L DISPLAY

The DSD utility QDSPLAY displays the content of a queued file listed in the queued file table (QFT).

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.

<u>Parameter</u>	<u>Description</u>
jsn	Job sequence name of the queued file you want to examine.

Figure 9-4 shows the initial QDSPLAY L display.

```
QDSPLAY. JSN = jsn. QFT = 0013. STATUS =  
THE FOLLOWING ARE VALID QUEUE DISPLAY COMMANDS -  
  
COMMAND                FUNCTION  
  
L.HELP.                LIST L DISPLAY COMMANDS FOR QDSPLAY UTILITY.  
L.DROP.                DROP FILE FROM QUEUE AND TERMINATE.  
L.END.                 RETURN FILE TO QUEUE AND TERMINATE.  
L.LINE.                SET FILE DISPLAY MODE TO LINE FORMAT.  
L.OCTAL.               SET FILE DISPLAY MODE TO OCTAL FORMAT.  
L.SS.                  DISPLAY IN OCTAL FROM BEGINNING OF SYSTEM SECTOR.  
L.*.                   RETURN FROM HELP/SS TO PREVIOUS FILE DISPLAY.  
L.+                     ADVANCE DISPLAY FORWARD (CIRCULAR FOR SS).  
L.BOI.                 POSITION FILE TO BOI AND DISPLAY DATA.  
L.DAY.                 POSITION PRINT FILE TO DAYFILE AND DISPLAY DATA.  
L.SL.                  SKIP 1 LINE FORWARD IN FILE AND DISPLAY DATA.  
L.SL,N.                SKIP N LINES FORWARD IN FILE AND DISPLAY DATA.  
L.SP.                  SKIP 1 PAGE FORWARD IN FILE AND DISPLAY DATA.  
L.SP,N.                SKIP N PAGES FORWARD IN FILE AND DISPLAY DATA.  
L.SR.                  SKIP 1 RECORD FORWARD IN FILE AND DISPLAY DATA.  
L.SR,N.                SKIP N RECORDS FORWARD IN FILE AND DISPLAY DATA.  
L.SR,*                 SKIP TO EOI IN FILE.  
L.BR.                  SKIP 1 RECORD BACKWARD IN FILE AND DISPLAY DATA.  
L.BR,N.                SKIP N RECORDS BACKWARD IN FILE AND DISPLAY DATA.  
  
SKIP COMMANDS ARE NOT ALLOWED UNDER HELP/SS DISPLAY.  
SKIP LINE/PAGE COMMANDS ARE NOT ALLOWED FROM OCTAL DISPLAY.
```

Figure 9-4. Initial QDSPLAY L Display

Valid QDSPLAY commands are:

<u>Command</u>	<u>Description</u>
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).
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 l 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.

<u>Command</u>	<u>Description</u>
L.SS.	Displays the data from the beginning of the system sector in octal display code format.
L.*.	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.+.	Advances the display forward to the next screen of data. This command is ignored under the L.HELP command where all information is presented on one screen. For the L.SS command, after the last screen of data, the first screen of the system sector is displayed again.

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,L0=option.

or

SCTD,outfile,option.

<u>Parameter</u>	<u>Description</u>
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. <ul style="list-style-type: none"> 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 9-4 shows the SCTD L display.

```

SERVICE CLASS CONTROL TABLE DISPLAY.

VALIDATED SERVICE CLASSES FOR EACH ORIGIN TYPE.

SYSTEM ORIGIN SERVICE CLASS(ES).
SY BC RB TS DI NS SS MA CT IO
I1 I2 I3

LOCAL BATCH ORIGIN SERVICE CLASS(ES).
BC

REMOTE BATCH ORIGIN SERVICE CLASS(ES).
RB

INTERACTIVE ORIGIN SERVICE CLASS(ES).
TS DI

PRIORITY ORDERED SERVICE CLASSES.
PO P1 P2 P3 P4 P5 P6 P7
TS I1 I2 IO DI I3
```

Figure 9-4. SCTD L Display

The SCTD L display shows the various origin types along with the service classes valid for each origin type.

At the bottom of the screen, 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.

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,L0=option.

or

SUBSYST,outfile,option.

<u>Parameter</u>	<u>Description</u>
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. <ul style="list-style-type: none">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 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
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 the LIDs Using the LIDOU Display	10-7
RHF Configuration Files	10-8
Network Configuration Statements	10-9
LNDR Statement	10-9
APPL Statement	10-10
NPID Statement	10-11
PATH Statement	10-11
LNAD Statement	10-12
RNAD Statement	10-13
DEBUG Statement	10-14
CHARGE Statement	10-14
Examples	10-15

This section contains information about LID configuration files and RHF configuration files.

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 3777778. 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.

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
```

id is a 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. For 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.

<u>Parameter</u>	<u>Description</u>
PID=pid	Physical identifier; pid is the 3-character physical identifier for the mainframe. 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).

<u>at</u>	<u>Description</u>
VALID	Specifies that USER command prevalidation is required.
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.

The default attributes are: USER command prevalidation is not required and PID is listable using 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).
------------	--

<u>sub</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
LID=lid	Logical identifier; lid is the 3-character logical identifier for the mainframe identified by the last NPID statement. lid can 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).

<u>at</u>	<u>Description</u>
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.
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.

Example

Figure 10-1 shows an equipment configuration consisting of six mainframes. Following the figure are the LIDCMid files for the three NOS mainframes in the configuration: LIDCM64, LIDCM42, and LIDCM05.

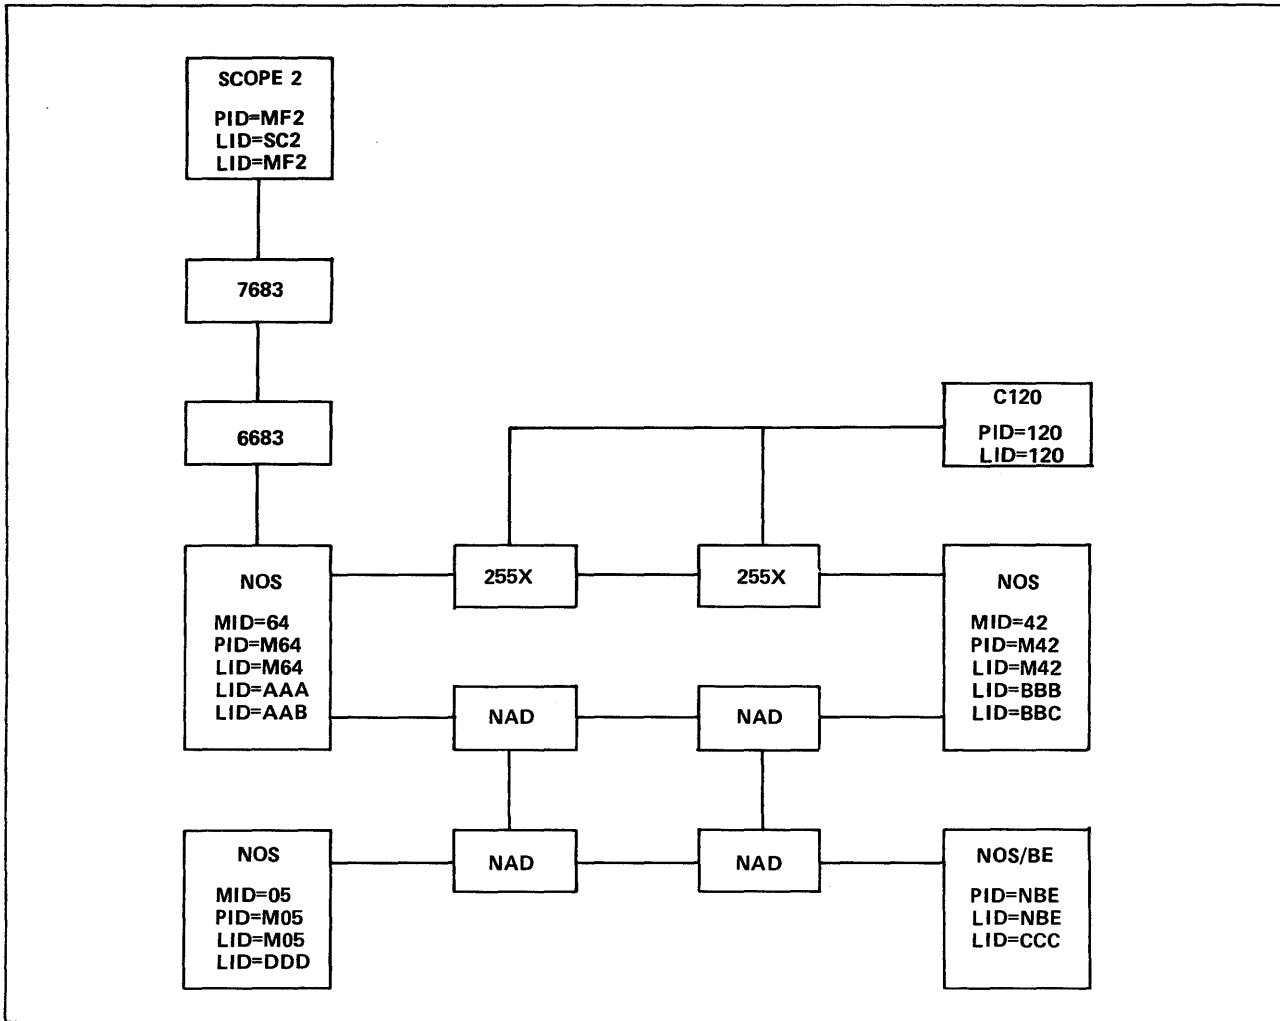


Figure 10-1. Equipment Configuration

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=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=LOOPB.
 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.

(Continued)

(Continued)

```
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=LOOPB.
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.
```

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.
```

nnnn is the number of central memory words allocated to the LDT (default=1, minimum=0, maximum=11008). You can calculate the value for nnnn by using this formula:

```
(3+lid)*pid
```

lid is the maximum number of LIDs allowed per PID.

pid is the 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.

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=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=LOOPB.
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.
```

(Continued)

(Continued)

```
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=LOOPB.
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.
```

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.
```

nnnn is the number of central memory words allocated to the LDT (default=1, minimum=0, maximum=11008). You can calculate the value for nnnn by using this formula:

```
(3+lid)*pid
```

lid is the maximum number of LIDs allowed per PID.
pid is the 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,LID=lid,PID=pid,L=list.
```

<u>Parameter</u>	<u>Description</u>
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 THE 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.

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.

The RCFGEN utility reads network configuration statements to create a direct access permanent file under user name SYSTEMX. RHF then uses the permanent file as the network description and can properly access the network.

When RHF is started, it expects to find a local configuration file named RCFILE. If it does not find this file, RHF will attach the RHF configuration file using file name RCFMid, where id is the machine identifier.

The following command calls the RCFGEN utility.

```
RCFGEN,I=input,L=list,O=output.
```

The parameters are order independent and optional; if omitted, the defaults are used.

<u>Parameter</u>	<u>Description</u>
I=input	Specifies the local file from which the network configuration statements are to be read. The default is file INPUT.
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.

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.

<u>Statement</u>	<u>Description</u>
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.

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:
 - NPID statement.
 - 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

<u>Parameter</u>	<u>Description</u>
------------------	--------------------

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 $\text{mxcopys} * (5 + 3 * \text{mxcons})$ central memory words (mxcopys and mxcons are defined below). To define an RHF application, enter:

```
APPL  NAME=name,ENABLED=status,MXCONS=mxcons,MXCOPYS=mxcopys,SVR=status,ASTART=status
```

<u>Parameter</u>	<u>Description</u>
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=mxcopys	Maximum number of simultaneously active copies of this application that are allowed. The maximum value for mxcopys is 127. Default is 1.
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.

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, ITF, and MLTF. The maximum simultaneous connections should be set to one (default) for QTFS, PTF, PTFS, and MLTF. The maximum number of simultaneously active copies should be set to one (default) for QTF, ITF, and MLTF. 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, ITF, or MLTF. ASTART=YES should be specified for both QTF and MLTF.

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.

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
```

<u>Parameter</u>	<u>Description</u>
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
```

<u>Parameter</u>	<u>Description</u>
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 two words in RHF's field length. To define a local NAD, enter:

```
laddr LNAD CH=ch,MAXNDRS=n,DEDICATE=status,CMPATHS=nn,CMBUFFS=nn.
```

<u>Parameter</u>	<u>Description</u>
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.

The software relies on the NAD to determine the actual number of paths and buffers reserved for code and conversion. The algorithm used by the NAM is $\text{Buffers Reserved} = (\text{the smaller of paths requested}) \text{ or } (\text{buffers reserved} * 2/3)$.

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

	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 6 buffers	5 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 one word in RHF's field length. To define a remote NAD, enter:

```
raddr RNAD ND=nn,DD=d,LOG=status
```

<u>Parameter</u>	<u>Description</u>
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 MLTF in the mainframe's error log. status may be YES or NO. Default is NO.

DEBUG Statement

The DEBUG statement controls the manner in which RHF uses queue entries.

DEBUG TRACE=status

<u>Parameter</u>	<u>Description</u>
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

<u>Parameter</u>	<u>Description</u>						
TYPE=type	Type of call for which the charge is being specified. <table><thead><tr><th><u>type</u></th><th><u>Description</u></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>	<u>type</u>	<u>Description</u>	1	Requires small amount of RHF processing time.	2	Requires large amount of RHF processing time.
<u>type</u>	<u>Description</u>						
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.

Examples

Figure 10-2 illustrates a sample LCN network and is the basis of the network configuration statements in examples 1 and 2.

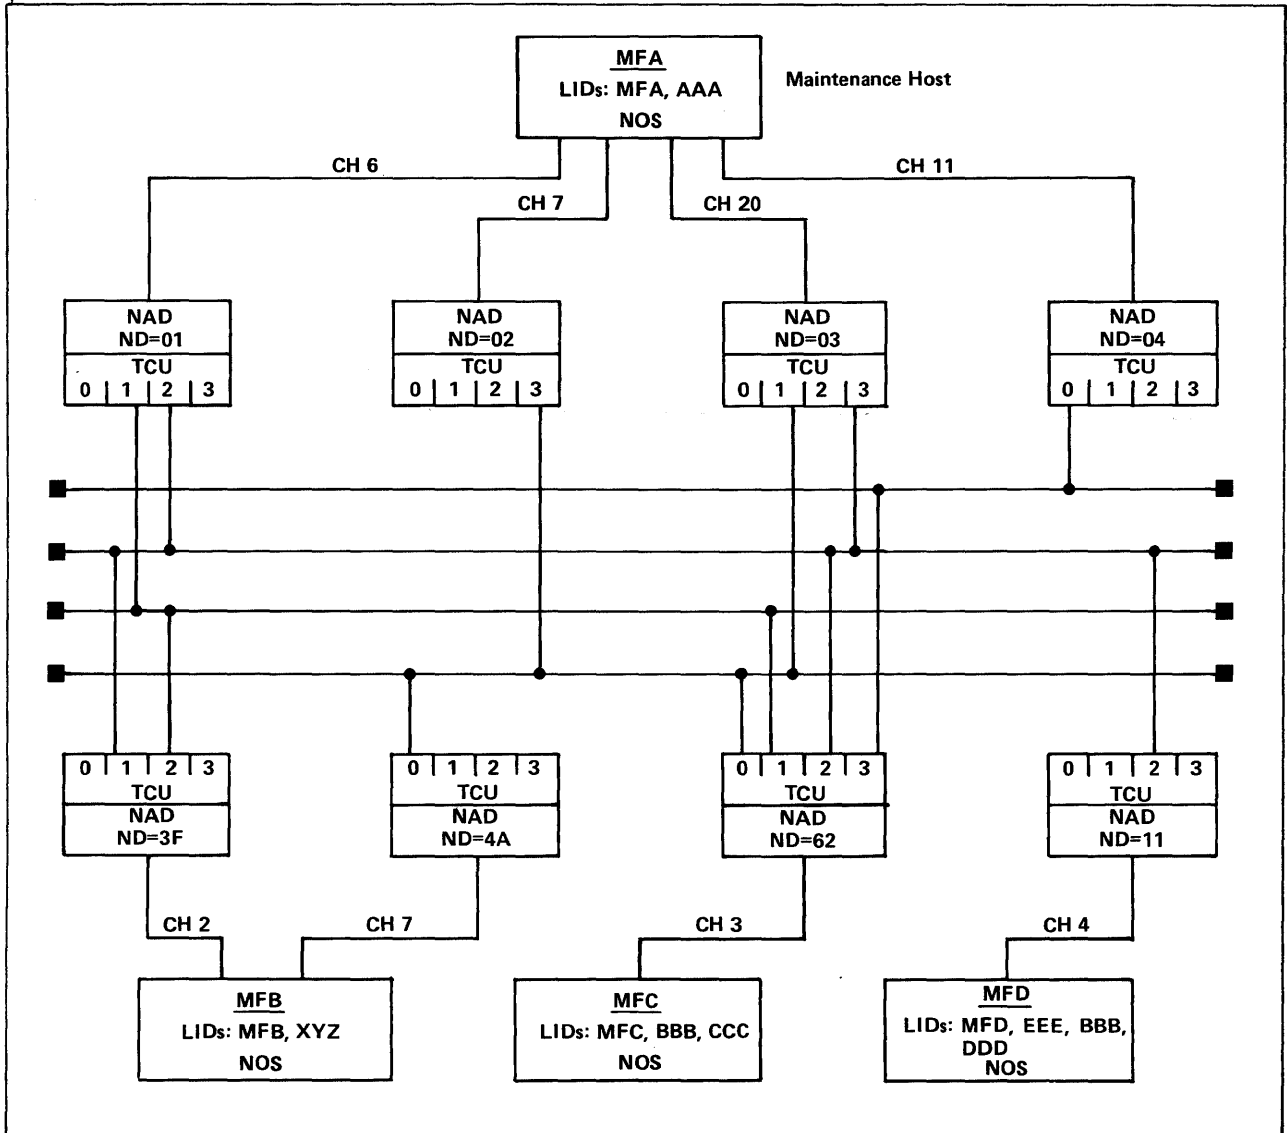


Figure 10-2. Sample LCN Network

Example 1.

Following are the network configuration statements for the RHF on mainframe MFA in figure 10-2.

```
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=MLTF,ASTART=YES
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=FOFO
PATH    ENABLED=YES,LT=0001,RT=1000,RNAD=RN1,LNAD=LN2,AC=FOFO
PATH    ENABLED=NO,LT=0001,RT=0100,RNAD=RN2,LNAD=LN3,AC=FOFO
PATH    ENABLED=YES,LT=0100,RT=1000,RNAD=RN1,LNAD=LN3,AC=FOFO

NPID    PID=MFC,MFTYPE=NOS2
PATH    ENABLED=NO,LT=0110,RT=0110,RNAD=RN3,LNAD=LN1,AC=FOFO
PATH    ENABLED=NO,LT=0001,RT=1000,RNAD=RN3,LNAD=LN2,AC=FOFO
PATH    LT=0101,RT=1010,RNAD=RN3,LNAD=LN3,AC=FOFO
PATH    LT=1000,RT=0001,RNAD=RN3,LNAD=LN4,AC=FOFO

NPID    PID=MFD,MFTYPE=NOS2
PATH    LT=0010,RT=0010,RNAD=RN4,LNAD=LN1,AC=FOFO
PATH    LT=0001,RT=0010,RNAD=RN4,LNAD=LN3,AC=FOFO

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
```

Example 2.

Following are the network configuration statements for the RHF on mainframe MFD in figure 10-2.

```
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=MLTF,ASTART=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=FOFO
PATH  RNAD=RN03,LNAD=LN1,LT=0010,RT=0001,AC=FOFO

NPID  PID=MFB,MFTYPE=NOS2
PATH  RNAD=RN3F,LNAD=LN1,LT=0010,RT=0100,AC=FOFO

NPID  PID=MFC,MFTYPE=NOS2
PATH  RNAD=RN62,LNAD=LN1,LT=0010,RT=0010,AC=FOFO

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
```


Introduction	11-1
Utility Description Terminology	11-2
Cartridge	11-2
SM Map	11-2
7990 Catalog	11-4
7990 Files	11-10
SSEXEC	11-10
Subfamily	11-11
Group	11-11
Volume	11-11
MSAS Utilities	11-11
SSDEF	11-12
SSLABEL	11-13
SSLABEL Directives	11-14
Parameters for SSLABEL Directives	11-17
SSLABEL Update Sequence	11-21
Restrictions to SSLABEL	11-21
Example	11-22
SSMOVE	11-22
SSMOVE Directives	11-24
Excluding Destage/Release Processing	11-28
Selection Algorithms	11-29
Disk Space/Dump Tape Management	11-31
Restrictions to SSMOVE	11-31
Output Example	11-31
SSVAL	11-37
Release Processing	11-38
Problem Reporting	11-39
Error Detection and Classification	11-40
Release Processing and Problem Fixing	11-42
Validation Report	11-42
Typical SSVAl Runs	11-44
Restrictions to SSVAl	11-46
SSUSE	11-46
Basic Usage Report	11-48
Optional Report A	11-50
Optional Report B	11-51
Optional Report C	11-52
Optional Report D	11-53
SSDEBUG	11-53
SSDEBUG Directives	11-54
Parameters for the SSDEBUG Directives	11-55
Restrictions to SSDEBUG	11-58
y,z Relationship With SM Map Ordinal	11-58
Format of Disk Files Read From 7990 by SSDEBUG	11-59

(Continued)

(Continued)

SSBLD	11-62
SSBLD Statements	11-62
7990 Equipment Configuration Constraints	11-64
SSALTER	11-65
Console Input	11-65
MSAS Operational Procedures	11-68
Initialization	11-68
Disk Space Management	11-68
7990 Space Management	11-69
Backup and Recovery	11-70
File Dumping	11-70
File Reloading	11-73
Error Conditions and Corrective Actions	11-73
Clearing Permanent File Error Flags	11-74
Permanent File Recovery	11-74
Cartridge Restoration and Reuse	11-75
Removal of Faulty or Missing Cartridges	11-75
7990 Catalog/SM Map/PFC Problems	11-76

INTRODUCTION

The Mass Storage Archival Subsystem (MSAS) 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 conventional 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. You can use both the 7990 and magnetic tapes to protect files from hardware and system failures.

The 7990 hardware is composed of the following components.

<u>Component</u>	<u>Description</u>
Cartridge	A plastic housing that encloses magnetic tape on which data is stored under program control. This cartridge is compatible with an IBM 3850 cartridge.
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 MSAS software supports up to two IBM channel interfaces per CYBER mainframe.

The basic function of MSAS 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 MSAS utilities are introduced. Detailed information about the call and operation of each MSAS utility, listed next, is included in this section.

<u>Utility</u>	<u>Description</u>
SSDEF	Creates the system files (SM maps and 7990 catalogs) necessary for MSAS 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.

<u>Utility</u>	<u>Description</u>
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 MSAS 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 MSAS 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 MSAS 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	
ℓ	sub-family	id	unused	
cm		csn		

<u>Field</u>	<u>Description</u>																
code	Number from 0 to 6 indicating how the cubicle is assigned.																
	<table border="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>code</u></th> <th style="text-align: left;"><u>Description</u></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.†</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>	<u>code</u>	<u>Description</u>	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.†	5	Assigned to a subfamily.	6	No cubicle exists at these coordinates.
<u>code</u>	<u>Description</u>																
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.†																
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.																
ℓ	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.																

†The pool is an area of the SM that stores scratch cartridges managed by the SSLABEL utility.

<u>Field</u>	<u>Description</u>						
cm	Cartridge manufacturer code.						
	<table border="0"> <thead> <tr> <th style="text-align: left;"><u>cm</u></th> <th style="text-align: left;"><u>Description</u></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>	<u>cm</u>	<u>Description</u>	A-	Cartridge manufacturer is IBM.	B-	Reserved for future cartridge manufacturer.
<u>cm</u>	<u>Description</u>						
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.

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*ig*, where *i* is the subfamily identifier (a number from 0 to 7). For example, file SFMCAT3 and user index 377763*g* 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:

59	17	11	2	0
familyname		sub-family	/ / / / / / / /	id
unused				
unused				

<u>Field</u>	<u>Description</u>
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:

59	53	47	41	35	26	17	11	0
id	length	FCT loc		AST loc		unused		
unused			date time					
AUs small		AUs large		AUs group		unused		

<u>Field</u>	<u>Description</u>
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.

flags One of the following flags.

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.

<u>Field</u>	<u>Description</u>												
	<table border="1"> <thead> <tr> <th><u>Bit</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <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>	<u>Bit</u>	<u>Description</u>	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.
<u>Bit</u>	<u>Description</u>												
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.												
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.												

<u>Field</u>	<u>Description</u>
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.
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.

<u>Bit(s)</u>	<u>Description</u>
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.

<u>Field</u>	<u>Description</u>
<u>Bit(s)</u>	<u>Description</u>
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 MSAS 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 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 MSAS activities. The MSAS 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 MSAS system files, the SM maps are created on the master device under user index 3777608 (subfamily 0) and one 7990 catalog is created under each user index 37776i8 (subfamily i), i=0,1,2,...,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 MSAS 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).

MSAS UTILITIES

MSAS utilities are described next. All of these utilities must be run from system origin jobs. In a multiframe 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 MSAS 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, P_1 , P_2 .

<u>P_1</u>	<u>Description</u>
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 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.

<u>P_i</u>	<u>Description</u>
I=filename	File containing the directives to SSLABEL.
I	Same as I=INPUT.
I omitted	Same as I=INPUT.

<u>P_i</u>	<u>Description</u>
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,...,P_n.

Example:

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:

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 MSAS 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 the exit tray (PT=D), or it removes 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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>						
CC=loc	Location into which the cartridge is to be stored.						
	<table border="0" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th style="text-align: left;"><u>loc</u></th> <th style="text-align: left;"><u>Description</u></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>	<u>loc</u>	<u>Description</u>	A	Store cartridge into location (0,0).	B	Store cartridge into location (0,15).
<u>loc</u>	<u>Description</u>						
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.						
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.						

<u>P_i</u>	<u>Description</u>
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.

<u>pkloc</u>	<u>Description</u>
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.
PT=ptloc	Location into which the cartridge or cubicle is to be put.

<u>ptloc</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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; $0 \leq y_1 \leq 21$, $0 \leq z_1 \leq 15$; 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).
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 MSAS system files, relabel the cartridge, and add the new assignment information to the MSAS 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.

Example

The following output 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.***
```

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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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.

<u>xxx</u>	<u>Description</u>
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.

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.

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.

<u>Parameter</u>	<u>Description</u>
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.

<u>process</u>	<u>Description</u>
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 2-character permanent filename beginning with Z from user index 002622g 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).

<u>directive</u>	<u>Description</u>
	Enter a 2-character value specifier directive keyword.
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.

<u>qualifier(s)</u>	<u>Description</u>
	Enter 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.
IA	Keyword for indirect access files.
DA	Keyword for direct access files.
DS	Keyword for destage decisions.
RL	Keyword for release decisions.

<u>parameter(s)</u>	<u>Description</u>
	Enter a 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.

BR Directive
Parameters

N=wf	wf is the weight factor to be used as the backup requirement value for decisions involving files with a BR=N attribute (no backup).†
Y=wf	wf is the weight factor to be used as the backup requirement value for decisions involving files with a BR=Y attribute (tape backup).†
MD=wf	wf is the 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).†

†The file owner specifies the backup requirement attribute using the BR parameter on the DEFINE or CHANGE command (refer to the NOS 2 Reference Set, Volume 3).

<u>parameter(s)</u>	<u>Description</u>
<u>FR Directive Parameters</u>	
MN=min	min is the minimum file size in PRUs.
MX=max	max is the maximum file size in PRUs.
DD=dd	dd is the number of days since the file was last accessed or updated.
TH=th	th is the 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.
<u>PR Directive Parameters</u>	
L=wf	wf is the weight factor to be used as the preferred residence value for decisions involving files with a PR=L attribute (locked to disk preference).†
D=wf	wf is the weight factor to be used as the preferred residence value for decisions involving files with a PR=D attribute (disk preference).†
M=wf	wf is the weight factor to be used as the preferred residence value for decisions involving files with a PR=M attribute (7990 preference).†
N=wf	wf is the weight factor to be used as the preferred residence value for decisions involving files with a PR=N attribute (no preference).†
<u>SM Directive Parameters</u>	
MG=mg	mg is the master goal specifying the percent of master device disk space not to be exceeded.
SG=sg	sg is the secondary goal specifying the percent of secondary device disk space not to be exceeded.
P1	Available for site use.
P2	Available for site use.
<u>WA Directive Parameters</u>	
AG=wa	wa is the weight adder to be added to the weighted file age.
LN=wa	wa is the weight adder to be added to the weighted file length.
AC=wa	wa is the weight adder to be added to the weighted file access count.
DV=sf	sf is the scaling factor given to the divisor when calculating destage and release values.

† The file owner specifies the preferred residence attribute using the PR parameter on the DEFINE or CHANGE command (refer to the NOS 2 Reference Set, Volume 3).

parameter(s)

Description

WM Directive
Parameters

AG=wf	wf is the weight factor given to the file age (days since the file was last accessed or updated).
LN=wf	wf is the weight factor given to the file length.
AC=wf	wf is the 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).

$$\text{value} = \frac{(a_1 + a_2 * \text{length}) * (b_1 + b_2 * \text{age}) * (\text{res}) * (\text{bkup})}{(c_1 + c_2 * \text{access}) * \text{sf}}$$

<u>Variable</u>	<u>Description</u>
a ₁	Site-defined length weight adder specified by the LN parameter of the WA directive.
a ₂	Site-defined length weight factor specified by the LN parameter of the WM directive.
length	Length of the file in PRUs.
b ₁	Site-defined age weight adder specified by the AG parameter of the WA directive.
b ₂	Site-defined age weight factor specified by the AG parameter of the WM directive.
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.

<u>Variable</u>	<u>Description</u>
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.
c1	Site-defined access count weight adder specified by the AC parameter of the WA directive.
c2	Site-defined access count weight factor specified by the AC parameter of the WM directive.
access	Number of times the file was accessed.
sf	Site-defined scaling factor specified by the DV parameter of the WA directive.

If a file's destage/release value exceeds a site-defined threshold (TH parameter of the FR directive), 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

Only one copy of SSMOVE per family can be run at a time. A second SSMOVE aborts if the first one has not completed.

SSMOVE and PFDUMP should not be run at the same time. The concurrent execution of PFDUMP and SSMOVE to destage files may result in errors in the PFC entries.

Output Example

The SSMOVE report consists of the following five sections.

- A list of directives from the input directive file plus any diagnostic messages.
- A list of run-time parameter values 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.
- An optional list of the files selected for processing. This list shows the file name, file type, user index, length, date, access count, processing action, and destage and release values.
- A device status report 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.
- A destage abandonment report 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.

SSMOVE REPORT.

SSMOVE,I=BIGO,LO=F,FM=MTST.

- 1 FR,MN=0,DD=0.
- 2 SM,MG=0,SG=0.

(Continued)

(Continued)

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

(Continued)

(Continued)

yy/mm/dd. hh.mm.ss. PAGE 4

SSMOVE REPORT.

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

(Continued)

(Continued)

yy/mm/dd. hh.mm.ss. PAGE 6

SSMOVE REPORT.

		(BEFORE)	DEVICE STATUS	(AFTER)	PERCENTS					
EO	DN	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	DIR.	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			
	NUMBER	DIRECT	PRU	INDIRECT	NUMBER	DIRECT	PRU	INDIRECT	NUMBER	DIRECT	PRU	INDIRECT
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	27919	69	56298
7	0	0	0	0	0	0	0	0	0	0	0	0

(Continued)

(Continued)

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
10	0	CARTRIDGE LOST
11	0	CLOSED DESTAGE

(BEFORE) DEVICE STATUS (AFTER) PERCENTS

EO	DN	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	DIR.	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				
	NUMBER	DIRECT	PRU	INDIRECT	PRU	NUMBER	DIRECT	PRU	INDIRECT	PRU	NUMBER	DIRECT	PRU	INDIRECT
0	0	0	0	0	0	38	6363	20	12190	38	6363	20	12190	
1	0	0	0	0	0	37	6352	20	12190	37	6352	20	12190	
2	0	0	0	0	0	37	10357	2	1201	37	10357	2	1201	
3	0	0	0	0	0	42	12362	5	3010	42	12362	5	3010	
4	0	0	0	0	0	37	6357	5	3010	37	6357	5	3010	
5	0	0	0	0	0	37	6357	5	3010	37	6357	5	3010	
6	0	0	0	0	0	19	27915	69	56298	19	27919	69	56298	
7	0	0	0	0	0	0	0	0	0	0	0	0	0	

SSVAL

SSVAL either performs release processing or reports on problems with the current MSAS 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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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 MSAS 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 MSAS 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 MSAS 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.

<u>Chain</u>	<u>Description</u>
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 MSAS 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

The following output 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

FAMILY = SYSTST

SSVAL,FM=SYSTST,SB=0,SM=A.

L = OUTPUT
RF = 0
AM = 0
SM = A
FM = SYSTST
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

REPORT COMPLETE

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 MSAS system files and/or PFC entries. The following examples show typical SSVAl runs that may be used periodically for these purposes.

Example:

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=DUMPl,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 MSAS system files and PFC entries. This is accomplished using the command:

```
SSVAL,FM=familyname,AM.
```

familyname is that on the release data file, DUMPl.

Example:

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,...,nh,AM,FX=n.
```

familyname and n₁,n₂,...,n_h 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=377760g), 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 MSAS processing, these periodic SSVAl runs can be done less frequently.

Restrictions to SSVAl

- 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:

<u>Report</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>										
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.										
OP=op	Type of report to be generated. Multiple options can be specified (for example, OP=AB).										
	<table border="0"> <thead> <tr> <th style="text-align: center;"><u>op</u></th> <th style="text-align: center;"><u>Description</u></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>	<u>op</u>	<u>Description</u>	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.
<u>op</u>	<u>Description</u>										
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.

The following output 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
 ** = SUBFAMILY TOTAL I = INHIBIT ALLOCATION FU = FLAWED AND UNALLOCATED
 - = UNASSIGNED GROUP F = FREE CARTRIDGE SF = START OF FRAGMENT
 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	A	1	2	2	1167	2662	0	0	0	2	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
1	A	-	3	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
1	A	**	5	2	1168	2662	0	0	0	2	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
2	A	1	2	2	1168	2606	0	0	0	2	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
3	A	-	3	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
3	A	**	5	2	1168	2586	0	0	0	2	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
4	A	1	2	2	1168	2662	0	0	0	2	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
5	A	-	3	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
5	A	**	5	2	1168	2662	0	0	0	2	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
6	A	1	3	3	5400	7	0	2	0	1	0	0	0	0	0	0	3	0
6	A	**	5	3	5400	7	0	2	0	1	0	0	0	0	0	0	3	0
7	A	-	3	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
7	A	**	5	2	1200	2662	0	0	0	2	0	0	0	0	0	0	0	0

(Continued)

(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	3	0
SUBFAMILY TOTALS FOR SPECIFIED SM-S																	
0	**	**	4	2	1167	2662	0	0	0	2	0	0	0	0	0	0	0
1	**	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0
2	**	**	5	2	1168	2606	0	0	0	2	0	0	0	0	0	0	0
3	**	**	5	2	1168	2586	0	0	0	2	0	0	0	0	0	0	0
4	**	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0
5	**	**	5	2	1168	2662	0	0	0	2	0	0	0	0	0	0	0
6	**	**	5	3	5400	7	0	2	0	1	0	0	0	0	0	3	0
7	**	**	5	2	1200	2662	0	0	0	2	0	0	0	0	0	0	0

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.

The following output 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      S = RESERVED FOR SYSTEM USE      C = RESERVED FOR CUSTOMER ENGINEERING
        F = NOT ASSIGNED TO THIS FAMILY    E = ERROR FLAG                  N = (Y,Z) DOES NOT EXIST IN SMMAP
        R = RESERVED FOR ALTERNATE SMMAP  GPORD = ORDINAL IN GROUP

  Y   Z   CM  CSN   FAMILY   SUBFAMILY   GROUP   GPORD   FLAGS
  0   0
  1   0
  2   0
  3   0
  4   0
  5   0
  6   0
  7   0
  8   0
  9   0
 10   0
 11   0
 12   0
 13   0
 14   0
 15   0
 16   0
 17   0
 18   0
 19   0  A-51680502  MTST      0         1         1
 20   0  A-13727600  MTST      0         1         0
 21   0

```

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.

The following output 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
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

```

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.

The following output 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
                FIRST  FIRST  DIVISION  OCL
                SMALL  LARGE  POINT
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

```

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).

The following output 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 0040611131 0040601131 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

```

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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>																		
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.																		
	<table border="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>flag</u></th> <th style="text-align: left;"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>AU conflict flag.</td> </tr> <tr> <td>EW</td> <td>Excessive write parity error flag.</td> </tr> <tr> <td>FC</td> <td>Frozen chain flag.</td> </tr> <tr> <td>FE</td> <td>Linkage error flag (in 7990 catalog FCT).</td> </tr> <tr> <td>IB</td> <td>Inhibit allocation flag.</td> </tr> <tr> <td>LC</td> <td>Lost cartridge flag.</td> </tr> <tr> <td>ME</td> <td>Linkage error flag (in SM map).</td> </tr> <tr> <td>SF</td> <td>Start of fragment flag.</td> </tr> </tbody> </table>	<u>flag</u>	<u>Description</u>	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.
<u>flag</u>	<u>Description</u>																		
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.																		
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=MMMBUG.																		
PF omitted	Same as PF=MMMBUG.																		
SB=subfamily	Subfamily to be used; $0 \leq \text{subfamily} \leq 7$.																		

<u>P_i</u>	<u>Description</u>
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 ≤ i ≤ 1931; i ≤ 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.
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.

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 (y,z) 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.

$$y = \frac{(352-\text{ord})}{16}$$

$$z = 352-\text{ord}-16y$$

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 11-1 shows the format of a disk file created when SSDEBUG reads a full stripe; figure 11-2 shows the format of a disk file created when SSDEBUG reads a partial stripe.

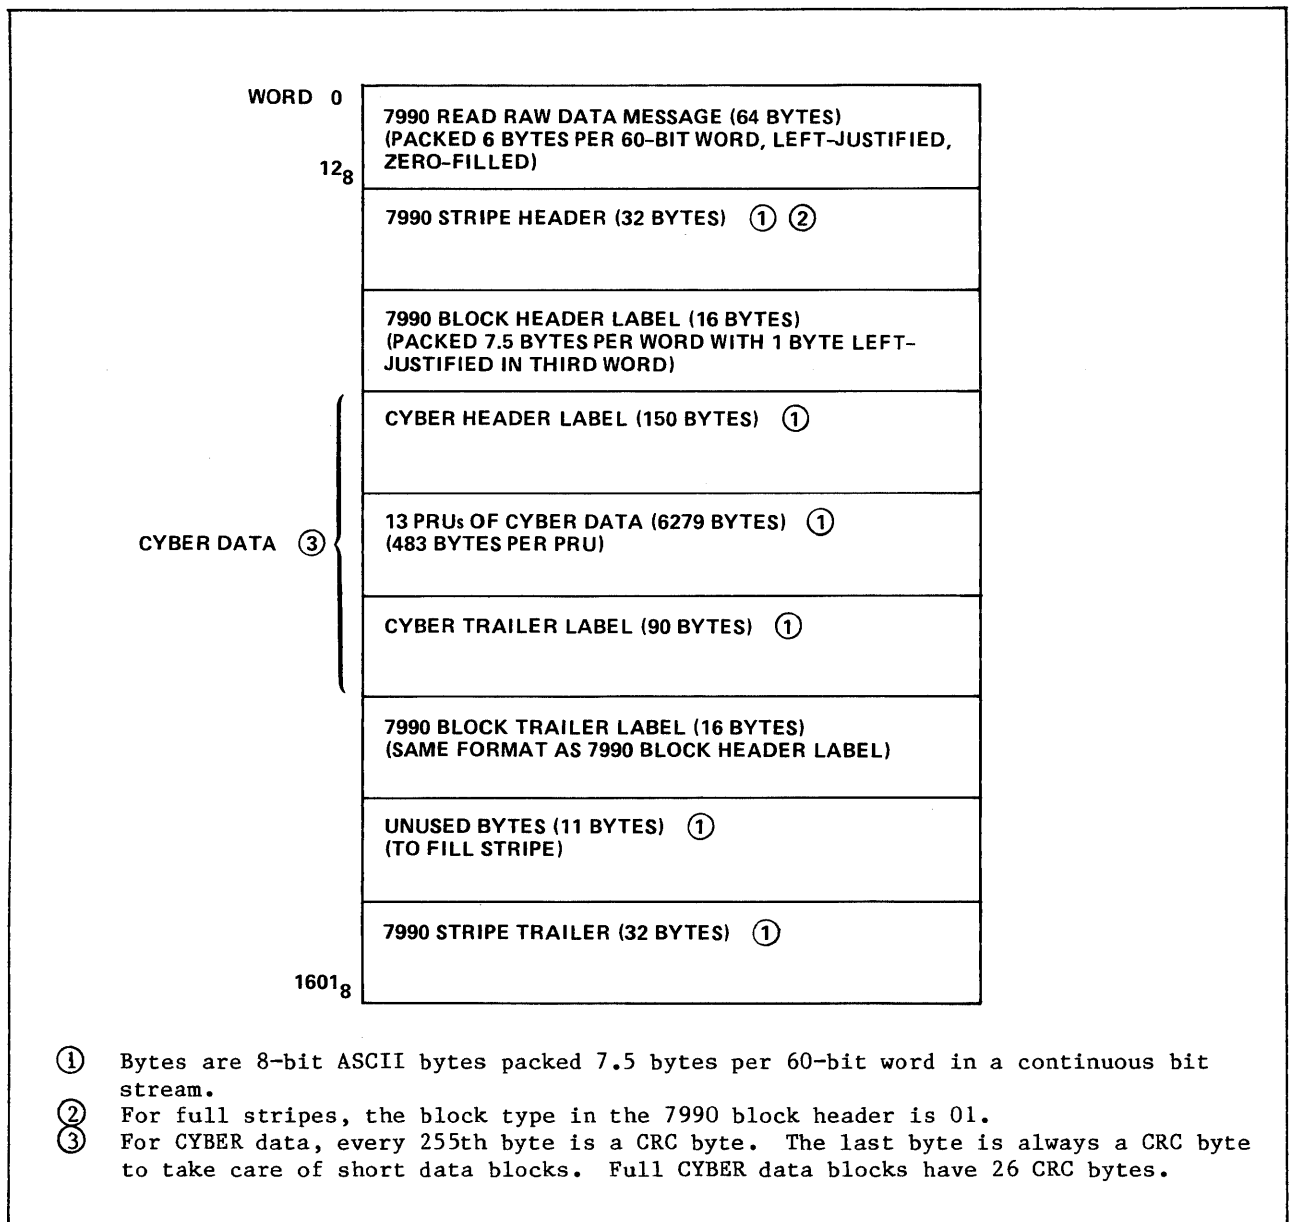
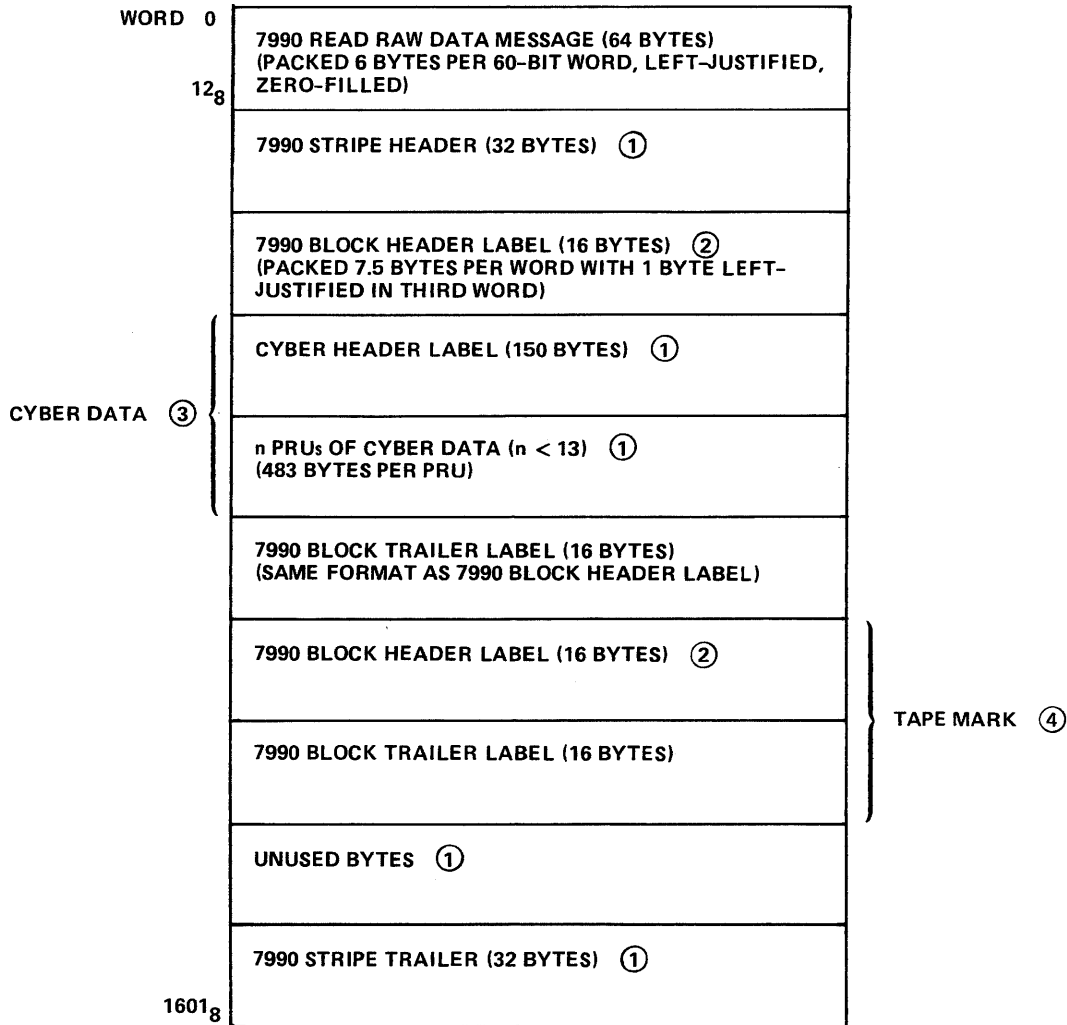


Figure 11-1. Full Stripe Disk File Format



- ① Bytes are 8-bit ASCII bytes packed 7.5 bytes per 60-bit word in a continuous bit stream.
- ② For partial stripes, the block type in the 7990 block header is 01 for data blocks and 10₁₆ for tape mark blocks.
- ③ For CYBER data, every 255th byte is a CRC byte. The last byte is always a CRC byte to take care of short data blocks.
- ④ A tape mark has a 7990 block header and trailer, but no data block.

Figure 11-2. Partial Stripe Disk File Format

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 11-1 shows the format of a disk file created when SSDEBUG reads a full stripe; figure 11-2 shows the format of a disk file created when SSDEBUG reads a partial stripe.

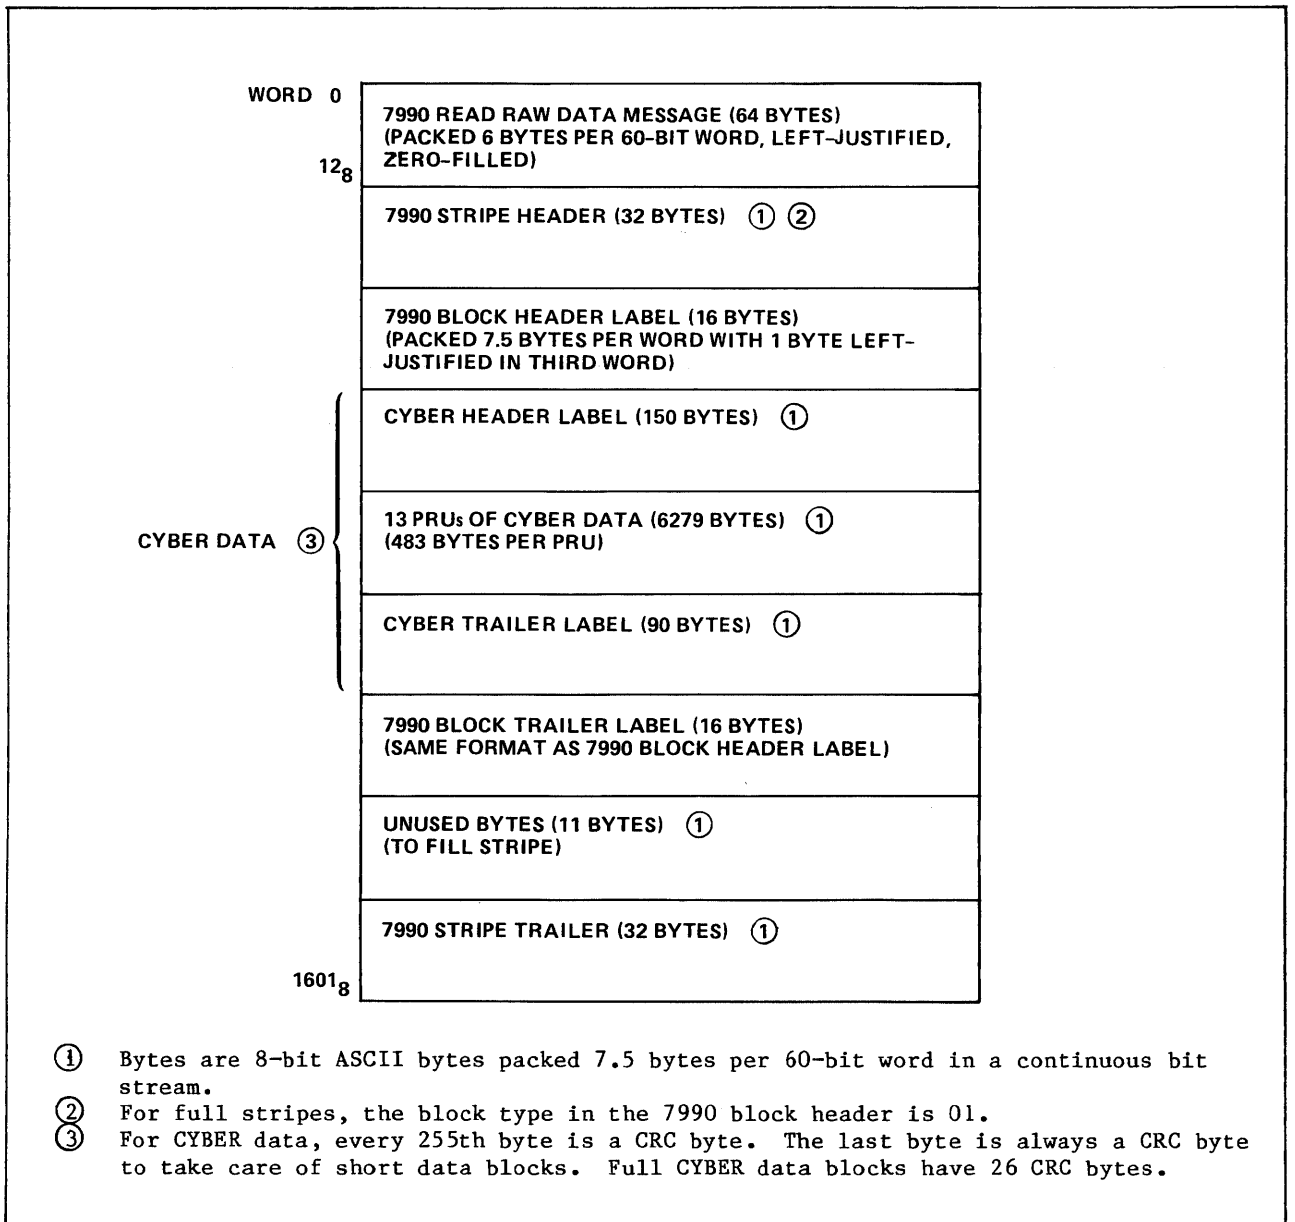
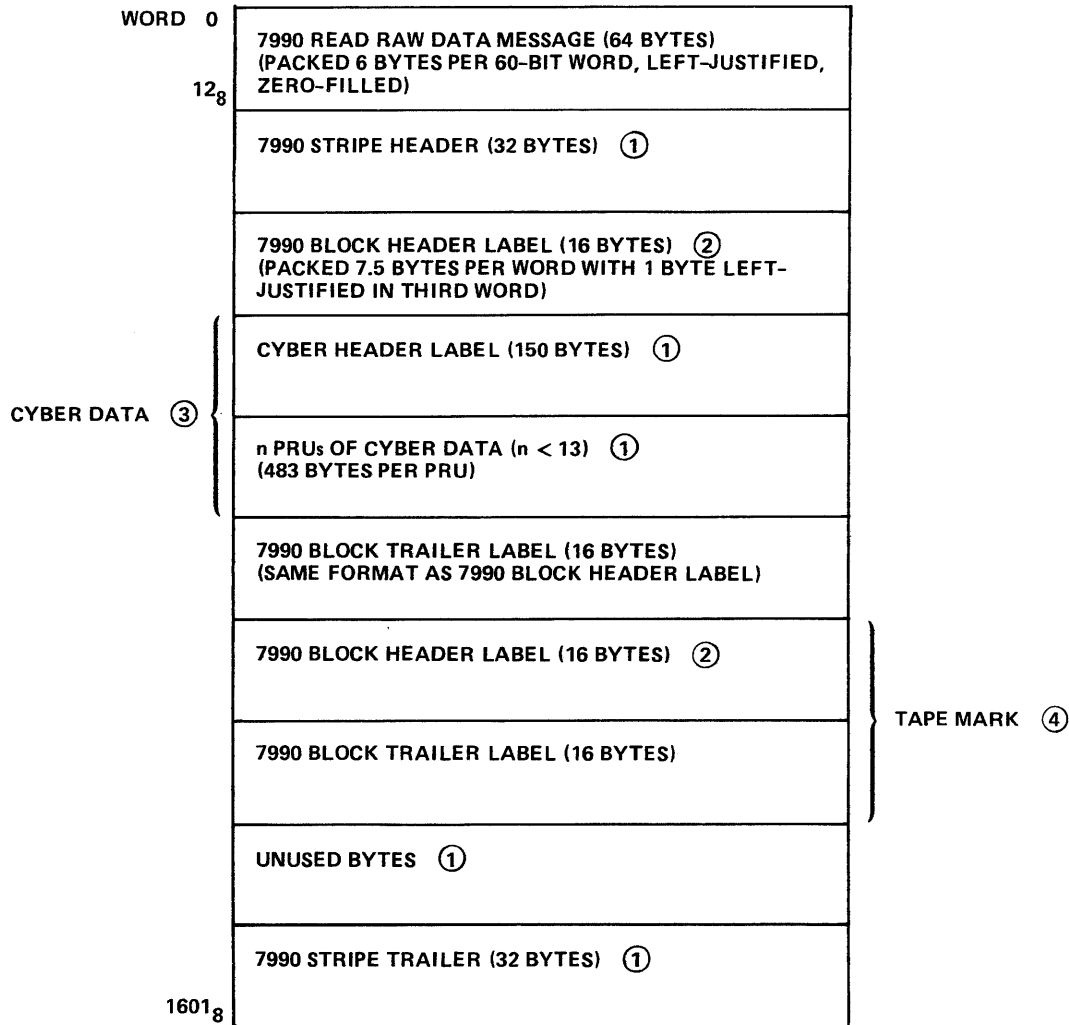


Figure 11-1. Full Stripe Disk File Format



- ① Bytes are 8-bit ASCII bytes packed 7.5 bytes per 60-bit word in a continuous bit stream.
- ② For partial stripes, the block type in the 7990 block header is 01 for data blocks and 10₁₆ for tape mark blocks.
- ③ For CYBER data, every 255th byte is a CRC byte. The last byte is always a CRC byte to take care of short data blocks.
- ④ A tape mark has a 7990 block header and trailer, but no data block.

Figure 11-2. 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.

<u>Bytes</u>	<u>Description</u>
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.

<u>Byte(s)</u>	<u>Description</u>												
0-1	Block type. The following block types are valid.												
	<table><thead><tr><th><u>Block type</u></th><th><u>Description</u></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>	<u>Block type</u>	<u>Description</u>	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.
<u>Block type</u>	<u>Description</u>												
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.												

<u>Byte(s)</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
CF=confile	Configuration file on which SSBLD statements are written under user index 377760g.
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 377760g.
BF	Same as BF=BUDT.
BF omitted	Same as BF=BUDT.

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 is the EST ordinal of the CU.
da is the device address of CUest.
st is the 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 is the device address of CIF (0 through 3).
DTI is the data transfer input.
DTO is the data transfer output.
st is the data transfer status (ON or OFF).

3. DTIn=DIFO=st,DIF1=st.

n is the device address of DTI (0 or 1).
DIF is the device interface.
st is the specified device interface status (ON or OFF).

4. DTOn=DIFO=st,DIF1=st.

n is the device address of DTO (0 or 1).
st is the device interface status (ON or OFF).

5. DIFn=DRC0=st,DRC1=st.

n is the device address of DIF (0 or 1).
DRC is the data recording controller.
st is the data recording controller status (ON or OFF).

6. DRCn=DRD0=st,...,DRD7=st.

n is the device address of DRC (0 or 1).
DRD is the data recording driver.
st is the data recording driver status (ON or OFF).

7. AIFn=AC0=st,AC1=st,AC2=st,AC3=st.

AIF is the accessor interface.
n is the device address of AIF.
AC is the accessor.
st is the accessor status (ON or OFF).

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.
```

id is the SM identifier (A through H).

st is the status (ON or OFF).

n is the device address of AC (0 through 3).

est is the EST ordinal of the CU(s) connected to SMid.

a is the device address of the DRD(s) connected to ACn.

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.

Example of SSBLD Statements

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=DTIO=ON,,DT00=ON,,.
DTIO=DIFO=ON,DIF1=ON.
DT00=DIFO=ON,DIF1=ON.
DIFO=DRCO=ON,DRC1=ON.
DIF1=DRCO=ON,DRC1=ON.
DRCO=DRDO=ON,DRD1=ON,,,,.
DRC1=DRDO=ON,DRD1=ON,,,,.
AIFO=ACO=ON,,,,.
SMA=ON,ACO,,CU41,DRDO=ON,DRD1=ON.
*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

Controls 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:

<u>Accessor Number</u>	<u>DRD Numbers</u>
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 BUDT) under user index 377760g 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,0 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. (jsn is the job sequence name of SSALTER as specified on the B,0 display)

3. Select the desired K display by entering

K.SM=id. (id is the SM identifier, A through H)

to select the SM K display (figure 11-3), which shows the 7990 configuration of the storage module and DRDs complete with status information; or

K.CUest. (est is the EST ordinal of the CU)

to select the CU K display (figure 11-4), which shows the status of various paths/nodes in the specified control unit.

 7990 SUBSYSTEM - SSALTER

		PCU ST	SCU ST	SM/DRD 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 11-3. 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.

MSAS OPERATIONAL PROCEDURES

Next is a summary of the operational procedures used in an MSAS environment. These procedures include modifications to existing procedures and new MSAS 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

MSAS 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 MSAS environment than in a non-MSAS 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.

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, run 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 MSAS 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 MSAS 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 MSAS environment are similar to those in a non-MSAS 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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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 and time 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.

<u>Parameter</u>	<u>Description</u>
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-MSAS environment, it is possible to make all previous dump tapes available for use after a full dump is completed. However, in an MSAS 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 the OP=M, BD, and BT parameters, but not specifying the OP=S parameter.† 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.

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 MSAS 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 377760g 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).

†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.

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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
OP=M	Selects incremental dumps.
SD	Updates the date/time field on each master device dumped, as previously described for full dumps.
AD=yyymmdd and AT=hhmmss	Selects the date and time 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.

<u>Parameter</u>	<u>Description</u>
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-MSAS environment, it is possible to make all previous dump tapes available for use after a full dump is completed. However, in an MSAS 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 the OP=M, BD, and BT parameters, but not specifying the OP=S parameter.† 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.

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 MSAS 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 377760g 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).

†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.

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 MSAS 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.

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 MSAS system file, or by abnormal termination of SSEXEC or MSAS utility processing, which can interrupt the updating of an MSAS 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 PFLOAD 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 and then 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=RS 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.

Introduction	12-1
Utility Description Terminology	12-2
Cartridge	12-2
CSU Map	12-3
MSF Catalog	12-5
MSF Files	12-9
MSSEXEC	12-10
Subfamily	12-10
MSS Utilities	12-10
ASDEF	12-11
ASLABEL	12-12
ASLABEL Directives	12-13
Parameters for ASLABEL Directives	12-16
ASLABEL Update Sequence	12-18
Restrictions to ASLABEL	12-19
Example	12-19
ASMOVE	12-20
Selection Algorithms	12-23
Disk Space/Dump Tape Management	12-24
Restriction to ASMOVE	12-24
Example	12-24
ASVAL	12-26
Release Processing	12-27
Problem Reporting	12-28
Error Detection and Classification	12-28
Release Processing and Problem Fixing	12-30
Validation Report	12-31
Typical ASVAL Runs	12-33
Restrictions to ASVAL	12-34
ASUSE	12-34
Basic Usage Report	12-36
Optional Reports A and B	12-37
Optional Report C	12-39
Optional Reports D and E	12-39
ASDEBUG	12-40
ASDEBUG Directives	12-41
Parameters for the ASDEBUG Directives	12-42
Restrictions to ASDEBUG	12-44
MSS Operational Procedures	12-45
Initialization	12-45
Disk Space Management	12-45
MSF Space Management	12-46
Backup and Recovery	12-47
File Dumping	12-47
File Reloading	12-50
Error Conditions and Corrective Actions	12-50
Clearing Permanent File Error Flags	12-50
Permanent File Recovery	12-51
Cartridge Restoration and Reuse	12-51
Removal of Faulty or Missing Cartridges	12-52
MSF Catalog/CSU Map/PFC Problems	12-52
Pseudorelease	12-54

Introduction	12-1
Utility Description Terminology	12-2
Cartridge	12-2
CSU Map	12-3
MSF Catalog	12-5
MSF Files	12-9
MSSEXEC	12-10
Subfamily	12-10
MSS Utilities	12-10
ASDEF	12-11
ASLABEL	12-12
ASLABEL Directives	12-13
Parameters for ASLABEL Directives	12-16
ASLABEL Update Sequence	12-18
Restrictions to ASLABEL	12-19
Example	12-19
ASMOVE	12-20
Selection Algorithms	12-23
Disk Space/Dump Tape Management	12-24
Restriction to ASMOVE	12-24
Example	12-24
ASVAL	12-26
Release Processing	12-27
Problem Reporting	12-28
Error Detection and Classification	12-28
Release Processing and Problem Fixing	12-30
Validation Report	12-31
Typical ASVAL Runs	12-33
Restrictions to ASVAL	12-34
ASUSE	12-34
Basic Usage Report	12-36
Optional Reports A and B	12-37
Optional Report C	12-39
Optional Reports D and E	12-39
ASDEBUG	12-40
ASDEBUG Directives	12-41
Parameters for the ASDEBUG Directives	12-42
Restrictions to ASDEBUG	12-44
MSS Operational Procedures	12-45
Initialization	12-45
Disk Space Management	12-45
MSF Space Management	12-46
Backup and Recovery	12-47
File Dumping	12-47
File Reloading	12-50
Error Conditions and Corrective Actions	12-50
Clearing Permanent File Error Flags	12-50
Permanent File Recovery	12-51
Cartridge Restoration and Reuse	12-51
Removal of Faulty or Missing Cartridges	12-52
MSF Catalog/CSU Map/PFC Problems	12-52
Pseudorelease	12-54

INTRODUCTION

The MSF hardware product is a large-capacity online mass storage device, which is a cost-effective extension to the disk file storage system and an alternative to conventional magnetic tape storage. Storing files on MSF 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. You can use both MSF and magnetic tapes to protect files from hardware and system failures.

MSF is composed of the following components.

<u>Component</u>	<u>Description</u>
Cartridge	A plastic housing that encloses 2540 mm (100 in) of magnetic tape on which data is stored under program control.
Coupler	The interface between the peripheral processor (PP) and the mass storage adapter (MSA), which includes a buffer that contains data going to or coming from the mass storage transport.
CSU	Cartridge storage unit (CSU), which includes storage cells (cubicles) for 2052 cartridges† and a selector that moves cartridges among the mass storage transport, the cubicles, and the input/output (I/O) drawers of the CSU. Each CSU has two I/O drawers; the top one is the input drawer and the bottom one is the output drawer. Each drawer has eight vertically aligned drawer slots, numbered from 0 (top slot) to 7 (bottom slot). There can be up to 13 CSUs in an MSS configuration. Each CSU in a configuration is associated with a letter from A to M, which is the CSU identifier.
MSA	Mass storage adapter, which interfaces between the coupler and the mass storage transport or the CSU.
MST	Mass storage transport, which includes storage positions for five cartridges: one being read or written, two queued for reading or writing, and two queued for storage by the selector into the CSU.

The MSS is the product consisting of the MSF hardware, the CYBER coupler, the diagnostics, and the operational software. The basic function of MSS is to store data on MSF and move it to disk upon request for access by an authorized user. Control of file movement between disk and MSF 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 MSS utilities are introduced.

†2000 cartridges are available for file data storage and the rest are reserved for customer engineer and system use.

Detailed information about the call and operation of each MSS utility is included in this section.

<u>Utility</u>	<u>Description</u>
ASDEF	Creates the system files (CSU maps and MSF catalogs) necessary for MSS processing.
ASLABEL	Manages allocation of cubicles and assignment of cartridges in the CSU.
ASMOVE	Controls the destaging of files (creating MSF images) and the release of disk space.
ASVAL	Controls the release of MSF space and analyzes the CSU maps, MSF catalogs, and PFC entries to identify and flag discrepancies within these three components.
ASUSE	Provides reports on the assignment and availability of cartridges and cubicles within a CSU.
ASDEBUG	Corrects error conditions detected by ASVAL and recovers data from MSF cartridges.

UTILITY DESCRIPTION TERMINOLOGY

Review the following terms; they are defined as used in the descriptions of the MSAS utilities later in this section.

Cartridge

A cartridge is the MSS data storage component consisting of 16 streams, called allocation units (AUs). Streams that have data written on them (allocated streams) are chained together in the MSF catalog (definition following) to identify the sequence of streams that must be accessed in order to read a file on MSF. A head-of-chain (HOC) flag identifies the first AU in the chain, a link field identifies the next AU in the chain, and an end-of-chain (EOC) flag identifies the last AU in the chain.

The cartridge label contains information that characterizes one of the three types of cartridges that can be used in an MSS environment. Cartridges having label types other than the following cannot be processed by MSS except by the ASLABEL utility, which can rewrite the cartridge label, or by the ASDEBUG utility, which can write the data on the cartridge to a permanent file.

- A manufacturer's label contains the volume serial number (vsn) of the cartridge in a machine-readable format. A cartridge with a manufacturer's label can be added to a CSU and assigned to a subfamily or pool (refer to the ASLABEL utility later in this section).
- A scratch label contains the vsn of the cartridge and additional system information indicating that the cartridge is available for assignment by the ASLABEL utility. A scratch label also contains the usage record of the cartridge. A cartridge with a scratch label is called a scratch cartridge.
- A family label contains the vsn of the cartridge and additional system information indicating that the cartridge is assigned to a particular subfamily. It also indicates the CSU identifier and the coordinates x and y of the cubicle where the cartridge resides when it is not in use.

The addition, removal, and reassignment of cartridges are managed by the MSS utilities. The cartridge labels, and also the MSF catalogs and CSU maps (definitions following), contain information concerning the location and content of the cartridges. When a cartridge is labeled and assigned to a subfamily, a label is written on each stream of the cartridge, which contains the cartridge vsn, the CSU id, x and y coordinates, family, and subfamily to which it was assigned, and the stream number of the particular stream. If such a cartridge is accessed for relabeling or for reading or writing an MSF file, the stream label verification procedure is used to verify this stream label information against the data in the MSF catalogs and CSU maps. If a discrepancy is detected, an error message is issued and corrective action must be taken to update the cartridge labels and/or system files as described in a later section. Whenever the cartridge location or content is changed, the cartridge label, MSF catalogs, and CSU maps are updated to reflect the current status of the cartridge. The descriptions of the MSS utilities later in this section contain further information on cartridge management.

CSU Map

A CSU map is a direct access permanent file that contains information indicating how cubicles in a CSU are assigned to a family and identifying the cartridges that reside in the CSU. There is one CSU map for each CSU in the configuration. The permanent file name of the CSU map is CSMAP*i*, where *i* is the CSU identifier (a letter from A to M); its user index is 377760g; and its family is the default family on the mainframe on which MSSEXEC executes (refer to the definition of MSSEXEC later in this section).

A CSU map contains an entry for each possible coordinate pair (x,y) that identifies a cubicle in the CSU, from (0,0), which is bottom right, to (57,36), which is upper left. There are no cubicles at the positions where x=30 or y=18, but there are entries in the CSU map for such coordinate pairs and these entries indicate that no cubicle exists there. Thirty-one cubicles are reserved for customer engineer use; they have these coordinates: (57,36), (0,36), (57,0), and (26 through 0,0). Twenty-one cubicles are reserved for system use; they have these coordinates: (48 through 31,0), (29,0), (28,0), and (27,0). In all, there are 2146 entries in a CSU map and they are ordered (assigned a map entry ordinal) according to the positions in the CSU of the corresponding cubicles. That is, the following scheme is used to relate the coordinate pairs (x,y) and the map entry ordinals.

- Given the coordinates x and y, the ordinal is $2146-x-(y*58)$.
- Given the ordinal, y is the whole number quotient of $(2146-ordinal)/58$ and x is the remainder.
- Ordinal 0 does not represent a coordinate pair; it is used as a map header entry.

Thus, in a CSU map the zero entry is the map header entry; the first entry describes the cubicle at coordinates (57,36); the second entry describes the cubicle at coordinates (56,36); and so forth.

Each entry in a CSU map has this format:

	59 58	53	42 41	36	0
	code	ord	familyname		
	sub-family		vsn		

<u>Field</u>	<u>Description</u>																
code	Number from 1 to 7 indicating how the cubicle is assigned.																
	<table border="0"> <thead> <tr> <th><u>code</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <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 CSU map.</td> </tr> <tr> <td>4</td> <td>Assigned to the cartridge scratch pool.†</td> </tr> <tr> <td>5</td> <td>Assigned to a subfamily.</td> </tr> <tr> <td>6</td> <td>Unassigned.</td> </tr> <tr> <td>7</td> <td>No cubicle exists at these coordinates (x=30 or y=18).</td> </tr> </tbody> </table>	<u>code</u>	<u>Description</u>	1	Reserved for customer engineer.	2	Reserved for system use.	3	Reserved for a different CSU map.	4	Assigned to the cartridge scratch pool.†	5	Assigned to a subfamily.	6	Unassigned.	7	No cubicle exists at these coordinates (x=30 or y=18).
<u>code</u>	<u>Description</u>																
1	Reserved for customer engineer.																
2	Reserved for system use.																
3	Reserved for a different CSU map.																
4	Assigned to the cartridge scratch pool.†																
5	Assigned to a subfamily.																
6	Unassigned.																
7	No cubicle exists at these coordinates (x=30 or y=18).																
ord	Ordinal for this cubicle in the MSF catalog of the subfamily to which the cubicle is assigned. This ordinal is referred to as the FCT ordinal (refer to the definition of MSF catalog later, 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.																
ℓ	Linkage error flag that is set by the ASVAL utility when a CSU map entry of a cubicle assigned to a family has no corresponding entry in the MSF catalog.																
subfamily	Number from 0 to 7 identifying the subfamily to which the cubicle is assigned. This field is meaningful only if code=5.																
vsn	The 6-character volume serial number of the cartridge assigned to the cubicle. If no cartridge is assigned, this field contains spaces.																

The zero entry in the CSU map is the map header entry. In this entry the code field is 7, the leftmost 6 bits of the second word contain the CSU identifier, and the remaining bits are unused.

The CSU map is updated whenever the ASLABEL, ASVAL, or ASDEBUG utility causes a change in cubicle or cartridge assignment. It is recommended that the CSU map be backed up after every update to avoid problems such as:

- Mismatches between CSU map and MSF catalog entries.
- Lost CSU 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.

†The pool is an area of the CSU that stores scratch cartridges that are managed by the ASLABEL utility.


Thus, you should make a copy of the CSU map on tape or another device or family in order to retain the latest version of the CSU map. If a device containing CSU maps is reloaded, the latest version of the CSU maps should be recovered from the backup copy after the reload is completed. After recovering the CSU maps, the analyst should run the ASVAL utility to check that the entries in the CSU maps and MSF catalogs match. If there are inconsistencies, corrective action should be taken as described under Error Conditions and Corrective Actions later in this section.

MSF Catalog

An MSF catalog is a disk-resident direct access permanent file that contains information describing which streams of each cartridge assigned to a particular subfamily are allocated to MSF files and which streams are available for allocation. There is one MSF catalog for each subfamily of a family that can have MSF-resident files, and it resides on the master device for the subfamily. The permanent file name of the MSF catalog file is MSFCAT*i* and its user index is 37776*ig*, where *i* is the subfamily identifier (a number from 0 to 7). For example, file MSFCAT3 and user index 377763*g* identify the MSF catalog for subfamily 3.

An MSF catalog is partitioned into subcatalogs, one subcatalog for each CSU used by the subfamily. The maximum number of subcatalogs in an MSF catalog is 13 (the maximum number of CSUs in a configuration). Each subcatalog consists of two parts: the file and cartridge table (FCT) and the available stream table (AST). The FCT has an entry for each cubicle assigned to the subfamily from the given CSU. The maximum number of FCT entries in a subcatalog is 2000 (the maximum number of user cartridges in a CSU). The AST contains information used by the allocation algorithm to select the cartridges on which a file will reside.

The first part of the MSF catalog contains a header and at most 13 subcatalog entries. The header identifies the family and subfamily of the MSF catalog. Each subcatalog entry contains the length and location of its FCT and AST, the CSU identifier, unallocated stream count, and date of the last ASVAL run that resulted in the releasing of MSF space assigned to the particular CSU or the setting of flags in entries for cubicles in the CSU. The format of the header is:



<u>Field</u>	<u>Description</u>
familyname	The 7-character name in 6-bit display code of the family for this MSF catalog.
subfamily	Number from 0 to 7 identifying the subfamily for this MSF catalog.

The format of the subcatalog entry is:

59	53	41	35	29	17	0
id	length	FCT loc	AST loc	streams		
unused			date time			

<u>Field</u>	<u>Description</u>
id	CSU identifier for the subcatalog (a letter from A to M).
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.
streams	Number of unallocated streams for this subcatalog.
date time	Date and time of the last releasing of orphan files on the CSU identified by id, or the last time flags were set for cartridges or cubicles in the CSU (whichever is later).

The header and subcatalog information is arranged in this order:

1. First word of the header.
2. First word of each of the subcatalog entries.
3. Second word of the header.
4. Second word of each of the subcatalog entries.

Each cubicle from the given CSU assigned to the subfamily has an entry in the FCT of the subcatalog. This entry contains the coordinates (x,y) of the assigned cubicle. If a cartridge has been assigned to the cubicle, the FCT entry also contains the vsn of the cartridge, usage information, status flags, and information about each of the 16 streams of the assigned cartridge. The format of each FCT entry is:

59	57	53	47	41	35	33	29	27	20	0
bps	flags ₁	x	y	vsn						
unused	link ₁	load threshold			load count					
unused	link ₂	pass threshold			pass count					
unused	link ₃	error threshold			error count					
x ₁	x ₂	x ₃	unused							
y ₁	y ₂	y ₃	unused							
reserved for site (2 words)										
u	pru	flags ₂			chain	u	stream ₁ detail			
⋮										
reserved for CDC (8 words)										

<u>Field</u>	<u>Description</u>														
bps	Number of blocks per stream on the cartridge assigned to this cubicle.														
flags ₁	One of the following flags. <table border="1"> <thead> <tr> <th><u>Bit</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>53</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 by a directive to the ASLABEL utility.</td> </tr> <tr> <td>52</td> <td>Lost cartridge flag indicating that the cartridge assigned to this cubicle was not there the last time MSSEEXEC tried to pick it. This flag can be cleared by a directive to the ASLABEL utility.</td> </tr> <tr> <td>51</td> <td>Excessive write parity error flag. Space from such a cartridge will not be allocated. This flag can be cleared by a directive to the ASDEBUG utility.</td> </tr> <tr> <td>50</td> <td>Reserved.</td> </tr> <tr> <td>49</td> <td>Linkage error flag indicating that the CSU map entry for the location (x,y) is inconsistent with this FCT entry.</td> </tr> <tr> <td>48</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bit</u>	<u>Description</u>	53	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 by a directive to the ASLABEL utility.	52	Lost cartridge flag indicating that the cartridge assigned to this cubicle was not there the last time MSSEEXEC tried to pick it. This flag can be cleared by a directive to the ASLABEL utility.	51	Excessive write parity error flag. Space from such a cartridge will not be allocated. This flag can be cleared by a directive to the ASDEBUG utility.	50	Reserved.	49	Linkage error flag indicating that the CSU map entry for the location (x,y) is inconsistent with this FCT entry.	48	Reserved.
<u>Bit</u>	<u>Description</u>														
53	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 by a directive to the ASLABEL utility.														
52	Lost cartridge flag indicating that the cartridge assigned to this cubicle was not there the last time MSSEEXEC tried to pick it. This flag can be cleared by a directive to the ASLABEL utility.														
51	Excessive write parity error flag. Space from such a cartridge will not be allocated. This flag can be cleared by a directive to the ASDEBUG utility.														
50	Reserved.														
49	Linkage error flag indicating that the CSU map entry for the location (x,y) is inconsistent with this FCT entry.														
48	Reserved.														
x	x coordinate of this cubicle.														
y	y coordinate of this cubicle.														
vsn	The 6-character volume serial number of the cartridge assigned to this cubicle. If no cartridge is assigned, this field contains spaces.														
link _i	Ordinal of the FCT entry for the cubicle containing the next cartridge n 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 in one CSU. Also, since there are only three link fields, if a cartridge contains several files, only three can be contained on other cartridges.														
load threshold	If the number of loads of the cartridge exceeds this threshold value, streams from this cartridge will not be allocated to new files.														
load count	Number of times the cartridge assigned to this cubicle has been loaded to a mass storage transport read/write station.														
pass threshold	If the number of passes across the read/write heads exceeds this threshold value, streams from this cartridge will not be allocated to new files.														
pass count	Number of times the cartridge assigned to this cubicle has been passed across a read/write head.														
error threshold	If the number of errors recorded exceeds this threshold value, streams from this cartridge will not be allocated to new files.														

<u>Field</u>	<u>Description</u>
error count	Number of recovered read errors for the cartridge assigned to this cubicle.
x _i	x coordinate of the cubicle in which the cartridge referred to by the link _i field resides.
y _i	y coordinate of the cubicle in which the cartridge referred to by the link _i field resides.
u	Unused.
pru	Number of PRUs on the stream.
flags ₂	One of the following flags.

<u>Bit(s)</u>	<u>Description</u>
47	Start of fragment flag; set by the ASVAL utility.
46	Frozen chain flag indicating a problem with this allocation chain. Streams in this chain are not reused until a directive to the ASDEBUG utility clears the flag.
45	Stream conflict flag indicating an allocation conflict involving this stream. This flag is set by the ASVAL utility or by MSSEXEC.
44	Stream selected flag indicating that the allocation routine has selected this stream but the file data has not yet been written on the stream. This field is meaningful only if the free/busy flag indicates that the stream is free.
43	Free/busy flag indicating whether or not this stream is currently allocated to a file.
42	Reserved.
41-40	Off-cartridge link flag indicating that the next stream of the file is on another cartridge. The next cartridge is identified by one of the link _i fields. The value of the off-cartridge link flag (1, 2, or 3) specifies which link _i to use. The chain field specifies the stream number on the next cartridge. Only 3 of the 16 streams of a cartridge can have this flag set at any one time. This field is meaningful only if the chain control field indicates that this stream is the first or middle stream of the file.
39-38	Unused.
37-36	Chain control indicating whether the stream is a first (1), last (2), only (3), or middle (0) stream of the file. This field is meaningful only if the stream is allocated to a file.
35-34	Unused.

<u>Field</u>	<u>Description</u>
chain	Next stream number, if any, containing file data for the file. This field is meaningful only if the stream is allocated to a file.
stream ₁ detail	PRU, flag, and chain control information about stream 1. The upper 30 bits of this word contain information about stream 0. The lower 30 bits are divided into the same fields, and contain information about stream 1. The next seven words are divided similarly: the upper 30 bits contain information about stream 2, 4, 6, 8, 10, 12, or 14 and the lower 30 bits contain information about stream 3, 5, 7, 9, 11, 13, or 15.

The MSF catalog is updated whenever the ASLABEL, ASMOVE, ASVAL, or ASDEBUG utility causes a change in cartridge or cubicle assignment that affects the subfamily. Because the MSF 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 MSF catalogs are automatically recovered and all MSF-resident files as indicated in a recovered PFC will also have entries in the recovered MSF catalog. Thus, no special operational procedures are needed to back up an MSF catalog. It is possible, however, that the MSF catalog will be inconsistent with the CSU 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.

MSF Files

In an MSS environment, permanent files can be categorized according to whether or not they reside on MSF. A disk file is a permanent file that resides on disk but not on MSF. An MSF file is a permanent file that resides on MSF 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 MSF files can also have backup images on tape (refer to the NOS 2 Reference Set, Volume 3).

When a user defines a direct access file, initially it is a disk file. A disk file becomes an MSF file when it is destaged to MSF; that is, an image of the file is created on MSF. You can destage files through the ASMOVE utility, which is run periodically to manage disk space. When ASMOVE is run, files are destaged to MSF and/or their disk space released depending on certain file characteristics (refer to ASMOVE later in this section). Thus, after an ASMOVE run a file can reside on disk, on MSF, or on both. If the file does have an MSF image, the asa field in the PFC entry for the file indicates the location of the MSF copy. The obsolete (AFOBS) flag in the PFC entry indicates whether or not the MSF image is a current version of the file.

When a user attaches an MSF file, it is staged to disk from MSF (that is, a disk image is created) if the current version of the file is not on disk. If the file is attached in write mode, the MSF image is marked obsolete; that is, the AFOBS flag in the PFC entry for the file is set. This is because the disk image is immediately updated when the user makes changes to the file, but the MSF image is not updated until the ASMOVE utility is run again. Hence, the current version of the file resides on disk only. Setting the AFOBS flag ensures that the current version of the file will be copied to MSF and will replace the obsolete MSF file the next time ASMOVE is run. When a user attaches a file, however, it is always the current version of the file that the user accesses; an obsolete file cannot be accessed.

If the file is purged, its disk space, but not its MSF space, is immediately released. The ASVAL utility must be run to release MSF space allocated to purged files. Thus, because a purged file has no PFC entry linking to its MSF catalog entry, the MSF image that still exists before ASVAL is run is called an orphan file. However, a user can never access an orphan file.

When a file is destaged to MSF, information is written on each stream to which the file data is written. This information includes the file's creation date and time and user index from the PFC entry, the identity of the first stream to which the file is written, the identity of the stream immediately preceding this one, and the number of disk PRUs of data recorded on previous streams (for all but the first stream of multistream files).

When a file is staged back to disk, the file label verification procedure is used to verify this file label information against the stream label information (refer to the definition of cartridge earlier in this section) for each stream of the file. 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 MSF file could not be accessed, and the stream conflict error flag is set in the MSF catalog entry for the particular cartridge(s) and stream(s).

MSSEXEC

MSSEXEC is the main processing program that is responsible for controlling MSS activities. The MSS utilities issue requests to MSSEXEC to destage files from disk to MSF, purge unneeded MSF files, label or relabel cartridges, update CSU maps and MSF catalogs, and so forth. In a multiframe environment, two versions of this program exist: the mainframe to which the MSF device is physically connected (the master mainframe) has a program called MSSEXEC, and all other mainframes (the slave mainframes) have a program called MSSSLV. Refer to section 13, Multiframe 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 ASDEF utility is run to create MSS system files, the CSU maps are created under user index 377760g (subfamily 0) and one MSF catalog is created under each user index 37776ig (subfamily i), i=0,1,2,...,7. When the ASLABEL utility is used to assign a CSU, cubicle, or cartridge to a family, it is possible to specify assignment to particular subfamilies of the family.

MSS UTILITIES

MSS utilities are described next. All of these utilities must be run from system origin jobs. In a multiframe environment, the ASDEF and ASUSE 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 MSSEXEC executes.

ASDEF

ASDEF creates the system files (CSU maps and MSF catalogs) that are necessary for MSS processing. If a CSU is added to the MSF hardware configuration, ASDEF is used to create the CSU map for that CSU. If a family is to be permitted to have MSF-resident files, ASDEF is used to create the eight MSF catalogs for that family (one catalog for each subfamily).

NOTE

If the CS parameter is specified, ASDEF will create a CSU map for the specified CSU (refer to the definition of CSU map earlier in this section). It is recommended that the CSU map be copied on tape or on another device or family immediately after it is created.

The format of the ASDEF command is:

ASDEF,p₁,p₂.

<u>P_i</u>	<u>Description</u>
CS=id	CSU identifier of the CSU for which a CSU map is to be created; id is a letter from A to M.
CS	Same as CS=A.
CS omitted	No CSU map is to be created. FM=familyname or FM must be specified.
FM=familyname	Family for which MSF catalogs are to be created, one catalog for each subfamily.
FM	Same as FM=system default family.
FM omitted	No MSF catalogs are to be created. CS=id or CS must be specified.

Example:

ASDEF,CS=B.

CSMAPB, the CSU map for CSU B, is created, and its entries are as described in the definition of CSU map earlier in this section. However, since ASDEF does not assign cubicles, the entries for cubicles available for use initially have zeros or spaces in all but the code field, which indicates the cubicles are unassigned.

ASLABEL

ASLABEL manages cartridge assignment and cubicle allocation in a CSU. The following functions are performed through the use of directives to ASLABEL.

- Add a CSU to a subfamily (AC directive).
- Remove a CSU from a subfamily (RC 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 (RS directive).
- Repair a cartridge label or overwrite a family label (FX directive).
- Inhibit or allow further allocation of files to a cartridge (IB directive).

Input to ASLABEL is via a directive file. ASLABEL reads the appropriate CSU maps and MSF catalogs to determine how to process each directive and then issues requests to MSSEXEC to read and/or write cartridge labels and to update the CSU maps and MSF catalogs. ASLABEL 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 CSU map or MSF catalog, the cartridge label information is included on this report and the cartridge is put into the output drawer. It may be possible to restore such a cartridge as described under Cartridge Restoration and Reuse later in this section.

NOTE

ASLABEL updates the CSU map for the specified CSU (refer to the definition of CSU map earlier in this section). It is recommended that the CSU map be copied on tape or on another device or family immediately after each update.

The format of the ASLABEL command is:

ASLABEL,p₁,p₂.

<u>P_i</u>	<u>Description</u>
I=filename	File containing the directives to ASLABEL.
I	Same as I=COMPILE.
I omitted	Same as I=INPUT.

<u>P_i</u>	<u>Description</u>
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 ASLABEL command. The I parameter is ignored.
Z omitted	Directives are contained on the file specified by the I parameter.

ASLABEL Directives

The directives to ASLABEL can be specified on a separate file (specified by the I parameter) or after the ASLABEL command (Z specified). If on the input file, each directive must be specified on a separate line via 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,...,p_n.

Example:

ASLABEL,I=DIRFILE.

DIRFILE is the directive file; it contains the following directives.

OP=AM,N=4,PK=D.
OP=RM,V=VSN444,FM,SB=1.

Two directives to ASLABEL are specified. OP=AM adds four cartridges to the pool of CSU A (CS parameter not specified). OP=RM removes the cartridge whose vsn is VSN444 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:

ASLABEL,Z./OP=AM,N=4,PK=D./OP=RM,V=VSN444,FM,SB=1.

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 ASLABEL. Some of these directives cause cartridges to be physically moved to and from cubicles and the input or output drawer (refer to OP=AM, OP=RM, OP=RS, and OP=FX). The remaining directives cause only logical operations to occur, updating the MSS system files. Descriptions of the parameters for these directives follow the directives.

OP=PC - Add CSUs

OP=AC adds a CSU to a subfamily. The MSF catalog for the specified subfamily is updated to reflect that cartridges and permanent files for the subfamily can reside on the specified CSU. This directive, however, does not manipulate cartridges or cubicles.

OP=RC - Remove CSUs

OP=RC removes a CSU from a subfamily. The MSF catalog for the specified subfamily is updated to reflect that cartridges and permanent files for the subfamily cannot reside on the specified CSU. Before OP=RC can be specified, all cubicles in the specified CSU 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 a CSU to a subfamily (PT=F), the pool (PT=P), or the reserved area of the CSU (PT=R). More than one cubicle (N=n) can be added at a time. Specific cubicles (XI and YI parameters) can be added, but they must be currently unassigned. For PT=R, XI and YI must be used to add multiple cubicles; N=n is not valid. If cubicles are to be assigned to a subfamily, ASLABEL selects available cubicles closest to the top of the CSU. If cubicles are to be assigned to the pool, ASLABEL selects available cubicles closest to the bottom of the CSU. The CSU map is updated to reflect the new assignment of cubicles.

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 CSU (PK=R). More than one cubicle (N=n) can be removed at a time. ASLABEL reads the CSU map and selects the first empty cubicle assigned to the subfamily, pool, or reserved area to be removed. Specific cubicles (XI and YI parameters) can be removed, but they must be empty. The CSU map is updated to reflect that the cubicles are unassigned.

OP=AM - Add Cartridges

OP=AM adds a cartridge with a manufacturer's label or a scratch label to either a specific subfamily (PT=F) or the pool (PT=P). ASLABEL 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 (V=vsn) from the pool can be added to a subfamily. The CSU map, MSF catalog, and cartridge label are updated to reflect the new assignment of the cartridge.

<u>P_i</u>	<u>Description</u>
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 ASLABEL command. The I parameter is ignored.
Z omitted	Directives are contained on the file specified by the I parameter.

ASLABEL Directives

The directives to ASLABEL can be specified on a separate file (specified by the I parameter) or after the ASLABEL command (Z specified). If on the input file, each directive must be specified on a separate line via 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,...,P_n.

Example:

ASLABEL,I=DIRFILE.

DIRFILE is the directive file; it contains the following directives.

OP=AM,N=4,PK=D.
OP=RM,V=VSN444,FM,SB=1.

Two directives to ASLABEL are specified. OP=AM adds four cartridges to the pool of CSU A (CS parameter not specified). OP=RM removes the cartridge whose vsn is VSN444 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:

ASLABEL,Z./OP=AM,N=4,PK=D./OP=RM,V=VSN444,FM,SB=1.

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 ASLABEL. Some of these directives cause cartridges to be physically moved to and from cubicles and the input or output drawer (refer to OP=AM, OP=RM, OP=RS, and OP=FX). The remaining directives cause only logical operations to occur, updating the MSS system files. Descriptions of the parameters for these directives follow the directives.

OP=PC - Add CSUs

OP=AC adds a CSU to a subfamily. The MSF catalog for the specified subfamily is updated to reflect that cartridges and permanent files for the subfamily can reside on the specified CSU. This directive, however, does not manipulate cartridges or cubicles.

OP=RC - Remove CSUs

OP=RC removes a CSU from a subfamily. The MSF catalog for the specified subfamily is updated to reflect that cartridges and permanent files for the subfamily cannot reside on the specified CSU. Before OP=RC can be specified, all cubicles in the specified CSU 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 a CSU to a subfamily (PT=F), the pool (PT=P), or the reserved area of the CSU (PT=R). More than one cubicle (N=n) can be added at a time. Specific cubicles (XI and YI parameters) can be added, but they must be currently unassigned. For PT=R, XI and YI must be used to add multiple cubicles; N=n is not valid. If cubicles are to be assigned to a subfamily, ASLABEL selects available cubicles closest to the top of the CSU. If cubicles are to be assigned to the pool, ASLABEL selects available cubicles closest to the bottom of the CSU. The CSU map is updated to reflect the new assignment of cubicles.

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 CSU (PK=R). More than one cubicle (N=n) can be removed at a time. ASLABEL reads the CSU map and selects the first empty cubicle assigned to the subfamily, pool, or reserved area to be removed. Specific cubicles (XI and YI parameters) can be removed, but they must be empty. The CSU map is updated to reflect that the cubicles are unassigned.

OP=AM - Add Cartridges

OP=AM adds a cartridge with a manufacturer's label or a scratch label to either a specific subfamily (PT=F) or the pool (PT=P). ASLABEL 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 (V=vsn) from the pool can be added to a subfamily. The CSU map, MSF catalog, and cartridge label are updated to reflect the new assignment of the cartridge.

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 the output drawer (PT=D), or it removes any cartridge from the pool (PK=P) and moves it to the output drawer (PT=D). Any cartridge currently assigned to a subfamily cannot be removed unless it is empty; that is, all 16 streams must be unallocated. More than one cartridge (N=n) can be removed or a specific cartridge (V=vsn) from the pool can be removed to the output drawer. The CSU map, MSF catalog, and cartridge label are updated to reflect the change in location of the cartridge. The removed cartridge then has a scratch label, which allows it to be reassigned via the OP=AM directive.

If the cartridge specified by the V parameter is lost (does not reside in its assigned cubicle), the LT parameter should be specified. This allows the appropriate entries in the CSU map and MSF 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 ASLABEL aborts.

OP=RS - Restore Lost Cartridges

OP=RS restores to its proper cubicle a cartridge that was inadvertently removed from a CSU. If restoration is successful, the lost flag in the MSF catalog is cleared. If data recorded on the cartridge label does not agree with the information in the MSF catalog and the CSU 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 output drawer for use in further processing of the cartridge.

OP=FX - Fix Cartridge Labels

OP=FX writes a scratch label on a cartridge 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 via the OP=AM directive, but it is not feasible to first remove the cartridge normally via the OP=RM directive. For example, if ASLABEL is run to add a cartridge to a subfamily, and a system failure occurs before the MSF catalog and CSU map are updated but after the cartridge is relabeled, then the cartridge label does not match the corresponding entries in the MSF catalog and CSU 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 and SB parameters must identify the family and subfamily to which the cartridge was assigned. The CSU map and MSF catalog are updated to reflect the new cartridge label.

OP=IB - Control Cartridge Allocation

OP=IB sets or clears the inhibit allocation flag in the MSF catalog entry for the specific cartridge (V=vsn must be specified). If the flag is set (ON), MSSEXEC does not allocate new MSF files to this cartridge. If the flag is cleared (OF), allocation of files to this cartridge is enabled.

Parameters for ASLABEL Directives

The descriptions of the parameters to the ASLABEL directives follow. Not all parameters are valid with all directives, as indicated.

<u>P_i</u>	<u>Description</u>
CS=id	CSU identifier of the CSU to be used by ASLABEL; id is a letter from A to M.
CS	Same as CS=A.
CS omitted	Same as CS=A.
D=d	Input drawer slot from which ASLABEL picks the cartridge; valid only with OP=AM, OP=RS, or OP=FX.
D	First nonempty input drawer slot is to be used; valid only with OP=AM, OP=RS, or OP=FX.
D omitted	Same as D.
FM=familyname	Family to which ASLABEL adds or from which it removes a cartridge or CSU. 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.
LT	CSU map and MSF 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 ASLABEL aborts.
N=n	Number of cartridges or cubicles to be added, removed, or repaired; $1 \leq n \leq 2000$; not valid if PT=R is specified. If V=vsn is specified, n must be 1.
N	Same as N=1.
N omitted	Same as N=1.
OF	Inhibit allocation flag in the MSF catalog is to be cleared; valid only with OP=IB.
ON	Inhibit allocation flag in the MSF catalog is to be set; valid only with OP=IB.

<u>P_i</u>	<u>Description</u>										
PK=pkloc	Location from which the cartridge or cubicle is to be picked; not valid if V=vsn is specified.										
	<table> <thead> <tr> <th><u>pkloc</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>D</td> <td>Cartridge is to be picked from the specified input drawer slot (D=d). PK=D is valid only with OP=AM, OP=RS, or OP=FX.</td> </tr> <tr> <td>F</td> <td>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.</td> </tr> <tr> <td>P</td> <td>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.</td> </tr> <tr> <td>R</td> <td>Cubicle is to be picked from the reserved area of the CSU. PK=R is valid only with OP=RB.</td> </tr> </tbody> </table>	<u>pkloc</u>	<u>Description</u>	D	Cartridge is to be picked from the specified input drawer slot (D=d). PK=D is valid only with OP=AM, OP=RS, 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 CSU. PK=R is valid only with OP=RB.
<u>pkloc</u>	<u>Description</u>										
D	Cartridge is to be picked from the specified input drawer slot (D=d). PK=D is valid only with OP=AM, OP=RS, 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 CSU. PK=R is valid only with OP=RB.										

PK Same as PK=P.

PK omitted Same as PK=P.

PT=ptloc Location into which the cartridge or cubicle is to be put.

<u>ptloc</u>	<u>Description</u>
D	Cartridge is to be put into the first available output drawer slot. 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 CSU. PT=R is valid only with OP=AB.

PT Same as PT=P.

PT omitted Same as PT=P.

SB=subfamily Subfamily to which ASLABEL adds or from which it removes a cartridge or CSU; $0 \leq \text{sub} < 7$. 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.

<u>P_i</u>	<u>Description</u>
V=vsn	Volume serial number of the cartridge to be added, removed, or repaired; not valid if PK=pkloc is specified. If V=vsn is specified, n must be 1 if N=n is specified.
V	Volume serial number of the cartridge is not specified.
V omitted	Same as V.
XI=x ₁	Column of the CSU to be added or removed; $0 \leq x_1 \leq 57$, $x_1 \neq 30$; valid only with OP=AB or OP=RB.
YI=y ₁	Row of the CSU to be added or removed; $0 \leq y_1 \leq 36$, $y_1 \neq 18$; valid only with OP=AB or OP=RB.
XI=x ₁ , YI=y ₁	x and y coordinates of the cubicle to be added or removed; $0 \leq x_1 \leq 57$, $0 \leq y_1 \leq 36$, $x_1 \neq 30$, $y_1 \neq 18$; valid only with OP=AB or OP=RB.
XI=x ₁ , YI=y ₁ , XF=x ₂ , YF=y ₂	Rectangle of cubicles to be added or removed; cubicles with x coordinates between x ₁ and x ₂ and y coordinates between y ₁ and y ₂ are included; valid only with OP=AB or OP=RB. At most, 100 cubicles can be included in the rectangle. $x_1, x_2 < 57$, $x_1, x_2 \neq 30$; $y_1, y_2 < 36$, $y_1, y_2 \neq 18$; $x_1 < x_2$; $y_1 < y_2$. XF and YF must both be specified, if either is specified. XF and YF cannot be specified unless both XI and YI are specified.
XI and YI 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.

ASLABEL Update Sequence

The general result of each directive to ASLABEL is an update of the CSU maps, MSF catalogs, and cartridge labels, whichever are appropriate, to reflect the changes in cartridge, cubicle, or CSU assignment. Because the MSF 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 MSF catalogs immediately after an ASLABEL run. However, the backup and recovery of CSU maps do require special operational procedures, which should be performed immediately after an ASLABEL run (refer to the definition of CSU map earlier in this section).

When ASLABEL 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 MSS system files, relabel the cartridge, and add the new assignment information to the MSS system files. If an interruption such as a system failure, ASLABEL abort, or MSSEXEC abort prevents ASLABEL from completing the update sequence, the location of the affected cartridge and the status of the CSU maps and MSF catalogs depend on the point of interruption, as follows:

- If the cartridge label, MSF catalog, and CSU map do not all match, then the cartridge is put into the output drawer. The OP=RS 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, MSF catalog, and CSU map do match, then the cartridge may be returned to its original location, the new location, or the output drawer, depending on the exact point of interruption. If the cartridge is in the output drawer, OP=RS 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 output drawer.

Restrictions to ASLABEL

The following restrictions apply to the ASLABEL utility.

- MSSEXEC must be running when ASLABEL is run.
- Only one copy of ASLABEL can be run at a time.
- ASLABEL, ASVAL, and ASDEBUG cannot be run at the same time.

Example

The following output shows the format of an ASLABEL report. The cartridge label information is included on the report because of a mismatch with the CSU map. An error message is issued for the first directive on the command (in this case, the only directive).

ASLABEL REPORT FILE

ASLABEL,Z./OP=RS,PK=D,CS=B.

1 OP=RS,PK=D,CS=B.

1 OP=RS,PK=D,CS=B.

VSN = P66157
FAMILY = SYSTST
SUBFAMILY = 0
CSU = B
X = 2
Y = 6

*** ERROR 8 DIRECTIVE 1
UNEXPECTED CSU, X, Y, FAMILY OR SUBFAM.***

ASMOVE

ASMOVE manages disk and MSF residence. That is, ASMOVE determines which files should be left on disk, which files should be released from disk and moved to MSF, and which files should be resident both on disk and on MSF. The ASVAL utility, not ASMOVE, controls the release of MSF space.

The selection process includes these two algorithms that weigh certain file characteristics as follows:

- Files are selected for destaging to MSF 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.

ASMOVE 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 MSF images but is to reside only on MSF, ASMOVE releases the disk space for the file. For files that do not have an MSF image, ASMOVE creates entries on the ASMOVE/MSSEXEC communication file, MOVCOM, 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. MSSEXEC then processes each destage and destage/release request on MOVCOM.

ASMOVE generates an output report that lists the files released by ASMOVE and the files contained on MOVCOM. A report on the use of disk space before and after the ASMOVE run is included in the output file and the dayfile.

The format of the ASMOVE command is:

ASMOVE,P1,P2,...,Pn.

<u>P_i</u>	<u>Description</u>
FM=familyname	Family to be used by ASMOVE.
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.
L=0	No output file is to be generated.
L omitted	Same as L=OUTPUT.
NW	No wait. ASMOVE will not wait for completion of destage and release processing by MSSEXEC.
NW omitted	ASMOVE will wait for completion of destage and release processing by MSSEXEC.
RD=yymmdd	Last access date. All files not accessed after yymmdd are to be released from disk.

<u>P_i</u>	<u>Description</u>
RD omitted	No files are to be released.
RO	Report only. ASMOVE does not release files from disk and does not send requests to MSSEXEC to destage or destage/release files.
RO omitted	Disk space is to be released and requests are to be sent to MSSEXEC, if appropriate.
RT=hmmss	Last access time. All files not accessed after hmmss of the day specified by the RD parameter are to be released.
RT	Same as RT=000000 (midnight).
RT omitted	Same as RT=000000 (midnight).
TM=mode	Selects or deselects test mode (refer to Pseudorelease later in this section).

<u>mode</u>	<u>Description</u>
Y	Select test mode. Pseudorelease of all files selected for release from disk by this ASMOVE run is performed.
N	Deselect test mode. The pseudorelease flag is cleared and disk images for all files from the selected family that were previously pseudoreleased are really released from disk. Normal release processing is performed for all files selected for release from disk by this ASMOVE run.
TM omitted	Normal release processing is to be performed for files that do not have the pseudorelease flag set. Files with the pseudorelease flag set are treated as if they have already been released.
UI=userindex	Destage and release processing is restricted to files having user index userindex.
UI=0	All user indexes are processed.
UI omitted	All user indexes are processed.

The following parameters for ASMOVE redefine the values of the weight factors or thresholds (installation parameters) used in the algorithms that select files to be destaged or released. Use these parameters 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 ASMOVE from selecting small files for destaging to MSF. Unless otherwise stated, each of these parameters causes the installation-defined value to be multiplied by the integer value n ; $n \leq 0$.

<u>Parameter</u>	<u>Description</u>
DB= n	n times the installation-defined DB weight factor is to be used as the preferred residence value for destage decisions involving files with a PR=M attribute.†
DB	Same as DB=1.
DB omitted	Same as DB=1.
DC= n	n times the installation-defined weight factor is to be used as the preferred residence value for destage decisions involving files with a PR=N attribute.†
DC	Same as DC=1.
DC omitted	Same as DC=1.
DL= n	n times the installation-defined length weight factor is to be used as the length weight factor for destage decisions.
DL	Same as DL=1.
DL omitted	Same as DL=1.
DT= n	n times the installation-defined time weight factor is to be used as the time weight factor for destage decisions.
DT	Same as DT=1.
DT omitted	Same as DT=1.
DV= n	n times the installation-defined destage control value is to be used as the destage control value.
DV	Same as DV=1.
DV omitted	Same as DV=1.
MN= n	n times the installation-defined minimum length threshold is to be used as the minimum allowable size in disk PRUs (64 words) for MSF files.

†The file owner specifies the preferred residence attribute via the PR parameter on the DEFINE or CHANGE command (refer to of the NOS 2 Reference Set, Volume 3).

<u>Parameter</u>	<u>Description</u>
MN	Same as MN=1.
MN omitted	Same as MN=1.
MX=n	n times the installation-defined maximum length threshold is to be used as the maximum allowable size in disk PRUs for MSF files.
MX	Same as MX=1.
MX omitted	Same as MX=1.

Selection Algorithms

ASMOVE determines which files to destage and/or release according to the following algorithms. Files that reside only on MSF are not considered because they have been destaged and released previously. Also, indirect access files are excluded from consideration because they cannot reside on MSF. For all other files, ASMOVE 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 MSF, ASMOVE 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, ASMOVE uses the destage algorithm to determine whether or not to destage the file. If the file is to be destaged to MSF, ASMOVE also determines via the release algorithm whether or not to release the file's disk space.

Destage Algorithm

The destage algorithm is used to select for destaging certain files that do not have current MSF images. ASMOVE calculates the destage value for each eligible file according to the following equation. If the destage value exceeds the destage control value (refer back to the DV parameter), ASMOVE selects the file for destaging. Otherwise, the file remains on disk only.

$$\text{destage value} = (1+t*\text{time})*(1+l*\text{length})*(\text{res})$$

<u>Variable</u>	<u>Description</u>
t	Installation-defined time weight factor.
time	Number of days since the file was last modified.
l	Installation-defined length weight factor.
length	Length of the file in PRUs.
res	Installation-defined preferred residence value. This value depends on whether the file owner specified no preferred residence or MSF residence preferred for the file. Refer back to the DB and DC parameters.

Release Algorithm

The release algorithm is used to select files whose disk space is to be released. Any file with a backup required (BR=Y)† attribute and whose master device has not been dumped since the file was last modified is not considered for disk space release. The date and time the master device was last dumped was set by the SD parameter on PFDUMP. For all other files, ASMOVE checks the last access date and time in the PFC entry for each file. If the file was last accessed before the date and time specified by the RD and RT parameters, its disk space is to be released. If the file to be released has a current MSF image, ASMOVE releases the disk space. If the file to be released does not have a current MSF image, MSSEXEC creates one before releasing the disk space.

Disk Space/Dump Tape Management

As more disk-resident files are created and more MSF-resident files are staged to disk, it will be necessary to monitor the availability of disk space. It is recommended that ASMOVE be used as a periodic disk space management procedure to avoid frequent disk full conditions (refer to Disk Space Management later in this section). ASMOVE 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 MSF and/or releasing disk space of files that need not be on the dump tape (refer to Dump Tape Management later in this section).

Restriction to ASMOVE

Only one copy of ASMOVE per family can be run at a time. A second ASMOVE aborts if the first one has not completed.

Example

The following output shows the format of an ASMOVE report. The files that are to be destaged and destaged/released are listed. An account of the disk space and the destage for backup, release date, and release time values is given.

†The file owner specifies the backup requirement using the BR parameter on the DEFINE or CHANGE command (refer to the NOS 2 Reference Set, Volume 3).

ASMOVE REPORT.

ASMOVE,FM=SYSTST,L=MOV,RD=791109,RT=085547,MN=1,MX=9999,TM=Y

FILE	UI	LENGTH	
AAAAAI	172	22	RELEASE AND DESTAGE
AAAAAJ	172	22	RELEASE AND DESTAGE
AAAAAK	172	7	RELEASE AND DESTAGE
AAAAAL	172	29	RELEASE AND DESTAGE
AAAAAM	172	41	RELEASE AND DESTAGE
AAAAAN	172	50	RELEASE AND DESTAGE
AAAAAO	172	43	RELEASE AND DESTAGE
AAAAAP	172	31	RELEASE AND DESTAGE
AAAAAQ	172	27	RELEASE AND DESTAGE
AAAAAR	172	34	RELEASE AND DESTAGE
AAAAAS	172	16	RELEASE AND DESTAGE
DEF002	170	78	RELEASE AND DESTAGE
AAOAAA	172	19	RELEASE AND DESTAGE
AAEAAA	172	17	RELEASE AND DESTAGE
AAAPAA	172	20	RELEASE AND DESTAGE
MOV0040	170	6	RELEASE AND DESTAGE
MOV0005	170	6	RELEASE AND DESTAGE
DLF0006	170	4	RELEASE AND DESTAGE
DLF0012	170	70	RELEASE AND DESTAGE
DMP0012	170	160	RELEASE AND DESTAGE
RDF0012	170	59	RELEASE AND DESTAGE
CLF0012	170	5	RELEASE AND DESTAGE
DWK0013	170	1065	RELEASE AND DESTAGE
WLF0013	170	29	RELEASE AND DESTAGE
MOV0014	170	7	RELEASE AND DESTAGE
DEF0015	170	692	RELEASE AND DESTAGE
DLF0017	170	78	RELEASE AND DESTAGE
CLF0017	170	5	RELEASE AND DESTAGE
STIMABS	170	103	DESTAGE
GLOSARY	172	66	DESTAGE
AAAVAAA	172	14	DESTAGE
DMP0017	170	287	DESTAGE
RDF0017	170	59	DESTAGE
DWK0018	170	1089	DESTAGE
WLF0018	170	31	DESTAGE
RPT0019	170	119	DESTAGE
MOV0020	170	8	DESTAGE
DLF0021	170	107	DESTAGE
DMP0021	170	481	DESTAGE
RDF0021	170	60	DESTAGE
RPT0022	170	150	DESTAGE
DEF0023	170	687	DESTAGE
AAAXAAA	172	12	DESTAGE
AAAYAAB	172	13	DESTAGE
AAAYAAC	172	14	DESTAGE

ACTPFSPACE BEFORE PROCESSING = 6488
 ACTPFSPACE AFTER PROCESSING = 3846
 PF SPACE RELEASED = 2642
 RELEASE DATE USED = 82/02/10.
 RELEASE TIME USED = 08.55.47.

ASVAL

ASVAL either performs release processing or reports on problems with the current MSS system files. That is, it either makes available MSF space presently allocated to files that are no longer needed, or reports on irregularities or discrepancies found in the current MSF catalogs and PFC entries for the specified family and, optionally, in certain CSU maps. The function to be performed is determined by whether or not the RF option is specified, as described in the following paragraphs.

NOTE

ASVAL updates the CSU map for the specified CSU. (Refer to the definition of CSU map earlier in this section.) It is recommended that the CSU map be copied on tape or on another device or family immediately after every update of the CSU map.

The format of the ASVAL command is:

ASVAL,P1,P2,...,Pn.

<u>P_i</u>	<u>Description</u>
AM	The CSU map for the CSU specified by the CS parameter is to be analyzed in addition to the MSF catalogs; not valid if RF=filename or RF is specified.
AM=	Same as AM.
AM omitted	CSU maps are not to be analyzed.
CS=id	CSU identifier of the CSU to be used. Up to 13 CSUs can be selected by the letters A through M. For example, CS=ACJG selects CSU A, C, G, and J.
CS	Same as CS=ABCDEFGHIJKLM.
CS omitted	Same as CS=ABCDEFGHIJKLM.
FM=familyname	Family to be analyzed; not valid if the RF option is specified.
FM	Same as FM=system default family; not valid if the RF option is specified.
FM omitted	Same as FM=system default family if the RF option is not specified. The family on the release data file is used if the RF option 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.

<u>P_i</u>	<u>Description</u>
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.
RF=filename	File that contains the release data file.
RF	Same as RF=ZZZZRDF.
RF omitted	Current versions of the MSF 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.
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. The minimum number of cartridges is the quotient of (number of streams+15)/16; the remainder is ignored. For example, if the file is contained on 100 streams and 10 cartridges, it is scattered if n=1 but is not scattered if n=5.
ST	Same as ST=0. 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=0.

Release Processing

If RF=filename or RF is specified, ASVAL determines which MSF files are no longer needed and issues a request to MSSEXC to purge these files so their MSF space can be reused. The procedure is for ASVAL to analyze copies of the MSF catalogs and PFC entries for the specified family that are contained on the release data file (RDF) specified by the RF parameter. Those MSF files described in an MSF catalog but not having a PFC entry (that is, orphans) can be purged. During the analysis (refer to both MSF Catalog Analysis and PFC Analysis later in this section), ASVAL 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 and the last purge date in the MSF catalog is prior to the RDF file dump date. That is, trouble-free orphans are purged and their MSF space is made available for reuse. The current MSF 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 MSF space. If RL is not specified or the last purge date in the MSF catalog follows the RDF file dump date, 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 MSF space.

The RDF used for this analysis is a file produced during a previous PFDUMP run and it contains versions of the MSF 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 MSF space. For example, the analyst might run ASVAL 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 MSF Space Management later in this section). The capability to release trouble-free orphans by analyzing the current MSF catalogs and PFC entries is not provided in order to ensure that an MSF file is released from MSF only after it is no longer needed.

Problem Reporting

If RF=filename or RF is not specified, ASVAL reports on problems with the current MSF catalogs and PFC entries for the specified family. If AM is specified, problems with CSU maps are also included in the report. ASVAL examines the MSS system files and PFC entries and searches for problem chains and fragments, problem asa values, and CSU map/MSF catalog mismatches. The procedures ASVAL uses to detect and classify inconsistencies and discrepancies are described later in this section under Error Detection and Classification. ASVAL keeps track of the error conditions, if any, it discovers, and if the error count is less than or equal to the threshold specified by the FX parameter, problem fixing is performed. That is, ASVAL sets flags in the appropriate entries of the CSU map, MSF catalog, and/or PFC entries to prevent propagation of errors due to the inconsistencies or discrepancies found and to permit error recovery by the ASDEBUG utility. A count of the errors is recorded in the dayfile. A validation report is issued that lists the errors encountered, the number of trouble-free orphans, and the amount of releasable MSF space.

Error Detection and Classification

ASVAL detects and classifies errors according to the following procedures. During the CSU map analysis, ASVAL detects and classifies errors in the CSU map. During the MSF catalog analysis, ASVAL detects errors with chains of AUs; during the PFC analysis, ASVAL 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.

CSU Map Analysis

If the AM option is specified, ASVAL attempts to locate problems with CSU map entries by comparing the MSF catalogs and the CSU map. For each coordinate pair (x,y) in an MSF catalog entry, ASVAL locates the corresponding CSU map entry. A type 1 error exists if the code field in the CSU map entry is not 5 (assigned to a subfamily), or if the family, subfamily, or vsn fields in the CSU map entry do not match those in the corresponding MSF catalog entry. ASVAL also scans the CSU 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 MSF catalog entry.

MSF Catalog Analysis

The MSF 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 MSF files resides on a cartridge that is lost or has excessive write parity errors.

ASVAL scans the MSF 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, but rather an incorrect link value exists, an AU links to an unallocated stream, 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.

<u>Chain</u>	<u>Description</u>
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 plus the minimum number of cartridges needed for the file.

ASVAL 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 MSF catalog analysis. For each PFC entry with asa≠0 (the file has an MSF image and the asa value identifies the first AU in the chain containing the file), ASVAL classifies the following errors. Error type 3 exists if the asa value is incorrect. 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.

ASVAL 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.

ASVAL 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 ASVAL was run for release processing purposes (RF=filename or RF specified), ASVAL issues a request to MSSEXEC to release trouble-free orphans if RL was specified. If RL was not specified, no release processing is performed. If ASVAL was run for problem reporting purposes (RF omitted), the following action is taken.

- For error type 1, the linkage error flag is set in the MSF catalog entry.
- For error type 2, the linkage error flag is set in the CSU map entry.
- For error type 3, there are two 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 MSF 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 and that ASVAL then be rerun to clear the asa field in preparation for other corrective action.
- 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 MSF 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 ASDEBUG utility, but prevents these AUs from being overwritten until then. Thus, the ASDEBUG utility can be used to inspect or save data from the corresponding streams or cartridges.

Validation Report

The validation report consists of a report heading and a series of report groups for each subfamily and CSU being reported on. The heading identifies the subfamily, the CSU, whether or not there are any problems, and the last purge date and time for the CSU (the last time orphans on this CSU 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 in the following paragraphs. Each report group, however, contains the following items.

- Error type (a number from 1 to 7).
- Identification (refer to the particular error type described in the following paragraphs).
- Chain information (MSF catalog ordinal, stream 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 MSF space, the total number of errors detected, and whether or not the MSS system files were updated.

Error Types 1 and 2

Error types 1 and 2 identify mismatches between the CSU map and the MSF catalog. In the validation report, the identification field lists the MSF catalog ordinal, the x and y coordinates, and the vsn 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 ASUSE utility to produce a detailed report of the appropriate CSU map and MSF 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 MSF catalog ordinal and stream 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 MSF catalog ordinals and stream numbers of the intersecting chains. Therefore, when intersections are reported, there are two entries for the affected subfamily and CSU: one lists only the intersections and the other lists all the errors encountered.

The following output shows the format of a validation report. ASVAL was run for problem reporting purposes (RF not specified), and no errors were detected.

ASVAL - VALIDATION REPORT

ASVAL - VER 1.0

FAMILY = SYSTST

ASVAL,FM=SYSTST,SB=0,CS=J.

L = OUTPUT
RF = 0
AM = 0
CS = J
FM = SYSTST
FX = 0
RL = 0
SB = 0
ST = 0

SUBFAMILY = 0 CSU = J -- GOOD -- LPDT = ***** *****

RELEASABLE MSS FILES = 12

RELEASABLE MSS STREAMS = 16

TOTAL VALIDATION ERRORS = 0

CATALOGS NOT MODIFIED

REPORT COMPLETE

Typical ASVAL Runs

As described previously, ASVAL is run either to make MSF space available for reuse or to report on problems with the current MSS system files and/or PFC entries. The following examples show typical ASVAL runs that may be used periodically for these purposes.

Example:

The following ASVAL command causes MSF space to be made available for reuse.

```
ASVAL,RF=DUMPl,RL.
```

DUMPl is the release data file produced by a previous PFDUMP run from which ASVAL can identify all the MSF files that were orphans at the time of the dump. If no error conditions are detected in the MSF catalog and PFC entries contained on file DUMPl (FX=0 by default), the orphans are purged and the MSF space assigned to them is released. The last purge date and time field in the subcatalog (in the MSF catalog) for each CSU for each subfamily is updated to the time of the ASVAL run if any file from that CSU is purged. This is to ensure that a subsequent ASVAL run does not release the same space a second time, since the MSF 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 ASVAL with the following command to ensure that release processing is performed (n is the number of errors detected during the previous ASVAL run).

```
ASVAL,RF=DUMPl,RL,FX=n.
```

After this ASVAL run has completed, an analysis of the current version of the CSU map, MSF catalog, and PFC entries should be made to determine whether or not the errors detected in the first ASVAL run exist in the current MSS system files and PFC entries. This is accomplished through the following command.

```
ASVAL,FM=familyname,AM.
```

familyname is that on the release data file, DUMPl.

Example:

Upon completion of any device reload that includes recovery of an MSF catalog and/or CSU map, ASVAL should be run to determine whether any CSU map/MSF catalog mismatches exist. The following call accomplishes this.

```
ASVAL,FM=familyname,SB=n1,n2,...,nh,AM.
```

familyname is that for which recovery was done and n₁,n₂,...,n_h are the affected subfamilies if just some of the devices of the family were reloaded.

ASVAL analyzes the CSU map, MSF 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 ASLABEL 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 CSU map will reflect the results of the ASLABEL run but the MSF catalog will not. Such mismatches will be reported on the validation report.

If the following ASVAL command is then made, the errors detected during the previous run will be flagged as described previously under Release Processing and Problem Fixing.

ASVAL,FM=familyname,SB=n₁,n₂,...,n_h,AM,FX=n.

familyname and n₁,n₂,...,n_h are the same as in the previous ASVAL run and n is the number of errors detected during the previous run. After investigating the cause of these errors, the analyst can run the ASDEBUG utility to correct the error conditions (refer to ASDEBUG later in this section).

If a device reload includes recovery of all CSU maps (default family, user index=3777608), the latest CSU maps should be recovered from backup copies. Then each family that has MSF-resident files should be analyzed via the following command to detect CSU map/MSF catalog mismatches.

ASVAL,FM=familyname,AM.

It is recommended that the previous ASVAL run be made periodically to check whether any unexpected error conditions exist. As the site analyst becomes more familiar with MSS processing, these periodic ASVAL runs can be done less frequently.

Restrictions to ASVAL

- Only one copy of ASVAL can be run at a time.
- ASVAL, ASLABEL, and ASDEBUG cannot be run at the same time.

ASUSE

ASUSE reads data in the MSF catalogs and CSU maps and produces reports on the availability of space on MSF cartridges and the allocation of cubicle space within a CSU. The reports may not be completely up to date because the MSF catalogs and CSU maps can be updated while the reports are being generated.

The types of reports that ASUSE generates are:

<u>Report</u>	<u>Description</u>
Basic usage report	Lists general information about the use of each CSU in a subfamily.
Optional report A	Identifies cartridges with a specified number of streams available for assignment.†
Optional report B	Identifies cartridges with flags set in the MSF catalog.
Optional report C	Lists the contents of a CSU as described in the CSU map.

† A cartridge that has the lost cartridge flag, inhibit allocation flag, or excessive write parity errors flag set is considered to have zero streams available for allocation regardless of the number of unallocated streams on the cartridge.

<u>Report</u>	<u>Description</u>
Optional report D	Lists detailed cartridge status information on each entry in the MSF catalog.
Optional report E	Lists detailed cartridge and stream status information on each entry in the MSF catalog.

The format of the ASUSE command is:

ASUSE,p1,p2,...,pn.

<u>P_i</u>	<u>Description</u>
CS=id	CSU identifier of the CSU to be used. Up to 13 CSUs can be selected by the letters A through M. For example, CS=ACJG selects CSU A, C, G, and J.
CS	Same as CS=ABCDEFGHJKLM.
CS omitted	Same as CS=ABCDEFGHJKLM.
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.
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).

<u>op</u>	<u>Description</u>
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.
E	Optional report E and basic usage report.
OP	Basic usage report only is to be generated.
OP omitted	Same as OP.

<u>P_i</u>	<u>Description</u>
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.
SL=n	Minimum number of streams available for assignment; valid only with optional report A. Cartridges with n or more streams available are reported. $0 \leq n \leq 16$, $n \leq m$ (refer to SU=m).
SL	Same as SL=0.
SL omitted	Same as SL=0.
SU=m	Maximum number of streams available for assignment; valid only with optional report A. Cartridges with m or less streams available are reported. $0 \leq m \leq 16$, $n \leq m$ (refer to SL=n).
SU	Same as SU=16.
SU omitted	Same as SU=16.

Basic Usage Report

The basic usage report includes a title line, which identifies the subfamily, family, and CSU being reported on. The statistics in this report contain separate totals for CSUs and subfamilies. The following items are listed.

- Number of CSU locations reserved for the subfamily.
- Number of cartridges in the subfamily.
- Number of cartridges with 0 streams available for assignment.
- Number of cartridges with 1 stream available for assignment.
- .
- .
- Number of cartridges with 16 streams available for assignment.
- Number of unassigned and unflagged streams.
- Number of assigned and unflagged streams.
- Number of flagged streams.
- Number of cartridges with the inhibit allocation flag set.
- Number of available off-cartridge links.
- Number of cartridges with available streams and no off-cartridge links.

The following output shows the format of a basic usage report.

ASUSE REPORT FILE

SUBFAMILY = 3 CSU = J FM = SYSTST

BASIC REPORT

COUNT OF CSU LOCATIONS RESERVED FOR SUB-FAMILY =	10
NUMBER OF CARTRIDGES IN THE SUB-FAMILY =	4
NUMBER OF CARTRIDGES WITH 0 STREAMS AVAILABLE =	4
NUMBER OF CARTRIDGES WITH 1 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 2 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 3 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 4 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 5 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 6 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 7 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 8 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 9 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 10 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 11 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 12 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 13 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 14 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 15 STREAMS AVAILABLE =	0
NUMBER OF CARTRIDGES WITH 16 STREAMS AVAILABLE =	0
NUMBER OF FREE AND UNFLAGGED STREAMS =	0
NUMBER OF ASSIGNED AND UNFLAGGED STREAMS =	64
NUMBER OF FLAGGED STREAMS =	0
NUMBER OF CARTRIDGES WITH INHIBIT FLAG SET =	0
NUMBER OF AVAILABLE OFF-MSC LINKS =	3
NUMBER OF CARTRIDGES WITH AVAILABLE STREAMS AND NO OFF CARTRIDGE LINKS =	0

Optional Reports A and B

Optional reports A and B both include a title line, which identifies the subfamily, family, and CSU being reported on, and list the following information about cartridges. Optional report A includes a cartridge only if the number of streams available for assignment is within the range specified by the SL and SU parameters. Optional report B includes only those cartridges for which flags are set in the MSF catalog.

- Volume serial number of the cartridge.
- x and y coordinates of the cartridge location.
- Number of streams available on the cartridge.
- Number of off-cartridge links available on the cartridge.

In addition to the preceding items, optional report B lists one or more of the following messages depending on the flags set in the MSF catalog. Refer to the definition of MSF catalog earlier in this section for more information about these flags.

- EXCESSIVE PARITY ERRORS.
- LOST.
- INHIBIT SET.
- STREAM n START OF FRAGMENT.
- STREAM n FROZEN CHAIN.
- STREAM n STREAM CONFLICT.

The following output shows the format of optional report A.

ASUSE REPORT FILE

SUBFAMILY = 3 CSU =J FM = SYSTST

OPTIONAL REPORT(S) AB

VSN = V99004

X = 45

Y = 36

AVAILABLE STREAMS = 0

OFF CARTRIDGE LINKS = 0

VSN =V99005

X = 44

Y = 36

AVAILABLE STREAMS = 0

OFF CARTRIDGE LINKS = 0

VSN = V99006

X = 43

Y = 36

AVAILABLE STREAMS = 0

OFF CARTRIDGE LINKS = 0

VSN = V99007

X = 42

Y = 36

AVAILABLE STREAMS = 0

OFF CARTRIDGE LINKS = 3

END OF REPORT(S) AB

Optional Report C

Optional report C includes a title line, which identifies the CSU being reported on. The CSU number and y coordinate are printed at the top of the page. There is a new y coordinate on each page, from y=36 to y=0. x coordinates are written two to a line, from x=57 to x=0. The information listed in this report includes the following items.

- x coordinate
- code
- vsn
- family
- subfamily
- FCT link
- flags

Optional Reports D and E

Optional reports D and E both include a title line, which identifies the subfamily, family, and CSU being reported on. Both reports list the following cartridge information about entries in the MSF catalog.

- FCT ordinal of the MSF catalog entry.
- x and y coordinates of the location of the cartridge in the cubicle described by this entry.
- Volume serial number of the cartridge in the cubicle described by this entry.

In addition to the preceding information, optional report E lists the following information about the streams on the cartridges previously described.

- Off-cartridge links. Also referred to as Off-MSC (mass storage cartridge) links.
- Two words reserved for site use.
- Stream detail in octal.

Refer to the definition of MSF catalog earlier in this section for more information about the catalog entries.

The following output shows the format of optional report E.

ASUSE REPORT FILE

OPTIONAL REPORT E MSFCATALOG FOR SUBFAMILY 3, CSU J FM = SYSTST

ORD	X	Y	VSN	LINK1=	LINK2=	LINK3=	SITE=
1	45	36	V99004	2	2	2	000000000000000000/000000000000000000
			STREAM				
			00-03	20106	20300	22100	24101
			04-07	26102	20300	20007	20010
			10-13	20011	20012	20013	20014
			14-17	20015	20200	20300	20300
2	44	36	V99005	3	3	3	000000000000000000/000000000000000000
			STREAM				
			00-03	20006	20014	20200	22101
			04-07	24102	26103	20007	20010
			10-13	20011	20012	20200	20300
			14-17	20015	20016	20200	20300
3	43	36	V99006	4	4	4	000000000000000000/000000000000000000
			STREAM				
			00-03	20106	20011	20013	22001
			04-07	24102	26103	20007	20200
			10-13	20300	20200	20300	20014
			14-17	20015	20016	20017	20200
4	42	36	V99007	0	0	0	000000000000000000/000000000000000000
			STREAM				
			00-03	20106	20014	20016	20200
			04-07	20300	20300	20007	20010
			10-13	20011	20012	20200	20300
			14-17	20200	20300	20017	20200

ASDEBUG

ASDEBUG allows the analyst to update appropriate entries in the CSU maps and/or MSF catalogs and thereby resolve inconsistencies reported by the ASVAL utility. ASDEBUG can also copy data from selected MSF files or cartridges to disk. This function of ASDEBUG is intended for use when errors such as unrecoverable read errors do not allow successful staging of an MSF file.

NOTE

ASDEBUG updates the CSU map for the specified CSU. (Refer to the definition of CSU map earlier in this section.) It is recommended that the CSU map be copied on tape or on another device or family immediately after every update of the CSU map.

Input to ASDEBUG is via a directive file. Refer to the ASLABEL utility earlier in this section for a description of the directive file format.

The format of the ASDEBUG command is:

ASDEBUG,P1,P2,P3.

<u>P_i</u>	<u>Description</u>
I=filename	File on which directives are written.
I	Same as I=COMPILE.
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 ASDEBUG command. The I parameter is ignored.
Z omitted	Directives are contained on the file specified by the I parameter.

ASDEBUG Directives

The following directives are available with ASDEBUG. The descriptions of the parameters for these directives follow this section.

OP=RS - Read Streams

OP=RS reads selected streams of a cartridge in a specified drawer (D=d) or a cartridge identified by its vsn (V=vs_n) or x and y coordinates (XI and YI parameters). The CS parameter specifies the CSU where the cartridge resides. The range of streams to be read is specified by the SL and SU parameters. The streams are written to the file specified by the PF parameter.

OP=RF - Read File

OP=RF reads the file whose alternate storage address is specified by the CS, FO, and ST parameters. The file is written to the file specified by the PF parameter.

OP=RP - Release Frozen Chain Space

OP=RP clears flags in the MSF catalog and releases MSF space for the chain whose alternate storage address is specified by the CS, FO, and ST parameters.

OP=RL - Remove Cartridge Entry from MSF Catalog

OP=RL removes an MSF catalog entry that is not linked properly to the CSU map. The CS and FO parameters identify the MSF catalog entry to be removed.

OP=RC - Remove Cartridge Entry from CSU Map

OP=RC removes a CSU map entry that does not have a corresponding FCT entry in the MSF catalog. The CS, XI, and YI parameters identify the CSU map entry to be removed. The cartridge at the particular location (x,y) is put into the output drawer and the CSU map entry is changed to unassigned. The CSU map should be backed up immediately after the ASDEBUG run to ensure consistency with cartridge labels and MSF catalogs (refer to the definition of CSU map earlier in this section).

OP=CF - Change Flag in CSU Map or MSF Catalog

OP=CF sets or clears flags in the CSU map or MSF catalog. The FL and ON or OF parameters indicate the flag type and the action to be taken. The appropriate FCT or CSU map entry is determined from the FO, V, or XI and YI parameters. For FCT stream detail flags (FL=SF, FC, or SC), the user can specify either a single stream via the ST parameter or a range of streams via the SL and SU parameters.

Parameters for the ASDEBUG Directives

The descriptions of the parameters to the ASDEBUG directives follow. Not all parameters are valid with all directives, as indicated.

<u>P_i</u>	<u>Description</u>
CS=id	CSU identifier of the CSU to be used by ASDEBUG; id is a letter from A to M.
CS	Same as CS=A.
CS omitted	Same as CS=A.
D=d	Input drawer slot to be used; $0 \leq d \leq 7$. Not valid for OP=RS if V=vsn or XI=x is specified.
D	First available input drawer slot is to be used.
D omitted	V=vsn or XI=x and YI=y must be specified for OP=RS.

<u>P_i</u>	<u>Description</u>																		
FL=flag	Name of flag in CSU map or MSF catalog to be set or cleared. Valid only for OP=CF.																		
	<table border="0" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th style="text-align: center;"><u>flag</u></th> <th style="text-align: center;"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>EW</td> <td>Excessive write parity error flag.</td> </tr> <tr> <td>FC</td> <td>Frozen chain flag.</td> </tr> <tr> <td>FE</td> <td>Linkage error flag (in MSF catalog FCT).</td> </tr> <tr> <td>IB</td> <td>Inhibit allocation flag.</td> </tr> <tr> <td>LC</td> <td>Lost cartridge flag.</td> </tr> <tr> <td>ME</td> <td>Linkage error flag (in CSU map).</td> </tr> <tr> <td>SC</td> <td>Stream conflict flag.</td> </tr> <tr> <td>SF</td> <td>Start of fragment flag.</td> </tr> </tbody> </table>	<u>flag</u>	<u>Description</u>	EW	Excessive write parity error flag.	FC	Frozen chain flag.	FE	Linkage error flag (in MSF catalog FCT).	IB	Inhibit allocation flag.	LC	Lost cartridge flag.	ME	Linkage error flag (in CSU map).	SC	Stream conflict flag.	SF	Start of fragment flag.
<u>flag</u>	<u>Description</u>																		
EW	Excessive write parity error flag.																		
FC	Frozen chain flag.																		
FE	Linkage error flag (in MSF catalog FCT).																		
IB	Inhibit allocation flag.																		
LC	Lost cartridge flag.																		
ME	Linkage error flag (in CSU map).																		
SC	Stream conflict flag.																		
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.																		
FO=ord	MSF catalog ordinal indicating the file to be read or the chain whose space is to be released. Not valid for OP=CF if V=vsn or XI=x is specified.																		
FO omitted	FO=ord must be specified for OP=RF, OP=RP, and OP=RL. V=vsn or XI=x and YI=y 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 MSF image (streams or file) is to be copied. Each stream copied is separated by an end-of-record. This file is defined under the user's current family and user index.																		
PF	Same as PF=ZZZZBUG.																		
PF omitted	Same as PF=ZZZZBUG.																		
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	Stream where OP=RS begins copying or OP=CF begins changing flags; $0 \leq i \leq 15$; $i \leq j$ (refer to SU=j).																		

<u>P_i</u>	<u>Description</u>
SL	Same as SL=0.
SL omitted	Same as SL=0.
ST=s	Stream where OP=RF begins reading or OP=RP begins releasing. If specified with OP=CF, ST=s indicates a single stream detail to be changed and takes precedence over the SL and SU parameters.
ST omitted	ST=s must be specified for OP=RF and OP=RP.
SU=j	Stream where OP=RS ends copying or OP=CF ends changing flags; $0 \leq j \leq 15$; $i \leq j$ (refer to SL=i).
SU	Same as SU=15.
SU omitted	Same as SU=15.
V=vsn	Volume serial number of the cartridge to be used; not valid for OP=RS if D=d, D, XI=x, or YI=y is specified; and not valid for OP=CF if FO=ord, XI=x, or YI=y is specified.
V omitted	D=d or D, or XI=x and YI=y must be specified for OP=RS. FO=ord or XI=x and YI=y must be specified for OP=CF.
XI=x	x coordinate of the cubicle where the cartridge to be read resides; $0 \leq x \leq 57$ and $x \neq 30$. YI=y must also be specified. D=d, D, or V=vsn must not be specified with OP=RS. FO=ord or V=vsn must not be specified with OP=CF.
XI omitted	D=d, D, or V=vsn must be specified for OP=RS. FO=ord or V=vsn must be specified for OP=CF. XI=x and YI=y must be specified for OP=RC.
YI=y	y coordinate of the cubicle where the cartridge to be read resides; $0 \leq y \leq 36$ and $y \neq 18$. XI=x must also be specified. D=d, D, or V=vsn must not be specified with OP=RS. FO=ord or V=vsn must not be specified with OP=CF.
YI omitted	D=d, D, or V=vsn must be specified for OP=RS. FO=ord or V=vsn must be specified for OP=CF. XI=x and YI=y must be specified for OP=RC.

Restrictions to ASDEBUG

The following restrictions apply to the ASDEBUG utility.

- MSSEEXEC must be running when ASDEBUG is run.
- Only one copy of ASDEBUG can be run at one time.
- ASDEBUG, ASVAL, and ASLABEL cannot be run at the same time.

MSS OPERATIONAL PROCEDURES

The operational procedures used in an MSS environment are summarized next. These include modifications to existing procedures and also new MSS 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

MSS processing is initiated when the ASDEF utility is run to create the system files necessary to allow storage of files on MSF. ASDEF creates both a CSU map for each CSU that is to store file data on cartridges and eight MSF catalogs (one for each subfamily) for each family that is to have MSF-resident files. Once these system files have been created, the ASLABEL utility can then be run to allocate portions of the CSU to each subfamily and to initialize cartridges for use by files belonging to these subfamilies. When permanent files are defined, they reside on disk until the ASMOVE utility is run. Depending on the parameters specified, ASMOVE can destage files to MSF 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 MSS environment than in a non-MSS environment. This is necessary because MSF-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 ASMOVE utility is used periodically to release disk space.

The ASMOVE utility destages selected files to MSF and, optionally, releases their disk space. It is recommended that the site analyst use ASMOVE 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, under File Dumping, 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 ASMOVE utility with the RD and RT parameters specified should be run. ASMOVE reads the PFC entries for the specified family, selects certain files to reside on MSF, and decides which files should be released from disk.

Files can be destaged to MSF without their disk space being released. This destage-only feature can be used prior to a routine backup dump to create MSF 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, under File Dumping, later in this section.) Having an MSF 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 ASMOVE is run again. That is, ASMOVE can directly release disk space of files that have both disk and MSF images but are to reside on MSF only. However, for files that have no MSF image, there is a delay before the disk space can be released because these files must first be destaged to MSF.

MSF SPACE MANAGEMENT

The availability of MSF space must be managed because MSF 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 an MSF-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, MSF 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 MSF only and typically not on any dump tape.

An MSF-resident file is stored only on the cartridges within the cubicles of CSUs assigned to a subfamily. When destaging a file to MSF, MSSEXEC selects the CSU with the most space available. If no CSU has enough free space to store the file being destaged, MSSEXEC abandons that destage request, issues a message to the dayfile, and continues to destage other files that fit on the available MSF space. To avoid this problem, use the ASVAL utility periodically to release MSF space for purged files, thereby making it available for reuse.

When you want to reuse MSF space currently allocated to files that have been purged, run ASVAL 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 MSF catalogs and PFC entries that were current at the time of the dump. ASVAL examines the RDF to identify unneeded MSF files (those without PFC entries) and their MSF 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 ASVAL to release MSF space, you should discard it and any other RDFs for the same set of devices. These RDFs cannot be used as input to ASVAL in any future run. The MSS software does not permit an RDF to be specified by the RF parameter on an ASVAL call, unless it was produced by PFDUMP after the last ASVAL run that released MSF space was completed. This prevents the inadvertent release of MSF 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 MSF). If the following call to ASVAL is made, the MSF space assigned to AFILE is released and made available for reuse, if no errors were encountered.

```
ASVAL,RF=RDF1,RL.
```

Now, suppose that a new file, BFILE, which did not exist when RDF1 and RDF2 were produced, is assigned to the MSF space on which AFILE has resided. If ASVAL were to accept either RDF1 or RDF2 on a subsequent ASVAL call, the MSF space allocated to file AFILE (which is now the space where BFILE resides) is released. However, any RDF created after the ASVAL run that purged AFILE would not have an entry for AFILE and, hence, ASVAL would not inadvertently release the MSF space presently allocated to BFILE.

After ASVAL uses a release data file to release MSF 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 MSF space may have been reassigned. The MSS 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 ASVAL call releases MSF space allocated to CFILE.

```
ASVAL,RF=RDF2,RL.
```

Now assume 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 MSF. 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 MSS environment are similar to those in a non-MSS 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 MSF-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 MSF, 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 MSF-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 MSF files without disk images are included in the dump. It is unnecessary and not feasible to dump the files that reside only on MSF. Specify OP=S and SD on the call to PFDUMP; use the RD=filename parameter at your own discretion.

<u>Parameter</u>	<u>Description</u>
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 ASMOVE is run again. Refer to Release Algorithm, under ASMOVE, earlier in this section.
RD=filename	Specifies that PFDUMP is to write a release data file to be used later by ASVAL. RD=filename is included or excluded on a call to PFDUMP depending on how often you want to purge unneeded MSF 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.

<u>Parameter</u>	<u>Description</u>
OP=M	Selects incremental dumps.
SD	Updates the date/time field on each master device dumped, as previously described for full dumps.
AD=yyymmdd and AT=hhmmss	Selects the date and time 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.

<u>Parameter</u>	<u>Description</u>
OP=S	Selects no staging. If the SD option was inadvertently used on any recent partial dump, OP=S should not be specified so that any file with a BR=Y attribute that had its disk space inadvertently released will be in the dump.
RD=filename	Specifies that PFDUMP is to write a release data file to be used later by ASVAL, as previously described for full dumps.

Dump Tape Management

In a non-MSS environment, it is possible to make all previous dump tapes available for use after a full dump is completed. However, in an MSS environment, such dump tape recycling cannot occur because files that have resided only on MSF 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 the OP=M, BD, and BT parameters, but not specifying the OP=S parameter. † 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 MSF; 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.

The number of dump tapes created can be reduced by using ASMOVE to reduce the amount of data written on these dump tapes. If ASMOVE is run prior to an incremental dump, files with the BR=MD attribute can be destaged to MSF and excluded from the dump. That is, the incremental dump will exclude the files with a BR=MD attribute that are on both MSF and disk. If ASMOVE 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 ASMOVE runs can be used for efficiency.

1. Run ASMOVE to destage files with the BR=MD attribute.
2. Take an incremental dump to back up files with the BR=Y attribute.
3. Run ASMOVE 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.

CSU Map/MSF Catalog Backup

A CSU map is created for each CSU in an MSS environment and is updated by the ASLABEL, ASVAL, and ASDEBUG utilities. (Refer to the definition of CSU map earlier in this section.) It is recommended that the CSU map be copied on tape or on another device or family immediately after every update of the CSU map. If the device containing subfamily 0 of the default family is reloaded, the latest CSU maps should be recovered from the backup copies. After recovering the CSU maps, run the ASVAL utility with the AM and FM parameters specified to ensure that the CSU map and MSF catalog entries match. If they do not, corrective action should be taken as described under Error Conditions and Corrective Actions later in this section. Failure to retain or recover the latest CSU maps can cause problems with future ASLABEL runs (refer to the definition of CSU map earlier in this section).

An MSF catalog is created for each subfamily of a family that can have MSF-resident files and is updated by the ASLABEL, ASMOVE, ASVAL, and ASDEBUG utilities. When PFDUMP copies permanent files to tape for a master device or all devices in a family, it ensures that corresponding MSF catalogs are also dumped. (If MSSEXEC has an MSF catalog attached in modify mode, PFDUMP interfaces with MSSEXEC so that the MSF catalog is included on the dump tape. While PFDUMP is accessing the MSF catalog, any file destage requests for that subfamily are delayed.) Consequently, if PFLoad is used to reload all files, the MSF catalogs are automatically recovered and entries for all MSF-resident files reflected in recovered PFC entries will exist in the recovered MSF catalog.

† 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 back to the parameters for PFDUMP). This performs a dump of all files last dumped during the interval specified (including MSF-resident files) and allows your site to recycle all dump tapes created during this interval.

If PFLOAD is run after the ASLABEL utility has updated the MSF catalog and CSU map but before the MSF catalog has been backed up, the recovered MSF catalog will not match exactly the corresponding CSU map. In this case, recovery procedures (as described in MSF Catalog/CSU Map Mismatch later in this section) must be taken to bring the CSU map, MSF 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 MSF-resident only files, PFLOAD reloads only the PFC and permit information (if PFDUMP dumped only this information). The following situations involve PFLOAD in an MSS 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 an MSF 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 an MSF image), the ASCVER 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 an MSF catalog and/or CSU map is made, ASVAL should be run with the AM and FM parameters specified. ASVAL examines the PFC entries for MSF-resident files, the MSF catalog, and the CSU map to detect any error conditions that exist.

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 MSF catalogs, CSU maps, PFC entries, or cartridge labels. The ASVAL utility can analyze the MSF catalogs, CSU maps, and PFC entries to detect these errors, and the ASLABEL 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 MSS system file, or by abnormal termination of MSSEXEC or MSS utility processing, which can interrupt the updating of an MSS system file.

Clearing Permanent File Error Flags

The permanent file error flags that have been set by MSSEXEC 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 ASVAL 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 MSF, the file data will be available on disk and can be copied to another file. ASVAL should then be run as described in step 2.
2. If the file cannot be attached, reload the file from a dump tape. To identify the PFDUMP run that saved a backup copy of the file, check the output from ASVAL 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, ASVAL should be run as follows:

```
ASVAL,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. ASVAL 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 MSF catalog or CSU map, but you can analyze and then correct the error by using the ASDEBUG utility. Before using ASDEBUG to correct the error, you should take an incremental dump so the old PFC entry with an asa value pointing to a problem MSF-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 ASDEBUG might recover the file data.
4. If the file cannot be recovered, purge it so there is no PFC entry pointing to a problem MSF-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 drawer of the CSU to which it is assumed to belong, and run ASLABEL with the OP=RS directive. If the cartridge label information agrees with the CSU map entry for the cubicle with the coordinates (x,y) written on the cartridge label, the cartridge will be restored to that cubicle and the MSF catalog entry for the cartridge is updated to clear the lost flag. Otherwise, the cartridge is put in the output drawer and the cartridge label information is included in the output from ASLABEL.
2. If the cartridge was not restored to its CSU cubicle, examine the label information reported by ASLABEL. If the cartridge is assigned to a different CSU, repeat step 1 using the correct CSU. If the family name is blank, the cartridge has a scratch label and can be assigned normally with the OP=AM directive to ASLABEL.

3. If steps 1 and 2 fail, examine the CSU map entry for the cubicle with the coordinates (x,y) 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 ASVAL to ensure the CSU map and MSF 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 ASLABEL to rewrite the cartridge label and thereby make a scratch cartridge. It might be useful to first run ASDEBUG with the OP=RS directive to preserve the data on the cartridge.

Removal of Faulty or Missing Cartridges

The following steps should be taken to remove a cartridge that has unrecoverable errors, or to release the CSU map entry for a cartridge that is missing.

For unallocated cartridges:

1. Attempt to remove the cartridge with the ASLABEL 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 ASLABEL OP=IB directive to inhibit further allocation of the cartridge.
2. Run ASVAL specifying the AM parameter. Examine the report for possible discrepancies between the CSU map, MSF catalog, and PFC entries related to this cartridge.
3. If no MSF catalog entry corresponds to the CSU map entry, use the OP=RC directive to ASDEBUG to remove the map entry.
4. If no discrepancy is reported by ASVAL, the affected files should be attached in write mode (if possible), reloaded using PFLOAD with the OP=Z parameter, or purged. Then ASVAL release processing should be done to release the MSF space on the cartridge, and the cartridge can be removed with the ASLABEL OP=RM directive.

MSF Catalog/CSU Map/PFC Problems

If error conditions exist in the MSF catalogs or CSU maps, the recovery procedure is to correct the appropriate entry rather than reload the entire file. Reloading other than the latest copy of the CSU 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 MSF 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 MSF catalogs, CSU maps, or PFC entries.

MSF Catalog/CSU Map Mismatch

The following examples of MSF catalog/CSU map mismatches can appear as error type 1 or 2 on the ASVAL 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 CSU map or MSF catalog.

- If a cartridge has an MSF catalog entry indicating that all AUs on the cartridge are available for allocation but is not identified in any CSU map entry (type 1 mismatch), use the OP=RL directive to ASDEBUG to remove the MSF catalog entry for the cartridge.
- If a cartridge is identified in a CSU map entry but has no MSF catalog entry (type 2 mismatch), use the OP=RC directive to ASDEBUG to update the CSU 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 output drawer.
- If a cartridge has an MSF 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 ASVAL to purge unneeded MSF files and thereby make the cartridge unallocated. Use the OP=RL directive to ASDEBUG to remove the MSF catalog entry for the cartridge.

NOTE

It sometimes takes two ASVAL/ASDEBUG runs to completely clear up type 1 or type 2 errors. Therefore, it is recommended that ASVAL be rerun with the AM parameter after an ASDEBUG run using OP=RC or OP=RL.

MSF Catalog Chain Problems

AUs that are allocated to an MSF-resident file are chained together in the MSF 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, an MSF-resident file can reside on several cartridges, and several disk accesses may be necessary to update its chain. If MSSEXEC is interrupted during such an update, the resultant chain will be incomplete. ASVAL 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 an MSF-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 MSF, 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 procedure for any of these bad chain conditions follows.

1. Run ASVAL with the FM parameter specified and use a sufficiently large FX=n value so that ASVAL updates the MSF catalog entries affected by the problem chains.
2. Recover the permanent files, if any, that have MSF images potentially affected by these error conditions. (The ASVAL 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 MSF chains from being recovered if a device reload is necessary.
4. Run ASDEBUG 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/MSF Catalog Mismatch

It is possible that an old version of an MSF 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 MSF catalog or can point to a problem chain. If the correct version of the MSF 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, ASVAL 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 ASVAL to clear the asa value.

PSEUDORELEASE

The pseudorelease capability allows the site analyst to experiment with MSS operational procedures in order to determine which procedures should be used normally, and to become adept at recognizing error conditions and performing appropriate recovery procedures. While employing the pseudorelease feature, the analyst runs no risk of losing any permanent files used in this checkout phase, as explained in the following paragraphs.

Pseudorelease mode is initiated by an ASMOVE run with TM=Y specified. ASMOVE selects files to be destaged and/or released from disk by using the destage and release algorithms as usual. The destage to MSF process is the same whether or not pseudorelease mode is in effect (refer to the description of ASMOVE earlier in this section for details on destaging). However, the procedures for disk space release vary depending on whether or not pseudorelease mode is enabled. Disk space is not actually released in pseudorelease mode. Rather, the AFPDR flag is set in the PFC entry for a file that has been selected for disk space release. This flag indicates that PFM should ignore the disk image for this file. Therefore, if the user attaches such a file, the disk image is ignored and the file is staged to disk from MSF. The normal staging process is followed except that the disk image is created on a scratch device instead of a permanent file device. Thus, in actuality, two disk images of the file exist. However, once the file is staged to disk, whether or not staging errors are detected and reported by MSSEXEC, the AFPDR flag is cleared so that the user accesses the original disk image. Also, the file staged to the scratch disk is released so the scratch space is available for reuse. If MSS processing becomes disabled, the original disk image can still be accessed using an ATTACH request.

There are two methods for terminating pseudorelease mode. One is to perform a full dump followed by a reload with OP=Z specified on the PFLoad call. This clears the asa field in the PFC entry for each file, thereby indicating that no MSF images exist and that the files reside on disk only. The other method is to run ASMOVE with TM=N specified. This causes the disk images for pseudoreleased files to be really released from disk, thereby leaving the file with only the MSF image.

You can use the following examples after pseudorelease mode is enabled to force different error conditions to occur. You can implement these incorrect operational procedures and then perform error recovery action without endangering any of the permanent files involved. These examples assume that ASDEF has been run to create a CSU map and the MSF catalogs for the CSU and family being tested.

Example 1:

1. Run ASLABEL to add cartridges (for example, A, B, C) to a subfamily.
2. Dump the CSU map file to tape.
3. Run ASLABEL to remove cartridge A and add another cartridge (for example, D).
4. Reload the CSU map file dumped in step 2. This creates a mismatch between the CSU map and MSF catalog, because the MSF catalog indicates that cartridge D resides in the CSU but cartridge A does not.
5. Run ASVAL to detect the mismatch.
6. Run ASDEBUG to correct the mismatch. The resultant CSU map and MSF catalog should contain entries for cartridges B and C, but not A or D.

Example 2:

1. Run ASLABEL to add several cartridges to a subfamily.
2. Create some permanent files with different backup requirements and preferred residence requirements.
3. Run ASMOVE to destage and/or release the files created in step 2.
4. Physically remove one of the cartridges from the CSU; do not use ASLABEL to remove it.
5. Attach the files whose disk space was released in step 3.
6. Run ASVAL to verify that only files resident on the removed cartridge are affected.
7. Run ASLABEL to restore the cartridge removed in step 4.
8. Attach the files reported in step 6. There should be no problems reported.

Example 3:

1. Run ASLABEL to add several cartridges to a subfamily. Make a backup copy of the CSU map and MSF catalog files.
2. Create several files with the BR=Y or BR=MD attribute. Run ASMOVE to destage all of them and release those with the BR=MD requirement.
3. Take a backup dump of the family.
4. Reload the CSU map and MSF catalog backed up in step 1.
5. Create several new files with the BR=MD attribute. Run ASMOVE to destage them and to release the disk space of the files with the BR=Y attribute in step 2.
6. Attach the files created in step 2. They should have been overwritten by the files destaged in step 5.
7. Recover the files with the BR=Y attribute from the dump tape made in step 3.

Example 4:

Try to interrupt MSSEXEC while it is destaging new files or releasing unused MSF files to create incomplete chains. ASVAL reports such errors and ASDEBUG fixes them.

Introduction	13-1
Linked SDM Operation	13-2
Deadstart	13-3
Shared Mass Storage	13-6
Configuring Shared Mass Storage	13-6
Device Access Tables	13-6
Mass Storage Recovery Tables	13-7
Device Initialization	13-7
Device Unload	13-8
Device Recovery	13-8
Device Checkpoint	13-9
MSS Processing	13-9
MSAS Processing	13-10
Independent SDM Operation	13-10
Deadstart	13-11
Shared Mass Storage	13-12
Configuring Shared Mass Storage	13-13
Device Index Tables	13-13
Mass Storage Recovery Tables	13-13
Device Initialization	13-14
Device Unload	13-14
Device Recovery	13-15
Device Checkpoint	13-15
MSS Processing	13-15
MSAS Processing	13-15

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), and independent SDM. 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 four computers can access mass storage devices. Any combination of 1 to 4 model 71, 72, 73, 74, 171, 172, 173, 174, 175, 720, 730, 740, 750, 760, 865, 875 mainframes or 6000 Computer Systems mainframes can be used. Each mainframe requires one CPU port connected to extended memory. The presence of a DDP on a CPU port decreases by one the total number of mainframes that can run together. Shared mass storage devices supported are: 844 (single and double density), 885-11/12, and extended memory.

Extended memory is the means for controlling shared mass storage and intermainframe communications. Each mainframe has a CPU 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 MSS is used, one mainframe in the linked SDM complex is physically connected to MSF and is called the master mainframe; all other mainframes in the complex can access MSF files and are called slave mainframes. MSSEXEC runs on the master mainframe and is the main processing program that controls MSS activity. MSSSLV runs on each slave mainframe and communicates with MSSEXEC to retrieve files from MSF in response to ATTACH requests by jobs running on the slave mainframes.

If MSAS 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 MSAS 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.

<u>CMRDECK Entry</u>	<u>Description</u>
MID=id.	The 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.
<u>EQPDECK Entry</u>	<u>Description</u>
SHARE=est ₁ ,est ₂ ,...,est _n .	est _i are the 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.
PRESET,n.	You can specify this entry only at the system console. n is the total number of devices to be shared by all mainframes in the complex. You must specify the PRESET entry on a level 0 deadstart on the first machine to be deadstarted in the linked SDM complex. This entry initializes the flag register and extended memory resident tables. If n is omitted, the default value is the number of shared devices defined for the particular machine.

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 multiframe 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 four 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. The indices are 1, 2, 3, and 4. The masks are 1, 2, 4, and 8.

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.

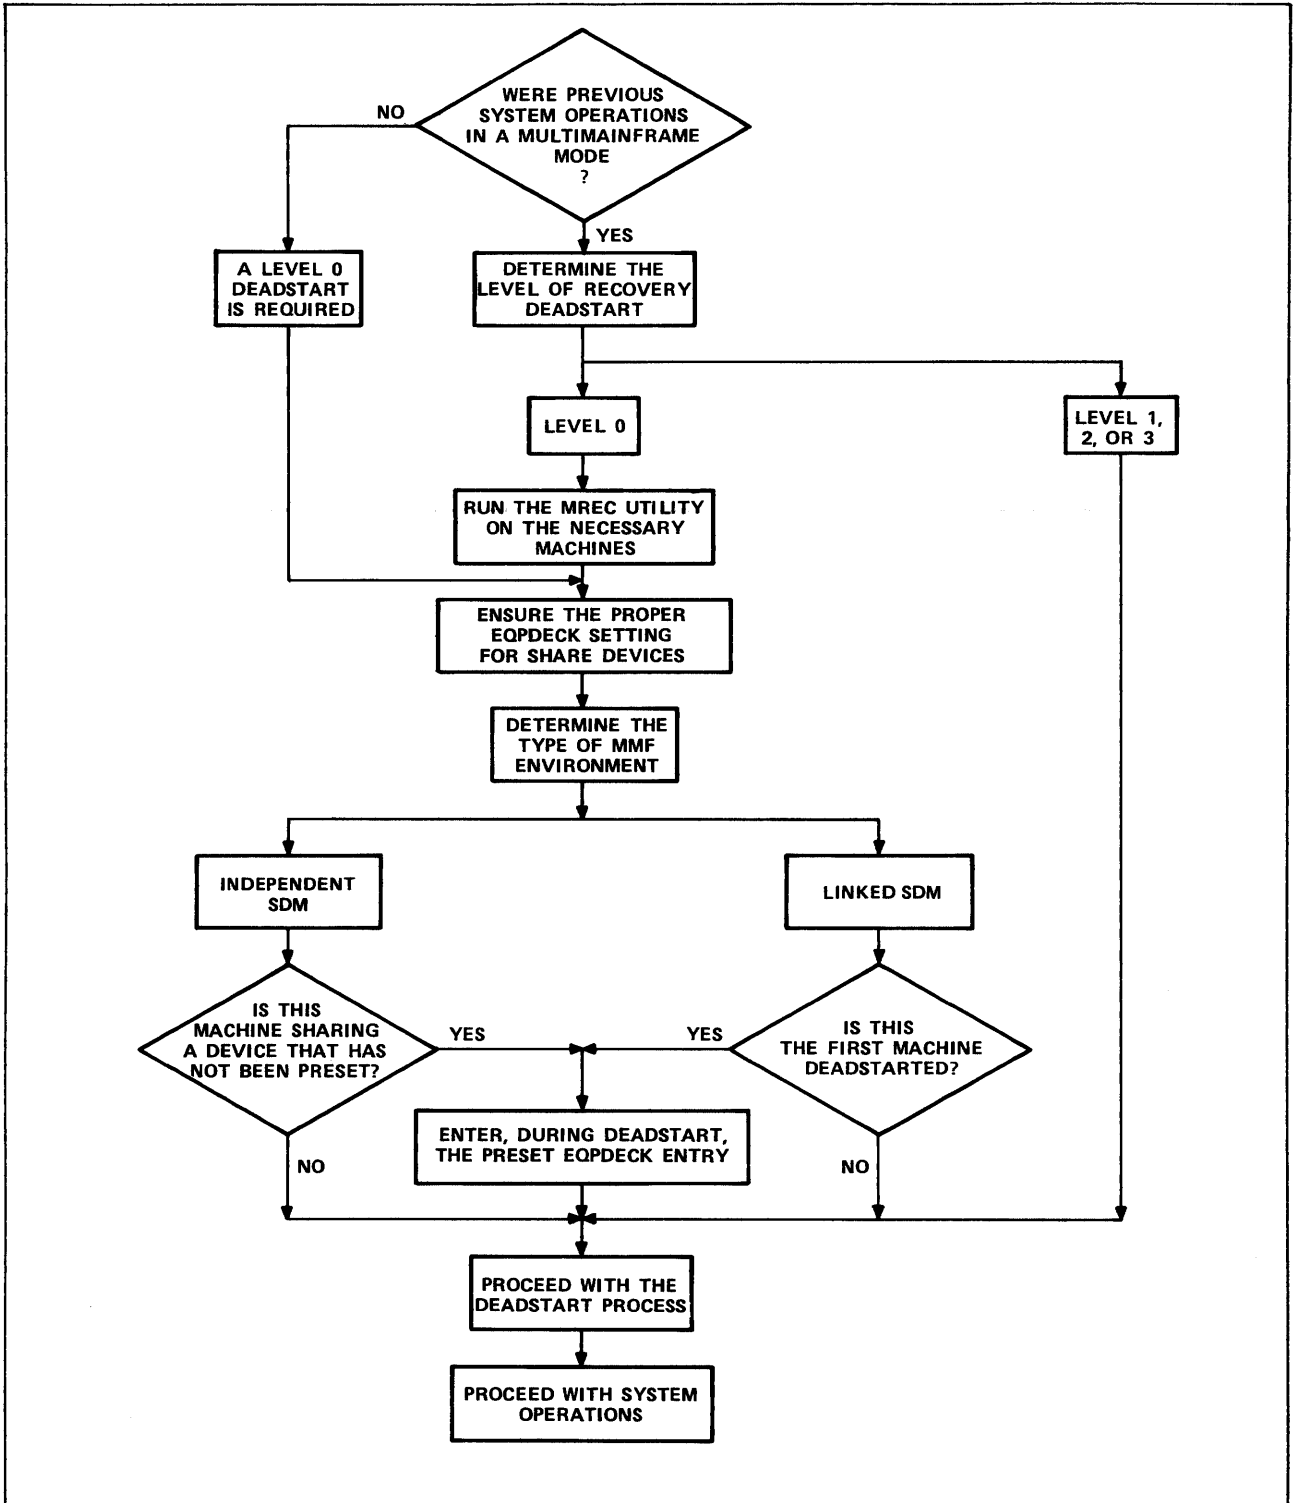


Figure 13-1. Deadstarting a Machine in a Multiframe Complex

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 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 (100g 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 for each machine resides in extended memory.

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 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 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 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 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.

MSS PROCESSING

Unlike most other mass storage devices, the MSF is not shared by mainframes. Rather, MSF 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 MSF-resident files are staged must be shared devices and accessible by all mainframes.

MSSEXEC runs on the master mainframe and is the main program that controls MSS processing activities. You can activate MSSEXEC by entering the following DSD commands at the system console.

```
L.SUBSYST.  
ENABLE,MASTER MSS.  
L.END.  
MSSffff.
```

MSSSLV runs on each slave mainframe and is the program that supports file retrieval from MSF for jobs running on a slave mainframe. You can activate MSSSLV by entering the following DSD commands at the system console.

```
L.SUBSYST.  
DISABLE,MASTER MSS.  
L.END.  
MSSffff.
```

When a job running on a slave mainframe attaches an MSF-resident file that has no disk image, MSSSLV is requested to stage the file to disk. However, MSSSLV cannot access MSF directly. MSSSLV must communicate with MSSEXEC, which stages the file to shared disk and then notifies MSSSLV when the file is available. MSSSLV then causes the requesting job to be resumed so that the file ATTACH can complete.

The permanent files used for communication between MSSEXEC and MSSSLV are initialized and recovered automatically and require no backup or recovery processing. The name of the permanent file used by MSSSLV running on mainframe *i* (*i*=1, 2, 3, 4) to send requests to MSSEXEC is STOMNO*i*. The name of the permanent file used by MSSEXEC to respond to all requests by all MSSSLVs is MTOSPFN. These files belong to the family that contains extended memory and are saved under the user index 377760g.

MSAS 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 MSAS processing activities. You can activate SSEXEC by entering the following DSD commands at the system console.

```
L.SUBSYST.  
ENABLE,MASTER MAS.  
L.END.  
MASffff.
```

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 MAS.  
L.END.  
MASffff.
```

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 377760g.

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 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 multi-mainframe 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 MSS or MSAS is used, one mainframe in the independent SDM complex is physically connected to MSF or 7990 and only that mainframe can assess 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.

<u>CMRDECK Entry</u>	<u>Description</u>
MID=id.	The 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.

<u>EQPDECK Entry</u>	<u>Description</u>
ISHARE=est ₁ ,est ₂ ,...,est _n .	est _i are the 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 .	You can specify this entry only at the system console. est _i is the EST ordinal of the shared mass storage device to be preset. 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.
- 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 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 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 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 multiframe 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.

MSS PROCESSING

Unlike most other mass storage devices, the MSF is not shared by mainframes. Rather, MSF is physically connected to and driven by one mainframe. The only machine that can access the MSF-resident files is the physically connected independent SDM. MSS processing is the same in an independent SDM environment as it is in a single mainframe environment. Refer to section 12, Mass Storage Subsystem.

MSAS 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. MSAS processing is the same in an independent SDM environment as it is in a single mainframe environment. Refer to section 11, Mass Storage Archival Subsystem.

NAD MAINTENANCE UTILITIES

14

Maintenance Logging Transfer Facility (MLTF)
Dump NAD Memory (DMPNAD)

14-1
14-1

This section describes utilities that are used to maintain network access devices (NADs).

MAINTENANCE LOGGING TRANSFER FACILITY (MLTF)

MLTF is a utility program that controls logging NAD error logs into the binary maintenance log. At regular intervals MLTF maintains the network description file on disk, calls the NETLOG program, and rolls out after execution. The time interval is determined by an input parameter to MLTF (refer to the NOS 2 Installation Handbook).

NETLOG uses the network configuration file to determine which remote NADs should be logged. Each time NETLOG runs it receives a copy of the network configuration file from RHF and saves the file on disk. If RHF is not active, NETLOG uses the saved copy of the network configuration file from disk to determine which NADs to log. Local NADs are always logged. For each NAD that is logged, NETLOG enters the NAD error log into the binary maintenance log.

The MLTF utility is started automatically by RHF when RHF is initialized. The mainframe identified as the maintenance host must have the MLTF utility running periodically even if RHF is not active.

DUMP NAD MEMORY (DMPNAD)

DMPNAD is a utility program that reads the NAD memory and formats the data into an output file. The dump is formatted so that each NAD word is printed in both hexadecimal and 7-bit ASCII code representations.

The format of the command is:

DMPNAD,p₁,p₂,...,p_n.

<u>P_i</u>	<u>Description</u>
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=cc	The 1- or 2-digit octal number specifying the channel number of the local NAD defined in the EST entry associated with the NAD.
I=infile	The 1- to 7-character name of the file on which the binary input is written. The default is no binary input file.

<u>Pi</u>	<u>Description</u>
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 online NAD is to be dumped. The dump can also be saved as a binary output file (B parameter) which can later be formatted for printing by using the file as a binary input file (I parameter) and no access to the NAD is required. No special privileges are required when using the I parameter.

For a local NAD dump, 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). 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 Loading Utility) 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.

Network Organization	15-1
NAM Startup	15-3
Initiating NAM Without Operator Intervention	15-3
Altering the Memory File	15-4
NAM Shutdown	15-5
Network Control by HOP	15-6
Disabling and Enabling of Couplers	15-6
Status Message Formats for HOP	15-6
Unsolicited Status Reports for HOP	15-7
HOP Commands	15-7
NS Control Commands	15-7
Cancel Alternate NPU Load File Command	15-8
Change NPU Load File Command	15-8
NPU Load Status Command	15-8
NS Recent History Command	15-11
NVF Control Commands	15-11
Disable Host Element Command	15-11
Enable Application Command	15-12
Idle Host Element Command	15-12
NVF Recent History Command	15-13
Status Host Element Command	15-14
Network Control by NOP	15-16
Becoming a NOP	15-16
Status Message Formats for NOP	15-18
Unsolicited Status Reports for NOP	15-20
NOP Commands	15-20
ALERT Command	15-20
CONTROL Command	15-21
Disable Network Element Command	15-22
Enable Network Element Command	15-24
NPU Go Command	15-26
NPU Load Command	15-27
NPU Memory Dump Command	15-28
Send Message Command	15-29
Network Control by DOP	15-30
Becoming a DOP	15-30
Status Message Formats for DOP	15-31
Unsolicited Status Reports for DOP	15-33
DOP Commands	15-34
CS Information Command	15-34
NPU Diagnostic Test Command	15-35
TST Drop Option	15-35
TST Msg Option	15-35
TST Request Option	15-36

(Continued)

(Continued)

Recent History Command	15-36
Report Unsolicited Status Command	15-37
Send Message Command	15-38
Status Network Element Command	15-39
Summary of Network Operation Commands	15-42
Control Commands for HOP	15-42
NAM K Display	15-42
NAM Startup	15-43
NAM Shutdown	15-43
NAM Mode Commands	15-43
NS Control Commands	15-44
NVF Control Commands	15-44
Control Commands for HOP and NOP	15-44
Control Commands for HOP, NOP, and DOP	15-45

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 (NPUs), trunks, lines, terminals, and application programs under network supervision.

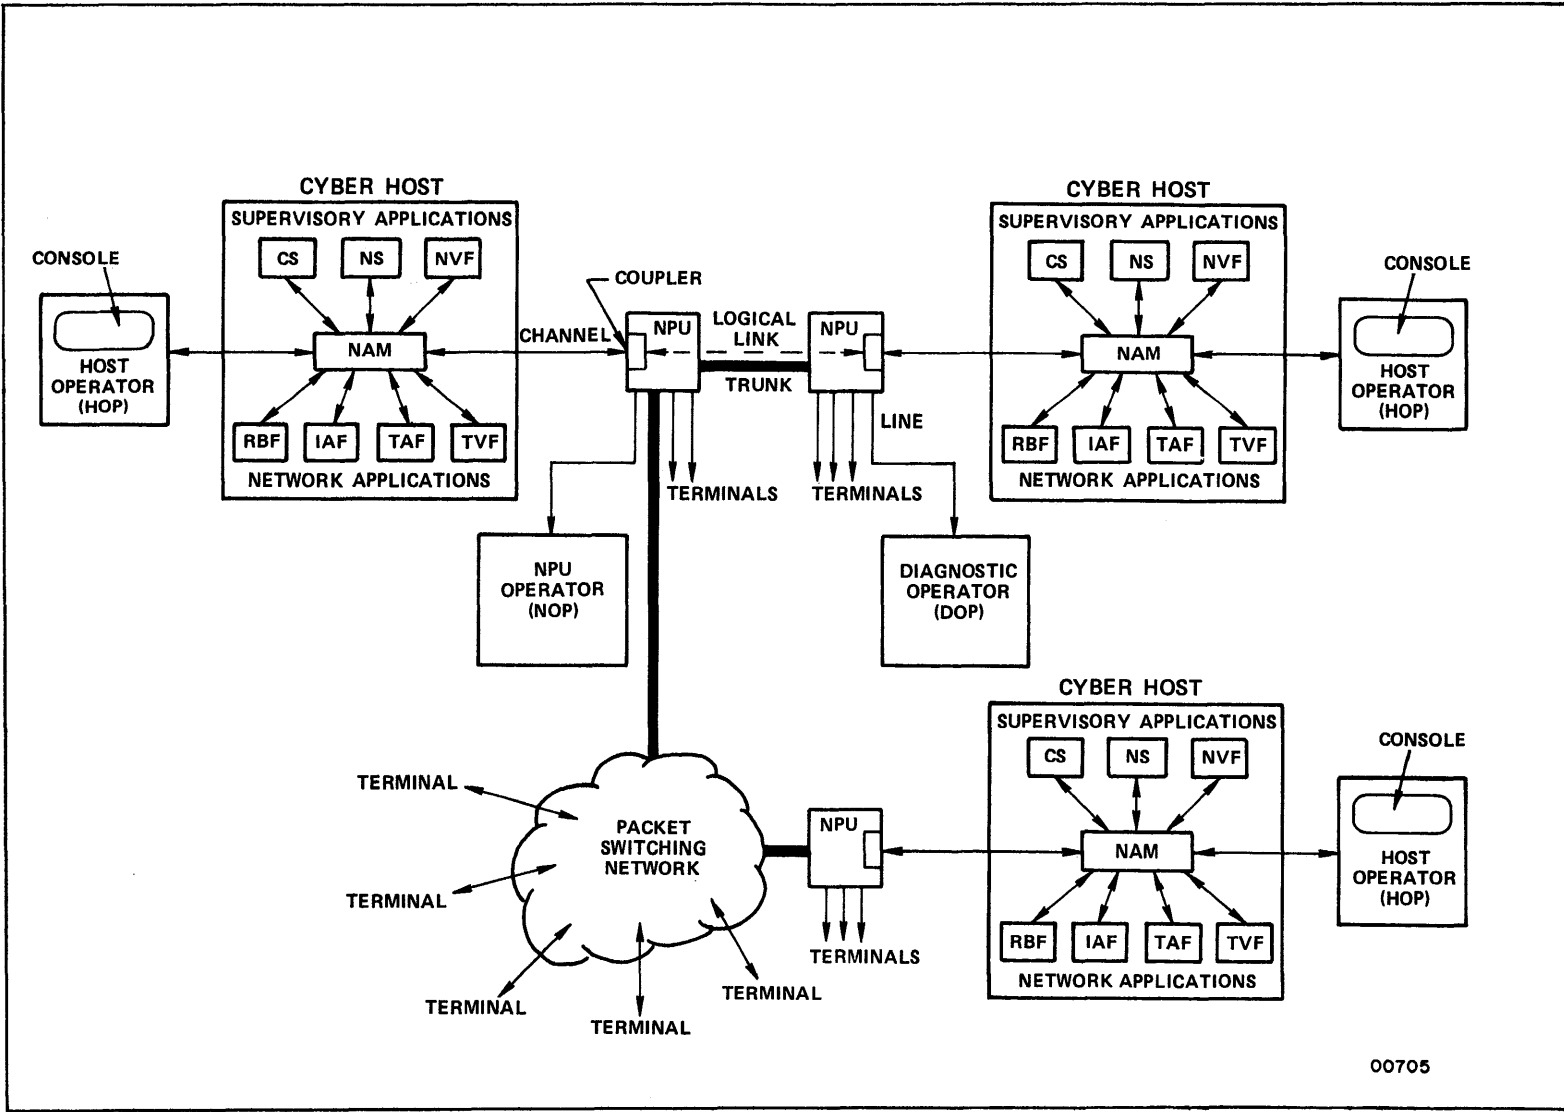
NETWORK ORGANIZATION

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 NPUs and all of their communication lines and terminals. NS loads and dumps all of the NPUs 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. 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.

The 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.



00705

Figure 15-1. Sample Network

The network recognizes three types of network operators: the host operator, the NPU operator, and the diagnostic operator. A host operator (HOP) resides at the system console, initiates NAM, and can control network elements being supervised by this host. An NPU operator (NOP) resides at a terminal and controls communication 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.

The HOP can always perform NOP and DOP functions regardless of the existence of 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, or NVF) executing in its host. 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.

NAM 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 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.†

NAM is brought up without further operator action. This results in starting the network using the previously specified master file and the parameter record RESTRT.

† 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.keyword₁=param₁,keyword₂=param₂,...,keyword_n=param_n.

The following keyword and parameter combinations are allowed on the CFO command.

<u>Parameter</u>	<u>Description</u>
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.
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.
RN=recordname †	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 from which dumps/traces are collected. If this parameter is not specified, the current network invocation number minus 1 is used.

† Specifying a record name changes recordname only for this initiation.

Table 15-1. Parameter Record Feature Descriptions

Feature Descriptions	Parameter Record Name					
	INIT	RESTRT	RECOVR	MULTI	MINIT	MRECOV
This feature collects dumps/traces on tape before startup.	N †	Y ††	N	Y	N	N
This feature stops local NPUs at startup.	Y	N	N	N	Y	N
This feature stops local NPUs at termination.	Y	Y	N	N	N	N
This feature takes host dumps when NPU fails.	Y	Y	Y	N	N	N
This feature takes NPU dumps before initial load.	N	Y	N	Y	N	N

† N indicates that the parameter record name causes the action not to occur.
†† Y indicates that the parameter record name causes the action to occur.

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.

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 the NPU 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 a 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 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.

- 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.

<u>Variable</u>	<u>Description</u>
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.

<u>Variable</u>	<u>Description</u>										
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="0"> <thead> <tr> <th><u>status</u></th> <th><u>Description</u></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>	<u>status</u>	<u>Description</u>	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.
<u>status</u>	<u>Description</u>										
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.

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.

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.

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.

<u>Parameter</u>	<u>Description</u>
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 specified username.

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.

The first form requests the current load/dump status of a specific NPU.

The command format is:

STATUS,NPU=npuname.

or

ST,NP=npuname.

The second form requests the current load/dump status of all NPUs.

The command format is:

STATUS,NPUS.

or

ST,NPS.

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
```

<u>Status</u>	<u>Description</u>								
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><thead><tr><th><u>dflag</u></th><th><u>Description</u></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>	<u>dflag</u>	<u>Description</u>	DUMP	Dump NPU before a load.	NODUMP	No NPU dump before a load.	blank	NPU is not currently being loaded.
<u>dflag</u>	<u>Description</u>								
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).								
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.								

Status

Description

message

Message is the NS activity and its associated status.

<u>Activity</u>	<u>Status</u>
SAMLOAD	STARTED COMPLETED ABORTED - error
DUMP/LOAD DUMP	REQUESTED STARTED COMPLETED ABORTED-error
LOAD	modname COMPLETED ABORTED - error

modname Name of the system module being loaded.

error One of the following:

<u>error</u>	<u>Description</u>
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:

operation

Keyword Abbreviation

APPL=applicationname

AP

HOST

HO

applicationname

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.

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 network.
- If the application is active at the time the command is entered, NAM requests the application to immediately terminate access to the 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 the network.
- When all applications (including CS and NS) terminate their access to the 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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>		
operation	One of the following:		
		<u>operation</u>	<u>Keyword Abbreviation</u>
		APPL=applicationname	AP
		HOST	HO
		applicationname	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.

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 the network are rejected.
- New application to application connections are rejected.
- New terminal connections to the host 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.
- NS ignores any new dump or load requests from NPUs. Any ongoing dump and load processes are completed normally before NS terminates.
- When all applications (including NS and CS) 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.

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:

element

Keyword Abbreviation

APPL=elementname

AP

TERM=elementname

TE

UNAME=elementname

UN

elementname

Name of the element whose status is desired. If 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.

The second form requests the status of all network applications in the host.

The command format is:

STATUS,APPLS.

or

ST,APS.

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.

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.

<u>Parameter</u>	<u>Description</u>
termelement	One of the following:
	<u>termelement</u>
	APPL=elementname
	UNAME=elementname
	elementname
	<u>Keyword Abbreviation</u>
	AP
	UN
	Name of the desired application or user.

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.

<u>Parameter</u>	<u>Description</u>
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.

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 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 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,† 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

† It is suggested that a hardcopy device be used 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.

- 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/cc,/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.

<u>Variable</u>	<u>Description</u>
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).

<u>Variable</u>	<u>Description</u>												
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.												
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><u>status</u></th> <th><u>Description</u></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>	<u>status</u>	<u>Description</u>	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.
<u>status</u>	<u>Description</u>												
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.

<u>Parameter</u>	<u>Description</u>
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:

<u>element</u>	<u>Keyword Abbreviation</u>
NPU=npuname	NP
NPUS	NPS
npuname	Name of the NPU supervised by your CS.

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.

<u>Parameter</u>	<u>Description</u>								
option	One of the following:								
	<table><thead><tr><th><u>option</u></th><th><u>Keyword Abbreviation</u></th></tr></thead><tbody><tr><td>NPU=npuname</td><td>NP</td></tr><tr><td>NPUS</td><td>NPS</td></tr><tr><td>AUTO</td><td>AU</td></tr></tbody></table>	<u>option</u>	<u>Keyword Abbreviation</u>	NPU=npuname	NP	NPUS	NPS	AUTO	AU
<u>option</u>	<u>Keyword Abbreviation</u>								
NPU=npuname	NP								
NPUS	NPS								
AUTO	AU								
	npuname Name of the NPU.								
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 two forms of the DISABLE command. Refer to Status Message Formats for NOP, earlier in this section, for the format of the status line returned.

The first form disables a single element in a network.

The command format is:

DISABLE,element.

or

DI,element.

<u>Parameter</u>	<u>Description</u>
element	One of the following:

<u>element</u>	<u>Keyword Abbreviation</u>
LINE=elementname	LI
LLINK=elementname	LL
TERM=elementname	TE
TRUNK=elementname	TR

elementname	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.

The second form disables archetype 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.

<u>Parameter</u>	<u>Description</u>
atname	Archetype SVC name.
numcir	Number of SVCs to change from an enabled state to a disabled state.

The second 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.

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.

The NPUs 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.

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 five 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.

The first form enables a single element in the network.

The command format is:

ENABLE,element.

or

EN,element.

Parameter

Description

element

One of the following:

element

Keyword Abbreviation

LINE=elementname

LI

LLINK=elementname

LL

TERM=elementname

TE

TRUNK=elementname

TR

elementname

Name of the element to be enabled.

The second 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.

The third form enables all disabled elements of the specified type on all NPUs.

The command format is:

ENABLE,element,NPUS.

or

EN,element,NPS.

<u>Parameter</u>	<u>Description</u>										
operation	One of the following:										
	<table><thead><tr><th><u>element</u></th><th><u>Keyword Abbreviation</u></th></tr></thead><tbody><tr><td>LINES</td><td>LIS</td></tr><tr><td>LLINKS</td><td>LLS</td></tr><tr><td>TERMS</td><td>TES</td></tr><tr><td>TRUNKS</td><td>TRS</td></tr></tbody></table>	<u>element</u>	<u>Keyword Abbreviation</u>	LINES	LIS	LLINKS	LLS	TERMS	TES	TRUNKS	TRS
<u>element</u>	<u>Keyword Abbreviation</u>										
LINES	LIS										
LLINKS	LLS										
TERMS	TES										
TRUNKS	TRS										

The fourth form enables all disabled terminals on a line.

The command format is:

ENABLE,TERMS,LINE=termline.

or

EN,TES,LI=termline.

<u>Parameter</u>	<u>Description</u>
termline	Name of the line to which terminals are connected.

The fifth form enables SVCs that are in a disabled state.

The command format is:

ENABLE,SVC=atname,NCIR=numcir.

or

EN,SV=atname,NC=numcir.

<u>Parameter</u>	<u>Description</u>
atname	Archetype SVC name.
numcir	Number of SVCs to change from a disabled state to an enabled state.

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.

<u>Parameter</u>	<u>Description</u>
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.

The first form causes the specified NPU to be stopped, reloaded, and reactivated.

The command format is:

LOAD,NPU=npuname,dumppoption.

or

LO,NP=npuname,dumppoption.

<u>Parameter</u>	<u>Description</u>
npuname	Name of the NPU.
dumppoption	Optional dump parameter; one of the following:

<u>dumppoption</u>	<u>Keyword Abbreviation</u>
DUMP	DU
NDUMP	ND

If the DUMP option is specified, a memory dump of the NPU is taken prior to reloading. If the NDUMP 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 NDUMP 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.

The second form causes all active NPUS to be stopped, reloaded, and reactivated.

The command format is:

LOAD,NPUS,dumppoption.

or

LO,NPS,dumppoption.

<u>Parameter</u>	<u>Description</u>
dumppoption	Optional parameter; one of the following:

<u>dumppoption</u>	<u>Keyword Abbreviation</u>
DUMP	DU
NDUMP	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 NDUMP 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 NDUMP 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.

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.

<u>Parameter</u>	<u>Description</u>
npuname	Name of the NPU.
status	ON or OFF.

The second form causes a memory dump of all NPUs to be taken.

The command format is:

DUMP,NPUS,status.

or

DU,NPS,status.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
------------------	--------------------

element	One of the following:
---------	-----------------------

<u>element</u>	<u>Keyword Abbreviation</u>
----------------	-----------------------------

LINE=elementname	LI
------------------	----

LLINK=elementname	LL
-------------------	----

NPU=elementname	NP
-----------------	----

TERM=elementname	TE
------------------	----

HOST	HO
------	----

NPUS	NPS
------	-----

elementname	Name of the element to which the message is addressed.
-------------	--

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, in order 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. 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 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 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,† 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

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.

† It is suggested that a hardcopy device be used for DOP terminals; this prevents the paging problem and provides a permanent log of network events.

The following are the general formats for these status messages.

- 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/cc,/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.

<u>Variable</u>	<u>Description</u>
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.

<u>Variable</u>	<u>Description</u>
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.
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.

<u>status</u>	<u>Description</u>
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.

<u>Variable</u>	<u>Description</u>
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 commands 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.

<u>Parameter</u>	<u>Description</u>								
npuname	Name of the NPU containing the diagnostic program with which the operator wants to communicate.								
option	One of the following:								
	<table><thead><tr><th><u>option</u></th><th><u>Keyword Abbreviation</u></th></tr></thead><tbody><tr><td>DROP</td><td>DR</td></tr><tr><td>MSG=message</td><td>MS</td></tr><tr><td>REQUEST</td><td>RE</td></tr></tbody></table>	<u>option</u>	<u>Keyword Abbreviation</u>	DROP	DR	MSG=message	MS	REQUEST	RE
<u>option</u>	<u>Keyword Abbreviation</u>								
DROP	DR								
MSG=message	MS								
REQUEST	RE								
message	Command that controls the execution of the diagnostic program in the NPU.								

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 Msg 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

    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.

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

<u>Parameter</u>	<u>Description</u>								
element	One of the following:								
	<table><thead><tr><th><u>element</u></th><th><u>Keyword Abbreviation</u></th></tr></thead><tbody><tr><td>NPU=npuname</td><td>NP</td></tr><tr><td>NPUS</td><td>NPS</td></tr><tr><td>npuname</td><td>Name of the NPU.</td></tr></tbody></table>	<u>element</u>	<u>Keyword Abbreviation</u>	NPU=npuname	NP	NPUS	NPS	npuname	Name of the NPU.
<u>element</u>	<u>Keyword Abbreviation</u>								
NPU=npuname	NP								
NPUS	NPS								
npuname	Name of the NPU.								
option	ON or OFF. If option is not specified, ON is assumed. If option is ON, CS reports all unsolicited status messages until you turn reporting off.								

By default, HOPs receive reports for all NPUs supervised by a CS, and 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

<u>Parameter</u>	<u>Description</u>						
element	One of the following:						
	<table><thead><tr><th><u>element</u></th><th><u>Keyword Abbreviation</u></th></tr></thead><tbody><tr><td>TERM=termname</td><td>TE</td></tr><tr><td>HOST</td><td>HO</td></tr></tbody></table>	<u>element</u>	<u>Keyword Abbreviation</u>	TERM=termname	TE	HOST	HO
<u>element</u>	<u>Keyword Abbreviation</u>						
TERM=termname	TE						
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 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.

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:

element

Keyword Abbreviation

CPLER=elementname CP

LINE=elementname LI

LLINK=elementname LL

NPU=elementname NP

SVC=elementname SV

TERM=elementname TE

TRUNK=elementname TR

elementname

Name of the element whose status is desired.

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.

<u>Parameter</u>	<u>Description</u>														
linkelement	One of the following:														
	<table border="0"> <thead> <tr> <th><u>linkelement</u></th> <th><u>Keyword Abbreviation</u></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>	<u>linkelement</u>	<u>Keyword Abbreviation</u>	CPLERS	CPS	LINES	LIS	LLINKS	LLS	SVCS	SVS	TERMS	TES	TRUNKS	TRS
<u>linkelement</u>	<u>Keyword Abbreviation</u>														
CPLERS	CPS														
LINES	LIS														
LLINKS	LLS														
SVCS	SVS														
TERMS	TES														
TRUNKS	TRS														
npuelement	One of the following:														
	<table border="0"> <thead> <tr> <th><u>npuelement</u></th> <th><u>Keyword Abbreviation</u></th> </tr> </thead> <tbody> <tr> <td>NPU=npuname</td> <td>NP</td> </tr> <tr> <td>NPUS</td> <td>NPS</td> </tr> </tbody> </table> <p style="margin-left: 100px;">npuname Name of the NPU.</p>	<u>npuelement</u>	<u>Keyword Abbreviation</u>	NPU=npuname	NP	NPUS	NPS								
<u>npuelement</u>	<u>Keyword Abbreviation</u>														
NPU=npuname	NP														
NPUS	NPS														
statuscode	Optional. If specified, only the elements with that status are reported.														
	<table border="0"> <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														

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.

<u>Parameter</u>	<u>Description</u>										
mode	One of the following:										
	<table><thead><tr><th><u>mode</u></th><th><u>Keyword Abbreviation</u></th></tr></thead><tbody><tr><td>LINE=elementname</td><td>LI</td></tr><tr><td>LLINK=elementname</td><td>LL</td></tr><tr><td>elementname</td><td>Name of the specified element.</td></tr></tbody></table>	<u>mode</u>	<u>Keyword Abbreviation</u>	LINE=elementname	LI	LLINK=elementname	LL	elementname	Name of the specified element.		
<u>mode</u>	<u>Keyword Abbreviation</u>										
LINE=elementname	LI										
LLINK=elementname	LL										
elementname	Name of the specified element.										
statuscode	Optional. If specified only the elements with that status are reported.										
	<table><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										

The fourth form requests the status of all NPUs supervised by the CS.

The command format is:

STATUS,NPUS.

or

ST,NPUS.

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.

<u>Parameter</u>	<u>Description</u>
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 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):

<u>Command</u>	<u>Description</u>
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.*	Assigns the NAM K display back to NAM mode after it has been assigned to CS, NS, or NVF.

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.

<u>Command</u>	<u>Description</u>
DB	Activates the inline debug code.
DE	Deactivates the inline debug code.
DU	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	Begins logging of the network traffic on the debug log file.
LE	Ends logging of the network traffic on the debug log file.
LR	Releases the debug log file.
RS	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.

NS Control Commands

The HOP assigns the NAM K display to NS before entering any of the following commands.

<u>Command</u> †	<u>Description</u>
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.

<u>Command</u> †	<u>Description</u>
DISABLE	Disables an application or all host network operations.
ENABLE	Enables an application.
IDLE	Terminates an application or all host network operations.
HISTORY	Requests a display of the recent unsolicited status reports.
STATUS	Requests the status of network elements.

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 commands by logging into CS.

<u>Command</u> †	<u>Description</u>
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.

† 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, 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 commands by logging into CS.

<u>Command</u> †	<u>Description</u>
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.

†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)	16-4
Load File Generator (LFG)	16-6
Input Directives	16-7
Network Load File	16-8
Summary Listing	16-8
NPU Dump Analyzer (NDA)	16-9
Input Directives	16-10
0 - Comments	16-10
1 - File 1 Register	16-10
2 - Macromemory Word	16-10
3 - Contiguous Block Data Structure	16-10
4 - Circular Buffer Data Structure	16-11
7 - Page Registers	16-11
8 - File 1 Registers	16-11
9 - Macromemory Words	16-12
A - LCB and TCB/CCB	16-12
B - Port Tables and MUXLCB	16-13
NDA Output File Content	16-24
NPU Dump Files	16-25
Dump Interpretation	16-25
PIP Dump Analyzer (LISTPPM)	16-25
AIP Trace With IAF	16-26
IAFTM	16-26
IAFTR	16-27
Network Description Files	16-28
Stimulator	16-29
STIMULA	16-30
Session File Construction	16-30
Dynamic Login	16-34
Optional Think Time Specification	16-35
Task Definitions	16-36
Session File Processing	16-37
Multiple Sessions	16-42
Parameter Output Format	16-42
Postprocessing	16-42

This section describes the following utilities and files involved in the operation of the network and interactive subsystems of NOS.

- NAM Initialization (NAMI). This utility controls the network software startup process.
- Network Dump Collector (COLLECT). This utility produces permanent files from the various NPU and host dump, trace, statistics, and list files resulting from any abnormal termination of the 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 (NPU's).
- 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.
- Stimulator. A stimulator is a collection of programs that artificially loads the system to analyze the effects of such a load on system performance and reliability.

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 (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 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 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 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,...,pn.

p_i is a parameter consisting of a keyword or a keyword equated to a value. All parameters are order independent.

<u>p_i</u>	<u>Description</u>
GO	This parameter initiates the network startup process without operator intervention.
MFN=filename †	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 †	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.
RN=recordname †	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.

<u>recordname</u>	<u>Description</u>
INIT	Purge existing dumps/traces, stop but do not dump local NPUs at network startup, take host dumps when an NPU fails, and stop local NPUs at 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, do not stop local NPUs at host termination, and do not take host dumps when an NPU fails. Use MINIT for the initial startup of multihost networks.
MRECOV	Purge existing dumps/traces, do not take NPU dumps before initial loading of NPUs, do not stop local NPUs at network startup or host termination, and do not take host dumps when an NPU fails. Use MRECOV for recovery startup of multihost networks after a failure.

†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.

P₁

Description

<u>recordname</u>	<u>Description</u>
MULTI	Collect dumps/traces on tape before network startup, take NPU dumps before initial loading of NPUs, do not stop local NPUs at network startup or host termination, and do not take host dumps when an NPU fails. Use MULTI for the normal restart of multihost networks.
RECOVR	Purge existing dumps/traces, take host dumps when an NPU fails, 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, take host dumps when an NPU fails, stop local NPUs at host termination, and do not stop local NPUs before network startup. Use RESTRT for the normal restart of single host networks.

RS=applname The name of the application to start up. This parameter allows a 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.

UN=username † 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.

† 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 two permanent files: one of the various host dump, trace, statistics, and list files that result from any abnormal termination of the network, and one of any NPU dump files. The collected local files may be copied to the two permanent files and then purged, or copied and not purged. The permanent files can then be copied to tapes. There is also an option to purge the local files without copying them to the permanent files. The format of the COLLECT command is:

```
COLLECT,NIN=nnn,NOPURGE,NOSAVE.
```

<u>Parameter</u>	<u>Description</u>
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 permanent 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 permanent files. If NOPURGE is omitted, the collected files are purged after they are copied to the permanent 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

pp is a 2-character product prefix.

<u>pp</u>	<u>Description</u>
CS	Communication Supervisor
IA	Interactive Facility
NI	Network Interface Package
NP	Network Processing Unit
NS	Network Supervisor
NV	Network Validation Facility
RB	Remote Batch Facility
TV	Terminal Verification Facility

t is the type of file.

<u>t</u>	<u>Description</u>
D	Dump file
L	List file
S	Statistics file
T	Trace file

s is the subtype 0, 1, 2, or 3 for the file.

nnn is the 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.

The permanent files produced by the COLLECT utility each have one file with many records. Each dump, trace, statistics, or list file is copied to a permanent 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 permanent file to show which collected files were copied to the permanent file.

The COLLECT utility assigns each permanent file a file name of the form:

DUMPnnn for host dump, trace, statistics, or list files.

or

DUNPnnn for NPU dump files.

nnn is the network invocation number from the COLLECT command.

The permanent files DUMPnnn and DUNPnnn are attached and copied to magnetic tape (or other operator-assigned equipment) if the RESTRT or MULTI startup procedure is selected. Once the files are successfully copied, they are purged.

The COLLECT utility purges all dump and trace files without collecting them on a permanent 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, p_1, p_2, \dots, p_n .

p_i is a parameter consisting of a keyword or a keyword equated to a value. All parameters are optional and order independent.

<u>p_i</u>	<u>Description</u>
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	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. 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). For example, (slant used for separator): LFG,Z./*FILE,PFILE1,PFILE2.
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 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,CFILE1,...,CFILEn.  
--EOI--
```

} file INPUT

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:

<u>NO.</u>	<u>PARTITION NAME</u>	<u>INPUT FILE</u>	<u>RECORDS</u>
1	PN1	AFILE1	3
2	PN2	AFILE2	7
3	PN3	AFILE3	2

<u>Field</u>	<u>Description</u>
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,...,P_n.

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.

<u>P_i</u>	<u>Description</u>
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.
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.

<u>opt</u>	<u>Description</u>
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.

<u>Columns</u>	<u>Field Name</u>	<u>Description</u>
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.

<u>Columns</u>	<u>Field Name</u>	<u>Description</u>
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.

<u>Columns</u>	<u>Field Name</u>	<u>Description</u>
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.

<u>Columns</u>	<u>Field Name</u>	<u>Description</u>
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 I 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
      MEANING OF BIT POSITIONS IN THE INITIALIZATION TABLE OR
      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
00150  BYWLCB      WORKLIST CONTROL BLOCK
00151  BYCURWLP   CURRENT ACTIVE MONITOR WORKLIST CONTRL BLK
00152  BITCB      INTERNAL PROCESSING TCB
3      CONTIGUOUS DATA STRUCTURE
00161 00020 00001 00001 00003  MLIA, CONSOLE, COUPLER, TCB
00165 00002 00001 00001 00001  INITIALIZATION COMPLETE TABLE
00150 00008 00008 00008 00019  WORKLIST CONTROL BLOCK NUMBER
4      CIRCULAR INPUT BUFFER DUMP FROM CEXX TO CEXX PATTERN
0003F 00040 00042 00030 00E00 0FF00 00000
A      PRINT ALL LCBS WITH THEIR TCBS -- SKIP LCBO
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
0015F 0001F 00015 00010 0000E 00000 00017  LCB, CCB, AND SO FORTH
B      PRINT PORT TABLE AND MUXLCBS, TEST FOR BIT 12, PORT WORD 1
000F1 00008 00001 00010 01000 01000 00001 000F3
7      DUMP PAGE REGISTERS
8      DUMP FILE 1 REGISTER PER COMMAND LO= PARAMETER VALUE
9      DUMP MACROMEMORY PER COMMAND BA=,EA=,LO= VALUES

```

PAGE 1

Input directive file

Figure 16-1. Sample NDA Output File (Sheet 1 of 10)

NPU DUMP FILE = NP04458 NDA,I=NDAINF,L=NDAOUTF,NDF=NP04458.

PAGE 1

TIME	02.22.37	} Header information
DATE	84/03/23	
NPU NAME	LOC3	
NPU NODE	21	
HALT CODE	0	
P REGISTER	0	

THIS IS A HEADING LINE FOR THE FIRST COMMENT PAGE.
 THIS IS A SAMPLE OF AN INPUT DIRECTIVE FILE SPECIFIED BY
 THE I 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
 MEANING OF BIT POSITIONS IN THE INITIALIZATION TABLE OR
 THE NAME ASSOCIATED WITH WORKLIST CONTROL BLOCK 4 CAN BE
 PRINTED.

PAGE 2

FILE 1 REGISTER INFORMATION
 (LOC 000F1) 950F FIRST PORT TABLE
 (LOC 000F3) 008A MAXIMUM NUMBER OF LCBS
 (LOC 00040) 938A NEXT FRAME POINTER

PAGE 3

MACROMEMORY INTERPRETATION
 (LOC 00150) 0663 BYWLCB WORKLIST CONTROL BLOCK
 (LOC 00151) 0753 BYCURWLP CURRENT ACTIVE MONITOR WORKLIST CONTRL BLK
 (LOC 00152) 05A4 BITCB INTERNAL PROCESSING TCB

PAGE 4

Figure 16-1. Sample NDA Output File (Sheet 2 of 10)

CONTIGUOUS DATA STRUCTURE				PAGE												
MLIA, CONSOLE, COUPLER, TCB				5												
OFFSET	1	2	3													
0	0000	0000	995F													
1	AA9F	AABE	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	89B4													
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	ECE9	0000	0000	0000	0000	0000	0000	CC30	0000	B770	ACA0	0000	BD60	0000
2	AC90	0000	DFE6	0000	AD10	0000	0000	0000	0000	CC30	0000	B770	ACA0	F5BC	BD60	0000
3	8003	0003	8003	4102	4001	4001	4001	8003	4102	8001	8005	8003	8001	0003	8003	4304
4	4D58	0000	5354	0000	0000	0000	0000	584D	434F	4250	4D4C	534D	5453	5248	4C49	4F4C
5	0000	0000	0000	0000	0000	0000	0000	06AB	06B3	06BB	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	20D1	3FC4	213D	8944	3885	80DA	366C

Figure 16-1. Sample NDA Output File (Sheet 3 of 10)

CONTIGUOUS DATA STRUCTURE		
WORKLIST CONTROL BLOCK NUMBER		
OFFSET	18	19
0	0000	0000
1	ABFO	AE40
2	ABFO	AE40
3	8003	8003
4	4850	4C50
5	06EB	06F3
6	0426	102C
7	3A73	274B

Figure 16-1. Sample NDA Output File (Sheet 4 of 10)

CIRCULAR INPUT BUFFER DUMP FROM CEXX TO CEXX PATTERN									PAGE	7
LOC										
938B	0E03	CC1C	0E03							
6	0000	0000	0000	0000	0000	0000	0002	0002		
7	0000	0000	0000	0000	0000	0000	00A0	00A0		
8	0000	0000	0000	0000	0000	0000	C000	3000		
9	0009	0000	000E	0000	0000	0000	0000	0000		
A	0004	0000	000F	0000	0000	0000	0000	0000		
B	0003	0000	0006	0000	0000	0000	0000	0000		
C	0001	0000	000E	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		
10	91A2	9177	91CD	91F8	0000	0000	0000	0000		
11	0000	0000	0000	0000	0000	0000	0000	0000		
12	0000	0000	0000	0000	0000	0000	0003	0003		
13	0303	0000	05C2	0000	0000	0000	0000	0000		
14	0000	0000	0000	0000	0000	0000	0000	0000		
15	0000	0000	0000	0000	0000	0000	0000	0000		
16	0000	0000	0000	0000	0000	0000	0000	0000		
17	0000	0000	0000	0000	0000	0000	0000	0000		
18	0000	0000	0000	0000	0000	0000	0000	0000		
19	0000	0000	0000	0000	0000	0000	0000	0000		
1A	0000	0000	0000	0000	0000	0000	0000	0000		
1B	0000	0000	0000	0000	0000	0000	0000	0000		
1C	0000	0000	0000	0000	0000	0000	0000	0000		
1D	0000	0000	0000	0000	0000	0000	0000	0000		
1E	0000	0000	0000	0000	0000	0000	0000	0000		
1F										
20										
21										
22										
23										
24										
25										
26										
27										
28										
29										
2A										
2B										
2C										
2D										
2E										
2F										

Figure 16-1. Sample NDA Output File (Sheet 5 of 10)

PRINT ALL LCBS WITH THEIR TCBS -- SKIP LCBO
LCB, TCB, AND SO FORTH

PAGE 9

OFFSET	LCB009	TCB000	LCB00A	TCB000	LCB00B	TCB000	LCB00C	TCB000	LCB00D	TCB000	LCB00E	TCB000	LCB00F	TCB000	LCB010	TCB000
LOC..	9B00		9B1F		9B3E		9B5D		9B7C		9B9B		9BBA		9BD9	
0	0000		0A00		0B00		0C00		0D00		0E00		0F00		1000	
1	0000		9B3E		9B5D		9B7C		9B9B		9BBA		9BD9		9BF8	
2	0000		00E7		0005		00E7		8003		0003		8003		8003	
3	0000		0000		0000		0000		0000		0000		0000		0000	
4	0000		0000		0000		0000		0000		0000		0000		0000	
5	0000		0A60		0AA0		0A60		0180		0280		0180		0180	
6	0000		0004		0100		0004		0002		0002		0002		0002	
7	0000		00A0		0041		00A0		80A0		00A0		00A0		80A0	
8	0000		4000		4000		4000		5000		F000		F000		6200	
9	0000		0000		001B		0000		0000		0000		0000		0000	
A	0000		0000		0013		0000		0000		0000		0000		0000	
B	0000		0000		0143		0000		0000		0000		0000		0000	
C	0000		0000		0003		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	
10	0000		0000		C10Q		0000		0000		0000		0000		0000	
11	0000		0000		0000		0000		0000		0000		0000		0000	
12	0000		0003		0D03		0003		0003		0003		0003		0003	
13	0000		0000		0547		0000		0000		0000		0000		0000	
14	0000		0000		0000		0000		0000		0000		0000		0000	
15	0000		AB1B		AB4B		AB7B		0000		0000		0000		0000	
16	0000		0000		0000		0000		0000		0000		0000		0000	
17	0000		0000		0000		0000		0000		0000		0000		0000	
18	0000		0000		0000		0000		0000		0000		0000		0000	
19	0000		0000		0000		0000		0000		0000		0000		0000	
1A	0000		0000		0000		0000		0000		0000		0000		0000	
1B	0000		0000		0000		0000		0000		0000		0000		0000	
1C	0000		0000		0000		0000		0000		0000		0000		0000	
1D	0000		0000		0000		0000		0000		0000		0000		0000	
1E	0000		0000		0000		0000		0000		0000		0000		0000	
1F																
20																
21																
22																
23																
24																
25																
26																
27																
28																
29																

Figure 16-1. Sample NDA Output File (Sheet 6 of 10)

A NONZERO COLUMN 1 FORCES A NEW PAGE FOR THE CCB CHAIN		PAGE	28
LCB, CCB, AND SO FORTH			
OFFSET	LCB017	TCB000	
LOC..	9CB2		
0	0000		
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 (Sheet 7 of 10)

PRINT PORT TABLE AND MUXLCBS, TEST FOR BIT 12, PORT WORD 1											PAGE 29	
OFFSET	PORT000 MUX	PORT001 MUX	PORT002 MUX	PORT003 MUX	PORT004 MUX	PORT005 MUX	PORT006 MUX	PORT007 MUX				
LOC..	950F	9517	951F	9527	952F	9537	953F	9547				
0	0000	AA81	0280	AA81	0280	0000	0000	28C0				
1	0000	ECF0	0000	ACD0	0000	0000	0000	AE60				
2	0000	8807	0001	8807	0001	0000	0000	8803				
3	0000	41FC	41FC	41FC	41FC	0000	0000	416C				
4	0000	8AA7	0000	8AA7	0000	0000	0000	8827				
5	0000	0000	0000	0000	0000	0000	0000	3200				
6	0000	E00F	0000	E00F	0000	0000	0000	0000				
7	0000	0000	0000	0000	0000	0000	0000	0000				
8												
9												
A												
B												
C												
D												
E												
F												
OFFSET	PORT008 MUX	PORT009 MUX	PORT00A MUX	PORT00B MUX	PORT00C MUX	PORT00D MUX	PORT00E MUX	PORT00F MUX				
LOC..	954F	9557	955F	9567	B290	956F	9577	B2C0	957F	B2D0	9587	B2A0
0	28C0	0000	0280	AA80	A180	0280	2840	0000	2880	0000	2840	0000
1	AE70	0000	0000	B290	4189	0000	B2C0	2000	B2D0	2000	B2A0	2000
2	8803	0000	0001	8807	DE44	0001	0803	0000	0803	0000	0803	0000
3	416C	0000	41FC	41FC	0004	41FC	4126	0000	4126	0000	4126	0000
4	8827	0000	0000	8AA7	0000	0000	0827	0000	0827	0000	0827	0000
5	3200	0000	0000	0000	4DA0	0000	3200	420E	3200	499F	3200	499F
6	0000	0000	0000	E00F	0000	0000	0000	0000	0000	0000	0000	0000
7	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
8					0008			000A		0008		0008
9					0000			0000		0000		0000
A					0000			0000		0100		0100
B					0000			0000		0100		0100
C					0000			0000		4EE8		4EE8
D					0000			0000		0000		0000
E					0023			0000		0000		0000
F					0000			0000		0000		0000

Figure 16-1. Sample NDA Output File (Sheet 8 of 10)

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	0008	0009	000A	000B	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	0D08	0013	CCBF	DBBF	1D08	0003	B81F	C91F	3D08	00F8	EF7F	EABF	7D08
000010	319D	0522	9520	DBB0	0000	0000	01C7	037E	03C7	03D7	0040	0000	0000	0000	ACA0	06C3
000020	000F	06C3	ACA0	0001	0001	0000	0000	0000	0000	8001	8B79	06E3	06DB	B300	0102	0184
000030	0185	0000	B300	959F	07C0	B300	0000	0002	A0B7	0002	0026	260A	0032	0026	0003	9310
000040	938A	007A	950F	E00F	41FC	0007	0000	8B7A	0000	9517	0000	AA81	995F	9971	06CC	0001
000050	0CC1	05E0	CD0F	0001	0000	0085	0086	0087	0801	0000	0028	0000	00C1	0000	0000	0000
000060	0000	0000	0000	0000	3994	0000	811E	2527	846F	DB3C	7806	000A	0003	0000	2C07	3899
000070	0007	38EF	0045	0045	0045	DB3C	0000	0000	0000	0000	00C1	00C1	0004	ECF3	8D14	ECF3
000080	0000	8003	0001	0720	0000	0000	0000	0000	0000	0000	0000	0000	0000	0548	0000	0006
000090	0000	040D	0004	5030	0000	0000	0000	0000	0000	0000	0663	2100	EAB6	0102	0300	3000
0000A0	0508	0500	CC00	0C00	0560	0F00	0D00	0001	0002	0003	0004	0005	0006	0007	0008	000A
0000B0	000E	000F	001F	0037	003F	0050	007F	0080	00C0	00F0	00FE	00F8	00FF	0200	07E0	0C00
0000C0	0F00	2000	2180	4000	C000	401F	8000	A180	CD00	0000	0000	0000	F000	00E0	0012	D280
0000D0	001F	0071	E640	000D	0018	0000	0000	0D0A	0D0A	1918	0000	0000	0001	0D0A	0000	0000
0000E0	DB3C	0071	2C87	AF6B	1918	7806	DB3C	001F	DD86	008D	848A	F0C6	0000	0D0D	DD86	000D
0000F0	0000	950F	01B8	008A	844D	0008	0004	0004	0000	0000	0000	0000	001B	1502	0614	8B8A
000100	--															

Figure 16-1. Sample NDA Output File (Sheet 9 of 10)

DUMP MACROMEMORY PER COMMAND BA=,EA=,LO= VALUES																PAGE 40															
ADDRESS	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	T	T	T	M	T	L									
000000	5400	5400	5400	4B09	0B00	0B00	0B00	0B00	0B00	0B00	0B00	0B00	0B00	0B00	5400	8B6C	T	T	T	M		T	L								
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															
000040*	0000	0044	00FE	0000	0040	0000	0001	8934	0A76	0000	2E37	002D	0030	001D	0001	001D	D	@		4	V	.7	-	0							
000050	0001	0002	FFFE	0091	0001	0002	000F	0001	000F	0001	CF82	0044	0000	0001	5E27	F200					D										
000060	FFFE	004D	0056	BD60	0001	BD60	004D	0000	005B	0000	0001	5C6F	004D	FFFE	0001	BD60	M	V		M	[O	M							
000070	F200	0052	0068	0000	0001	6203	07BA	065D	0091	0001	0002	000F	0001	000F	0001	5B11	R	H		B]			[
000080	0001	FFFE	0091	0001	0002	0019	0001	0019	0001	0780	007F	2020	0108	0118	0128	0001								(
000090	0001	0001	4101	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	A														
0000A0	0000	0000	0A00	0000	0000	0000	0000	0000	0000	0000	024C	024C	024C	024C	024C	024C						L	L	L	L	L	L	L			
0000B0	024C	024C	0240	0500	0000	010B	0000	0000	0000	0000	0000	04F8	0000	0000	0000	0000	L	L	@												
0000C0	7F4F	FF4F	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	O	O													
0000D0	0140	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	@														
0000E0	0000	0000	0000	2020	2020	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	1FAC	0000	3C45	1400	1FD5	0000	0000	1400	81F9	0000	3FAA	1400	821A	0000			<	E								?			
000110	0000	5400	828C	0E10	2945	5400	6A46	0E14	0000	5400	6A46	0E18	0000	5400	828C	0E1C	T)	E	T	J	F	T	J	F	T				
000120	8BC0	1400	8254	0000	0000	5400	828C	0E24	0000	5400	828C	0E28	0000	5400	828C	0E2C			T		T	\$	T	(T		,				
000130	0000	5400	828C	0E30	0000	5400	828C	0E34	0000	5400	828C	0E38	0000	5400	828C	0E3C	T		0	T	4	T	8	T	<						
000140	1400	0000	1400	8B4D	1400	F02C	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000			M	,											
000150	0663	0753	05A4	05A5	079B	019F	0A6E	07BA	3EA9	2222	8BC5	2787	950F	9310	995F	99E9	C	S		N	>	"	"	"	"	"	"	"			
000160	AA9F	0A9F	0AE7	0C1C	9177	0E0E	ABCO	0036	1916	0614	3128	803F	0000	019E	9251				W		6		1	(?	Q					
000170	020C	0A02	020C	0A02	0000	0000	0400	AFBF	0003	950F	4920	AAA1	BAE0	8610	8E30	9E70											I	O	P		
000180	4510	4D30	5D70	7DF0	6DB0	0000	03F0	0412	0434	0456	0478	049A	04BC	04DE	0500	0522	E	M	O]P	M		4	V	X	"					
000190	0544	0566	0000	9310	01FF	000A	0C15	0200	8A00	0000	0000	0000	0000	0000	0011	0000	D	F													
0001A0	0001	0005	000D	001D	007D	001D	001D	00FF	01FF	03FF	07FF	0FFF	1FFF	3FFF	7FEF	FFFF														?	
0001B0	8F59	0000	CAB0	80F0	CA09	CA1A	0000	065D	0000	0BFD	0000	0BF7	0BF9	0BF9	0BF9	0BF9	Y														
0001C0	0000	0BFB	0BFB	0BFB	0BFB	0BFB	0BFB	0BFB	0BFB	0000	0000	0000	0000	0000	0000	0000															
0001D0	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000															
000310*	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000															
000320	3FFF	CC00	0000	00CF	0000	FF00	0000	0000	0000	FC00	FF00	00FF	FFFC	FFFF	FFFF	FFFF														?	
000330	CFF3	FFFF	FFFF	FFFF	FFFF	FFFF	0000	0F00	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF															
000340	FFFF	FFFF	0000	0000	0000	0000	FF00	0000	0000	FF00	0000	0000	0000	00FF	00FF	0000															
000350	00FF	FF00	0000	0000	0000	0000	0000	0000	0000	0300	0000	0000	0000	0000	0000	FF00															
000360	FCFC	FF00	FF00	0000	0000	0000	00FF	0000	FC00	0000	0000	0000	0000	0000	0000	FC00															
000370	0000	0000	0000	0000	00FC	0000	0000	0000	0000	0000	0000	FF00	FFFF	0000	0000	FC00															
000380	0000	0000	0000	0000	FF00	00FF	00FF	0000	0000	0000	0000	0000	0000	0000	0000	0000															
000390	0000	0000	FC00	0000	0000	0000	FF00	0000	0000	0000	0000	0000	00FF	0000	0000	0000															
0003A0	000F	0000	0000	0000	FF00	0000	0000	00FF	0000	0000	0000	0000	0000	0000	0000	0000															
0003B0	FC00	0000	FC00	0000	00FC	0000	0000	FFFF	0000	0000	00FF	0000	0000	00FF	0000	0000															
0003C0	0000	0000	FF00	0000	0000	0000	FF00	8000	4000	4000	8000	2000	1000	1000	2000	0800						@	@								
0003D0	0400	0400	0800	0200	0100	0100	0200	C000	3000	3000	C000	0C00	0300	0300	0C00	0032														2	
0003E0	ABCO	FFFC	003F	FFFD	0607	0608	0606	0605	0609	095A	0600	0601	0602	060A	060E	060B															
0003F0	0000	0000	4121	0000	0000	FFFF	0020	FFFF	8840	FFFF	0200	FFFF	0100	FFFF	8800	FFFF			?				Z								
000400	0100	FDFD	0000	FDFD	0000	FEFF	0000	0400	0000	0400	0000	FFFF	0000	FFFF	0000	FFFF			A!			@									
000410	0000	FFFF	2440	BFBF	4126	0000	0000	7FFF	0020	FFFF	0840	FFFF	4200	7EFF	8100	FDFD			\$	@	A	&		@	B						
000420	0800	FFFF	C100	FDFD	0000	FDFD	0000	7EFF	0000	0400	0000	0000	0000	FFFF	0000	FFFF															
000430	0000	FFFF	0000	FFFF	2E40	BFBF	4126	0000	0000	7FFF	0020	FFFF	0840	FFFF	4200	FEFF			.	@	A	&		@	B						
000440	8100	FDFD	0800	FFFF	C100	FDFD	0000	FDFD	0000	7EFF	0000	0400	0000	0400	0000	FFFF															
000450	0000	FFFF	0000	FFFF	0000	FFFF	0A40	FFFF	416C	0000	0000	FFFF	0020	FFFF	8840	FFFF							@	AL		@					
000460	4200	FFFF	0100	FFFF	8800	FFFF	4100	FDFD	0000	FDFD	0000	FEFF	0000	0400	0000	0400	B					A									
000470	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	0000	0000	0000	0000	FFFF	0000	FFFF															
000480	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF															
000490	0000	0400	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	0000	0000	FFFF															
0004A0	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF															
0004B0	0000	0400	0000	0400	0000	FFFF	0000	FFFF	0000	FFFF	0000	FFFF	0021	FFBF	41B7	0000														!	A

Figure 16-1. Sample NDA Output File (Sheet 10 of 10)

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

<u>Header</u>	<u>Description</u>
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 30₁₆. If the content of that location is nonzero, a halt has occurred; the halt code value is contained in that location.

If memory location 30₁₆ equals zero after a halt during initialization, find the address table that begins at fixed memory address 150₁₆. 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.

<u>Parameter</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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.

<u>File Name</u>	<u>Description</u>
LCF and NCF	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†	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 LCF and NCF to be permanent, direct access public files stored under the NETOPS user index (377772g). LCF and NCF 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 (377777g).

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.

The following format shows the terminal definition directive. No embedded blanks are allowed.

```
/terminalname,p1,p2,...,pn.
```

<u>Parameter</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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.

†i is the network file identifier as specified in the TAF configuration file (refer to the TAF Reference Manual).

<u>Parameter</u>	<u>Description</u>
<u>P_i</u>	<u>Description</u>
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 77777778). 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.

STIMULATOR

A 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 STIMULA stimulator consists of a central memory program (STIMULA) and a peripheral processor program (ITS).

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 stimulator sessions are described under Multiple Sessions later in this section. Errors that can occur while the stimulator software is in use are described in the NOS 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.

<u>Parameter</u>	<u>Description</u>
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 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 familyname entry, a username entry, and a password entry.
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
PERFDEV
MATH,MATH
CHARGE(7159,73KDF)A
> 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
> >
RUN
2.0
0.0
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)
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

```

} Source line input

} Program input

Figure 16-3. Sample Session for Use by STIMULA


```

ANSWERBAC
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

```

ENTER STIMULATION FILE PARAMETERS IN FORM -

P = VALUE

F =          SESSION FILE NAME.
UI = 0B      USER INDEX (OCTAL BASE ASSUMED.)

N = 1        NUMBER OF FILE ON SESSION FILE.
GO           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 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\$\$\$

<u>Parameter</u>	<u>Description</u>
username	User name; 1 to 4 alphabetic characters.
password	Password; 1 to 4 alphabetic characters.
\$	Special symbol 12/8/6 (6-bit display code 76g), 1 (6-bit display code 34g) 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 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:

> xxx

<u>Parameter</u>	<u>Description</u>
xxx	Think time (octal): xxx \leq 1778.

To return the terminal to the initial think time (refer to figure 16-6), enter this command:

>>

ENTER STIMULATION PARAMETERS IN FORM -

P = VALUE	(NUMERIC VALUES ARE ASSUMED DECIMAL BASE.) (SEE RIGHT SCREEN FOR SPECIAL PARAMETERS.)
NT = 310B	NUMBER OF TERMINALS TO STIMULATE. (1 - 100B)
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	UPPER BOUND OF RANDOM THINK TIME INCREMENT. 0 - 64 (MUST BE A POWER OF 2.) THINK TIME WILL VARY BETWEEN TT AND TT+TI.
AC = 4	ACTIVATION COUNT. (1 - 64) NUMBER OF TERMINALS TO ACTIVATE EVERY AD SECONDS.
AD = 10	ACTIVATION DELAY IN SECONDS. (0 - 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	LOOP ON SESSION FILE (YES OR NO)

Figure 16-6. STIMULA Session File Parameters

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 70g) in the session calls the next most eligible task in the session file, which is determined by the following equation.

$$f_i = \frac{\text{Total initiations of task}_i}{(\text{Total initiations of all tasks})(\text{Task usage parameter for task}_i)}$$

For each task_i defined in the session file, the 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 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 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 stimulator fails to calculate the next most eligible task on time.

The first line in a task record must be:

\$task nnn

<u>Parameter</u>	<u>Description</u>
\$	Denotes the beginning of a new task (character position 1).
task	The 1- to 6-character task name.
nnn	Task usage percentage; $0 \leq \text{nnn} \leq 100$ (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 71g), 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 stimulator software is used to process the interactive session file following this procedure:

1. Ensure that an IAF stimulator entry is present in the EST and that the stimulator has an ON status (refer to the section 3, Deadstart Decks). The stimulator entry must be on a channel that is not used by any other equipment that is currently ON in the EST.
2. Use DSD to activate the 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 (37777g). The procedure file STM is called when you enter STM. The stimulator is invoked by a STIMULA command within the STM procedure file. The 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.

<u>Format</u>	<u>Default</u>	<u>Description</u>
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 37777g.
K.GO	None	Indicates to the 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 special parameters, enter:

KK.

Figure 16-7 appears on the right screen.

TO ENTER MIXED MODE OF PARAMETER ASSIGNMENT -
K.MIXED.

Figure 16-7. STIMULA Special Parameters


6. Enter the stimulator parameters 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,G0.

Enter any special parameters (refer to figure 16-9) before the G0. entry.

<u>Format</u>	<u>Default</u>	<u>Description</u>
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 3 to 1000g.
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 (600g, for example) are run, the maximum realizable line speed is determined by the overhead of lTD and lTS. The worst case can be determined by checking the lTD cycle time dayfile message, which is issued 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 the 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 n (the value of NT). The AC and AD parameters provide a method to stagger the I/O commands and data to the IAF.
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 127 seconds.

<u>Format</u>	<u>Default</u>	<u>Description</u>
K.RO=0	ON	When RO=ON, the output from the IAF is recovered on file STIMOUT. The format of this file is:



Terminal
number.

Each byte is an ASCII character if less than 4000g; otherwise, it is the running-second time.

When RO=OFF, the output is not recovered.

K.LF=0	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 special 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.GO	None	Indicates to the stimulator that the parameters are entered. Stimulation is initiated.
K.MX= s ₁ -s ₂ , tys,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.)
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 one page.
K.-	None	Page matrix display backward one page.
K.GO	None	Indicates to the stimulator that the parameters are entered. Stimulation is initiated.

STIMULATION PARAMETER MATRIX

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

NUMBER OF TTYS ASSIGNED = 200

MAXIMUM TTYS PERMITTED = 200

CF = MATRIX CHANGE FILE

Figure 16-8. Parameter Matrix (Left Screen)

DEFINITION OF PARAMETERS

SCRIPT - SCRIPT(S) WHICH THE PARAMS WILL AFFECT
(SCRIPT(N) OR SCRIPT(N)-SCRIPT(N+X))

TTYS - NUMBER OF TTYS ASSIGNED TO SCRIPT(S)

LS - LINE SPEED ASSIGNED TO SCRIPT(S)

IS - INPUT SPEED ASSIGNED TO SCRIPT(S)

TT - THINK TIME ASSIGNED TO SCRIPT(S)

LD - LOGOUT DELAY ASSIGNED TO SCRIPT(S)

RC - REPEAT COUNT ASSIGNED TO SCRIPT(S)

*NOTE - NULL PARAMETER WILL KEEP CURRENT VALUE

TO MAKE ENTRIES, FOLLOW THE FORMAT -

K.MX=SCRIPT,TTYS,LS,IS,TT,LD,RC.

TO CHANGE MATRIX VIA PREDEFINED PARAMETERS-

K.CF=XXX WHERE XXX IS PF CONTAINING CHANGE
DIRECTIVES IN FOLLOWING FORMAT-
MX=SCRIPT,TTYS,LS,IS,TT,LD,RC.

TO ADVANCE DISPLAY ENTER K.+

TO BACKUP DISPLAY ENTER K.-

Figure 16-9. Mixed Mode Parameters (Right Screen)

Example:

`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.

MULTIPLE SESSIONS

Each session on the session file is allocated to each terminal line. When each session has been allocated, the stimulator goes to the first session and reallocates the sessions 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=OFF, the session currently assigned to the terminal is repeated. If LF=ON, 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=OFF), or the next four sessions on the session file (with repetition, if necessary) are executed (LF=ON).

PARAMETER OUTPUT FORMAT

The parameter information (shown in figure 16-10) is placed on the output file for printing upon completion of the stimulator run (refer to Session File Processing earlier in this section).

POSTPROCESSING

The postprocessing of the interactive 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 4000g 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.

```

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

An example of output from DEMUX:

```
ANSWERBAC
*****TTT.TTT (Optional output)
*****TTT.TTT (Optional output)
WELCOME TO THE NOS SOFTWARE SYSTEM.
COPYRIGHT CONTROL DATA 1978, 1983.
```

```
yy/mm/dd.hh.mm.ss.
```

```
.
.
.
.
.
```

```
END OF THE LINE _____ n
```

(Page Eject)

```
ANSWERBAC
```

```
.
.
```

The n in the previous listing is the decimal encounter number.

The format of the DEMUX command is:

```
DEMUX,p1,p2,...,pn.
```

p_i is either a keyword or a keyword equated to a value.

<u>P_i</u>	<u>Description</u>
I=filename	Name of the file containing 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.

Introduction	17-1
Utility Description Terminology	17-2
Archive File	17-2
User Index	17-3
Device Masks	17-4
Mass Storage Table	17-7
Catalog Track	17-7
Permanent File Catalog Entry	17-9
Track Reservation Table (TRT)	17-16
Family	17-19
Calling the Utilities	17-23
Console Input	17-23
DIS Input	17-28
Batch Input	17-28
Terminal Input	17-29
Procedure Files	17-29
Parameters for the Permanent File Utilities	17-30
Description of Permanent File Utility Parameters	17-31
Preassigning the Archive File	17-37
Permanent File Utility Routines	17-38
Catalog Archive File (PFATC)	17-41
Catalog Permanent File Device (PFCAT)	17-41
Copy Archive File (PFCOPY)	17-45
Dump Permanent File (PFDUMP)	17-45
PFDUMP in a Secured System	17-45
Incremental Permanent File System Dump	17-46
Partial Permanent File System Dump	17-46
Full Permanent File Dump	17-46
PFDUMP in an MSS/MSAS Environment	17-47
PFDUMP Protected Files	17-47
Load Permanent File (PFLOAD)	17-47
Catalog Image Record	17-49
PFLOAD Selection Process	17-49
Family Reconfiguration	17-52
PFLOAD in an MSS/MSAS Environment	17-52
PFDUMP/PFLOAD Example	17-52

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, and copying archived files to a job as local files.

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.

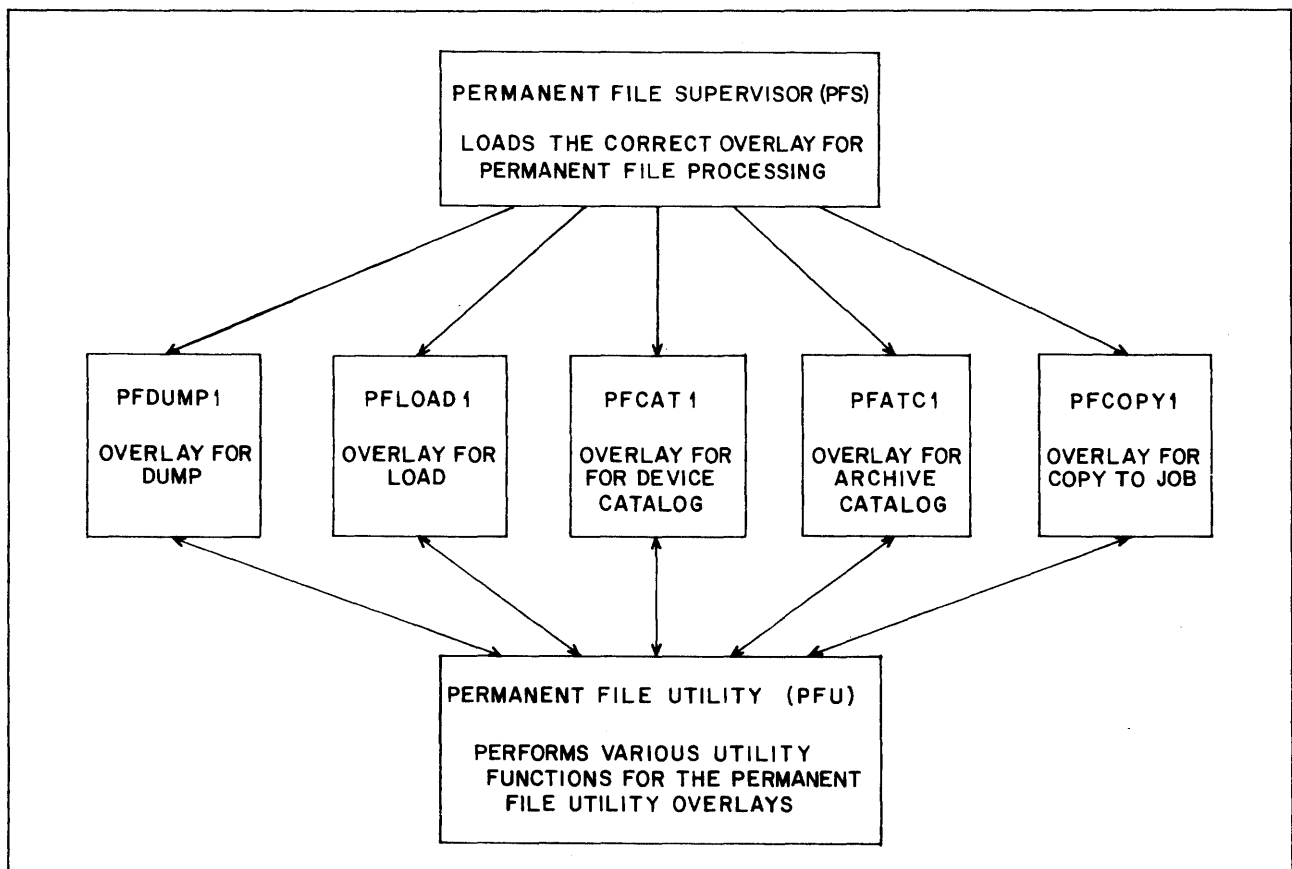


Figure 17-1. Functional Overview of the Permanent File Utility Overlays

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.

<u>Utility</u>	<u>Description</u>
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.

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 multiframe 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.

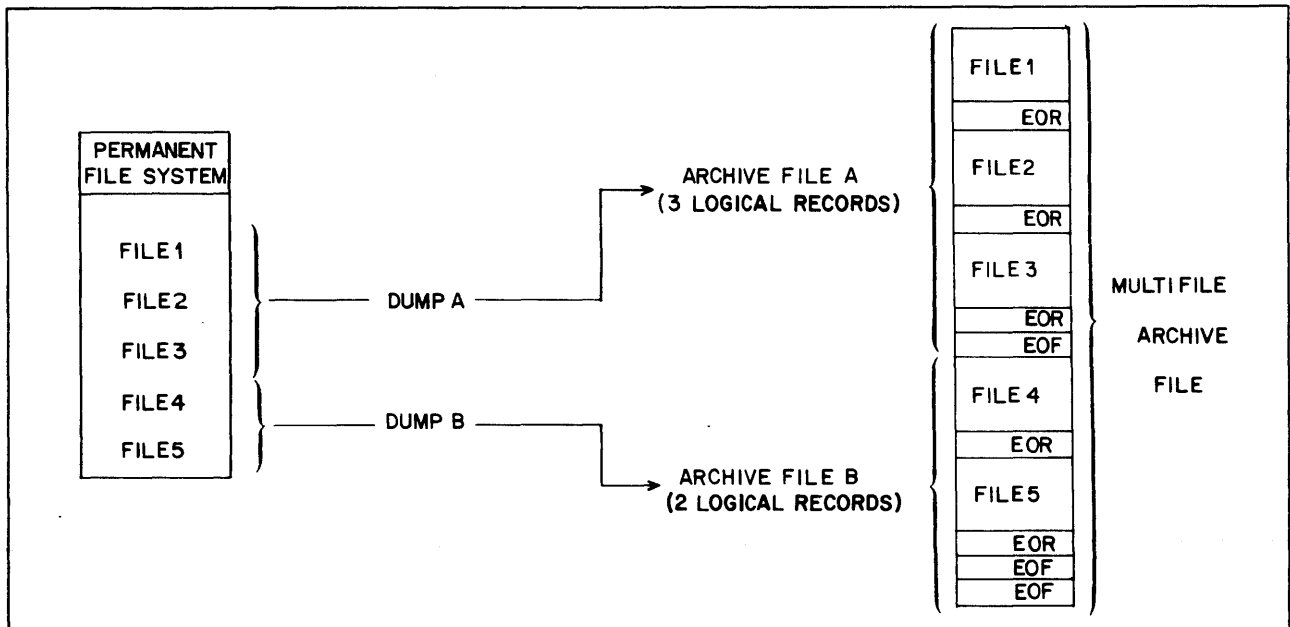


Figure 17-2. Example of Multifile Archive File Structure

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 Subsystem (MSS) environment or a Mass Storage Archival Subsystem (MSAS) 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:

The device identified by the mask 261₈ 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 377₈.
- Each bit position (0 through 7) must be set exactly once for the masks of the devices in one family.

Example:

For a family of three devices, 221₈, 042₈, and 114₈ are valid device masks because their sum is 377₈ 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 377₈, 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:

For a family of four devices, 142₈, 020₈, 010₈, and 204₈ are not valid device masks because their sum is less than 377₈. Any user whose index ends in 0 would have no master device and could not create a permanent file.

```
0 1 1 0 0 0 1 0 = 1428
0 0 0 1 0 0 0 0 = 0208
0 0 0 0 1 0 0 0 = 0108
1 0 0 0 0 1 0 0 = 2048
-----
1 1 1 1 1 1 1 0 = 3768
```

If the sum of the device masks is greater than 377₈, 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:

For a family of four devices, 212₈, 106₈, 040₈, and 021₈ are not valid device masks because their sum is greater than 377₈. Bit 1 is set twice, which indicates that any user whose index ends in 1 would have two master devices, one with mask 212₈ and one with mask 106₈.

```
1 0 0 0 1 0 1 0 = 2128
0 1 0 0 0 1 1 0 = 1068
0 0 1 0 0 0 0 0 = 0408
0 0 0 1 0 0 0 1 = 0218
-----
1 1 1 1 1 1 1 1 4018
      1
```

It is a necessary but not sufficient condition that all masks for one family total 377₈. It is possible, but not allowable, that in one family the device masks total 377₈ but some bit is set in more than one device mask.

Example:

For a family of three devices, 261₈, 115₈, and 001₈ 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.


```
1 0 1 1 0 0 0 1 = 2618
0 1 0 0 1 1 0 1 = 1158
0 0 0 0 0 0 0 1 = 0018
-----
1 1 1 1 1 1 0 1 3778
      1
      1
```

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:


A device whose secondary mask is 072₈ 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:

The following situations occur for a device whose mask is 261₈ and secondary mask is 072₈.



- 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 377700₈ (AUIMX). However, PFLoad will enforce the secondary mask check when loading these files. Direct access files for the Mass Storage Facility (MSF/7990) subfamily user indexes (377760 through 377767₈) 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 377₈.

Example:

For a family of three devices, 000, 131g, and 326g 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 131g. Any user whose index ends in 1, 2, 4, 6, or 7 can place direct access files on the device whose secondary mask is 326g. 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.

```
0 0 0 0 0 0 0 = 000g
0 1 0 1 1 0 0 1 = 131g
1 1 0 1 0 1 1 0 = 326g
-----
1 1 0 1 1 1 1 1
  1   1
```

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 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:

<u>Type of Device</u>	<u>Default Number of Catalog Tracks</u>
Extended memory (DE)	4
834 (DD)	40g
844 (DI, DJ, DK, DL)	40g
885-11/12 (DM, DQ)	10g
885-42 (DB)	10g
819 (DV, DW)	10g
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 200g. Therefore, the possible numbers of catalog tracks on a device are:

1, 2, 4, 10g, 20g, 40g, 100g, 200g

Reducing these numbers by 1 produces the following numbers, referred to as track masks.

0, 1, 3, 7, 17g, 37g, 77g, 177g

The track masks in binary form are:

```

      0
      1
     11
    111
   1111
  11111
 111111
1111111
1111111

```

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.

<u>Device</u>	<u>Device Mask</u>	<u>Number of Catalog Tracks</u>	<u>Corresponding Track Mask</u>
1	221g or 10 010 001	40g	40g-1=37g or 11 111
2	156g or 01 101 110 377g 11 111 111	20g	20g-1=17g or 1 111

A user whose index is 14224g is assigned device 1 as his master device because bit 4 (last digit in 14224g) is set in the device mask for device 1. Therefore, device 1 contains this user's catalog track.

The binary form of 14224g is 001 100 010 010 100.

Therefore, the catalog track number for this user is 22g (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 22g

```

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.

	59	53	47	44	41	35	23	17	11	0	
word 0	filename						userindex				FCFN, FCUI
1	file length							track	sector		FCLF, FCBT, FCBS
2	random index					creation date and time					FCRI, FCCD
3	access count					data modification date and time					FCAC, FCMD
4	ct	mode	ef	ec	dn	last access date and time					FCCT, FCAM, FCEF, FCEC, FCEO, FCAD
5						control modification date and time					FCKD
6	pr	br	ss	ac			utility control date and time				FCRS, FCBR, FCFS, FCAP, FCUD
7	file password						expiration date				FCPW, FCXD
10	aflags				alt	asa					FCAF, FCAT, FCAA
11					al	access categories					FCAL, FCFC
12	charge number										FCCN
13	project number										FCP1
14	project number										FCP2
15											
16	user control word										FCUC
17	reserved for installation										FCIW

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>
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.
random index	2	59-36	Random disk address of first permit sector.
creation date and time	2	35-0	yyymmddhhmmss 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.

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>
data modification date and time	3	35-0	yyymmddhhmmss 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: <ul style="list-style-type: none"> 0 Private. 1 Semiprivate. 2 Public.
mode	4	53-48	Mode of access for semiprivate and public files: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 0 No error. 1 EOI changed by recovery.
ec	4	44-42	Error code: <ul style="list-style-type: none"> 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 Reserved. 6 Reserved. 7 Reserved.

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>
dn	4	41-36	Device number (1 through 778) 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.
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.
pr	6	59-57	Preferred residence † for this file: <ul 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 MSF/7990 residence preferred (PR=M specified by file owner). 4 No preferred residence (PR=N specified by file owner).
br	6	56-54	Backup requirement † for this file: <ul style="list-style-type: none"> 1 Backup required on dump tape (BR=Y specified by file owner). 2 Backup required on dump tape only if the current version of the file does not have an MSF/7990 image (BR=MD specified by file owner). 3 Backup on dump tape is not to be done (BR=N specified by the file owner).
ss	6	53-48	Subsystem code for this file: <ul style="list-style-type: none"> 0 Null Subsystem. 1 BASIC Subsystem. 2 FORTRAN Subsystem. 3 FTNLS Subsystem. 4 Execute Subsystem. 5 Batch Subsystem.

† Refer to Permanent File Commands in the NOS 2 Reference Set, Volume 3 for details on the PR, BR, and AC parameters.

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>
ac	6	47	Alternate CATLIST permission: † 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 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.
aflags	10	59-48	Alternate storage flags; refer to section 11, MSAS, or section 12, MSS, for details.

<u>Flag</u>	<u>Bit</u>	<u>Description</u>
AFLOK	59	Lock to disk flag: 0 No indirect access permanent file is being destaged to MSF/7990. 1 Indirect access permanent file is being destaged to MSF/7990.
AFFRE	54	Free up cartridge flag: 0 Do not remove file from cartridge. 1 Remove file from cartridge.

† Refer to Permanent File Commands in the NOS 2 Reference Set, Volume 3 for details on the PR, BR, and AC parameters.

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>
<u>Flag</u>	<u>Bit</u>		<u>Description</u>
AFTMP	53		Temporary error flag:
		0	No temporary error conditions have been detected.
		1	A temporary error condition prevented the MSF/7990 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	If alternate storage address asa≠0, a successful stage/destage operation has occurred and system control errors should not be encountered.
		1	The PFC entry with asa≠0 was reloaded and a stage attempt has not yet been made. If a system control error is detected, the probable cause is that the asa value is obsolete, probably because an obsolete dump tape was used during the reload.
AFPDR	51		Pseudorelease flag:
		0	The file can be attached immediately if it has a disk image.
		1	The file has a disk image that cannot be attached until the MSF/7990 image of the file is staged to disk. After the stage attempt (whether or not it is successful), this flag is cleared.

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>																														
			<table border="1"> <thead> <tr> <th><u>Flag</u></th> <th><u>Bit</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>AFPDE</td> <td>50</td> <td>Data error flag:</td> </tr> <tr> <td></td> <td>0</td> <td>No unrecoverable read errors have been detected.</td> </tr> <tr> <td></td> <td>1</td> <td>An unrecoverable read error prevented the MSF/7990 image of the file from being staged to disk.</td> </tr> <tr> <td>AFPSE</td> <td>49</td> <td>System control error flag:</td> </tr> <tr> <td></td> <td>0</td> <td>No system error conditions have been detected.</td> </tr> <tr> <td></td> <td>1</td> <td>A system error condition prevented the MSF/7990 image of the file from being staged to disk.</td> </tr> <tr> <td>AFOBS</td> <td>48</td> <td>Obsolete flag; set to 1 whenever the file is attached in write mode.</td> </tr> <tr> <td></td> <td>0</td> <td>MSF/7990 image (if asa#0) is a current version of the file.</td> </tr> <tr> <td></td> <td>1</td> <td>Disk image is the current version of the file; the MSF/7990 image (if asa#0) is obsolete.</td> </tr> </tbody> </table>	<u>Flag</u>	<u>Bit</u>	<u>Description</u>	AFPDE	50	Data error flag:		0	No unrecoverable read errors have been detected.		1	An unrecoverable read error prevented the MSF/7990 image of the file from being staged to disk.	AFPSE	49	System control error flag:		0	No system error conditions have been detected.		1	A system error condition prevented the MSF/7990 image of the file from being staged to disk.	AFOBS	48	Obsolete flag; set to 1 whenever the file is attached in write mode.		0	MSF/7990 image (if asa#0) is a current version of the file.		1	Disk image is the current version of the file; the MSF/7990 image (if asa#0) is obsolete.
<u>Flag</u>	<u>Bit</u>	<u>Description</u>																															
AFPDE	50	Data error flag:																															
	0	No unrecoverable read errors have been detected.																															
	1	An unrecoverable read error prevented the MSF/7990 image of the file from being staged to disk.																															
AFPSE	49	System control error flag:																															
	0	No system error conditions have been detected.																															
	1	A system error condition prevented the MSF/7990 image of the file from being staged to disk.																															
AFOBS	48	Obsolete flag; set to 1 whenever the file is attached in write mode.																															
	0	MSF/7990 image (if asa#0) is a current version of the file.																															
	1	Disk image is the current version of the file; the MSF/7990 image (if asa#0) is obsolete.																															
alt	10	41-36	Alternate storage type:																														
		0	File resides on disk only.																														
		1	File has an MSF image.																														
		2	File has a 7990 image.																														
asa	10	35-0	Alternate storage address:																														
		0	An MSF/7990 image of the file does not exist.																														
		Non-zero	Specifies the location of the beginning of the file on MSF/7990. The AFOBS flag indicates whether or not this MSF/7990 image is current.																														

<u>Field</u>	<u>Word</u>	<u>Bit(s)</u>	<u>Description</u>
al†	11	38-36	File access level: 0 Access level 0. 1 Access level 1. . . 7 Access level 7.
access categories	11	31-0	File 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.
user control word	16	59-0	User control information.

† 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.

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.

<u>Entry</u>	<u>Description</u>
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.

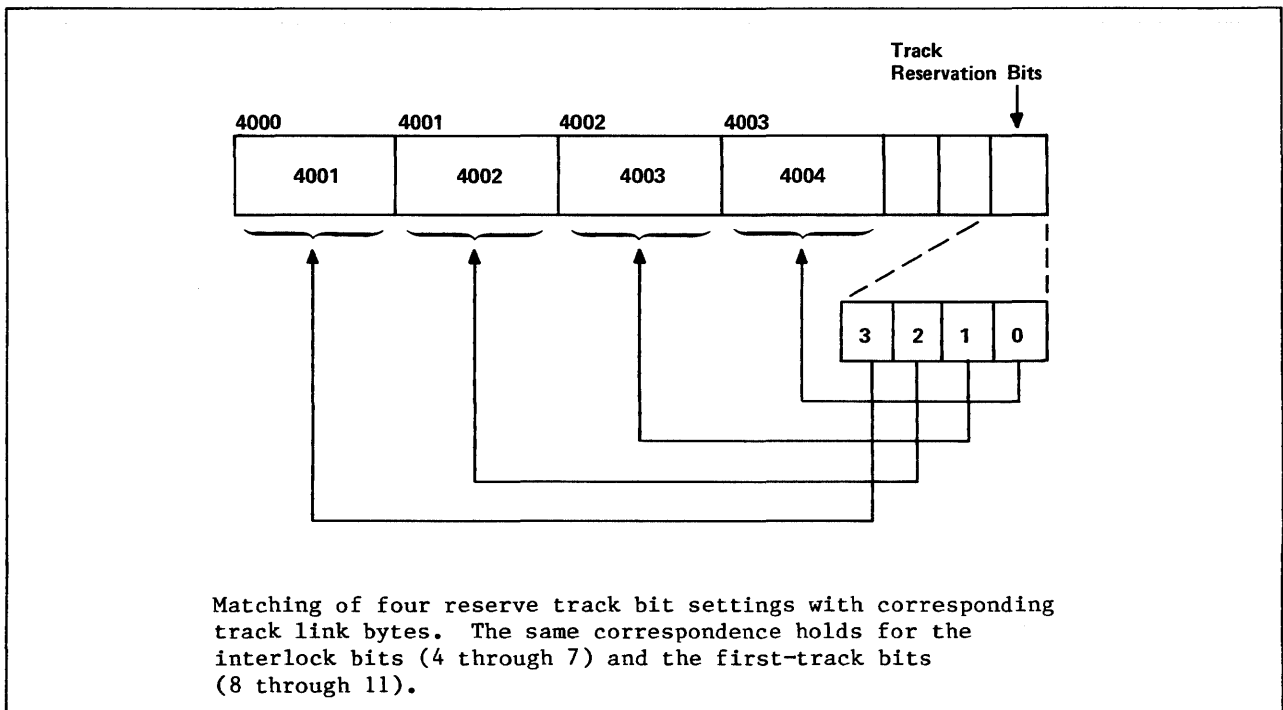



Figure 17-4. Bit Settings for Track Link Bytes



EXAMPLE : THIS ILLUSTRATES THE SEQUENTIAL LINKAGE

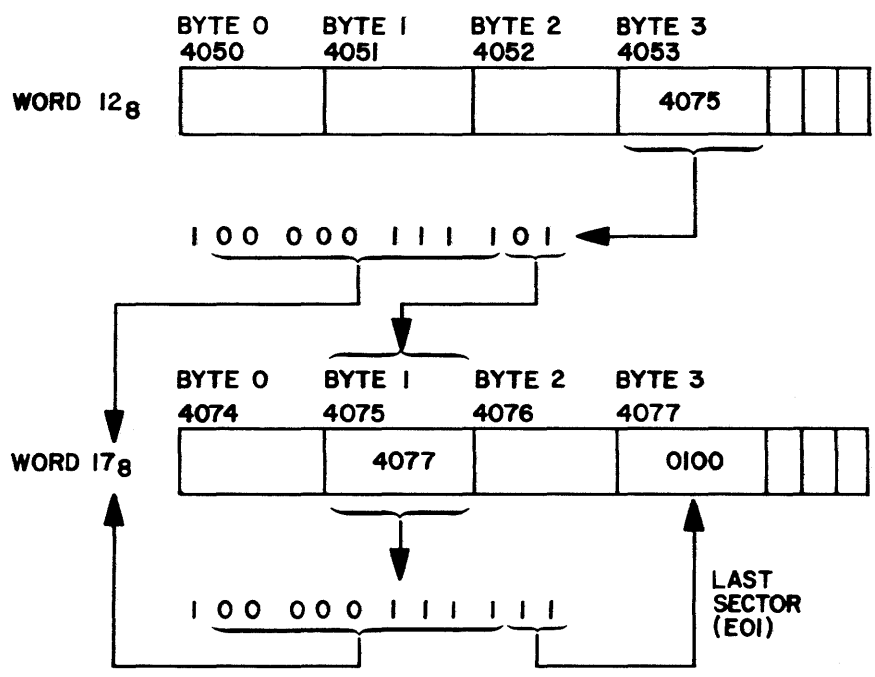


Figure 17-5. Linkage Format With Example

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. † Any job can change its associated family name by using the USER command to specify the new family name. †† 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.

<u>Parameter</u>	<u>Description</u>
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.

† The default family name is specified at deadstart time with the EQPDECK FAMILY entry and cannot be changed dynamically.

†† 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 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.

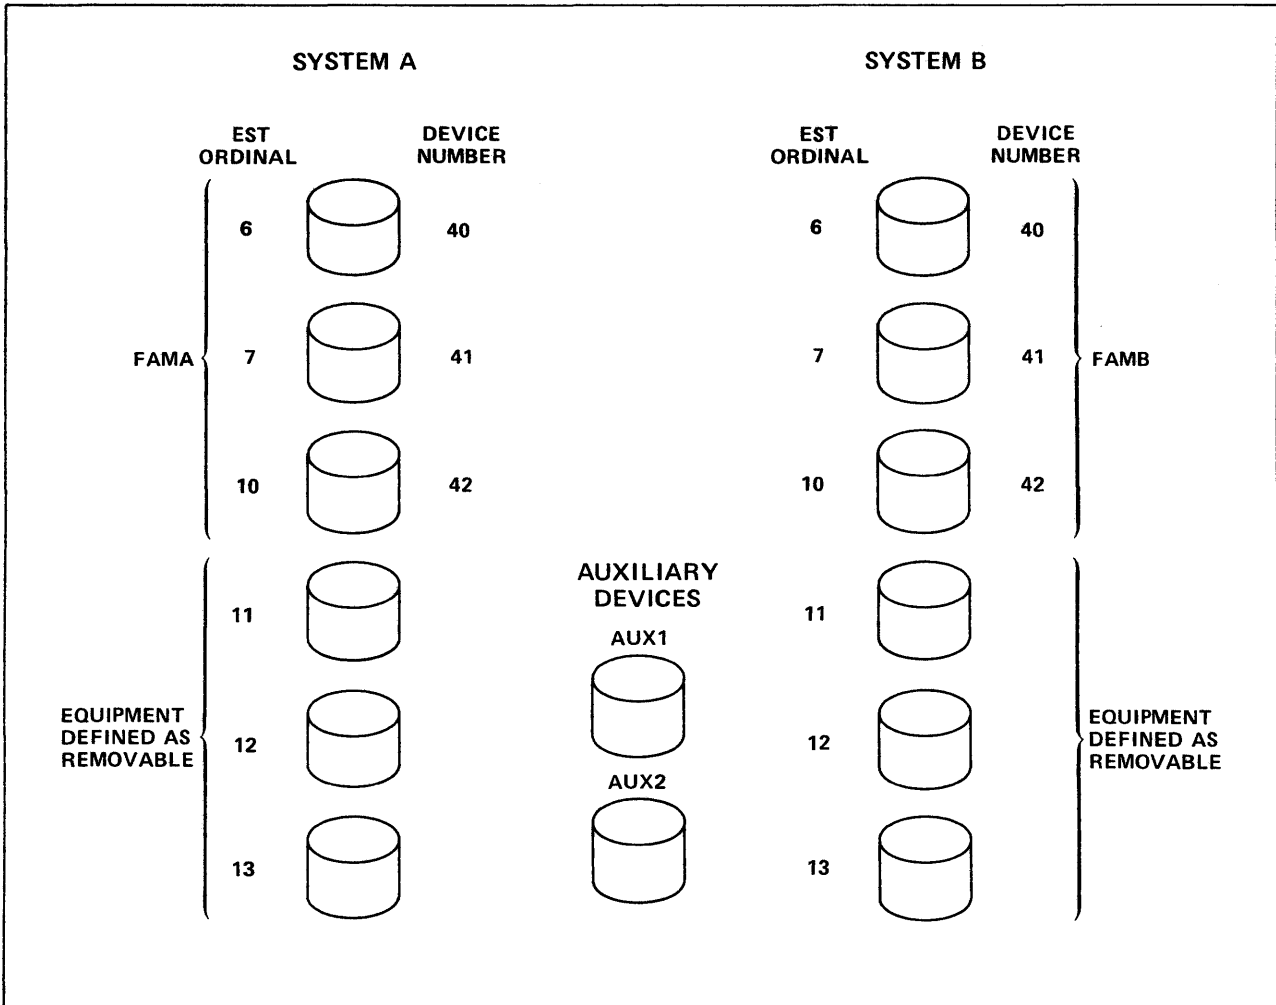


Figure 17-6. Example of Families in Two Configurations

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

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 is the 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 at the bottom of the display describe how to select the desired utility.

*** UTILITY OPTIONS ***

CURRENT OPTION VALUE	DESCRIPTION
-------------------------	-------------

ENTER K.UO. WHERE *UO* IS

AT - TAPE CAT	CA - DEVICE CAT
CP - TAPE COPY	DU - DUMP LD - LOAD

Figure 17-8. Initial Permanent File Utilities K Display (Left Screen)

3. Activate the right K display by entering:

KK.

The display shown in figure 17-9 appears on the right screen. A description of the permanent file utility options is shown.

*** PERMANENT FILE UTILITY OPTION DESCRIPTION ***

OPTION	DESCRIPTION
UT *UTILITY*	*LD* - PFLOAD *DU* - PFDUMP *CA* - PFCAT *AT* - PFATC *CP* - PFCOPY
LO *LIST OPTION*	T - FILES PROCESSED CATALOG C - PERMANENT FILE DEVICES CATALOG E - ERRORS S - SUMMARY
OP *UTILITY OPTIONS*	C - CREATION* A - LAST ACCESS* M - LAST MODIFICATION* I - INDIRECT ACCESS** D - DIRECT ACCESS** L - LOAD LEVELING P - PURGE AFTER DUMP R - REPLACE Q - ADD CATALOG AND PERMIT RECORDS E - EXTRACT CATALOG IMAGE RECORD O - OMIT CATALOG IMAGE RECORD S - SUPPRESS MSF STAGING Z - ZERO OUT ASA LINKAGE

NOTE - OPTIONS ENTERED AS FOLLOWS -

LO = TCS OR
OP = CIP

* ONLY ONE TYPE OF DATE MAY BE SPECIFIED
** ONLY ONE ACCESS OPTION MAY BE SPECIFIED

Figure 17-9. Permanent File Utilities K Display (Right Screen)

4. Select the desired utility by entering:

K.uo.

uo is one of the following:

<u>uo</u>	<u>Description</u>
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.

The parameter options available under the chosen utility appear on the left screen. Figure 17-10 shows the left screen after the DU utility has been called ①. The TCE appearing after the LIST OPTIONS VALID ② 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 CAMIDPS ③ appearing after the OPTIONS VALID indicates that

C	Creation
A	Last access
M	Last modification
I	Indirect access files
D	Direct access files
P	Purge after dump
S	Suppress staging of files

are the options available for the dump utility. Different combinations of utility options are available with the other utilities.

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).

6. Initiate execution by entering:

K.GO.

The lower lines of the left display ④ disappear and are replaced by:

DEVICE	MASK
DATE	TIME

Values appear with these identifiers as processing continues.

*** PFDUMP OPTIONS ①†***

CURRENT OPTION VALUE	DESCRIPTION
FM = 0	FAMILY NAME
PN = 0	PACK NAME
DN = 0	DEVICE NUMBER
TD = 0	TRUE DEVICE NUMBER
UN = 0	USER NUMBER
UI = 0	USER INDEX
DD = -- N/A --	DESTINATION DEVICE NUMBER
DI = -- N/A --	DESTINATION USER INDEX
LO = 0	LIST OPTIONS VALID - TCE ②
L = OUTPUT	OUTPUT FILE NAME
OP = 0	OPTIONS VALID - CAMIDPS ③
EO PROCESS	ERROR OPTION
SD NO SETTING	DISK SPACE RELEASE DATE
UD -- N/A --	UTILITY CONTROL DATE
AD = 0	AFTER DATE YYMMDD
AT = 0	AFTER TIME HHMMSS
BD = 0	BEFORE DATE YYMMDD
BT = 0	BEFORE TIME HHMMSS
PF = NAME	PERMANENT FILE NAMES - 24 ALLOWED
PF =	
PF =	
PF =	
PF =	
PF =	

NOTE - N/A DENOTES INVALID PARAMETER ④

*** PFDUMP OPTIONS ***

CURRENT OPTION VALUE	DESCRIPTION
T = TAPE	ARCHIVE FILE NAME
SF = 0	NUMBER OF FILES TO SKIP
N = -- N/A --	NUMBER OF FILES TO PROCESS
VF = 0	VERIFY FILE NAME
MF = -- N/A --	MASTER FILE NAME
RD = 0	RELEASE DATA FILE NAME
LA = 0	LOWER ACCESS LEVEL
UA = 0	UPPER ACCESS LEVEL

NOTE - N/A DENOTES INVALID PARAMETER

† The circled numbers are identified in the text preceding the figure.

Figure 17-10. PFDUMP Left K Display

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.

uo is one of the following utility options, and opt_i=val_i is a desired parameter option and value for the selected utility (refer to Parameters for the Permanent File Utilities later in this section).

<u>uo</u>	<u>Description</u>
ATC	Catalog archive file.
CAT	Catalog permanent files.
COPY	Copy archive file to job as local files.
DUMP	Dump permanent files.
LOAD	Load permanent files.

If an error in parameters is detected at initiation of a utility, control is returned to PFS and the correct parameters can then be entered 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 377778), 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 DUMPS.
```

```
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	Utility				
	PFATC	PFCAT	PFCOPY	PFDUMP	PFLoad
AD	X	X	X	X	X
AT	X	X	X	X	X
BD	X	X	X	X	X
BT	X	X	X	X	X
DD					X
DI					X
DN		X		X	X
EO				X	X
FM		X		X	X
L	X	X	X	X	X
LA	X	X	X	X	X
LO	X	X	X	X	X
MF			X		
N	X		X		X
OP†	X	X	X	X	X
PF	X	X	X	X	X
PN		X		X	X
RD				X	
SD				X	
SF	X		X	X	X
T	X		X	X	X
TD				X	X
UA	X	X	X	X	X
UD					X
UI	X	X	X	X	X
UN	X	X	X	X	X
VF				X	

† 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

<u>Parameter</u>	<u>Description</u>
AD=yymmdd †	After date in the form year month day. Files with a last access date (if OP=A), a creation date (if OP=C), or a utility control date (if OP=M) more recent than this date are to be processed. If AT=hmmss is specified, the default is the current date. If neither AD nor AT are specified, after date and time are not used as selection criteria.
AT=hmmss †	After time in the form hour minute second. Files meeting the AD=yymmdd criterion and with a last access time (if OP=A), a creation time (if OP=C), or a utility control time (if OP=M) more recent than this time are to be processed. If AD=yymmdd is specified, the default is 000000 (midnight). If neither AD nor AT are specified, after date and time are not used as selection criteria.
BD=yymmdd †	Before date in the form year month day. Files with a last access date (if OP=A), a creation date (if OP=C), or a utility control date (if OP=M) prior to this date are to be processed. If BT=hmmss is specified, the default is the current date. If neither BD nor BT are specified, before date and time are not used as selection criteria.
BT=hmmss †	Before time in the form hour minute second. Files meeting the BD=yymmdd criterion and with a last access time (if OP=A), a creation time (if OP=C), or a utility control time (if OP=M) prior to this time are to be processed. If BD=yymmdd is specified, the default is 000000 (midnight). If neither BD nor BT are specified, before date and time are not used as selection criteria.
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 under which PFLOAD loads all files being processed. Default is 0 (the destination user index is the user index from which the file was 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 should always be specified for PFCAT (no default).

† When entering dates and times, always enter six characters. For example, 3:30 a.m. should be entered as 033000.

<u>Parameter</u>	<u>Description</u>										
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.										
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).										
	<table border="0"> <thead> <tr> <th><u>opt</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>C</td> <td>List all files in the catalog image record. This option is used only with PFATC, PFDUMP, and PFLOAD.</td> </tr> <tr> <td>E</td> <td>List errors.</td> </tr> <tr> <td>S</td> <td>List cumulative statistics. This option is used only with PFCAT.</td> </tr> <tr> <td>T</td> <td>List PFC data for all files processed.</td> </tr> </tbody> </table>	<u>opt</u>	<u>Description</u>	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.
<u>opt</u>	<u>Description</u>										
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.										
MF=filename	<p>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.</p> <p>This option allows you to extract a file from an archive file and change its name as a local file.</p>										
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).										

Parameter

Description

Only one of options A, C, or M 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), or a utility control date/time (if OP=M) 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), or a utility control date/time (if OP=M) 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), or a utility control date/time (if OP=M) 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.

See the descriptions of the AD, AT, BD, and BT parameters for more information.

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. |

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.

Only one of the following two options can be used at a time. They can be used in conjunction with the A, C, or M option.

opt

Description

- | | |
|---|------------------------------------|
| D | Select direct access files only. |
| I | Select indirect access files only. |

Parameter

Description

The following options are used only with PFDUMP. Only one of them can be used at a time.

<u>opt</u>	<u>Description</u>
P †	Purge after dump. All files for which file data is included in the dump are purged after the dump is completed. 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 files with media-dependent backup requirement (BR=MD) are copied to disk and fully dumped.
S ††	Suppress staging of files. If a file to be dumped resides on MSF/7990 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 MSF/7990 file is copied to disk and included in the dump. This option should normally be specified for full dumps.
Y †	This option causes files with backup requirement BR=MD to be processed the same as files with backup requirement BR=Y.

The following five options are used only with PFLOAD.

<u>opt</u>	<u>Description</u>
R	Load with the replace option. If R is specified, PFLOAD loads all selected files from the archive file. 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.
Z ††	Zero the asa field in the PFC entry when the PFC entry is loaded, if the asa field points to an MSF/7990 image. Also suppresses the loading of PFC only files. Default is to leave the asa field 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.

† Refer to Permanent File Catalog Entry earlier in this section for a description of the backup requirement field.

†† Refer to section 11, MSAS, or to section 12, MSS, for more information.

Parameter

Description

The following option is used only with PFCOPY.

<u>opt</u>	<u>Description</u>
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	The 1- to 7-character name of a permanent file to be processed. The operator can enter up to 24 permanent file names. The selected file names appear on the bottom of the first page of the left K-display screen (refer to figure 17-10). To delete a permanent file name that has been entered, enter the same file name again. When used from one of the utility commands, up to 10 file names can be entered. The format of the parameter is PF=filename ₁ ,PF=filename ₂ ,...,PF=filename _n for both the K-display and the utility commands. PF is associated with UI. Default is that permanent file names are not selection criteria.
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.
RD=filename†	The 1- to 7-character name of the release data file (RDF) to be created by PFDUMP, which identifies those MSF/7990-resident files that are pointed to by PFC entries at the time of the dump. The MSS/MSAS utility ASVAL/SSVAL uses this file to identify unneeded MSF/7990-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.
SD	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 the MSS/MSAS utility ASMOVE/SSMOVE to inhibit the releasing of the disk image of any file with a BR=Y attribute which was last modified after this date/time. The disk space can be released after the next incremental or full dump is taken.
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.
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.
TD=dn	The 1- or 2-digit octal number that specifies the device within the family to be dumped or loaded. 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).

†Refer to section 11, MSAS, or to section 12, MSS, for more information.

<u>Parameter</u>	<u>Description</u>
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. Refer to PFDUMP in a Secured System later in this section.
UD †	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.
UI=userindex	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.
UN=username	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.
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 2 Reference Set, Volume 3 for a complete description of VERIFY.)

<u>Command</u>	<u>Description</u>
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.

† Refer to section 11, MSAS, or to section 12, MSS, for more information.

Parameter

Description

Command

Description

VERIFY,AA,BB,N=0,A,R.

A binary comparison of AA and BB is performed. If words do not match, this command lists:

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, 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=PFDDUMP,NT,D=1600.

or

ASSIGN,NT,TAPE.

If no archive file is present, PFDDUMP 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.

NOTE

The archive file is not rewound or unloaded before or after processing by any of the utilities.

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	ucdate
password	count	userindex	mode	subsystem	time	time	time	time	time
exp date	level	pr br rs ac				flags		at asa	amsg
charge	project								

Figure 17-11 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.

<u>Header Information</u>	<u>Description</u>
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 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, or EXECUTE. These modes are explained in detail in the NOS 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.
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.

Header InformationDescription

(14)	cmdate time	Date and time of the last control information modification made in the file's PFC entry or permit entries.
(15)	update 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 specified by the file owner. This field can be D (disk residence preferred), L (file locked to disk), M (MSF/7990 residence preferred), or N (no preferred residence).
(19)	br	Backup requirement for the file specified by the file owner. This field can be either Y (backup required on dump tape), MD (backup required only if current version of file does not reside on MSF/7990), or N (no backup is to be performed).
(20)	rs	Current residence of the file. This field can be D (disk only), A (MSF/7990 only), or B (both disk and MSF/7990).
(21)	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).
(22)	flags	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:

<u>Field</u>	<u>Description</u>
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)	at	Alternate storage type. This field can be NONE (no MSF/7990 image), MSF (file resides on MSF), or 7990 (file resides on 7990).
(24)	asa	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.

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	
EXPIRES	LEVEL	PR	BR	RS	AC	ASA	FLAG	CODES	
CHARGE NO.	PROJECT	NUMBER						TYPE	
								ASA	
								VALUE	
301	<u>BIN</u> ①	IND. PRIVATE		13	83/07/08.	83/07/08.	83/07/08.	83/07/08.	83/07/08.
		0 3006	READ		15.32.46.	15.32.46.	15.32.46.	15.32.46.	15.32.46.
		LVLO	N Y D N						NONE
	SYSTEM								
302	ISDWRK ②	<u>IND.</u> <u>PRIVATE</u> ③		12	83/07/08.	83/07/08.	83/07/08.	83/07/08.	83/07/08.
		0 3006	WRITE		15.32.46.	15.32.46.	15.32.46.	15.32.46.	15.32.46.
		LVLO	N Y D N						NONE
	SYSTEM								
303	DATA83	IND. PRIVATE		17	83/07/08.	83/07/08.	83/07/08.	83/07/08.	83/07/08.
		0 3006	WRITE		15.32.47.	15.32.47.	15.32.47.	15.32.47.	15.32.47.
		LVLO	N Y D N						NONE
	SYSTEM								
304	MODSETS	DIR. PUBLIC		236 * ⑤	83/07/08.	83/07/08.	83/07/08.	83/07/08.	83/07/08.
		0 3006	READ		15.32.48.	15.32.48.	15.32.48.	15.32.48.	15.32.48.
		LVLO	N Y D Y						NONE
	SYSTEM								
305	IBLKFR	IND. PRIVATE		6	83/07/08.	83/07/12.	83/07/12.	83/07/08.	83/07/12.
		⑦ ⑪ ⑧ ①	READ ⑨		10.01.25.	12.26.49.	12.26.40.	10.01.25.	12.26.40.
		LVLO	N Y D N						NONE
	SYSTEM								
306	SCBIN	IND. PRIVATE		2	83/01/27.	83/01/27.	83/01/27.	83/01/27.	83/01/27.
	⑥ <u>PASSW</u>	0 11	READ <u>BATCH</u>		14.56.09.	14.56.09.	14.56.09.	14.56.09.	14.56.09.
	⑬ <u>83/07/20.</u>	LVLO	N Y D N	⑩					NONE
	SYSTEM								
307	PF254A	IND. PRIVATE		1	⑪ 83/07/08.	⑫ 83/07/08.	⑬ 83/07/08.	⑭ 83/07/08.	⑮ 83/07/08.
		0 13	WRITE		⑮ 12.55.12.	⑯ 12.55.12.	⑰ 12.55.12.	⑱ 12.55.12.	⑲ 12.55.12.
		⑰ LVLO	N Y D N						NONE
	SYSTEM								
308	A	IND. PRIVATE		0	78/11/06.	78/11/06.	78/11/06.	78/11/06.	78/11/06.
		0 13	WRITE		15.48.31.	15.48.31.	15.48.31.	15.48.31.	15.48.31.
		LVLO	N Y D N		⑳			㉓	NONE
	SYSTEM								
309	G	IND. PRIVATE		1	79/08/21.	79/09/10.	79/08/21.	79/08/21.	79/08/21.
		2 15	WRITE		13.20.23.	12.26.44.	13.20.23.	13.20.23.	13.20.23.
		LVL2	N Y D N						NONE ㉔
	SYSTEM								㉕
310	X	IND. PUBLIC		1	79/02/05.	79/02/05.	79/02/05.	79/02/05.	79/02/05.
		1 13	WRITE		16.51.07.	16.55.59.	16.51.07.	16.55.59.	16.51.07.
		LVLO	N Y D Y						NONE
	⑲ 1234								
		⑲ 1ABC2345		㉗					

Figure 17-11. 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 17-11. 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 17-16.

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 17-12. 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 17-13 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 17-14 lists a page of the information given for each user index. Figure 17-15 lists the general device information.

```

CATALOG FILE
DIRECTORY OF PERMANENT FILE DEVICE 40 SYS606          YY/MM/DD. HH.MM.SS.
CATALOG OF USER INDEX 4054                          PAGE 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 RS AC ASA FLAG CODES TYPE ASA VALUE
CHARGE NO. PROJECT NUMBER

1 VERTICAL IND. PUBLIC 1 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 WRITE 13.01.32. 13.01.32. 13.01.32. 13.01.32. 13.01.32.
LVLO N Y D Y --- --- --- --- NONE
54321 2AAA3BBB

2 AUCTEST IND. PRIVATE 2 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 WRITE 13.01.29. 13.01.29. 13.01.29. 13.01.29. 13.01.29.
LVLO N Y D N --- --- --- --- NONE
SYSTEM

3 SETB DIR. PRIVATE 56 * 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 READ 13.01.34. 13.01.34. 13.01.34. 13.01.34. 13.01.34.
LVLO N Y D N --- --- --- --- NONE
SYSTEM

4 CATLI7 IND. PUBLIC 5 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 WRITE 13.01.29. 13.01.29. 13.01.29. 13.01.29. 13.01.29.
LVLO N Y D Y --- --- --- --- NONE
SYSTEM

5 TESTINP IND. PRIVATE 6 83/03/02. 83/03/30. 83/03/02. 83/03/02. 83/03/02.
6 WRITE 15.29.09. 10.30.17. 15.29.09. 15.29.09. 15.29.09.
LVLO N Y D N --- --- --- --- NONE
SYSTEM

6 HISTORY IND. PRIVATE 6 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 READ 13.01.32. 13.01.32. 13.01.32. 13.01.32. 13.01.32.
LVLO N Y D N --- --- --- --- NONE
SYSTEM

7 WRKTEST IND. PRIVATE 3 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 READ 13.01.30. 13.01.30. 13.01.30. 13.01.30. 13.01.30.
LVLO N Y D N --- --- --- --- NONE
SYSTEM

8 NS2104B DIR. PRIVATE 231 * 83/07/12. 83/07/12. 83/07/12. 83/07/12. 83/07/12.
0 READ 13.01.36. 13.01.36. 13.01.36. 13.01.36. 13.01.36.
LVLO N Y D N --- --- --- --- NONE
SYSTEM

TOTALS 8 FILE(S), 310 SECTORS

```

Figure 17-12. Sample Directory Produced by PFCAT,LO=T,DN=40,UI=4054.

```

MASS STORAGE TABLE REPORT
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 17-13. Mass Storage Table Report

SUMMARY REPORT
 DIRECTORY OF PERMANENT FILE DEVICE 40 SYS964

82/01/11. 12.48.50.
 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 17-14. Cumulative Statistics by User Index PFCAT,LO=S,DN=40.

INDIRECT ACCESS FILE HOLES

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

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 17-15. Cumulative Statistics for Entire Device PFCAT,LOS=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.

Incremental Dump OP=M is specified and BD=yyymmdd and BT=hmmss are not specified.

Partial Dump Any dump that is not incremental and not full.

Full Dump OP=0 is specified or OP is not 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 (K.GO), 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. †

† 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 2 Operations Handbook for display formats.

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=yyymmdd and BT=hhmmss are not specified. Other options specified on the PFDUMP call can restrict the dump even further.

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 PFDUMP 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. 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 (PFDUMP) later in this section.

Partial Permanent File System Dump

A partial dump copies permanent files according to any specified option(s), 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. 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 a certain device (DN or TD option) or can dump all files in the system (DN and TD not specified) that meet 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. That is, the OP option is not specified (or OP=0), and either a particular device is specified (device dump) or no device (DN=0 or TD=0) is specified (system dump).

PFDUMP interlocks a catalog track during the time it takes to dump all files whose PFC entries reside on that catalog track. No files whose PFC entries reside there are accessible until PFDUMP begins processing the next catalog track. 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 in an MSS/MSAS Environment

In an MSS/MSAS environment, the data for a permanent file to be dumped may reside on MSF/7990 and not on disk. However, if the file is to be copied to a dump file, it must have a disk image. The analyst can either stage the file to disk from MSF/7990 in order to include its data on the dump file (OP=S omitted), or copy to the dump file only the PFC and permit entries for the file and not the file data (OP=S specified). Typically, OP=S is specified for full dumps; otherwise, all MSF/7990 files would be included in the dump, which is not feasible because of the time involved to dump such a large permanent file base.

NOTE

PFDUMP will dump only the PFC and permit entries for a direct access permanent file if that file has an alternate storage copy (MSF/7990 image) and is attached in write mode at dump time. The alternate storage obsolete flag (AFOBS) is cleared in the PFC written to the dump file.

As described in the NOS 2 Reference Set, Volume 3, the file owner can supply a backup requirement for a file using the BR parameter on the DEFINE, CHANGE, or SAVE command. Options include backup on a dump file, backup only if the current version of the file is not on MSF/7990, or no backup on a dump file. This partially determines whether or not a file is included in a particular PFDUMP run. Refer to section 12, MSS, or to section 11, MSAS, for further information on how the PFDUMP utility is used in an MSS/MSAS environment.

PFDUMP Protected Files

When PFDUMP specifies purging of files (OP=P), files whose user indexes are greater than PGUI (default value is 300000g) 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 17-11. If LO=C (list all files in the catalog image record) is specified, the listing produced is similar to that shown in figure 17-16.

PFATC CATALOG OF CATALOG IMAGES					PAGE	4
ON	YY/MM/DD.	AT	HH.MM.SS.			
FILE NAME	USER INDEX	ACCESS COUNT	LAST ACCESS DATE	LAST ACCESS TIME	DEVICE NUMBER	
76	SYSB	6000	11	79/11/04.	12.42.09.	
77	IBL2	1	6	82/09/23.	03.12.33.	
78	IBL4	1	3	82/09/23.	03.12.33.	
79	SET2	3	0	83/02/08.	06.31.35.	
80	OVWO	3	0	83/02/08.	06.36.17.	
81	PROCFIL	1	24	82/02/22.	17.26.11.	
82	OVWRTLD	3	1	83/02/08.	06.43.56.	
83	PACBIN	3006	2	82/06/01.	00.34.46.	
84	PRB150	1	32	83/03/21.	12.50.37.	
85	FNTNOS	1000	0	81/10/21.	10.04.42.	
86	MODVMOD	1000	3	81/10/27.	13.01.14.	
87	LIST	1000	0	81/11/07.	13.14.07.	
88	FNT2MOD	1000	0	81/11/07.	13.17.22.	
89	OVWRTMZ	3	1	83/02/08.	06.44.08.	
90	OVWRTNL	3	1	83/02/08.	06.44.08.	
91	MOD	3006	1	83/06/16.	14.42.54.	
92	JOB	1	1	83/04/27.	08.10.28.	
93	IBL	1	3424	83/06/14.	05.18.48.	
94	IBLFNT	1	47	82/01/07.	13.41.20.	
95	XED	3006	0	83/07/08.	15.32.31.	
96	XEDITH	3006	0	83/07/08.	15.32.41.	
97	STM150	1	3	83/03/29.	18.56.16.	
98	IBLFNT3	1	0	82/01/06.	18.51.52.	
99	PRB200	1	15	82/06/03.	03.36.05.	
100	PACOPL	3006	2	82/07/15.	01.19.15.	

Figure 17-16. Catalog of Catalog Image Record

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.

Catalog Image Record

Each incremental dump (OP=M) 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 on the device(s) being dumped at the time of the incremental dump (not just files included in the dump).

The format of each entry is:

59	41	35	17	0
filename			userindex	
access count		dn	date time	

<u>Field</u>	<u>Description</u>
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 of the device being loaded. No other information can be changed over dumps and loads without updating the utility control date and time.

Figure 17-16 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 17-11 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 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 archive file mask (written in the header record by PFDUMP) and the load selection mask. 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.

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 is a master device, all files (direct and indirect) 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. If both the DN and TD parameters are specified, only files that meet the previously described criteria for both parameters are selected.

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 MSS/MSAS Environment

In an MSS/MSAS 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. Files that resided only on MSF/7990 at the time of the dump may or may not have been copied to the dump tape. If OP=S was specified on the PFDUMP call, then only the PFC and permit entries for the MSF/7990 files were included on the dump tape. 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 12, MSS, or to section 11, MSAS, for further information on how the PFLOAD utility is used in an MSS/MSAS 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.

<u>Device</u>	<u>Description</u>
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 17-17 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.

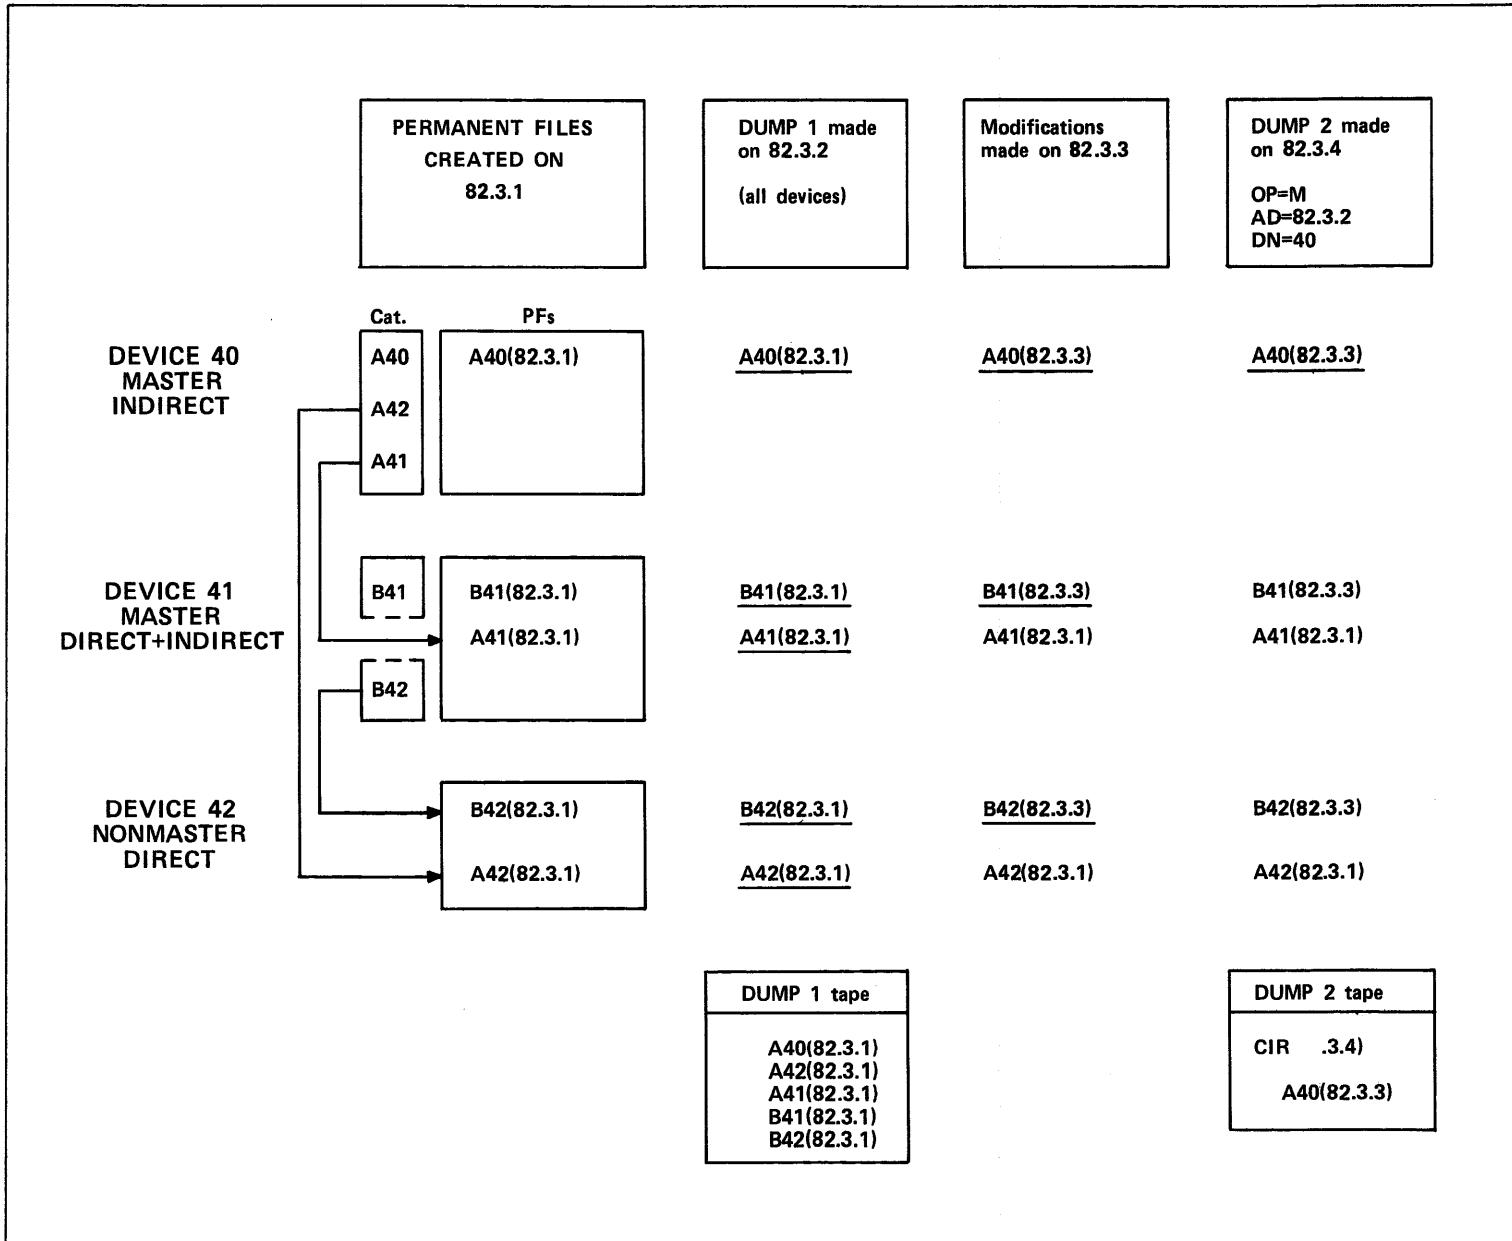


Figure 17-17. PFDUMP/PFLOAD Example (Sheet 1 of 3)

	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
DEVICE 40 MASTER INDIRECT	A40(82.3.3)	A40(82.3.3)	<u>A40(82.3.5)</u>	<u>A40(82.3.5)</u>	A40(82.3.5)
DEVICE 41 MASTER DIRECT+INDIRECT	<u>B41(82.3.3)</u> A41(82.3.1)	B41(82.3.3) A41(82.3.1)	<u>B41(82.3.5)</u> A41(82.3.1)	<u>B41(82.3.5)</u> A41(82.3.1)	B41(82.3.5) A41(82.3.1)
DEVICE 42 NONMASTER DIRECT	<u>B42(82.3.3)</u> A42(82.3.1)	<u>B42(82.3.3)</u> A42(82.3.1)	<u>B42(82.3.5)</u> A42(82.3.1)	<u>B42(82.3.5)</u> A42(82.3.1)	<u>B42(82.3.5)</u> A42(82.3.1)
	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 17-17. PFDUMP/PFLOAD Example (Sheet 2 of 3)

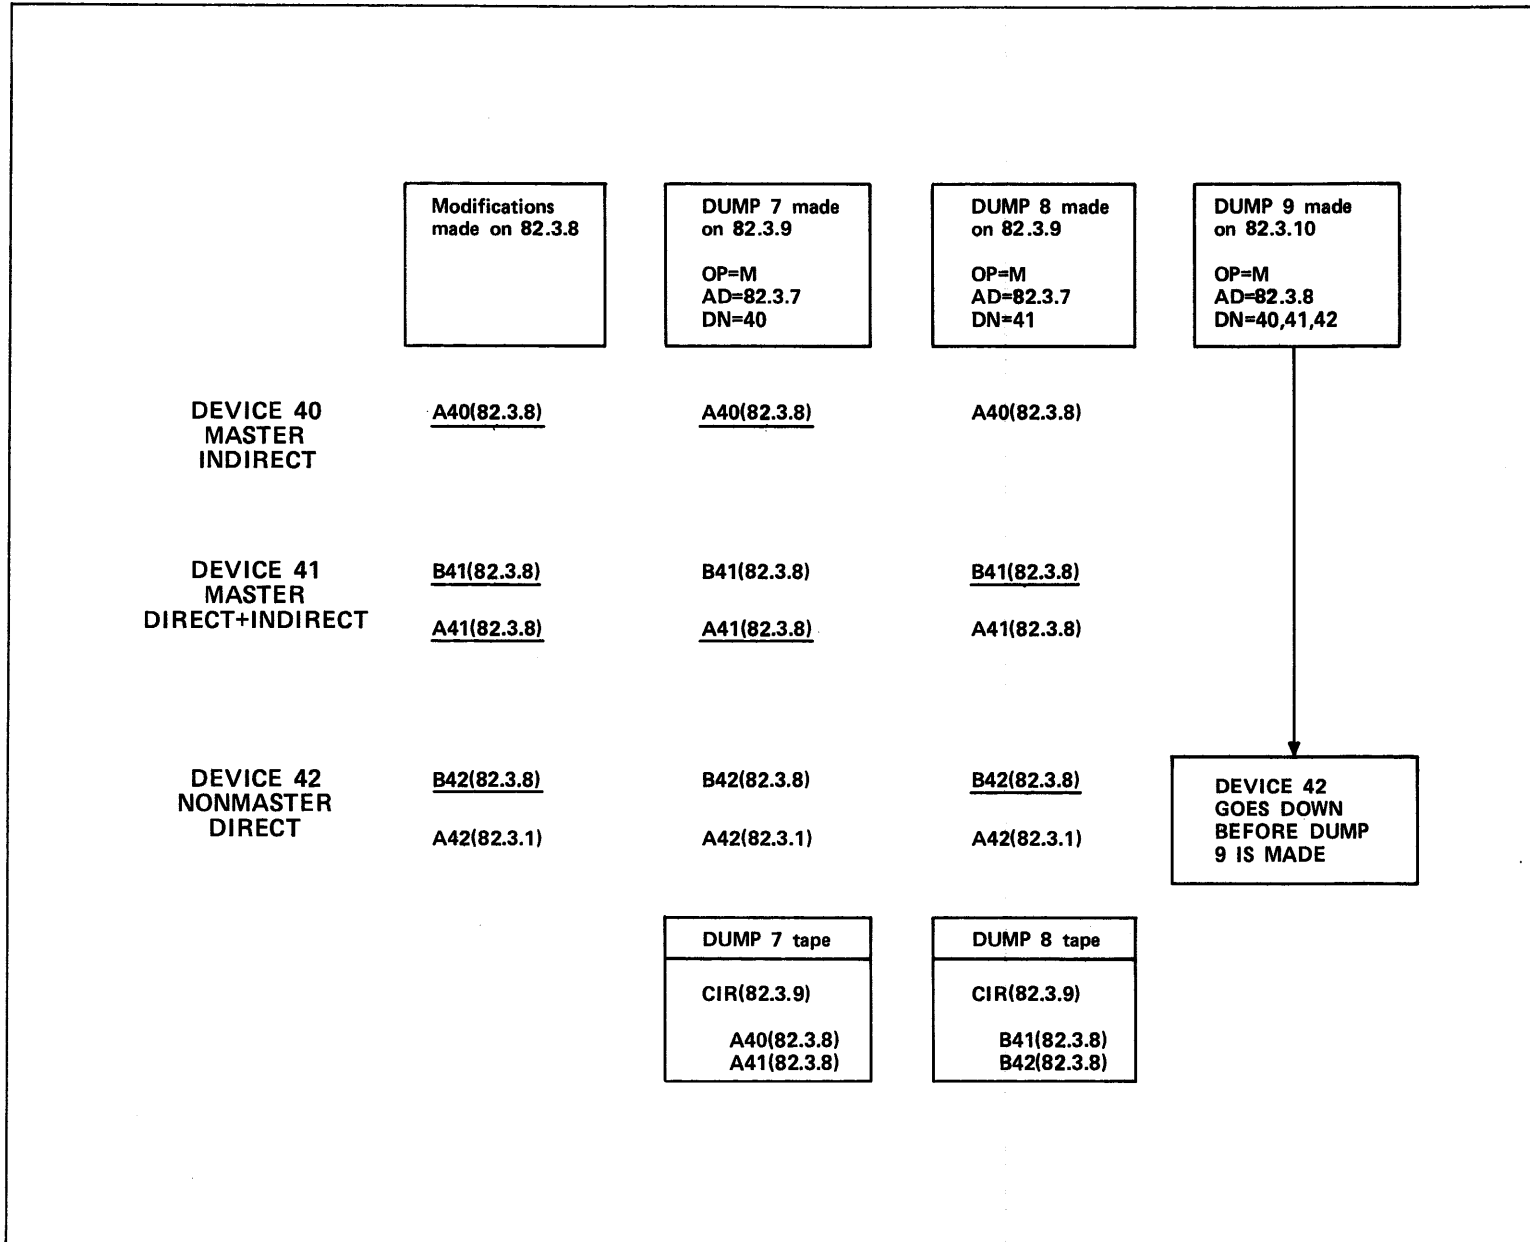


Figure 17-17. PFDUMP/PFLOAD Example (Sheet 3 of 3)

Actions during the 10 days in March in which the dumps and modifications were made run as follows:

<u>Date</u>	<u>Action</u>
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.
82.3.9	DUMP 7. An incremental dump of device 40 is made with AD=82.3.7. Two files qualify.

<u>Date</u>	<u>Action</u>
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:

<u>Archive File</u>	<u>Action</u>
DUMP8 tape	The analyst assigns DUMP8 tape and enters the command:
CIR(82.3.9)	PFLOAD,T=DUMP8,DN=42.
B41(82.3.8) B42(82.3.8)	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:

B42	userindex
-----	-----------

The utility then requests the next reel.

DUMP7 tape	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.
CIR(82.3.9)	
A40(82.3.8) A41(82.3.8)	
DUMP6 tape	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.
CIR(82.3.7)	
B42(82.3.5)	

Archive File

Action

DUMP5 tape
CIR(82.3.6)
A40(82.3.5)
B41(82.3.5)
B42(82.3.5)

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.

DUMP4 tape
CIR(82.3.4)
B42(82.3.3)

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.

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.

Queue File Utilities	18-1
Calling the Queue File Utilities	18-2
Console Input	18-2
Batch Input	18-9
Terminal Input	18-9
Parameters for the Queue File Utilities	18-9
Description of Queue File Utility Parameters	18-11
File Preassignment	18-20
Active and Inactive Queued Files	18-20
Queue File Utility Routines	18-20
DFLIST	18-20
DFTERM	18-21
LDLIST	18-22
QALTER	18-23
QDUMP	18-27
QFTLIST	18-31
QLIST	18-32
QLOAD	18-32
QMOVE	18-34
QREC	18-38
Dayfile Dumping Utilities	18-43

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 queue file utilities are 10 utilities under the control of the queue 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.

QUEUE FILE UTILITIES

The names and functions of the queue file utilities follow. The rest of this section includes detailed information about the call and operation of each utility.

<u>Utility</u>	<u>Description</u>
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 QUEUE FILE UTILITIES

All queue file utility operations, except DFLIST, LDLIST, and QLIST, can be initiated through console input to a K display. All queue 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.

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,0 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. (jsn is the job sequence name of QFSP as specified on the B,0 display.)

The display shown in figure 18-1 appears on the left screen. Instructions at the bottom of the display describe how to select the desired utility.

*** QUEUED FILE SUPERVISOR ***

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.

SELECT DESIRED UTILITY BY ENTERING K.OPTION
WHERE OPTION IS ONE OF THE UTILITIES LISTED
ABOVE.

Figure 18-1. QFSP K Display (Left Screen)

3. Select the desired utility by entering:

K.utility. (utility is 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 18-3, 18-9, 18-12, 18-14, and 18-17 exhibit initial K displays for the various utilities. Different parameters are available with the different utilities. (Refer to Description of Queue File Utility Parameters 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. (utility is one of the utilities shown in figure 18-1.)

The B,0 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. (jsn is the job sequence name as specified on the B,0 display.)

4. Several commands are available to aid in the use of the utilities and direct processing. A right screen display lists and defines each available command. Activate this K display by entering:

KK.

Figure 18-2 shows the resultant display.

Following is a list of the queue file utility commands included on the display and a description of each.

<u>Command</u>	<u>Description</u>
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.
STOP	Terminates the active utility and ends the K-display interaction.
RESET	Resets all options displayed on the left screen K display to their default values.
LIST	If the DFTERM utility is active, this command displays a list of all permanent dayfiles (refer to figure 18-4). The list of files displayed is also written on the output file specified by the DFTERM L parameter.

<u>Command</u>	<u>Description</u>
	<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 18-6). 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 18-18). 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 18-7) 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.
+	Used in conjunction with the LIST command; displays succeeding pages (screens) of the information listed. This command applies only to the QREC, QALTER, and QFTLIST utilities.
*	Enables you to page the left screen K display. This command applies only to the QREC, QDUMP, QLOAD, QMOVE, QFTLIST, and QALTER utilities.
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.
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.
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.
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.

<u>Command</u>	<u>Description</u>
REWIND	Rewinds the load or dump file specified by the FN parameter before processing. This command applies only to the QLOAD and QDUMP utilities.
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>

*** QUEUE FILE UTILITY COMMANDS *** PAGE 1 OF 2

THE FOLLOWING COMMANDS ARE PROVIDED-

- GO - PROCEED WITH PROCESSING (FM MUST BE SPECIFIED).
(ALL UTILITIES)
- STOP - TERMINATE INPUT, END RUN. (ALL UTILITIES)
- RESET - RESET PARAMETERS TO DEFAULT VALUES.
(ALL UTILITIES)
- LIST - DISPLAY JOB NAMES OF ACTIVE/INACTIVE QUEUED FILES
ON RIGHT SCREEN. (QREC, DFTERM, QALTER, QFTLIST).
- + - PAGE RIGHT SCREEN. (QREC, QALTER, QFTLIST, QFSP).
- * - PAGE LEFT SCREEN.
(QREC, QDUMP, QLOAD, QMOVE, QFTLIST, QALTER).
- CLEAR - SET BOTH SCREENS TO FIRST PAGE OF K-DISPLAY.
(QREC, DFTERM, QFTLIST, QALTER).
- OUT - RELEASE FILE SPECIFIED BY L PARAMETER TO
PRINTER. (ALL UTILITIES)
- 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).
- REWIND- IF ENTERED, DUMP/LOAD FILE WILL BE REWOUND WHEN
GO COMMAND IS ENTERED. (QDUMP AND QLOAD ONLY).
- ERROR - TOGGLE STATUS OF ERRORED FILE PROCESSING.
(QDUMP, QLOAD, QMOVE ONLY).

Figure 18-2. Queue File Utility Commands K Display (Sheet 1 of 2)

*** QUEUE FILE DIRECTIVE FORMATS *** PAGE 2 OF 2

THE FOLLOWING DIRECTIVES MAY HAVE ONE OF SEVERAL FORMS-

DA=YMMDD PROCESS THIS DATE ONLY.
DA=YMMDD-YMMDD PROCESS ALL DATES IN SPECIFIED RANGE.
DA=ALL PROCESS ALL DATES.

UI=XXXXXX PROCESS THIS USER 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.
IF 77777B IS SPECIFIED, ALL SIZES LARGER ARE INCLUDED.

ID=XX PROCESS THIS ID ONLY.
ID=XX-YY PROCESS ALL ID,S IN SPECIFIED RANGE.

FC=FORMS CODE.
FC=F1/F2/F3 SPECIFY FORMS CODE LIST.
FC=F1-F2 SPECIFY FORMS CODE RANGE.

DS=DT FILE DESTINATION=DISPOSITION TYPE.
FILE DESTINATION TYPES INCLUDE THE FOLLOWING-
BC - BATCH RB - REMOTE BATCH
DISPOSITION TYPES INCLUDE THE FOLLOWING-
IN - INPUT PU - PUNCH
PL - PLOT PR - ANY PRINTER
P2 - 512 PRINTER LR - 580-12 PRINTER
LS - 580-16 PRINTER LT - 580-20 PRINTER
WT - WAIT QUEUE SF - SPECIAL QUEUE TYPES
ALL - ALL QUEUES PRINT - ALL PRINT QUEUES
NONE - NO QUEUES

Figure 18-2. Queue File Utility Commands K Display (Sheet 2 of 2)

5. Select the desired parameters by entering:

K.opt₁=val₁,opt₂=val₂,...,opt_n=val_n.

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 Queue File Utilities later in this section).

If the FM parameter is specified for QDUMP, QMOVE, or QREC (indicating the processing of queue files for only the specified family of devices), the initial left screen display (refer to figures 18-9, 18-14, and 18-17) is replaced by the secondary left screen display (refer to figures 18-10, 18-15, and 18-19). 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.

6. 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.

7. Terminate use of the chosen utility and end K-display interaction by entering:

K.STOP.

An output file is generated indicating the disposition of queued files or dayfiles processed. Examples of output files generated by each utility are shown in figures 18-4, 18-11, 18-13, 18-16, and 18-20.

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,...,opt_n=val_n,I=filename,PO=N)

<u>Parameter</u>	<u>Description</u>
utility	Any queue or dayfile utility.
opt _i =val _i	Parameters (refer to Parameters for the Queue 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 queue/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,0 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. The following sequence of commands is used to call a utility.

```
ujn.  
USER,username,password,familyname.  
utility,opt1=val1,opt2=val2,...,optn=valn,I=filename,PO=N.
```

<u>Parameter</u>	<u>Description</u>
utility	Any of the following utilities: QREC, QDUMP, QLOAD, QMOVE, QLIST, or LDLIST.
opt _i =val _i	Parameters (refer to Parameters for the Queue 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 queue 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 QUEUE FILE UTILITIES

Table 18-1 specifies the parameters accepted by each queue 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. Queue File Utility Parameters

Parameter	Utility								
	QREC	QDUMP	QLOAD	QMOVE	QALTER	LDLIST	QFTLIST	DFTERM	QLIST
BC	X	X	X	X	X	X	X		X
DA	X	X	X	X		X			X
DD			X	X		X			
DF			X	X	X	X	X		
DN	X	X		X	X		X	X	X
FC	X	X	X	X	X	X	X		X
FM	X	X		X	X		X	X	X
FN		X	X			X			
FS	X	X	X	X		X			X
FT								X	
FU	X	X	X	X		X			X
ID	X	X	X	X	X	X	X		X
JSN	X	X	X	X	X		X		X
L	X	X	X	X	X	X	X	X	X
LA	X	X	X	X	X	X	X		X
LD	X	X	X	X	X	X	X		X
LO					X		X		
MI	X	X	X	X		X			X
NAL					X				
NDC					X				
NDF					X				
NEC					X				
NF		X	X			X			
NFC					X				
NID					X				
NLD					X				
NM								X	
NPR					X				
NRC					X				
NUN					X				
OP	X		X	X	X			X	
PO	X	X	X	X	X		X	X	
RB	X	X	X	X	X	X	X		X
SC			X			X			
TF	X	X	X	X					
TP		X		X					
TUI	X	X	X	X		X			X
UA	X	X	X	X	X	X	X		X
UI	X	X	X	X		X			X
UJN					X		X		
UN					X		X		

Description of Queue File Utility Parameters

<u>Parameter</u>	<u>Description</u>
BC=dt	Queued files destined for local batch devices can be selected for processing by disposition type (dt).

<u>dt</u>	<u>Description</u>
ALL	Files of all disposition types.
IN	Input files.
LR	580-12 printer files.
LS	580-16 printer files.
LT	580-20 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.

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.
-----------	--

<u>Parameter</u>	<u>Description</u>
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 (1 through 77g). 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.
FC=fc	Forms code for output files considered for processing. This parameter does not deselect input files.

<u>fc</u>	<u>Description</u>
ALL	All files are considered for processing regardless of forms codes.
fc ₁ /fc ₂ /fc ₃	fc ₁ 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 ₁ 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.
---------------	--

<u>Parameter</u>	<u>Description</u>
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 < 777778. Default is to consider all files for processing regardless of file size.
FT=t	Type of dayfile to be terminated by the DFTERM utility.

<u>t</u>	<u>Description</u>
ACCOUNT	Account dayfile.
DAYFILE	System dayfile.
ERRLOG	Error log.
MAINLOG	Binary maintenance log.

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.

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 678) 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.
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.

<u>Parameter</u>	<u>Description</u>																
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.																
LO=opt	List option associated with the output file specified by the L parameter.																
	<table border="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>opt</u></th> <th style="text-align: left;"><u>Description</u></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>	<u>opt</u>	<u>Description</u>	F	All routing information for each file.	S	Condensed listing of each selected queued file.										
<u>opt</u>	<u>Description</u>																
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="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>dt</u></th> <th style="text-align: left;"><u>Description</u></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>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>	<u>dt</u>	<u>Description</u>	LP	Any printer.	LR	580-12 printer.	LS	580-16 printer.	LT	580-20 printer.	NONE	No device code specified.	PR	Any printer.	P2	512 printer.
<u>dt</u>	<u>Description</u>																
LP	Any printer.																
LR	580-12 printer.																
LS	580-16 printer.																
LT	580-20 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.																

<u>Parameter</u>	<u>Description</u>
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> <ul style="list-style-type: none"> A4 ASCII graphic 48-character set. A6 ASCII graphic 63/64-character set. A9 ASCII graphic 95-character set. B4 BCD graphic 48-character set. B6 BCD graphic 63/64-character set. <p>For punch files:</p> <ul style="list-style-type: none"> AS Punch ASCII. PB Punch system binary. PH Punch 026 mode. P8 Punch 80-column binary. P9 Punch 029 mode. <p>For plot files:</p> <ul style="list-style-type: none"> 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.
NID=id	The 1- or 2-digit number (0 through 67g) 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.

<u>Parameter</u>	<u>Description</u>
NM=filename	<p>The 1- to 5-character 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). Default is automatic naming of the file by DFTERM as follows:</p> <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	<p>The 1- to 4-digit number (0 through 7777g) 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.</p>
NRC=c	<p>The 1- or 2-digit number (0 through 37g) 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.</p>
NUN=username	<p>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.</p>
OP=opt	<p>Processing option specifying the function to be performed by each utility.</p>

With QREC, opt can be any of the following. Default is OP=RI.

<u>opt</u>	<u>Description</u>
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.

Parameter

Description

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.

<u>opt</u>	<u>Description</u>
BC	Change to local batch.
NC	Do not change file destination.
PR	Purge files.
RB	Change to remote batch.

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.

<u>opt</u>	<u>Description</u>
A	Active queues/dayfiles are specified.
I	Inactive queues/dayfiles are specified.

PO=opt

Processing option specifying the function to be performed.

<u>opt</u>	<u>Description</u>								
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.								
E	Errored file processing status is toggled. Default conditions for queue utilities include: <table><thead><tr><th><u>Default</u></th><th><u>Utility</u></th></tr></thead><tbody><tr><td>Dump file.</td><td>QDUMP</td></tr><tr><td>Do not load file.</td><td>QLOAD</td></tr><tr><td>Do not move file.</td><td>QMOVE</td></tr></tbody></table>	<u>Default</u>	<u>Utility</u>	Dump file.	QDUMP	Do not load file.	QLOAD	Do not move file.	QMOVE
<u>Default</u>	<u>Utility</u>								
Dump file.	QDUMP								
Do not load file.	QLOAD								
Do not move file.	QMOVE								
F	Entries are made in FOT for inactive families and the files are queued.								
L	Entries are made in the LID table for inactive LIDs and the files are queued.								

<u>Parameter</u>	<u>Description</u>
RB=dt	Queued files destined for remote batch devices can be selected for processing by disposition type (dt).

<u>dt</u>	<u>Description</u>
ALL	Files of all disposition types.
IN	Input files.
LR	580-12 printer files.
LS	580-16 printer files.
LT	580-20 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.

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.
---------------	--

<u>Parameter</u>	<u>Description</u>								
TP=t	Type of files to move or dump. Default is TP=ALL.								
	<table border="0" style="margin-left: 40px;"> <thead> <tr> <th style="text-align: left;"><u>t</u></th> <th style="text-align: left;"><u>Description</u></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>	<u>t</u>	<u>Description</u>	A	Active files.	ALL	Active and inactive files.	I	Inactive files.
<u>t</u>	<u>Description</u>								
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.								
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 queue utilities), nor do they appear on any DSD display.

Three queue 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.

QUEUE FILE UTILITY ROUTINES

The queue 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. The 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 18-4).

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 18-3).

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 dayfile list K display (refer to figure 18-4). 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 (page), 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 queue 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 18-4, which shows essentially the same format.) If the LIST command was not entered, no output file is generated.

*** DAYFILE TERMINATION ***

ACTIVE DAYFILES-

TYPE	FAMILY	DN
DAYFILE	SYST72	1B
ACCOUNT	SYST72	1B
ERRLOG	SYST72	1B
MAINLOG	SYST72	1B

INACTIVE DAYFILES-

TYPE	FAMILY	DN
------	--------	----

OPTIONS

DESCRIPTION

FT = DAYFILE	FILE TYPE TO BE TERMINATED. (DAYFILE, ACCOUNT, ERRLOG, MAINLOG)
FM = SYST72	FAMILY FOR NEW DEVICE NUMBER THAT NEW DAYFILE IS TO RESIDE ON (1-7 CHARACTERS).
DN = 1B	DEVICE NUMBER OF NEW DEVICE (1-77B). FM AND DN DENOTE RESIDENT DEVICE DAYFILE IS INACTIVE.
OP = A	ACTIVE OR INACTIVE FILE TO BE TERMINATED (A OR I).
NM =	NAME OF PERMANENT FILE. OVERRIDES AUTOMATIC NAMING (1-5 CHARACTERS).
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).

Figure 18-3. DFTERM Display

*** PERMANENT DAYFILE CATALOG ***

NO.	TYPE	FM/PN	DN	PFNAME	DATE	TIME	LENGTH
1.	DAYFILE	SYST72	01	DFA0331	80/03/31	13.35.33	1171
2.	ACCOUNT	SYST72	01	ACA0331	80/03/31	13.35.42	402
3.	ERRLOG	SYST72	01	ERA0331	80/03/31	13.35.49	55
4.	MAINLOG	SYST72	01	MLA0331	80/03/31	13.35.52	1032

Figure 18-4. DFTERM Dayfile List

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,...,opt_n=val_n.

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 18-13).

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.

- Access level (requires security administrator privileges).
- File destination.
- File identifier (local batch-destined files).
- Destination family (remote batch-destined files).
- Destination user name (remote batch-destined files).
- Forms code.
- Priority of a queued file.
- Repeat count.
- Disposition type.
- External characteristics.
- Destination LID.

When QALTER is selected at the system console, the QALTER K display 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 figure 18-5). Follow this procedure to alter active queued files:

1. Enter the selection criteria specifying which files are to be altered (parameters FM, DN, JSN, LD, UJN, ID, DF, UN, FC, LA, UA, BC, and RB).
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 18-6, 18-7, and 18-8). 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, NDC, NEC, NLD, NDF, NUN, NID, NFC, NPR, NRC, and NAL). 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.

*** LIST ACTIVE QUEUED FILES *** PAGE 1 OF 2

OPTIONS	DESCRIPTION
FM = ALL	FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL	DEVICE NUMBER (1-77B,ALL).
JSN=	JOB SEQUENCE NUMBER (1-4 CHARACTERS). JOB SEQUENCE NUMBER WILL BE CLEARED IF PREVIOUSLY ENTERED.
LD =	DESTINATION LID (3 CHARACTERS).
UJN=	USER JOB NAME (1-7 CHARACTERS).
ID = 0B 67B	ID OF FILE TO LIST (0-67B).
LA =	LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA =	UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
DF =	DESTINATION FAMILY (1-7 CHARACTERS).
UN =	DESTINATION USER (1-7 CHARACTERS).
FC =	FORMS CODE.
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
LO = S	LIST OPTION (S=SHORT,F=FULL).
OP = NC	OPTION (BC,RB,NC,PR).
NAL =	NEW ACCESS LEVEL.
NDC =	NEW DEVICE CODE.
NEC =	NEW EXTERNAL CHARACTERISTICS.
NLD =	NEW DESTINATION LID (3 CHARACTERS).
NDF =	NEW DESTINATION FAMILY (1-7 CHARACTERS).
NUN =	NEW DESTINATION USER (1-7 CHARACTERS).
NID =	NEW FILE IDENTIFIER (0-67B).
NFC =	NEW FORMS CODE (AA-99,**).
NPR =	NEW PRIORITY FOR OUTPUT (0-7760B).
NRC =	NEW REPEAT COUNT (0-37B).

*** LIST ACTIVE QUEUED FILES *** PAGE 2 OF 2

(DS=DT) SELECT BY FILE DESTINATION AND
DISPOSITION TYPE.
OPTION WILL BE CLEARED
IF PREVIOUSLY SELECTED.
* = OPTION SELECTED.

S	I	P	P	P	P	L	L	L	W
F	N	U	L	R	2	R	S	T	T
BC	*	*	*	*	*	*	*	*	*
RB	*	*	*	*	*	*	*	*	*

Figure 18-5. QALTER/QFTLIST Left K Display

*** ACTIVE QUEUE FILE LIST ***

ORD.	JSNS.	DT	LID	FAMILY	USERNAM	TUI/ID	FC	EXTCH	LENGTH
1.	AAAOS.	PL				55		T6	4
2.	AAAPB.	LS				22		A9	6
3.	AAQR.	LP		SYS964	GAK1234	1234	AX	LP	41
4.	AAARB.	LS				22		A9	5
5.	AAASB.	LS				22		A9	4
6.	AAATR.	LP		SYS964	GAK1234	1234	AX	LP	41
7.	AAAUR.	LP		SYS964	GAK1234	1234	AX	LP	41
10.	AAAVB.	LP				55		LP	7
11.	AAAWR.	LP		SYS964	GAK1234	1234	AX	LP	41
12.	AAAXB.	LS	M64			22		A9	5
13.	AAAYR.	LP		SYS964	GAK1234	1234	AX	LP	41
16.	AABBR.	LP	LBK	SYS964	GAK1234	1234	AX	LP	41
17.	AABCB.	LS				22		A9	3
20.	AABDB.	LS				22		A9	2

END OF DISPLAY.

Figure 18-6. QALTER/QFTLIST LIST Command Display

*** ACTIVE QUEUE LIST ***

JSN	= AADU	ORDINAL	= 7
SERVICE	= SYSTEM	QUEUE	= PRINT
DESTINATION		CREATION	
FAMILY	=	FAMILY	= SYS964
LID	= M42	LID	= M64
USER	=	USER	=
TUI/ID	= 0	USR INDX	= 0
FORMS	= AX	USR JOBNM	= DIS
DISP CODE	= LP	DATE	= 83/08/08
EXT.CHR.	= LP	LENGTH	= 2
INT.CHR.	= DIS	REPEAT	= 0
RESIDENCE		INTERRUPT	= NO
FAMILY	= SYST64	PRIORITY	= 7110
DEVICE	= 3	ACCESS LV	=

Figure 18-7. QALTER/QFTLIST LIST=qft Command Display

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--

CURRENT	DESTINATION	CREATION	RESIDENCE
*****	*****	*****	*****
1. JSN = AACG SERVICE = ICLASS2 QUEUE = PRINT ORDINAL = 3	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LP EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= NO REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 55 LENGTH = 13 ACCESS LV=
2. JSN = AACI SERVICE = TIMESH QUEUE = PRINT ORDINAL = 5	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LP EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= NO REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 55 LENGTH = 5 ACCESS LV=
3. JSN = AACJ SERVICE = REMOTE QUEUE = PRINT ORDINAL = 6	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LT EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= NO REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 56 LENGTH = 15 ACCESS LV=
4. JSN = AACW SERVICE = BATCH QUEUE = PRINT ORDINAL = 23	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LP EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= YES REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 56 LENGTH = 13 ACCESS LV=
5. JSN = AACY SERVICE = TIMESH QUEUE = PRINT ORDINAL = 25	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LP EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= YES REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 56 LENGTH = 5 ACCESS LV=
6. JSN = AACZ SERVICE = ICLASS2 QUEUE = PRINT ORDINAL = 26	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LP EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= YES REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 56 LENGTH = 13 ACCESS LV=
7. JSN = AADB SERVICE = REMOTE QUEUE = PRINT ORDINAL = 30	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LT EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= YES REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 56 LENGTH = 15 ACCESS LV=
8. JSN = AADF SERVICE = BATCH QUEUE = PRINT ORDINAL = 34	FAMILY = SYS964 USER = BMS4247 TUI/ID = 4247 DLID =	FORMS = DEVICE = LP EXT.CHR. = LP INT.CHR. = DIS	FAMILY = SYS964 USER = BMS4247 USR INDX = 4247 SLID = USR JOBNM= DIS DATE = 83/07/07 INTERRUPT= NO REPEAT = 0 FAMILY = SYST64 DEVICE = 2 PRIORITY = 57 LENGTH = 13 ACCESS LV=

Figure 18-8. QALTER Printer Listing

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.
```

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 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 figure 18-9).

*** QUEUE FILE DUMP *** PAGE 1 OF 2

INACTIVE QUEUES RESIDE ON FOLLOWING FAMILIES.
SYST64 SYS964 PACKV2

OPTIONS	DESCRIPTION	(ERRORED FILE DUMPING SET)
MI = 64	MACHINE ID (1-2 CHARACTERS,ALL).	
FM = ALL	FAMILY FOR DEVICES (1-7 CHARACTERS).	
DN = ALL	DEVICE NUMBER (1-77B,ALL).	
FU = ALL	FAMILY FOR USER INDEX (1-7 CHARACTERS).	
UI = 0B 377777B	USER INDEX RANGE (0-377777B).	
DA = 83/08/03. 83/08/08.	DUMP DATE RANGE (YYMMDD,ALL).	
FS = 0B 777777B	FILE SIZE RANGE IN PRUS (0-777777B).	
JSN=	JOB SEQUENCE NUMBER (1-4 CHARACTERS). JOB SEQUENCE NUMBER WILL BE CLEARED IF PREVIOUSLY ENTERED.	
LD =	DESTINATION LID (3 CHARACTERS).	
ID = 0B 67B	ID OF FILES TO DUMP (0-67B).	
LA =	LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).	
UA =	UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).	
TUI= 0B 377777B	DESTINATION USER INDEX (0-377777B).	
TF = ALL	FAMILY FOR TERMINAL (1-7 CHARACTERS).	
FC = ALL	FORMS CODE.	
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).	
FN = QFILES	DUMP FILE NAME (1-7 CHARACTERS).	
NF = 0D	NUMBER OF MEDIA FILES TO SKIP.	
TP = ALL	TYPE (A-ACTIVE,I-INACTIVE,ALL).	

*** QUEUE FILE DUMP *** PAGE 2 OF 2

(DS=DT)	DUMP BY FILE DESTINATION AND DISPOSITION TYPE.	OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.
S I P P P P L L L W		* = OPTION SELECTED.
F N U L R 2 R S T T		
BC * * * * * * * * *		
RB * * * * * * * * *		

Figure 18-9. Initial QDUMP Display

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 18-10). 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.

```

*** QUEUE FILE DUMP ***    PAGE 1 OF 2

INACTIVE QUEUES RESIDE ON FOLLOWING DEVICES IN FAMILY - SYST64
01 02 03

OPTIONS          DESCRIPTION      (ERRORED FILE DUMPING SET)
MI = 64          MACHINE ID (1-2 CHARACTERS,ALL).
FM = SYST64      FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL         DEVICE NUMBER (1-77B,ALL).
FU = ALL         FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = 0B         USER INDEX RANGE (0-377777B).
377777B
DA = 83/08/03.   DUMP DATE RANGE (YMMDD,ALL).
83/08/08.
FS = 0B         FILE SIZE RANGE IN PRUS (0-777777B).
777777B
JSN=            JOB SEQUENCE NUMBER (1-4 CHARACTERS).
                JOB SEQUENCE NUMBER WILL BE CLEARED IF
                PREVIOUSLY ENTERED.

LD =            DESTINATION LID (3 CHARACTERS).
ID = 0B         ID OF FILES TO DUMP (0-67B).
67B
LA =            LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA =            UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
TUI= 0B        DESTINATION USER INDEX (0-377777B).
377777B
TF = ALL        FAMILY FOR TERMINAL (1-7 CHARACTERS).
FC = ALL        FORMS CODE.
L = OUTPUT      FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
FN = QFILES     DUMP FILE NAME (1-7 CHARACTERS).
NF = 0D        NUMBER OF MEDIA FILES TO SKIP.
TP = ALL        TYPE ( A-ACTIVE,I-INACTIVE,ALL).

```

```

*** QUEUE FILE DUMP ***    PAGE 2 OF 2

(DS=DT)          DUMP BY FILE DESTINATION AND
                  DISPOSITION TYPE.
                  OPTION WILL BE CLEARED
                  IF PREVIOUSLY SELECTED.
                  * = OPTION SELECTED.
S I P P P P L L L W
F N U L R 2 R S T
BC * * * * *
RB * * * * *

```

Figure 18-10. Secondary QDUMP Display

After QDUMP processing has been completed and K-display interaction has been terminated by K.STOP.

an output file is generated, indicating all files dumped. Figure 18-11 is an example of this output.

QDUMP PROCESSOR.		83/07/07. 13.33.19. PAGE 1		*** DUMPED QUEUE FILE DISPOSITION ***			
NO.	JSNS MID FAMILY	DN DT DS	DATE	LENGTH (PRUS)	ACCESS 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 18-11. QDUMP Output File

QDUMP does not purge files upon completion. If you want to purge files, the following sequence of commands is recommended.

<u>Command</u>	<u>Description</u>
QREC,OP=DI.	Dequeues the files.
QDUMP,TP=I.	Dumps the inactive files. The dump file (default name = QFILES) should be preassigned before 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 figure 18-5). 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 18-18).

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 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 figure 18-12).

After QLOAD processing has been completed and K-display interaction has been terminated by

```
K.STOP.
```

an output file is generated, indicating all queued files which were contained on the dump file, and whether they were loaded or not. Figure 18-13 contains an example of this output.

*** QUEUE FILE LOAD *** PAGE 1 OF 2

OPTIONS	DESCRIPTION	(ERRORED FILE LOADING NOT SET)
MI = 64	MACHINE ID (1-2 CHARACTERS,ALL).	
DF = ALL	DESTINATION FAMILY (1-7 CHARACTERS).	
DD = ALL	DESTINATION DEVICE (1-77B).	
FU = ALL	FAMILY FOR USER INDEX (1-7 CHARACTERS).	
UI = 0B	USER INDEX RANGE (0-377777B).	
	377777B	
DA = 83/08/03.	LOAD DATE RANGE (YMMDD,ALL).	
	83/08/08.	
FS = 0B	FILE SIZE RANGE IN PRUS (0-777777B).	
	777777B	
JSN=	JOB SEQUENCE NUMBER (1-4 CHARACTERS).	
	JOB SEQUENCE NUMBER WILL BE CLEARED IF	
	PREVIOUSLY ENTERED.	
LD =	DESTINATION LID (3 CHARACTERS).	
ID = 0B	ID OF FILES TO LOAD (0-67B).	
	67B	
LA =	LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).	
UA =	UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).	
TUI= 0B	DESTINATION USER INDEX (0-377777B).	
	377777B	
TF = ALL	FAMILY FOR TERMINAL (1-7 CHARACTERS).	
FC = ALL	FORMS CODE.	
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).	
FN = QFILES	LOAD FILE NAME (1-7 CHARACTERS).	
NF = 0D	NUMBER OF MEDIA FILES TO SKIP.	
SC = 0D	NUMBER OF QUEUED FILES TO SKIP.	
OP = A	LOAD OPTION (1 CHARACTER).	
	A-LOAD/ACTIVATE. I-LOAD/INACTIVATE	

*** QUEUE FILE LOAD *** PAGE 2 OF 2

(DS=DT) LOAD BY FILE DESTINATION AND DISPOSITION TYPE.

	S I P P P P L L L W	OPTION WILL BE CLEARED
	F N U L R 2 R S T T	IF PREVIOUSLY SELECTED.
		* = OPTION SELECTED.
BC	*****	
RB	*****	

Figure 18-12. QLOAD Display

QLOAD/LDLIST PROCESSOR.			83/07/07. 13.33.38.				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 18-13. QLOAD Output File

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.
```

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 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 figure 18-14).

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 18-15). 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.

```

*** QUEUE FILE MOVE *** PAGE 1 OF 2
INACTIVE QUEUES RESIDE ON FOLLOWING FAMILIES.
  SYST64   SYS964   PACKV2

OPTIONS      DESCRIPTION      (ERRORED FILE MOVING NOT SET)
MI = 64      MACHINE ID (1-2 CHARACTERS,ALL).
FM = ALL     FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL     DEVICE NUMBER (1-77B, ALL).
DF =         DESTINATION FAMILY (1-7 CHARACTERS).
DD =         DESTINATION DEVICE (1-77B).
FU = ALL     FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = 0B      USER INDEX RANGE (0-377777B).
             377777B
DA = 83/08/03. DATE RANGE (YYMMDD,ALL).
             83/08/08.
FS = 0B      FILE SIZE RANGE IN PRUS (0-777777B).
             777777B
JSN=         JOB SEQUENCE NUMBER (1-4 CHARACTERS).
             JOB SEQUENCE NUMBER WILL BE CLEARED IF
             PREVIOUSLY ENTERED.

LD =         DESTINATION LID (3 CHARACTERS).
ID = 0B      ID OF FILE TO MOVE (0-67B).
             67B
LA =         LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA =         UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
TUI= 0B      DESTINATION USER INDEX (0-377777B).
             377777B
TF = ALL     FAMILY FOR TERMINAL (1-7 CHARACTERS).
FC = ALL     FORMS CODE.
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).

```

```

*** QUEUE FILE MOVE *** PAGE 2 OF 2
(DS=DT)      MOVE BY FILE DESTINATION AND
              DISPOSITION TYPE.
              OPTION WILL BE CLEARED
              IF PREVIOUSLY SELECTED.
              * = OPTION SELECTED.
S I P P P P L L L W
F N U L R 2 R S T T
BC * * * * *
RB * * * * *

```

Figure 18-14. Initial QMOVE Display

*** QUEUE FILE MOVE *** PAGE 1 OF 2
 INACTIVE QUEUES RESIDE ON FOLLOWING DEVICES IN FAMILY - SYST64
 01 02 03

OPTIONS	DESCRIPTION (ERRORED FILE MOVING NOT SET)
MI = 64	MACHINE ID (1-2 CHARACTERS,ALL).
FM = SYST64	FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL	DEVICE NUMBER (1-77B, ALL).
DF =	DESTINATION FAMILY (1-7 CHARACTERS).
DD =	DESTINATION DEVICE (1-77B).
FU = ALL	FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = 0B 377777B	USER INDEX RANGE (0-377777B).
DA = 83/08/03. 83/08/08.	DATE RANGE (YMMDD,ALL).
FS = 0B 777777B	FILE SIZE RANGE IN PRUS (0-777777B).
JSN=	JOB SEQUENCE NUMBER (1-4 CHARACTERS). JOB SEQUENCE NUMBER WILL BE CLEARED IF PREVIOUSLY ENTERED.
LD =	DESTINATION LID (3 CHARACTERS).
ID = 0B 67B	ID OF FILE TO MOVE (0-67B).
LA =	LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA =	UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
TUI= 0B 377777B	DESTINATION USER INDEX (0-377777B).
TF = ALL	FAMILY FOR TERMINAL (1-7 CHARACTERS).
FC = ALL	FORMS CODE.
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).

*** QUEUE FILE MOVE *** PAGE 2 OF 2

(DS=DT) MOVE BY FILE DESTINATION AND
 DISPOSITION TYPE.
 OPTION WILL BE CLEARED
 IF PREVIOUSLY SELECTED.
 * = OPTION SELECTED.

S	I	P	P	P	P	L	L	L	W
F	N	U	L	R	2	R	S	T	T
BC	*	*	*	*	*	*	*	*	*
RB	*	*	*	*	*	*	*	*	*

Figure 18-15. Secondary QMOVE Display

After QMOVE processing has been completed and K-display interaction has been terminated by K.STOP.

an output file is generated, indicating all files moved. Figure 18-16 contains an example of this output.

```
QMOVE PROCESSOR.                83/07/07. 13.34.15. PAGE 1
*** MOVED QUEUE FILE DISPOSITION ***
ACTIVE FILES MOVED TO FAMILY SYST64 DEVICE 2 AND LEFT ACTIVE
                                LENGTH ACCESS
                                (PRUS) LEVEL  DISPOSITION ERRORS
NO.  JSNS MID FAMILY  DN DT DS  DATE
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 18-16. QMOVE Output File

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.  
K,jsn.           (jsn is the job sequence name of QFSP)  
K.QMOVE.  
K.FM=SYSTEM, DN=45.  
K.DF=FAM1, DD=5.  
K.OP=A.         (the moved files are active)  
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.
```

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 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 figure 18-17).

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 18-18). 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.

*** 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 18-18. Inactive Queues List

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 18-19). 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.

After QREC processing has been completed and K-display interaction has been terminated by

K.STOP.

an output file is generated, indicating the disposition of all queued files processed. Figure 18-20 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=N)

*** QUEUED FILE REQUEUING *** PAGE 1 OF 2

INACTIVE QUEUES RESIDE ON FOLLOWING DEVICES IN FAMILY - SYS964
40

OPTIONS	DESCRIPTION
MI = 64	MACHINE ID (1-2 CHARACTERS,ALL).
FM = SYS964	FAMILY FOR DEVICES (1-7 CHARACTERS).
DN = ALL	DEVICE NUMBER (1-77B, ALL).
FU = ALL	FAMILY FOR USER INDEX (1-7 CHARACTERS).
UI = 0B 377777B	USER INDEX RANGE (0-377777B).
DA = 83/08/03. 83/08/08.	QUEUED DATE RANGE (YYMMDD,ALL).
FS = 0B 777777B	FILE SIZE RANGE IN PRUS (0-777777B).
JSN=	JOB SEQUENCE NUMBER (1-4 CHARACTERS). JOB SEQUENCE NUMBER WILL BE CLEARED IF PREVIOUSLY ENTERED.
LD =	DESTINATION LID (3 CHARACTERS).
ID = 0B 67B	ID OF FILES TO PROCESS (0-67B).
LA =	LOWER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
UA =	UPPER SECURITY ACCESS LEVEL (1-7 CHARACTERS).
TUI= 0B 377777B	DESTINATION TUI (0-377777B).
TF = ALL	FAMILY FOR TERMINAL (1-7 CHARACTERS).
FC = ALL	FORMS CODE.
L = OUTPUT	FILE TO RECEIVE OUTPUT (1-7 CHARACTERS).
OP = RI	PROCESS OPTION (2 CHARACTERS). PROCESS OPTION (RP/RI/PI/DI).

*** QUEUED FILE REQUEUING *** PAGE 2 OF 2

(DS= DT)	REQUEUE BY FILE DESTINATION AND DISPOSITION TYPE.
	OPTION WILL BE CLEARED IF PREVIOUSLY SELECTED.
	* = OPTION SELECTED.
S I P P P P L L L W	
F N U L R 2 R S T T	
BC * * * * * * * * *	
RB * * * * * * * * *	

Figure 18-19. Secondary QREC Display

*** QUEUE DISPOSITION ***

NO.	JSNS	MID	FAMILY	DN	DT	DS	DATE	LENGTH (PRUS)	ACCESS 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.	AABCB	64	SYST64	2	PR	RB	83/07/07.	13		ACTIVATE

Figure 18-20. QREC Output File

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 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/TUI pertain 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, and FC parameters restrict 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.
```

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.
```

<u>Parameter</u>	<u>Description</u>
utility	Utility called to perform the dayfile dump.

<u>Utility</u>	<u>Description</u>
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 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 2 Operations Handbook.

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.

FR=string 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.

OP=op 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.

<u>op</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>										
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="0"> <thead> <tr> <th><u>pd</u></th> <th><u>Default pl</u></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>	<u>pd</u>	<u>Default pl</u>	3	30	4	40	6	60	8	80
<u>pd</u>	<u>Default pl</u>										
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	19-2
*CM Directive	19-2
*MS Directive	19-3
*AD Directive	19-3
*PROC Directive	19-5
*SC Directive	19-5
*FL Directive	19-5
*/ Directive	19-6
*DELETE Directive	19-6
*FILE Directive	19-7
*IGNORE Directive	19-7
*PPSYN Directive	19-7
Record Types	19-8

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. 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.

<u>p_i</u>	<u>Description</u>
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; in that case, 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.

<u>P_i</u>	<u>Description</u>
L=0	No list output.
L omitted	Same as L=0.
NA	SYSEDIT 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 SYSEDIT. The directives *CM, *MS, *AD, *PROC, *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,...,ty_n/rec_n

<u>Parameter</u>	<u>Description</u>
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. The following programs will automatically reside in CMR.

- Mass storage drivers.
- Programs specified in the internal tables of SLL (SYSEDT service routine), such as IDD and ODF.
- 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 LMT, IMS, and PFM) reside in CMR.

Any addition to central memory allows less space for user jobs.

REL type records cannot reside in CMR.

***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,...,ty_n/rec_n

<u>Parameter</u>	<u>Description</u>
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/rec₂,...,ty_n/rec_n

Parameter

Description

xx

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.

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.

ty_i/rec_i

Record type and record name of the routine.

ty_i Record type of routine:

Record

Type

Description

ABS	CPU multiple entry point overlay.
OVL	CPU overlay.
PP	PP absolute.
REL	Relocatable CPU routine.

If a record type other than ABS, OVL, PP, or REL is specified, the message ILLEGAL CM/AD RESIDENCE. is issued to the output device; the run is aborted.

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, SYSEDIT 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.

***PROC DIRECTIVE**

The *PROC directive identifies a record as a procedure. It can be treated as any command, with parameters as required by the procedure itself. Further information on procedure creation and execution can be found in the NOS 2 Reference Set, Volumes 2 and 3.

The format of the directive is:

*PROC,rec₁,rec₂,...,rec_n

<u>Parameter</u>	<u>Description</u>
rec _i	Record name of routine to be defined as a procedure file.

***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 2 Reference Set, Volume 3).

The format of the directive is:

*SC,ty₁/rec₁,ty₂/rec₂,...,ty_n/rec_n

<u>Parameter</u>	<u>Description</u>
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,...,ty_n/rec_n-fl_n

<u>Parameter</u>	<u>Description</u>
ty _i /rec _i	Record type and record name of the routine.
fl _i	Field length divided by 1008 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:
 - 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.
 - 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

<u>Parameter</u>	<u>Description</u>
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,...,ty_n/rec_n

or

*D,ty₁/rec₁,...,ty_n/rec_n

<u>Parameter</u>	<u>Description</u>
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

<u>Parameter</u>	<u>Description</u>
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,...,ty_n/rec_n

<u>Parameter</u>	<u>Description</u>
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,...,name_n

<u>Parameter</u>	<u>Description</u>
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.

<u>Type</u>	<u>Description</u>
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.

ISF
Deadstart Sequencing

20-1
20-3

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.† A fast-attach file is a special direct-access file under user index 37777g 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 other jobs, and when the above conditions have been met [refer to the family status display (E,F) in the NOS 2 Operations Handbook], the IDLEFAMILY command must be used again to allow the ISF command to be entered.

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 multiframe network, additional information is also maintained in the fast-attach table (FAT) 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 77g exists on the number of files that can be entered as global fast attach in a multiframe 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=37777g) for files with the predefined names previously listed. They are defined in a table internal to ISF. The format of the ISF command is:

ISF,option,FM=familyname,SJ=filename,SP=filename.

†The resource files are generated and maintained separately for each machine id in a multiframe 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).

<u>Parameter</u>	<u>Description</u>
option	Specifies whether to initialize or release system files.

<u>option</u>	<u>Description</u>
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)

NOTE

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 (377777g). 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	Parameter	
	Initialize (E)	Release (R)
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=37777g), 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.

ISF,FM=0,SJ. (recovery deadstarts)
ISF,FM=0,SJ,SP. (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).

TRACER Utility	21-1
TRACER Commands	21-4
ICPD Command	21-4
ACPD Command	21-5
ENDCPD Command	21-6
Output File Format	21-7
Summary File Format	21-18
Data Items Reported by TRACER	21-18
Fast Loop Items	21-20
Medium Loop Items	21-25
Slow Loop Items	21-29
Statistical Summary	21-30
PROBE Utility	21-33

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.
- 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 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.

The TRACER utility includes the following programs:

<u>Program</u>	<u>Description</u>
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.

ICPD Command

ICPD defines a mass storage file to which CPD will write statistical data and then initiates system monitoring by CPD.

Format:

ICPD, p_1,p_2,\dots,p_n .

<u>P_i</u>	<u>Description</u>
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 EOF 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:

<u>mode</u>	<u>Description</u>
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).

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.

<u>P_i</u>	<u>Description</u>
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.

<u>P_i</u>	<u>Description</u>
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.
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.

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.

START DATE	83/05/25.	}	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	8.0	}	ICPD command parameters (FL, ML, SL, FW)
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 PPUS	20		
MEMORY SIZE / 100B	7777B		
USER EM / 1000B	200B		
MAGNETIC TAPE UNITS	6		
NUMBER OF EST ENTRIES	61B	}	Software configuration at beginning of sampling interval
NUMBER OF FMT ENTRIES	23B		
NUMBER OF EJT ENTRIES	1100B		
NUMBER OF QFT ENTRIES	620B		
NUMBER OF FOT ENTRIES	10B		
NUMBER OF CONTROL POINTS	34B		
CRSIZ / 100B	622B		
LIBDECK NUMBER	1		
RECOVERY LEVEL	0		
NUMBER OF IAF TERMINALS	200B		
MACHINE ID	64		
CPU USAGE CATEGORIES (CTCUL)	5B		
CHANNEL TABLE LENGTH (CTALL)	36B		
LENGTH OF AN EST ENTRY (ESTE)	2B		
LENGTH OF AN FMT ENTRY (FNTE)	2B		
LENGTH OF AN EJT ENTRY (EJTE)	4B		
LENGTH OF A QFT ENTRY (QFTE)	4B		
NUMBER OF SERVICE CLASSES (MXSC)	17B		
NUMBER OF JOB STATUSES (MXJS)	13B		
NUMBER OF CONNECTION STATUSES (MXCS)	3B		
NUMBER OF FILE TYPES (MXFT)	21B		
MAXIMUM ORIGIN TYPES (MXOT)	6B		
NUMBER OF MASS STORAGE DEVICES	14B	}	System delay parameters at beginning of sampling interval
LENGTH OF MSAL TABLE (MXRS)	14B		
CPU RECALL DELAY	30B		
PP AUTO RECALL DELAY	1750B		
MAXIMUM JOB SWITCH DELAY	20B		
MINIMUM JOB SWITCH DELAY	4B		
INPUT FILE SCHEDULING INTERVAL	2B	}	System buffers at beginning of sampling interval
SCHEDULER CYCLE INTERVAL	1B		
HIGH SPEED DISK BUFFERS	0B		
EM/PP BUFFERS	1B		


Figure 21-1. Example of TRACER Output (Sheet 1 of 10)

EQUIPMENT STATUS TABLE

EQUIPMENT NO.	TYPE	STAT	EQ	UN	CHANNELS	EST	ENTRY	TRACK	FILES
0	RD	ON	1	11	0	0000000000022041271	00000000000000000000		
1	DS	ON	7	0	10	0000001000004237000	00000000000000000000		
2	NE	ON	0	0	0	20000000000016050000	000000000000000070000		
3	TE	ON	0	0	0	20000000000024050000	000000000000000070000		
4	TT	ON	0	0	0	20000000000024240200	000000000000000070000		
5	DE	ON			0	40100000000004051274	00000000000000000000	770	-----
6	DQ	ON			13	62100013137704211277	00000000000000000000	3222	-XT-----N
7	DQ	ON			13	62100013137704211302	00000000000000000000	3222	-XT-----N
10	DJ	ON			26 32	42103226137704121305	00000000000000000000	3150	-----
11	DJ	ON	1	32	26	47002632137704121310	00000000000000000000	3150	-----
12	DJ	ON	2	26	32	47003226137704121313	00000000000000000000	3150	-----
13	DI	ON	3	32	26	47002632137704111316	00000000000000000000	3140	-----
14	DI	ON	4	26	32	47003226137704111321	00000000000000000000	3140	-----
15	DI	ON	6	26	32	47003226137704111324	00000000000000000000	3140	-----
16	DJ	ON			0 13	47000013137704121327	00000000000000000000	3150	-----
17	DJ	ON			1 13	47000013137704121332	00000000000000000000	3150	-----
20	DJ	ON			2 13	47000013137704121335	00000000000000000000	3150	-----
21	CR	ON	4	0	12	0000001200003224000	00000000000000000000		
22	CP	ON	5	0	12	0000001200003205000	000000000000000000014		
23	LT	ON	7	0	12	02000012000014247100	00000000000000000000		
24	LT	ON	6	0	12	00000012000014246000	00000000000000000000		
30	TT	OFF	6	10	6	2000006000064246310	00000000000000000000		
31	TT	OFF	6	6	6	2000006000064246226	00000000000000000000		
32	NP	ON	7	3	3	0000003000016207103	00000000000000000006		
34	NP	OFF	7	1	1	0000001000056207101	00000000000000000000		
35	NP	OFF	7	7	2	0000002000056207107	00000000000000000000		
36	NP	OFF	7	4	7	0000007000056207104	00000000000000000000		
43	CS	ON	0	6	0	0000000440003230106	00000000000000000000		
44	CT	ON	0	0	0	0000000440003240000	00000000000000000000		
45	CT	ON	0	2	0	0000000440203240202	00000000000000000000		
50	MT	ON	0	0	31 33	02003331000015240000	00000000000000000000		
51	NT	ON	0	1	31 33	02003331000016240001	00000000000000000000		
52	NT	ON	0	2	31 33	02003331000016240002	00000000000000000000		
53	NT	ON	0	3	31 33	02003331000016240003	00000000000000000000		
54	NT	ON	0	4	31 33	02003331000016240004	00000000000000000000		
55	NT	ON	0	5	31 33	02003331000016240005	00000000000000000000		
60	NC	ON	0	5	6	2000006000016030045	00000000000000000000		

770 -----
 3222 -XT-----N
 3222 -XT-----N
 3150 -----
 3150 -----
 3150 -----
 3140 -----
 3140 -----
 3140 -----
 3150 -----
 3150 -----
 3150 -----
 3150 -----

Track capacity
 for mass storage
 devices



① The letters to be used here are order dependent and indicate file types that can reside on a given device. The order is GXTIORDPLBSRRN, defined as follows.

Letter	File Type	Letter	File Type
G	Checkpoint	P	Primary
X	System	L	Local
T	Temporary	B	LGO
I	Input	S	Secondary rollout
O	Output	R	Reserved for CDC
R	Rollout	R	Reserved for CDC
D	Dayfile	N	Temporary on nonshared device

Figure 21-1. Example of TRACER Output (Sheet 2 of 10)

SYSTEM CONTROL INFORMATION

SERVICE CLASS	QUEUE	PRIORITIES					SERVICE LIMITS					
		IL	LP	UP	WF	IP	PR FL EC	CP AN EM	CM DS	NJ TP CS	TD FS	
SY	IN		7770	7776	0		30	100	20	7777	341	
	EX	4000	2000	7000	0	7000	3777	77777		4004		
	OT		7000	7776	0		3777	7777	0	0	0	
BC	IN		10	4000	0		30	400	200	7777	341	
	EX	2000	1000	4004	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
RB	IN		10	4000	0		30	400	200	7777	341	
	EX	2000	1000	4004	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
TS	IN		7000	7770	0		30	40	10	7777	341	
	EX	3770	3700	7000	0	4000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
DI	IN		10	4000	0		30	100	20	7777	341	
	EX	2000	1000	4000	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
NS	IN		7770	7776	0		74	400	200	7777	341	
	EX	7772	7770	7776	0	7772	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
SS	IN		7770	7776	0		70	100	20	7777	341	
	EX	7772	7770	7776	0	7772	3777	77777		4004		
	OT		7400	7776	0		3777	7777	0	0	0	
MA	IN		1	10	0		2	100	20	7777	341	
	EX	4	1	10	0	10	3777	77777		4004		
	OT		7000	7776	0		3777	7777	0	0	0	
CT	IN		7770	7776	0		30	400	200	7777	341	
	EX	4004	4000	7000	0	7000	3777	77777		4004		
	OT		7000	7776	0		3777	7777	0	0	0	
I0	IN		10	4000	0		30	400	200	7777	341	
	EX	2000	1000	4004	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
I1	IN		10	4000	0		30	400	200	7777	341	
	EX	2000	1000	4004	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
I2	IN		10	4000	0		30	400	200	7777	341	
	EX	2000	1000	4004	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	
I3	IN		10	4000	0		30	400	200	7777	341	
	EX	2000	1000	4004	0	2000	3777	77777		4004		
	OT		1	7000	0		3777	7777	0	0	0	

This page shows the status of the job control area at the beginning of the sampling interval. The values are set either by the IPRDECK entries QUEUE and SERVICE or by the DSD commands QUEUE and SERVICE.

Figure 21-1. Example of TRACER Output (Sheet 3 of 10)

A C P D - VER 8.0		CDC NETWORK OPERATING SYSTEM.				NOS 2		yy/mm/dd. hh.mm.ss.				PAGE 4	
2 MINS INTERVAL		(2) 23.52.13	23.54.13	23.56.13	23.58.12	00.00.12	00.02.13	00.04.12	00.06.13	00.08.13	00:18 HR	(5) SUBTOTAL	(6)
		INTERVAL	INTERVAL	INTERVAL	INTERVAL	INTERVAL	INTERVAL	INTERVAL	INTERVAL	INTERVAL			
(3) FAST LOOP SAMPLES		10480	10519	10525	10517	10528	10524	10516	10507	10517		94633	
PPUS ACTIVE	(4) AV	(4.236)	6.272	4.987	7.186	[7.708]	4.552	4.504	4.250	6.154		5.539	
	SD	(0.014)	0.139	0.118	0.151	[0.363]	0.238	0.034	0.021	0.143		1.261	
	PC	(21.180)	31.359	24.933	35.931	[38.538]	22.760	22.519	21.250	30.770		27.693	
MOVE REQUEST PENDING	SD	(0.000)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	[0.001]		0.001	
	PC	(0.000)	0.010	0.067	0.000	[0.124]	0.048	0.000	0.000	0.067		0.035	
NO PPU AVAILABLE	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	
EM TRANSFER IN PROG	SD	0.000	0.000	0.000	0.000	0.001	0.000	[0.001]	(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	
MTR CYCLE TIME	AV	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	
MONITOR MODE - CPU 0	SD	[0.013]	(7) 0.002	0.012	0.009	0.003	(0.002)	0.010	0.002	0.005		0.020	
	PC	(7.887)	12.358	8.608	[13.605]	10.961	8.200	9.870	8.376	9.186		9.895	
MONITOR MODE - CPU 1	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	
SCHEDULER ACTIVE	SD	(0.001)	(8) 0.001	0.006	0.041	0.022	0.019	0.001	0.001	[0.069]		0.057	
	PC	5.114	(5.067)	7.335	15.915	[19.136]	7.051	5.696	5.805	12.110		9.248	
CHANNEL RESERVED													
CH 12	SD	0.000	0.000	(0.000)	0.000	0.036	0.000	0.001	0.000	[0.068]		0.034	
	PC	0.105	0.076	0.076	0.095	3.776	0.095	(0.038)	[6.898]	1.250			
CH 13	SD	0.010	0.262	0.012	(0.003)	0.033	0.059	0.016	0.005	[0.275]		0.168	
	PC	7.781	[39.808]	11.135	16.810	12.914	19.718	25.217	(7.434)	33.625		19.383	

- (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 FL 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 (Sheet 4 of 10)

9	MEDIUM LOOP SAMPLES	1170	1175	1175	1173	1175	1175	1175	1173	1175	10566
	CPS IN W STATUS	AV (0.225)	0.339	0.889	1.868	[2.494]	1.699	2.085	1.282	1.278	1.351
		SD 0.043	0.144	[0.320]	0.113	0.293	(0.002)	0.122	0.134	0.137	0.747
		PC (0.805)	1.209	3.175	6.671	[8.907]	6.067	7.446	4.579	4.566	4.825
	CPS IN X STATUS	AV (5.787)	6.452	5.894	[6.802]	6.078	5.906	6.080	5.827	6.079	6.101
		SD 0.020	0.014	(0.003)	0.044	0.101	0.009	0.162	[0.190]	0.159	0.328
		PC (20.669)	23.043	21.052	[24.293]	21.705	21.091	21.714	20.811	21.712	21.788
	CPS IN I STATUS	AV (0.067)	0.130	0.083	[0.185]	0.177	0.091	0.095	0.075	0.141	0.116
		SD (0.001)	0.016	0.015	0.004	0.005	0.011	0.008	0.002	[0.075]	0.049
		PC (0.238)	0.465	0.295	[0.661]	0.632	0.325	0.340	0.268	0.505	0.414
	SAME MOVE REQUEST	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
	FL AVAILABLE	AV [43478]	42368	33638	26758	24718	30178	(24128)	33758	31408	32508
		SD (1.227)	7.470	22.167	82.904	[126.210]	5.241	60.961	9.916	117.505	337.642
		PC [61.730]	59.740	48.185	39.794	36.213	42.016	(34.957)	48.459	44.193	46.143
	USER EM AVAILABLE	AV 08	08	08	08	08	08	08	08	08	08
		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
	NON INTERACTIVE JOBS	AV (1.022)	2.096	2.668	4.243	[6.591]	2.603	2.351	2.191	3.598	3.040
		SD 0.136	(0.023)	0.456	[0.619]	0.139	0.462	0.258	0.033	0.178	1.556
		PC (0.177)	0.364	0.463	0.737	[1.144]	0.452	0.408	0.380	0.625	0.528
	DETACHED JOBS	AV 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 0.000	0.000	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	1.000	1.000
		SD 0.000	0.000	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	0.174	0.174
	PRE-INITIAL JOB STEP										
	SY	AV (1.000)	1.000	1.001	1.029	1.000	1.000	1.000	1.001	[1.344]	1.042
		SD (0.000)	0.000	0.001	0.029	0.000	0.000	0.000	0.001	[0.344]	0.157
		PC (0.174)	0.174	0.174	0.179	0.174	0.174	0.174	0.174	[0.233]	0.181
	BC	AV (0.025)	1.000	0.684	1.166	[1.665]	0.094	0.236	0.092	0.462	0.603
		SD 0.025	(0.000)	0.125	0.166	0.335	0.094	0.236	0.029	[0.462]	0.582
		PC (0.004)	0.174	0.119	0.202	[0.289]	0.016	0.041	0.016	0.080	0.105
	RB	AV 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 0.000	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
		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

9 The sampling frequency for medium loop items is specified by the ML 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 (Sheet 5 of 10)

⑩	SLOW LOOP SAMPLES	AV	119	119	119	119	119	119	119	119	1071
	IAF USERS	SD	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
	IAF POTS ALLOCATED	AV	455.000	455.000	455.000	455.000	455.000	455.000	455.000	455.000	455.000
		SD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	IAF POTS IN USE	AV	(193.000)	193.000	193.017	193.008	193.017	[197.297]	193.000	193.000	193.482
		SD	(0.000)	0.000	0.000	0.008	0.017	[4.297]	0.000	0.000	1.967
		PC	(42.418)	42.418	42.421	42.419	42.421	[43.362]	42.418	42.418	42.524
	QUEUE FILES ASSIGNED	AV	(0.000)	0.008	0.000	0.008	0.050	0.000	0.000	0.000	[0.092]
		SD	(0.000)	0.008	0.000	0.008	0.050	0.000	0.000	0.000	[0.092]
		PC	(0.000)	0.002	0.000	0.002	0.012	0.000	0.000	0.000	[0.023]
	INPUT FILES										
	SY	AV	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
		PC	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
		SD	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
	RB	AV	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
		PC	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
		SD	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
	PRINT FILES										
	SY	AV	(0.000)	0.000	0.000	0.000	[0.008]	0.000	0.000	0.000	0.001
		SD	(0.000)	0.000	0.000	0.000	[0.008]	0.000	0.000	0.000	0.004
		PC	(0.000)	0.000	0.000	0.000	[0.002]	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
		SD	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
	RB	AV	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
		PC	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
		SD	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
		PC	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 ICPD 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 (Sheet 6 of 10)

11

SNAP SHOT OF INSL

(IN0L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0200	0200	0200	0200	0200	0200	0200	0200	0200
(IN1L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
(IN2L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
(IN3L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
(IN4L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
(IN5L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
(IN6L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
(IN7L) =	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000
	0000	0000	0000	0000	0000	0000	0000	0000	0000

Installation area
from CMR

11

The sampling frequency for snapshot loop items is specified by the FW parameter of the ICPD command.

Figure 21-1. Example of TRACER Output (Sheet 7 of 10)

MTR MXN TIME	743	743	743	743	743	743	743	743	743	
WORST CASE MTR CYCLE TIME	6	6	6	6	6	6	6	6	6	
MISSED CLOCK UPDATES	0	0	0	0	0	0	0	0	0	
EM MOVES	0	0	0	0	0	0	0	0	0	
CM MOVES	243	243	276	276	315	344	344	344	377	
TOTAL ROLLOUTS	54	56	70	74	77	85	90	96	97	
SECONDARY ROLLOUTS	0	0	0	0	0	0	0	0	0	
TOTAL SECTORS ROLLED	6046	6100	7168	7240	7783	8862	10029	10119	10146	
SECONDARY SEC ROLLED	0	0	0	0	0	0	0	0	0	
ROLLOUTS/USER LIMITS	0	0	0	0	1	1	3	3	3	
TIME SLICES	20	22	40	45	46	49	51	51	52	
PP PRIORITY EXCHANGES	0	0	0	0	0	0	0	0	0	
NO COMM BUFFER AVAIL	0	0	0	0	0	0	0	0	0	
EJT SCANS	112	114	190	334	714	781	786	797	846	
SCHEDULABLE JOBS	1024	1171	1357	1498	1633	1795	1946	2101	2238	
JOBS PREEMPTED	0	0	0	0	0	0	0	0	0	
JOBS SCHEDULED	72	74	98	104	106	115	120	127	130	
SCHEDULED NO CONSTRAINTS	0	0	0	0	0	0	0	0	0	
INSUFFICIENT CM SCANS	0	0	0	0	0	0	0	0	0	
INSUFFICIENT EM SCANS	0	0	0	0	0	0	0	0	0	
NO CONTROL POINT SCANS	0	0	0	0	0	0	0	0	0	

Statistical data area
from CMR

Figure 21-1. Example of TRACER Output (Sheet 8 of 10)

SNAP SHOT OF CPTW - CPO	0000	0000	0000	0000	0000	0000	0000	0000	0000	} CP accumulator area from CMR
	0104	0104	0104	0104	0104	0104	0104	0104	0104	
	7475	7475	7475	7475	7475	7475	7475	7475	7475	
	0005	0005	0005	0005	0005	0005	0005	0005	0005	
	7750	7750	7750	7750	7750	7750	7750	7750	7750	
SNAP SHOT OF CPTW - CPN	0000	0000	0000	0000	0000	0000	0000	0000	0000	
	0000	0000	0000	0000	0000	0000	0000	0000	0000	
	1141	1342	1406	1616	1672	1752	2072	2120	2163	
	5056	1115	6704	4124	5014	4315	7132	4576	5100	
	2400	4000	1000	1400	2000	2400	7000	0400	6000	
SNAP SHOT OF RTCL	0040	0040	0040	0040	0040	0040	0040	0040	0040	} Real time clock area from CMR
	1551	1741	2131	2321	2511	2701	3071	3261	3451	
	0000	0000	0000	0000	0000	0000	0000	0000	0000	
	0325	0362	0417	0455	0512	0547	0605	0642	0677	
	3020	5320	7620	2120	4420	6721	1220	3520	6020	
SNAP SHOT OF PDTL	0000	0000	0000	0000	0000	0000	0000	0000	0000	} Packed date/time area from CMR
	0000	0000	0000	0000	0000	0000	0000	0000	0000	
	1505	1505	1505	1505	1505	1505	1505	1505	1505	
	3127	3127	3127	3127	3200	3200	3200	3200	3200	
	6415	6615	7015	7214	0014	0215	0414	0615	1015	

Figure 21-1. Example of TRACER Output (Sheet 9 of 10)

2 MINS INTERVAL

00:18 HR 83/05/25 TO 83/05/26

		(13) TOTAL	(14) *MAX*	(15) *MIN*
(12) FAST LOOP SAMPLES		94633		
PPUS ACTIVE	AV	5.539	7.708	4.236
	SD	1.261	0.363	0.014
	PC	27.693	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
EM TRANSFER IN PROG	SD	0.001	0.001	0.000
	PC	0.125	0.209	0.048
MTR CYCLE TIME	AV	0.392	0.436	0.364
	SD	0.029	0.023	0.000
MONITOR MODE - CPU 0	SD	0.020	0.013	0.002
	PC	9.895	13.605	7.887
MONITOR MODE - 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	0.034	0.068	0.000
	PC	1.250	6.898	0.038
CH 13	SD	0.168	0.275	0.003
	PC	19.383	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 (Sheet 10 of 10)

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 table 21-1. 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. Data Items (Sheet 1 of 2)

Fast Loop Items	Weighting Factor
Fast loop samples	None
PPs active	Number of PPs available
Move request pending	1 †
No PP available	1
EM transfer in progress	1
MTR cycle time	100 ††
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
†The average value will not be reported if the weighting factor is 1. ††The percentage value will not be reported if the weighting factor is 100.	

Table 21-1. Data Items (Sheet 2 of 2)

Medium Loop Items	Weighting Factor
Medium loop samples CPS in W status CPS in X status CPS in I status Same move request FL available User EM available Noninteractive jobs Detached jobs Online jobs Preinitial job step Executing Scheduler rollout SCP rollin SCP rollout Timed/event rollout Interactive rollout Disabled rollout Suspended rollout Rollout file error EJT entries in use FL at control points FL in rollout queue EM memory at CP EM in rollout queue Subsystem FL	None Number of control points Number of control points Number of control points 1 Available field length (machine size - CMR size) Available user extended memory field length Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Number of EJT entries Available field length Available field length User extended memory size/1000g User extended memory size/1000g Available field length
Slow Loop Items	Weighting Factor
Slow loop samples IAF users IAF pots allocated IAF pots in use Queue files assigned Input files Print files Punch files Other queue files QFT entries in use Tape drives in use Tracks available	None Number of terminals defined 100 IAF pots allocated Number of QFT entries Number of QFT entries Number of QFT entries Number of QFT entries Number of QFT entries Number of QFT entries Number of QFT entries Number of available tape drives Maximum number of tracks for monitored device

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.

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 (ISJ) 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 ISJ is called to process EJT entries or QFT entries, not both. For example, if ISJ 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 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 LHP.

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. It records which devices have at least 1/8 of their track capacity available for use.
- Second, the system picks the device with the least activity. If all potential devices have less than 1/8 of their tracks available, then the system chooses the least active device.
- 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 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 IHP.

This can be used for debugging purposes when developong 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 channel busy counter for each buffered I/O channel and increments a counter whenever a data transfer is in progress on that channel. TRACER uses the buffered I/O channel control table to determine the channel 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.

<u>Counter</u>	<u>CPU Use</u>
IDLE	CPU is not currently being used.
SYSTEM	CPU is being used by CPUMTR.
SUB-SYS	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.

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 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 takes less CPU overhead than rolling jobs to UEM; 2) with more memory for the users, NOS performs less storage moves.
- 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 SUB-SYS CPU usage data. TRACER maintains a CPU usage counter for every subsystem. Whenever the SUB-SYS 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 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 CP

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 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.

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).

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.

Extended Memory 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.

Total Rollouts

TRACER reports the total number of jobs rolled out by the job rollout routine (IRO).

Secondary Rollouts

TRACER reports the number of jobs rolled out to secondary rollout devices by the job rollout routine (IRO).

Total Sectors Rolled

TRACER reports the total number of sectors of mass storage used by the job rollout routine (IRO) for rollouts.

Secondary Sec Rolled

TRACER reports the number of sectors of mass storage used by the job rollout routine (IRO) 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 (ISJ).

Time Slices

TRACER reports the number of jobs whose scheduling is set to the lower bound for its service class due to an expired CPU or CM time slice.

PP Priority Exchanges

TRACER reports the number of priority exchange requests issued by PPs. 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.

EJT Scans

TRACER reports the number of times the job scheduler (ISJ) 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 EJT on each EJT scan.

Jobs Preempted

TRACER reports the number of jobs rolled out so the job scheduler (ISJ) 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:

- The number of times a PP routine was loaded.
- The number of CIO RA+1 requests by function number.
- The number of PP requests to CPUMTR by function number.
- The number of MTR requests to CPUMTR by function number.
- The statistical data accumulated in CMR includes items such as number of sectors rolled and number of rollouts.

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.

	<u>Description</u>
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.

<u>opt</u>	<u>Description</u>
A	Sort data in alphabetic order.
F	Sort data by frequency of loads.
R	Sort data by location and frequency of loads.

OP=opt Processing option. Default is P.

<u>opt</u>	<u>Description</u>
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.

PROBE VERSION 1.0.

yy/mm/dd. hh.mm.ss. PAGE 1

PROBE VERSION 1.0 84/09/15. 10.09.16.
START OF SAMPLE INTERVAL 84/09/18. 07.19.13.
END OF SAMPLE INTERVAL 84/09/18. 14.57.20.

MONITOR REQUESTS

	PROGRAM MODE	MONITOR MODE	
	CPU0	CPU0	CPU1
AFAM(21)	0	7177	0
DLKM(22)	0	0	0
DTKM(23)	131214	189955	0
RTCM(24)	129969	129969	0
STBM(25)	0	121555	0
VMSM(26)	9108	9108	0
ACTM(27)	1133	20059	0
BFMM(30)	0	0	0
CKSM(31)	435	435	0
CSTM(32)	34405	34670	0
ECSM(33)	0	0	0
PIOM(34)	0	0	0
RDCM(35)	3170	3170	0
(36)	0	0	0
ABTM(37)	0	76	0
B IOM(40)	0	0	0
BMI(41)	0	0	0
CCAM(42)	0	77321	0
CEFM(43)	0	832	0
DCPM(44)	0	25696	0
DECM(45)	0	41425	0
DFMM(46)	0	36827	0
DPPM(47)	0	1108876	0
EATM(50)	0	48694	0
JACH(51)	0	59846	0
LDAM(52)	0	844946	0
MTRM(53)	0	2495	0
PLFM(54)	0	76738	0
RCLM(55)	0	447037	0
RCPM(56)	0	1110	0
RECM(57)	0	46638	0
REQM(60)	0	41426	0
RJSM(61)	0	675	0
RLMM(62)	0	30666	0
ROCM(63)	0	15092	0
RPNM(64)	0	0	0
RPPM(65)	0	106311	0
RSJM(66)	0	15481	0
SCDM(67)	0	62387	0
SFBM(70)	0	70802	0
SJCM(71)	0	22237	0
SPLM(72)	0	1358228	0
TDAM(73)	0	17566	0
TGPM(74)	0	11532	0
T IOM(75)	0	21228	0
TSEM(76)	0	37099	0
UADM(77)	0	583735	0
UTEM(100)	0	1473766	0
VFLM(101)	0	0	0
VFPM(102)	0	14518	0
VSAM(103)	0	13180	0
(104)	0	0	0
(105)	0	0	0
SUBTOTAL		7230584	0
TOTAL	309434		7230584

Figure 21-2. Example of PROBE Output (Sheet 1 of 7)

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	0

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
COMMUNICATIONS BUFFER NOT AVAILABLE =	0
PP PRIORITY EXCHANGE REQUESTS =	0

1R0 PERFORMANCE PARAMETERS

NUMBER OF ROLLOUTS =	18393
ROLLOUTS TO SECONDARY DEVICES =	0
NUMBER OF SECTORS ROLLED =	2918049
SECTORS ROLLED TO SECONDARY DEVICES =	0

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

Figure 21-2. Example of PROBE Output (Sheet 2 of 7)

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 (Sheet 3 of 7)

PPU PROGRAM LOADS

NAME	LOC	LOADS	NAME	LOC	LOADS
CFM	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	DCI	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
DBF	CMR	37181	OSI	DSK	0
3QS	CMR	36988	OTD	DSK	0
4MC	CMR	31223	1FA	DSK	0
3NW	CMR	29968	1HP	DSK	0
3SB	CMR	27491	1IE	DSK	0
3SD	CMR	27491	1IP	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
3SC	CMR	8245	2IO	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
3PA	CMR	5122	2LD	DSK	0
3LC	CMR	5085	2MB	DSK	0
3AI	CMR	4722	2ME	DSK	0

Figure 21-2. Example of PROBE Output (Sheet 4 of 7)

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
QAV	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
OFA	CMR	1162	3BD	DSK	0
DSP	DSK	1161	3BE	DSK	0
3DB	DSK	1120	3BF	DSK	0
ODQ	DSK	1086	3CA	DSK	0
3RG	DSK	1029	3CV	DSK	0
2NW	DSK	995	3FB	DSK	0
3MI	CMR	981	3IC	DSK	0
OAU	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	3KB	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	3MH	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	3G1	DSK	0
4DD	DSK	426	3RF	DSK	0
3SA	DSK	425	3RU	DSK	0
3MO	CMR	383	3SU	DSK	0
3MM	DSK	382	3SZ	DSK	0
3ML	DSK	354	3TA	DSK	0
3MH	DSK	350	4DF	DSK	0

Figure 21-2. Example of PROBE Output (Sheet 5 of 7)

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
ORF	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
ORT	DSK	74	5BD	DSK	0
3MS	DSK	67	5CU	DSK	0
2NQ	DSK	65	5CV	DSK	0
T76	DSK	58	5D1	DSK	0
6DI	CMR	58	5D2	DSK	0
OPT	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 (Sheet 6 of 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	54I	DSK	0
OPI	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	58A	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
OST	DSK	2	9CA	CMR	0
3SY	DSK	2	9CB	CMR	0
OTI	DSK	1	9CC	CMR	0
OTJ	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
HFM	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 (Sheet 7 of 7)

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 SYSEDIT command (refer to section 19, SYSEDIT).

APPENDIXES

A. CHARACTER SETS	A-1
Character Set Anomalies	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-12
B. GLOSSARY	B-1
C. SCOPE 2 STATION FACILITY	C-1
Equipment Configuration	C-1
SSF File Transfers	C-2
Staged File Transfers	C-2
Spooled File Transfers	C-3
Error Logout	C-3
D. STATUS/CONTROL REGISTER SIMULATOR	D-1
Using the Simulator	D-1
Console Operation	D-2
Batch Input	D-2
Simulator Commands	D-5
Holding Register Commands	D-5
Cycle Commands	D-6
+ and STOP Commands	D-7
E. 580 PROGRAMMABLE FORMAT CONTROL	E-1
Building Programmable Format Control Arrays	E-1
Adding Programmable Format Control Arrays	E-2
F. 881/883 PACK REFORMATTING UTILITY	F-1
FORMAT Command	F-2
Input Formats	F-3
Control Directives	F-4
Data Statements	F-4
Output Formats	F-5
Accessing Disk Devices	F-9
Access for Read Operations	F-9
Access for Read and Write Operations	F-10

CHARACTER SETS

A

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 00g to 77g.

The 6/12-bit display code is a combination of 6-bit codes and 12-bit codes. The 6-bit codes are 00g through 77g, excluding 74g and 76g. (Refer to Character Set Anomalies, next, for the interpretation of the 00g and 63g codes.) The 12-bit codes begin with either 74g or 76g and are followed by a 6-bit code. Thus, 74g and 76g 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 7401g, 7402g, 7404g, 7407g, and 7601g through 7677g. All other 12-bit codes (74xxg and 7600g) 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 0000g code from the end-of-line byte. The 7-bit codes are 0001g through 0177g and 4000g.

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 63g code, and the 00g code is undefined. However, if the installation uses the 64-character set, output of 6/12-bit display codes 7404g or 6-bit display code 00g produces a colon. In ASCII mode for interactive jobs, a colon can be input only as a 7404g 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 00g 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 (026 and 029) 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 0040g through 0176g.

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.

<u>Character Set</u>	<u>Print Train</u>	<u>PSU Printer Print Band</u>
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 0040g through 0176g), 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 (Sheet 1 of 2)

ASCII Graphic (64 Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12- Bit Display Code	7-Bit ASCII Code	ASCII Graphic (64 Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12- Bit Display Code	7-Bit ASCII Code
: colon		00†			# number sign	# number sign	60	60	0043
†Display code 00 is undefined at sites using the 63-character set.					[opening bracket	[opening bracket	61	61	0133
A	A	01	01	0101] closing bracket] closing bracket	62	62	0135
B	B	02	02	0102	% percent sign	% percent sign	63	63†	0045
C	C	03	03	0103	† colon	† colon	64	64	0042
D	D	04	04	0104	" quote	" quote	65	65	0137
E	E	05	05	0105	_ underline	_ underline	66	66	0041
F	F	06	06	0106	! exclamation point	! exclamation point	67	67	0046
G	G	07	07	0107	& ampersand	& ampersand			
					' apostrophe	' apostrophe	70	70	0047
H	H	10	10	0110	? question mark	? question mark	71	71	0077
I	I	11	11	0111	< less than	< less than	72	72	0074
J	J	12	12	0112	> greater than	> greater than	73	73	0076
K	K	13	13	0113	@ commercial at	@ commercial at	74†	7401	0100
L	L	14	14	0114	\ reverse slant	\ reverse slant	75	75	0134
M	M	15	15	0115	ˆ circumflex	ˆ circumflex	76†	7402	0136
N	N	16	16	0116	; semicolon	; semicolon	77	77	0073
O	O	17	17	0117					
P	P	20	20	0120	: colon	: colon		7404	0072
Q	Q	21	21	0121	˘ grave accent	˘ grave accent	74†	7407	0140
R	R	22	22	0122					
S	S	23	23	0123	a	a		7601	0141
T	T	24	24	0124	b	b		7602	0142
U	U	25	25	0125	c	c		7603	0143
V	V	26	26	0126	d	d		7604	0144
W	W	27	27	0127	e	e		7605	0145
					f	f		7606	0146
X	X	30	30	0130	g	g		7607	0147
Y	Y	31	31	0131					
Z	Z	32	32	0132	h	h		7610	0150
0	0	33	33	0060	i	i		7611	0151
1	1	34	34	0061	j	j		7612	0152
2	2	35	35	0062	k	k		7613	0153
3	3	36	36	0063	l	l		7614	0154
4	4	37	37	0064	m	m		7615	0155
					n	n		7616	0156
5	5	40	40	0065	o	o		7617	0157
6	6	41	41	0066					
7	7	42	42	0067	p	p		7620	0160
8	8	43	43	0070	q	q		7621	0161
9	9	44	44	0071	r	r		7622	0162
+ plus	+ plus	45	45	0053	s	s		7623	0163
- dash	- dash	46	46	0055	t	t		7624	0164
* asterisk	* asterisk	47	47	0052	u	u		7625	0165
					v	v		7626	0166
/ slant	/ slant	50	50	0057	w	w		7627	0167
(opening parenthesis	(opening parenthesis	51	51	0050					
) closing parenthesis) closing parenthesis	52	52	0051	x	x		7630	0170
\$ dollar sign	\$ dollar sign	53	53	0044	y	y		7631	0171
= equal	= equal	54	54	0075	z	z		7632	0172
space	space	55	55	0040	{ opening brace	{ opening brace	61†	7633	0173
, comma	, comma	56	56	0054	vertical line	vertical line	75†	7634	0174
. period	. period	57	57	0056	} closing brace	} closing brace	62†	7635	0175
					~ tilde	~ tilde	76†	7636	0176
					DEL	DEL		7637	0177

†The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

Table A-1. Character Sets for Interactive Jobs (Sheet 2 of 2)

ASCII Graphic (64 Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12- Bit Display Code	7-Bit ASCII Code	ASCII Graphic (64 Character)	ASCII Character (128 Character)	6-Bit Display Code	6/12- Bit Display Code	7-Bit ASCII Code
	NUL		7640	4000		DLE		7660	0020
	SOH		7641	0001		DC1		7661	0021
	STX		7642	0002		DC2		7662	0022
	ETX		7643	0003		DC3		7663	0023
	EOT		7644	0004		DC4		7664	0024
	ENQ		7645	0005		NAK		7665	0025
	ACK		7646	0006		SYN		7666	0026
	BEL		7647	0007		ETB		7667	0027
	BS		7650	0010		CAN		7670	0030
	HT		7651	0011		EM		7671	0031
	LF		7652	0012		SUB		7672	0032
	VT		7653	0013		ESC		7673	0033
	FF		7654	0014		FS		7674	0034
	CR		7655	0015		GS		7675	0035
	SO		7656	0016		RS		7676	0036
	SI		7657	0017		US†		7677	0037

†Reserved for network use. Refer to Character Set Tables in this appendix.

Table A-2. Character Sets for Batch Jobs (Sheet 1 of 3)

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code	Punch Code	
						026	029
: colon†	: colon†		00†			8-2	8-2
Display code 00 is undefined at sites using the 45-character set.							
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
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

†The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

Table A-2. Character Sets for Batch Jobs (Sheet 2 of 3)

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code	Punch Code	
						026	029
/ 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†
] closing bracket] closing bracket] closing bracket	62	62	0135	0-8-2	11-8-2†
% percent sign†	% percent sign†	% percent sign†	63†	63†	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
⤵ concatenation	underline	underline	65	65	0137	0-8-5	0-8-5
∨ logical OR	! exclamation point	! exclamation point	66	66	0041	11-0	12-8-7
∧ logical AND	& ampersand	& ampersand	67	67	0046	0-8-7	12
↑ 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	@ 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	^ circumflex	76			12-8-6	11-8-7
; semicolon	; semicolon	; semicolon	77	77	0073	12-8-7	11-8-6
		@ commercial at	74†	7401	0100		
		^ circumflex	76†	7402	0136		
		: colon†		7404†	0072		
		% percent sign		7406	0045		
		grave accent	74†	7407	0140		
		a		7601	0141		
		b		7602	0142		
		c		7603	0143		
		d		7604	0144		
		e		7605	0145		
		f		7606	0146		
		g		7607	0147		

†The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies elsewhere in this appendix.

Table A-2. Character Sets for Batch Jobs (Sheet 3 of 3)

CDC Graphic (64 Character)	ASCII Graphic (64 Character)	ASCII Graphic (95 Character)	6-Bit Display Code	6/12-Bit Display Code	7-Bit ASCII Code	Punch Code	
						026	029
		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†	7633	0173		
		vertical line	75†	7634	0174		
		} closing brace	62†	7635	0175		
		~ tilde	76†	7636	0176		

†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 (Sheet 1 of 2)

ASCII Character (128 Character)	7-Bit ASCII Code		6/12-Bit Display Code	ASCII Character (128 Character)	7-Bit ASCII Code		6/12-Bit Display Code
	Octal	Hexadecimal			Octal	Hexadecimal	
NUL	4000	00	7640	0	0060	30	33
SOH	0001	01	7641	1	0061	31	34
STX	0002	02	7642	2	0062	32	35
ETX	0003	03	7643	3	0063	33	36
EOT	0004	04	7644	4	0064	34	37
ENQ	0005	05	7645	5	0065	35	40
ACK	0006	06	7646	6	0066	36	41
BEL	0007	07	7647	7	0067	37	42
BS	0010	08	7650	8	0070	38	43
HT	0011	09	7651	9	0071	39	44
LF	0012	0A	7652	: colon††	0072	3A	7404††
VT	0013	0B	7653	; semicolon	0073	3B	77
FF	0014	0C	7654	< less than	0074	3C	72
CR	0015	0D	7655	= equal	0075	3D	54
SO	0016	0E	7656	> greater than	0076	3E	73
SI	0017	0F	7657	? question mark	0077	3F	71
DLE	0020	10	7660	@ commercial at	0100	40	7401
DC1	0021	11	7661	A	0101	41	01
DC2	0022	12	7662	B	0102	42	02
DC3	0023	13	7663	C	0103	43	03
DC4	0024	14	7664	D	0104	44	04
NAK	0025	15	7665	E	0105	45	05
SYN	0026	16	7666	F	0106	46	06
ETB	0027	17	7667	G	0107	47	07
CAN	0030	18	7670	H	0110	48	10
EM	0031	19	7671	I	0111	49	11
SUB	0032	1A	7672	J	0112	4A	12
ESC	0033	1B	7673	K	0113	4B	13
FS	0034	1C	7674	L	0114	4C	14
GS	0035	1D	7675	M	0115	4D	15
RS	0036	1E	7676	N	0116	4E	16
US†	0037	1F	7677†	O	0117	4F	17
space	0040	20	55	P	0120	50	20
! exclamation point	0041	21	66	Q	0121	51	21
" quote	0042	22	64	R	0122	52	22
# number sign	0043	23	60	S	0123	53	23
\$ dollar sign	0044	24	53	T	0124	54	24
% percent sign††	0045	25	63††	U	0125	55	25
& ampersand	0046	26	67	V	0126	56	26
' apostrophe	0047	27	70	W	0127	57	27
(opening parenthesis	0050	28	51	X	0130	58	30
) closing parenthesis	0051	29	52	Y	0131	59	31
* asterisk	0052	2A	47	Z	0132	5A	32
+ plus	0053	2B	45	[opening bracket	0133	5B	61
, comma	0054	2C	56	\ reverse slant	0134	5C	75
- dash	0055	2D	46] closing bracket	0135	5D	62
. period	0056	2E	57	^ circumflex	0136	5E	7402
/ slant	0057	2F	50	_ underline	0137	5F	65

†Reserved for network use. Refer to Character Set Tables in this appendix.

††The interpretation of this character or code may depend on its context. Refer to Character Set Anomalies in this appendix.

Table A-3. ASCII to 6/12-Bit Display Code Conversion (Sheet 2 of 2)

ASCII Character (128 Character)	7-Bit ASCII Code		6/12-Bit Display Code	ASCII Character (128 Character)	7-Bit ASCII Code		6/12-Bit Display Code
	Octal	Hexadecimal			Octal	Hexadecimal	
` grave accent	0140	60	7407	p	0160	70	7620
a	0141	61	7601	q	0161	71	7621
b	0142	62	7602	r	0162	72	7622
c	0143	63	7603	s	0163	73	7623
d	0144	64	7604	t	0164	74	7624
e	0145	65	7605	u	0165	75	7625
f	0146	66	7606	v	0166	76	7626
g	0147	67	7607	w	0167	77	7627
h	0150	68	7610	x	0170	78	7630
i	0151	69	7611	y	0171	79	7631
j	0152	6A	7612	z	0172	7A	7632
k	0153	6B	7613	{ opening brace	0173	7B	7633
l	0154	6C	7614	vertical line	0174	7C	7634
m	0155	6D	7615	} closing brace	0175	7D	7635
n	0156	6E	7616	~ tilde	0176	7E	7636
o	0157	6F	7617	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>External BCD</u>	<u>6-Bit Display Code</u>
00	16 (Z)	00
33 (0)	12 (0)	33 (0)
63 (:)	12 (0)	33 (0)
	Output →	← Input
<u>64-Character Set</u>		
<u>6-Bit Display Code</u>	<u>External BCD</u>	<u>6-Bit Display Code</u>
00 (:)	12 (0)	33 (0)
33 (0)	12 (0)	33 (0)
63 (Z)	16 (Z)	63 (Z)
	Output →	← Input

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				6-Bit Display Code		7-bit ASCII				6-Bit Display Code	
Code (Hex)	Char†	Code (Hex)	Char††	Char	Code (Octal)	Code (Hex)	Char†	Code (Hex)	Char††	Char	Code (Octal)
20	space	00	NUL	space	55	3E	>	1E	RS	>	73
21	!	7D	}	!	66	3F	?	1F	US	?	71
22	"	02	STX	"	64	40	@	60	`	@	74
23	#	03	ETX	#	60	41	A	61	a	A	01
24	\$	04	EOT	\$	53	42	B	62	b	B	02
25	%	05	ENQ	%	63	43	C	63	c	C	03
26	&	06	ACK	&	67	44	D	64	d	D	04
27	'	07	BEL	'	70	45	E	65	e	E	05
28	(08	BS	(51	46	F	66	f	F	06
29)	09	HT)	52	47	G	67	g	G	07
2A	*	0A	LF	*	47	48	H	68	h	H	10
2B	+	0B	VT	+	45	49	I	69	i	I	11
2C	,	0C	FF	,	56	4A	J	6A	j	J	12
2D	-	0D	CR	-	46	4B	K	6B	k	K	13
2E	.	0E	SO	.	57	4C	L	6C	l	L	14
2F	/	0F	SI	/	50	4D	M	6D	m	M	15
30	0	10	DLE	0	33	4E	N	6E	n	N	16
31	1	11	DC1	1	34	4F	O	6F	o	O	17
32	2	12	DC2	2	35	50	P	70	p	P	20
33	3	13	DC3	3	36	51	Q	71	q	Q	21
34	4	14	DC4	4	37	52	R	72	r	R	22
35	5	15	NAK	5	40	53	S	73	s	S	23
36	6	16	SYN	6	41	54	T	74	t	T	24
37	7	17	ETB	7	42	55	U	75	u	U	25
38	8	18	CAN	8	43	56	V	76	v	V	26
39	9	19	EM	9	44	57	W	77	w	W	27
3A	:	1A	SUB	:	00	58	X	78	x	X	30
3B	;	1B	ESC	;	77	59	Y	79	y	Y	31
3C	<	7B	{	<	72	5A	Z	7A	z	Z	32
3D	=	1D	GS	=	54	5B	[1C	FS	[61
						5C	\	7C		\	75
						5D]	01	SOH]	62
						5E	^	7E	~	^	76
						5F	_	7F	DEL	_	65

†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.

††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.

†††A 6-bit display code space always translates to a 7-bit ASCII space.

Table A-5. Nine-Track EBCDIC Coded Tape Conversion

EBCDIC				6-Bit Display Code		EBCDIC				6-Bit Display Code	
Code (Hex)	Char†	Code (Hex)	Char††	Char††	Code (Octal)	Code (Hex)	Char†	Code (Hex)	Char††	Char	Code (Octal)
40	space	00	NUL	space	55	C4	D	84	d	D	04
4A	↑	1C	IFS	[61	C5	E	85	e	E	05
4B	.	0E	SO	.	57	C6	F	86	f	F	06
4C	<	C0	{	<	72	C7	G	87	g	G	07
4D	(16	BS	(51	C8	H	88	h	H	10
4E	+	0B	VT	+	45	C9	I	89	i	I	11
4F		DO	}	!	66	D1	J	91	j	J	12
50	&	2E	ACK	&	67	D2	K	92	k	K	13
5A	!	01	SOH]	62	D3	L	93	l	L	14
5B	\$	37	EOT	\$	53	D4	M	94	m	M	15
5C	*	25	LF	*	47	D5	N	95	n	N	16
5D)	05	HT)	52	D6	O	96	o	O	17
5E	;	27	ESC	;	77	D7	P	97	p	P	20
5F	-	A1	~	^	76	D8	Q	98	q	Q	21
60	/	0D	CR	-	46	D9	R	99	r	R	22
61	/	0F	SI	/	50	E0	\	6A		\	75
6B	,	0C	FF	,	56	E2	S	A2	s	S	23
6C	%	2D	ENQ	%	63	E3	T	A3	t	T	24
6D	~	2D	ENQ	space†††	55	E4	U	A4	u	U	25
6E	>	07	DEL	>	73	E5	V	A5	v	V	26
6F	?	1E	IRS	?	71	E6	W	A6	w	W	27
7A	:	1F	IUS	:	71	E7	X	A7	x	X	30
		3F	SUB	:	00	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

†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.

††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.

†††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.

Table A-6. Seven-Track Coded Tape Conversions

External BCD	ASCII Character	6-Bit Display Code (Octal)	External BCD	ASCII Character	6-Bit Display Code (Octal)
01	1	34	40	-	46
02	2	35	41	J	12
03	3	36	42	K	13
04	4	37	43	L	14
05	5	40	44	M	15
06	6	41	45	N	16
07	7	42	46	O	17
10	8	43	47	P	20
11	9	44	50	Q	21
12†	0	33	51	R	22
13	=	54	52	!	66
14	"	64	53	\$	53
15	@	74	54	*	47
16†	%	63	55	'	70
17	[61	56	?	71
20	space	55	57	>	73
21	/	50	60	+	45
22	S	23	61	A	01
23	T	24	62	B	02
24	U	25	63	C	03
25	V	26	64	D	04
26	W	27	65	E	05
27	X	30	66	F	06
30	Y	31	67	G	07
31	Z	32	70	H	10
32]	62	71	I	11
33	,	56	72	<	72
34	(51	73	.	57
35	(65	74)	52
36	#	60	75	\	75
37	&	67	76	^	76
			77	;	77

†As explained previously in this section, conversion of these codes depends on whether the tape is being read or written.

GLOSSARY

B

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.

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 MSAS environment.

Allocation Unit (AU)

In an MSS environment, an allocation unit is the name that is given to each of the 16 data streams that make up a cartridge.

In an MSAS environment, an allocation unit is the smallest allocatable portion of a cartridge. Each AU consists of 14 data strips; each cartridge has 1931 AUs.

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 CDC 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.

asa

The alternate storage address.

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.

ASDEBUG Utility

A utility that allows an analyst to update appropriate entries in the CSU maps and/or MSF catalogs and thereby resolve inconsistencies reported by the ASVAL utility. ASDEBUG can also be used to copy data from selected MSF files or cartridges to disk.

ASDEF Utility

A utility that creates the system files (CSU maps and MSF catalogs) that are necessary for MSS processing.

ASLABEL Utility

A utility that manages cartridge assignment and cubicle allocation in a cartridge storage unit.

ASMOVE Utility

A utility that manages disk and Mass Storage Facility residence. ASMOVE determines which files to leave on disk, which files to release from disk and move to MSF, and which files should reside both on disk and on MSF.

AST

See Allocation Summary Table and Available Stream Table.

ASUSE Utility

A utility that reads data in the MSF catalogs and CSU maps and produces reports on the availability of space on MSF cartridges and the allocation of cubicle space within a CSU.

ASVAL Utility

A utility that either performs release processing or reports on problems with the current MSS system files.

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.

Auxiliary Device

A disk device that is not part of a permanent file family. Auxiliary devices can contain direct or indirect access permanent files.

Available Stream Table (AST)

A table that contains information used by the allocation algorithm to select the cartridges on which a file will reside in an MSS environment.

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

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.

Cache

A high-speed memory that resides in the central processor.

Cartridge

A component of the MSF or 7990. The cartridge consists of a plastic housing that encloses a strip of magnetic tape on which data is stored under program control.

Cartridge Storage Unit (CSU)

An MSS device that includes storage cells (cubicles) for 2052 cartridges and a selector that moves cartridges among the mass storage transport, the cubicles, and the input/output drawers of the CSU.

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.

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

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.

CIR

See Catalog Image Record.

CLA

See Communications Line Adapter.

CMR

See Central Memory Resident.

CCP

See Communications Control Program.

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.

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.

Coldstart

Procedure used to deadstart if the tape or disk controller has not yet been loaded with controlware, or the controlware is not running.

Channel Number

The number of the data channel on which a peripheral device controller can be accessed.

Combined Input/Output (CIO)

A system routine that performs NOS I/O.

Character

Unless otherwise specified, references to characters in this manual are to 7-bit ASCII code.

Common Testing and Initialization (CTI)

Common deadstart process that resides on the deadstart file and the maintenance system library.

Charge Number

An alphanumeric identifier the installation uses to allocate charges to individual users for system usage.

Communication Line

A complete communication circuit between a terminal and its network processing unit.

Checkpoint

The process of writing a copy of your job's central memory, the system

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.

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.

Coupler

A component of the MSF. The coupler interfaces between the peripheral processor and the mass storage adapter.

CS

See Communications Supervisor.

CSU

See Cartridge Storage Unit.

CSU Map

A direct access permanent file that contains information indicating how cubicles in a CSU are assigned to a family and identifying the cartridges that reside in the CSU.

CTI

See Common Testing and Initialization.

CYBERLOG

A program initiated after deadstart that requests data about the previous service interruption.

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 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.

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 multimainframe complex.

Device Index Table (DIT)

A table that is used to determine device usage in an independent shared device multimainframe complex.

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.

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.

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.

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.

End-of-Chain Flag (EOC)

In an MSS environment, EOC is a flag in the MSF catalog that identifies the last allocation unit in a chain of allocated data streams.

In an MSAS 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

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.

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.

EQPDECK

The equipment description deck used during deadstart to define the system equipment configuration.

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.

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.

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)

ECS provides optional additional memory. ECS contains 60-bit words; it has a large amount of storage and fast transfer rates. ECS can be used only for program and data storage, not for program execution. Special hardware instructions exist for transferring data between central memory and ECS.

Extended Memory

An extension to central memory that is physically located outside of the machine.

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 CSU or 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

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.

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)

Flag that identifies the first allocation unit in a chain of allocated data streams in the MSF catalog.

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.

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

See Inactive Queued File Table.

ITF

See Interactive Transfer Facility.

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.

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 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

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.

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.

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 Adapter (MSA)

A component of the MSF. The adapter interfaces between the coupler and the mass storage transport or the cartridge storage unit. See Mass Storage Facility.

Mass Storage Archival Subsystem (MSAS)

MSAS is the product consisting of the 7990 hardware, the channel interface, the diagnostics, and the operational software. MSAS stores data on the 7990 and moves it to disk upon request for access by an authorized user.

Mass Storage Device

An extended memory or disk unit that has defined logical attributes such as family, file residency, and so on.

Mass Storage Facility (MSF)

A hardware product that is a large capacity online mass storage device. MSF is a cost-effective extension to the disk file storage system and an alternative to conventional magnetic tape storage.

Mass Storage Facility (MSF) Catalog

A disk-resident direct access permanent file that contains information describing which streams of each cartridge assigned to a particular subfamily are allocated to MSF files and which streams are available for allocation.

Mass Storage Subsystem (MSS)

MSS is the product consisting of the MSF hardware, the CYBER coupler, the diagnostics, and the operational software. MSS stores data on the MSF 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.

Mass Storage Transport

A component of the MSF. MST has storage positions for five cartridges: one being read or written, two queued for reading or writing, and two queued for storage by the selector into the CSU.

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.

MID

See Machine Identification.

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.

MSA

See Mass Storage Adapter.

MSAS

See Mass Storage Archival Subsystem.

MSF

See Mass Storage Facility.

MSS

See Mass Storage Subsystem.

MSSEEXEC Program

The main processing program that controls MSS activities, such as destaging files from disk to MSF, purging unneeded MSF files, labeling or relabeling cartridges, updating CSU maps, and updating MSF catalogs.

MSSSLV Program

A program that runs on each slave mainframe and communicates with the MSSEEXEC program to retrieve files from MSF in response to ATTACH requests by jobs running on the slave mainframes.

MST

See Mass Storage Table.

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.

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.

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.

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 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 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.

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.

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.

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.

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 (PFC)

A 16-word entry that the system maintains and uses to determine 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

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.

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.

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.

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.

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.

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.

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.

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 MSF/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.

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.

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.

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.

SSEXEC Program

The main processing program that controls MSAS 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.

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.

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 MSAS 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.

TAF

See Transaction Facility.

TCU

See Trunk Control Unit.

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.

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.

UDT

See Unit Device Table.

UEM

See Unified Extended Memory.

Unified Extended Memory (UEM)

A type of extended memory. 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.

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.

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).

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

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 and 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

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 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.

jsn The job sequence name of SCRSIM noted on the B,0 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.

command One of the commands shown in figure D-2.

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.

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)

SIMULATOR COMMANDS

COMMAND	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)

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.

<u>Command</u>	<u>Description</u>
AREA,a,m,y.	Set m bits in the holding register, from bit a to bit a+m-1, to the octal value y.

<u>Parameter</u>	<u>Description</u>
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.
---------------	--

<u>Parameter</u>	<u>Description</u>
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.
--	--

<u>Parameter</u>	<u>Description</u>
a _i	Bit number to be cleared; from 0 to 203. A decimal value is assumed, but a postradix of D or B may be included.

<u>Command</u>	<u>Description</u>
LINE,x,y.	Set line x of the holding register to the octal value y.

<u>Parameter</u>	<u>Description</u>
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.
--	--

<u>Parameter</u>	<u>Description</u>
a ₁	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

<u>Command</u>	<u>Description</u>
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.

<u>Parameter</u>	<u>Description</u>
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

<u>Command</u>	<u>Description</u>
+	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.

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 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 17g. 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 17g) 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 17g 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 legally 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.

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 is included in this entry. This value can be any 6-bit binary value. 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

<u>Variable</u>	<u>Description</u>
a, b, c, d	Programmable format control numbers for 6 or 8 lpi. These numbers can be any value from 0 to 12 and 178.

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 (Sheet 1 of 4)

Location	6 lpi Default †		6 lpi Alternate ††		8 lpi Default		8 lpi Alternate	
	Form Length		Form Length		Form Length		Form Length	
	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.
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
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

† Default (SC=0).
 †† Alternate (SC=1).

Table E-1. Released Programmable Format Control Arrays (Sheet 2 of 4)

Location	6 lpi Default †		6 lpi Alternate ††		8 lpi Default		8 lpi Alternate	
	Form Length		Form Length		Form Length		Form Length	
	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.
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
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

† Default (SC=0).
†† Alternate (SC=1).

Table E-1. Released Programmable Format Control Arrays (Sheet 3 of 4)

Location	6 lpi Default †		6 lpi Alternate ††		8 lpi Default		8 lpi Alternate	
	Form Length		Form Length		Form Length		Form Length	
	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.
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
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

† Default (SC=0).
†† Alternate (SC=1).

Table E-1. Released Programmable Format Control Arrays (Sheet 4 of 4)

Location	6 lpi Default †		6 lpi Alternate ††		8 lpi Default		8 lpi Alternate	
	Form Length		Form Length		Form Length		Form Length	
	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.	11 in.	8.5 in.
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	
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	

†Default (SC=0).
††Alternate (SC=1).



881/883 PACK REFORMATTING UTILITY

F

Each 881 disk pack used in the 844 disk contains factory-recorded flawing information on cylinder 632g (410), track 0, sectors 0, 1, and 2. Each 883 pack contains this information on cylinder 1466g (822), track 0, sectors 0, 1, and 2. The following information is included on the cylinders.

- Cylinder 632g (or 1466g 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 632g (or 1466g), track 0, sector 1 contains the factory map. This map contains a list of all factory-detected flaws, both correctable and uncorrectable.
- Cylinder 632g (or 1466g), 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 632g (or 1466g)] 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 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, online, or is the result of running FORMAT. 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, . . . , P_n.

Each P_i is a keyword or a keyword equated to a value.

<u>P_i</u>	<u>Description</u>
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.
MODE=mode	Operational mode for FORMAT.

<u>mode</u>	<u>Description</u>
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 632g (or 1466g), 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.

<u>P_i</u>	<u>Description</u>
MODE	Same as MODE=FETCH.
O=filename	Optional output file to contain the output extracted from the disk pack.

NOTE

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=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.

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.

<u>Directive</u>	<u>Description</u>
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

<u>Parameter</u>	<u>Description</u>						
x	Type of flaw to be set or cleared. Acceptable values are: <table><thead><tr><th><u>x</u></th><th><u>Description</u></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>	<u>x</u>	<u>Description</u>	S	Specifies that the SET or CLEAR directive applies to a sector.	T	Specifies that the SET or CLEAR directive applies to a track.
<u>x</u>	<u>Description</u>						
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 632g (or 1466g), track 0, sector 0.
- A listing of the factory flaw map contained on cylinder 632g (or 1466g), track 0, sector 1.
- A listing of the utility flaw map contained on cylinder 632g (or 1466g), 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 F-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 F-1. FORMAT Output, MODE=RESTORE

Example 2:

The flaws noted in the factory flaw map from example 1 (refer to figure F-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 F-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 F-2. FORMAT Output, MODE=ALTER

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 F-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 F-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, Multimainframe Operations.)
 - If the device is shared, a global unload should be set.
3. The device must be declared logically off. Use the OFF,est 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.

INDEX

- Access count 17-9
- Access level limits 3-61,106
- Account dayfile 3-10,66
- ACPD 21-5
- Active dayfiles 18-21
- Active queued files 18-20,23,31
- *AD LIBDECK entry 3-29
- Add a mass storage device 3-64
- Add cartridges 11-15
- Add cubicles 11-15
- Add SMS 11-15
- AFD 18-44
- AIP trace 16-1,26
- Allocation information 3-33
- Allocation summary table 11-4
- Allocation unit(s) 11-2; 12-2
- Alter memory file 15-4
- Alternate CATLIST permission 17-12
- Alternate permanent file family 3-31
- Alternate storage address 17-14,39
- Alternate storage facility 3-99
- Alternate storage flags
 - Data error 17-14,39
 - Free up cartridge 17-12,39
 - Lock(ed) to disk 17-12,39
 - Obsolete 17-14,39
 - Pseudorelease 17-13,39
 - System control error 17-14,39
 - Temporary error 17-13,39
 - Verification 17-13,39
- Alternate storage type 17-14,39
- Alternate system device 3-29,112
- Alternate system library device assignment entry 3-55
- Answerback 16-30
- Application interface program trace 16-1,26
- APRDECK
 - Description 3-69
 - Display 2-15
 - Entries
 - Refer to tab for section 3; 3-69
 - Modifying 2-15
- APRINST display 2-15
- Archive file
 - Description 17-1
 - Multifile 17-2
 - Preassignment 17-37
- asa 12-9; 17-14,39
- ASDEBUG
 - Description 12-2,40
 - Directives 12-41
 - Parameters 12-42
 - Restrictions 12-44
- ASDEF 12-2,11,45
- ASLABEL
 - Description 12-2,12
 - Directives 12-13
 - Example 12-19
 - Parameters 12-16
 - Restrictions 12-19
 - Update sequence 12-18
- ASMOVE
 - Description 12-2,9,20
 - Example 12-24
 - Restrictions 12-24
 - Selection algorithms 12-23
- ASSIGN 17-37; 18-20
- Associate job files with devices 3-50
- ASUSE
 - Basic usage report 12-34,36
 - Description 12-2,34
 - Optional report A 12-38
 - Optional reports 12-34
 - Optional reports A and B 12-37
 - Optional reports C, D, and E 12-39
- ASVAL
 - Description 12-2,26
 - Error detection and classification 12-28
 - Examples 12-33,46,47
 - Problem reporting 12-28,30,40,50
 - Release processing 12-20,27,30,46
 - Restrictions 12-34
 - Validation report 12-31,32,51,53
- Attach request 3-102
- AU conflict flag 11-8
- AUs
 - Available for large files 11-5
 - Available for small files 11-5
- Auto recall interval 3-98
- Automatic flawing F-1
- Automatic flawing of devices 3-70
- Automatic label checking for removable devices 3-105
- Auxiliary device 3-32,51; 17-19
- Auxiliary device pack name 3-52
- Auxiliary packs 3-95
- Available stream table 12-5

Backup requirement(s) 11-10,72;
 12-9,20,24,45,49,51; 17-11,39,46
 Batch I/O 3-94
 Becoming a DOP 15-30
 Becoming a NOP 15-16
 Binary maintenance log 3-12,41,66
 Dayfile dumping utility 18-43
 Block identification on 7021
 controller 3-20
 Buffered mass storage devices 3-41

Calling FLAW utility 8-11
 Calling MREC utility 8-21
 Calling the express deadstart dump
 interpreter 6-2
 Cancel alternate NPU load file command 15-8
 Card punch 3-16
 Card reader 3-16
 Carriage control tape E-1
 Cartridge
 Add 11-15
 Allocation to a file 12-5,7,15,34,46
 Assignment 11-2,21; 12-1,2,4,9,12,14,
 15,19
 Availability 12-2,34
 Description 11-1,2; 12-1,2
 File allocation 11-17
 Label(s)
 Information 11-2,21; 12-2,12,14,15,
 19,51
 Repairing 12-12,15
 Updating 12-2,14,18
 Location 12-17
 Lost, restore 11-16
 Manufacturer code 11-4,6
 Pool 11-17
 Recover data from 12-2,40
 Removal of 11-75; 12-52
 Remove 11-16
 Residence in CSU 12-3
 Restoration 11-75; 12-12,15,51
 Serial number 11-4
 Stream 12-2
 Cartridge accessor unit 11-1
 Cartridge storage unit 3-20
 Assignment to subfamily 12-12,14,45
 CSU identifier 12-1,11,16,26,35,42
 CSU map 12-3,4,49
 Analysis 12-28
 Problems 12-52
 Description 12-1
 Pool 12-4,14
 Reserved area 12-3,14
 Catalog entry 17-9
 Catalog image record
 Catalog of 17-48
 Description 17-49
 PF utility option 17-34
 Catalog information 3-33

Catalog track
 Chain 17-8
 Description 17-7
 Entry 17-9
 Number of 17-7
 Central memory
 Clearing 3-87
 Descriptions 3-2
 Reservation 3-5
 Resident 3-2
 Time slice 3-108
 Central processor multiplier 3-82
 Chain control field 11-9
 Chain of streams 12-2,9,29,30,31,53
 Change flag in CSU map 12-42
 Change flag in MSF catalog 12-42
 Change flag in SM map 11-55
 Change NPU load file command 15-8
 Changing LIDs using the LIDOU display 10-7
 Channel 15-1
 Channel control commands 5-50
 Character set(s) A-1
 Anomalies A-2
 Mode 3-84
 Tables A-2
 Charge number 17-15,39
 Check mass storage (CMS) routine 3-65
 Checkpoint file 3-55
 Checkpoint, multimainframe 13-9,15
 Clear all flaw reservations 3-70
 Clear logical flaws on any mass storage
 device 3-71
 Clear physical extended memory track or disk
 area flaws 3-72
 Clear tape unit assignment 3-18
 Clearing permanent file error flags 11-74;
 12-50
 Cluster control blocks 16-12
 CMRDECK
 Description 3-1
 Display 2-2
 Entries
 Refer to tab for section 3; 3-2
 Entry format 3-2
 Listing 3-1
 Modifying 2-2; 3-1
 CMRINST 3-1
 Display 2-2
 Code conversion parameters, suggested
 values 10-13
 Commands under DIS 4-13
 Common library table 3-2,81
 Common test and initialization software 7-1
 Common testing and initialization 3-7
 PPU CMRDECK relationship 3-7
 Communications control program files 16-6
 Communications supervisor (CS) 15-1
 Configuring shared mass storage devices
 Independent SDM 13-13
 Linked SDM 13-6

Console equipment EST entry 3-15
 Console operation under DIS 4-10
 Continuation AU flag 11-9
 Control cartridge allocation 11-17
 Control module
 Controlware 1-1
 EST entry 3-46
 Control of jobs 3-75
 Control point area 3-2
 Control point, job leaves 3-75
 Control points, number of 3-5
 Controller, 7021 block identification 3-20
 Controlware
 Loading utility 1-1
 Status display 2-11
 Type 1-3
 Version number 3-61
 Conversion mode, tape 3-92
 Coupler 15-1
 Disable or enable 15-6
 Off 15-6
 On 15-6
 Coupler, MSF 12-1
 CP multiplier 3-82
 CPU
 Priorities 3-81,107
 Recall delay 3-98
 Time slice 3-107
 Creating the LID configuration file 10-1
 CS
 Communications supervisor 15-1
 Information command 15-34
 Cubicle assignment 11-3
 Cubicles 11-1
 Current residence 17-39
 CYBERLOG
 Component 8-4
 Edit 8-9
 Elapsed time 8-7
 Event type 8-3
 Impact 8-8
 K display 8-2
 Lost time 8-7
 Reason 8-3
 Responding to displays 8-2
 CYBERPLUS ring port EST entry 3-28

 Data gathering facility 3-88
 Data recording drives 11-1
 Data storage display (F) 4-5
 Database control system, CYBER 3-94
 Date line 3-5
 Date/time entry 2-18

 Dayfile
 Account 3-10
 Description 3-10
 Error log 3-11
 Messages 3-99
 System 3-11
 Dayfile commands 5-2
 Dayfile dumping utilities 18-1,43
 Dayfile dumps 5-3
 Dayfile residence examples 3-67
 Dayfile termination 18-13,21
 DDP rollout path 3-98
 Deadstart
 DSD entry 3-85
 Error idle recovery 2-35
 Error processing 2-28
 Level 0 initial 2-28
 Level 1 recovery 2-26
 Level 2 recovery 2-27
 Level 3
 Left screen display 2-23
 Recovery 2-23
 Right screen display 2-24
 Mass storage device recovery 2-30
 Process 2-1
 Recovery
 Level 1 2-26
 Level 2 2-27
 Level 3 2-23
 Preparation 2-21
 Deadstart decks 3-1
 Deadstart dump interpreter
 Buffer controller directive 6-41
 Command 6-2
 Description 6-1
 Directive format 6-6
 Dump hardware registers 6-37
 Dump memory
 Central 6-11
 CMR 6-21
 Extended 6-11
 PP 6-16
 Dump subsystem 6-34
 Dump tape request 6-4
 Examples
 Interactive 6-47
 Printer output 6-49
 Features 6-1
 File manipulation and control
 directives 6-9
 Input directives 6-5
 Interactive use 6-42
 List control directives 6-7
 Parameters 6-2
 Terminal output directives 6-43

Deadstart file
 CMRDECK 3-1
 EQPDECK 3-8
 Deadstart sequencing 20-3
 Debug log file processor 16-26
 Debug mode 3-97
 Decision algorithm hierarchy 11-30
 Default family name assignment entry 3-57
 Delay values, record of 5-9
 Delete device 3-56
 DEMUX 16-42,44
 Destage 11-10; 12-2,9,20,23,24,45,46
 Destage algorithm 12-23
 Destage/release processing
 Algorithms 11-29
 Excluding 11-28
 Device
 Auxiliary 3-32
 Family 3-32
 Link 3-33
 Master 3-33
 Private 3-32
 Public 3-32
 Device access table 13-2,6
 Device configuration
 Independent SDM 13-11
 Linked SDM 13-3
 Device index table 13-10,13
 Device initialization options table 8-16
 Device mask 3-52; 17-4
 Examples 17-4,5
 DFD 18-44
 DFLLIST (refer also to Queue file utilities) 18-20
 DFTERM (refer also to Queue file utilities) 18-21
 Diagnostic operator (DOP) 15-3
 Direct access files 3-33,110
 Directory display (Z) 4-9
 DIS
 Commands 4-13
 Console operation 4-10
 Data storage display (F) 4-5
 Directory display (Z) 4-9
 Display identifiers 4-2
 Display selection commands 4-12
 Exchange package display (B) 4-4
 Job dayfile display (A) 4-3
 Job status display (J) 4-7
 Keys with special meaning 4-10
 Memory displays 4-5
 Memory entry commands 4-16
 Messages 4-11
 Operations 4-1
 PP call commands 4-18
 Program storage display (G) 4-6
 Disable host element command 15-11
 Disable network element command 15-22
 Disk controlware 1-1
 Disk drives
 844 3-34
 885 3-34
 Disk equipment EST entry 3-42
 Disk file format
 Full stripe 11-59
 Partial stripe 11-60
 Disk space
 Management 11-10,31,68; 12-9,20,24,45
 Release 11-10; 12-2,9,20,23,24,45
 Display console equipment EST entry 3-15
 Display identifiers (DIS) 4-2
 Display selection commands (DIS) 4-12
 Disposition type 18-11,14,18
 DMPNAD utility 14-1
 DOP
 Commands 15-34
 Diagnostic operator 15-3
 Network control 15-30
 Status message formats 15-31
 Unsolicited status reports 15-33
 DPFC macro E-3
 DS equipment EST entry 3-15
 DSD commands
 Categories of 5-1
 Description 5-1
 Refer to tab for section 5
 Dual-access unit 3-42
 Dump
 Dayfiles 5-3
 Full 11-71; 12-47,48; 17-45
 Incremental 11-71; 12-48; 17-45
 MSS/MSAS environment 17-47
 Partial 17-45
 Secured system 17-45
 Dump NAD memory
 Description 14-1
 Local NAD dumping 14-2
 Remote NAD dumping 14-2
 Dump tape management 11-31,72; 12-24,49
 ELD 18-44
 Elements of network 15-1
 Enable application command 15-12
 Enable subsystems 3-94
 End-of-chain 11-2,40,43; 12-2,29,31,53
 ENDCPD 21-6
 Engineering mode 3-98
 Entering flaws for a device 3-69
 Environment interface loading 3-85
 EQPDECK
 Description 3-8
 Display 2-3
 Entries 2-12; refer to tab for section 3; 3-10
 Entry format 3-10
 Listing 3-8
 Modifying 2-3; 3-8

- Range of EST ordinals 3-42
- Restrictions 3-9
- Selection 3-3
- EQPINST 3-8
 - Display 2-3
- Equipment assignments
 - Clearing 3-13
 - Description of entries 3-12
 - Mass storage 3-28
 - Nonmass storage 3-14
- Equipment configuration example 10-4
- Equipment status display 2-4
- Error code 17-10
- Error flag 17-10
- Error idle recovery 2-35
- Error log 3-11,66
- Error processing
 - Deadstart from disk 2-29
 - Deadstart from tape 2-28
 - Error idle recovery 2-35
- EST entry
 - Control module 3-46
 - CYBERPLUS ring port 3-28
 - Disk equipment 3-42
 - Extended memory 3-47
 - Magnetic tape equipment 3-18
 - MAP equipment 3-27
 - Mass storage archival subsystem 3-22
 - Mass storage facility
 - CSU 3-20
 - Mass storage transport 3-21
 - Network access device 3-26
 - Network processing unit 3-24
 - Nonstandard equipment 3-14
 - Pseudoequipment 3-14
 - Stimulator equipment 3-23
 - System display console equipment 3-15
 - Two-port multiplexer 3-26
 - Unit record equipment 3-15
 - 6683 satellite coupler 3-27
- Excessive write parity error flag 11-7
- Exchange package display (B) 4-4
- Executing job table 3-3
- Expander addressing map 3-35
 - EST entry 3-34
- Express deadstart dump interpreter (DSDI) 6-1
- Extended memory
 - As a link device 3-33
 - Clearing 3-87
 - EST entry 3-47
 - Field length 3-108
 - Flag register commands 5-51
 - Maintenance mode 3-64
 - Reserve for user access or buffer 3-61,62
 - User 3-111
- Extended semiconductor memory (ESM) (also refer to extended memory)
 - As a link device 3-33
 - EST entry 3-47
- External characteristics code 18-15
- Family 17-19
- Family, adding a new 3-53
- Family device 3-32,51
- Family name to family ordinal relationship 3-4
- Family ordinal table 3-4; 18-4
- Family reconfiguration 17-52
- Fast-attach files 20-1
- Fast-attach table 20-1
- FCT ordinal 11-3,7; 12-4,7,39
- Field length 3-108
- File access categories 17-15
- File access level 17-15,39
 - Lower limit 17-32
 - Upper limit 17-36
- File and cartridge table 11-4; 12-5,42
- File category 17-10,38
- File dumping 11-70
- File label verification procedure 12-10
- File length field 11-9
- File name and volume serial number association 3-15
- File name table 3-3
- File reloading 11-73
- File staging 3-99
- File type to move or dump 18-19
- Files
 - Direct access 3-33,110
 - Indirect access 3-33,109
- Fix cartridge labels 11-17
- Fixed-length data structures 16-11
- Flaw data, disk pack F-1
- Flaw entries 8-12
 - CLF 8-12
 - CPF 8-12
 - SLF 8-12
 - SPF 8-12
- FLAW K display 8-10
- Flaw reservations lost 3-64
- FLAW utility 8-10
- Flawed area 8-10
- Flawed AU flag 11-9
- Flawing
 - Mass storage 3-69
- Flawing information for disks and EM 8-14
- Flawing tracks on mass storage devices 8-10
- Format of CMRDECK 3-2
- Format of EQPDECK 3-10
- FORMAT utility 3-98

Forms code 3-17; 18-12
FOTD L display 9-2
Fragment 11-41,43; 12-8,29,30,31,53
Free/busy flag 11-8
Free cartridge files 11-17
Free cartridge flag 11-7
Frozen chain 12-8,30
Frozen chain flag 11-8
Full dump 17-46

General job control commands 5-4
GETLOG 18-44
Glossary B-1
Group 11-11
Group-encoded tapes 3-19

Hardware diagnostic information 3-12
Hash table 3-45; 6-28
Head-of-chain 11-2,40,43; 12-2,29,31,53
Highest priority 3-74,104
HOP
 Commands 15-7
 Host operator 15-3
 Network control 15-6
 Status message formats 15-6
 Unsolicited status reports 15-7
Host operator (HOP) 15-3

IAFEX 16-26
IAFTM 16-26
IAFTR 16-27
ICPD 21-4
Idle host element command 15-12
Inactive dayfiles 18-21
Inactive queued files 18-20,32
Incremental dump 17-46
Incremental load 17-50
Independent SDM operation
 CMRDECK entry 13-11
 Deadstart 13-11
 Description 13-10
 Device checkpoint 13-15
 Device configuration 13-11
 Device recovery 13-15
 EQPDECK entry 13-11
 MSAS processing 13-15
 MSS processing 13-15
Independent shared device 3-68
Independent shared device entry 3-58
Index, user 17-3
Indirect access files 3-33,109
Inhibit allocation flag 11-6
INIT startup procedure 16-2,5
Initial DSD commands 3-85
Initialize and flaw tracks procedure 8-15
INITIALIZE K display 8-14
Initializing and releasing fast-attach files 20-3
Initiating
 Job processing 2-19
 NAM 15-1
Initiating NAM 15-3
 Without operator intervention 15-3
INSTALL 7-1
Instruction-stack purging 3-86
Interactive facility (IAF) 3-94
 Number of network terminals 3-88
Interactive job control commands 5-6
Interactive stimulator 3-23
Intersections 11-43; 12-29,32
IPRDECK
 Deadstart entries and DSD commands 3-93
 Deadstart only entries 3-81
 Description 3-73
 Display 2-16
 Entries
 Refer to tab for section 3; 3-73
 Listing 3-73
 Modifying 2-16; 3-73
 Selection 3-4
 Subsystem entries 3-94
IPRINST 3-73
 Display 2-16
ISF 20-1

Job
 Maximum extended memory length 3-108
 Maximum field length 3-108
 Maximum number 3-108
Job communication commands 5-7
Job control 3-75,77
Job control information 3-73
Job dayfile display (A) 4-3
Job flow commands 5-8
Job processing
 Control commands 5-3
 Initiating 2-19
Job queue type 3-104
Job scheduling 3-74
Job scheduling delay 3-98
Job status display (J) 4-7
Job switch delay 3-98
Jobs
 Batch A-3
 Interactive A-2
 Using line printers A-3
 Using magnetic tapes A-11

K-display utilities 8-1
 Keypunch mode 3-86
 Keys with special meaning for DIS 4-10

L-display utilities 9-1
 FOTD 9-2
 LIDOU 9-3
 QDSPLAY 9-6
 SCTD 9-8
 SUBSYST 9-10
 LABEL 17-37; 18-20
 Label, cartridge
 Description 12-2
 File label verification procedure 12-10
 Repairing 12-3,12,15
 Stream label verification procedure 12-3
 Updating 12-3,14,15,18
 Label checking 3-105
 Label sector 3-63
 Label track 17-16
 Large central memory extended (LCME) (refer also to extended memory) EST entry 3-47
 LDLIST (refer also to Queue file utilities) 18-22
 Level 0 initial deadstart 2-28
 Level 1 recovery deadstart 2-26
 Level 2 recovery deadstart 2-27
 Level 3 recovery deadstart 2-23
 Levels of recovery deadstart 2-22
 LIBDECK
 *AD entry 3-29
 Entries 3-112
 Selection 3-4
 LID configuration file 10-1
 LID/RHF configuration files 10-1
 LID table
 Building or rebuilding 10-6
 Listing 10-7
 Specifying size 10-6
 LIDOU L display (HELP) 9-5
 LIDOU L display (LID) 9-4
 Line 15-1
 Line control blocks 16-12
 Line printer 3-16
 Link device 3-33
 Link field 11-9
 Linkage byte 17-16,18
 Linkage error flag 11-7
 Linkage statements
 Node/path 11-63
 SM/CU 11-64
 Linked SDM operation
 CMRDECK entry 13-3
 Deadstart 13-3
 Description 13-2
 Device checkpoint 13-9
 Device configuration 13-3
 Device recovery 13-8
 EQPDECK entry 13-3
 MSAS processing 13-10
 MSS processing 13-9
 Using MSAS 13-2
 Using MSS 13-2
 List option 18-14
 Load buffer controllers entry 3-59
 Load file generator
 Description 16-1,6
 Input directives 16-7
 Network load file 16-6,8
 Parameters 16-6
 Summary listing 16-8
 Load partition 16-8
 Load (refer also to PFLOAD or QLOAD)
 Incremental 17-50
 MSS/MSAS environment 17-52
 Nonincremental 17-50
 LOADBC 1-1,2
 Loading microcode 3-87
 Loading the system 2-17
 Local configuration file 16-28
 Logical device removal entry 3-56
 Logical identifier (LID) 10-1; 18-4
 Table length 3-4
 Logical link 15-1
 Logical mapping 11-62
 Logically remove a device 3-56
 Looping chain 11-40; 12-29
 Lost cartridge 11-41; 12-7,16,30
 Lost cartridge flag 11-7
 Lower bound priority 3-104
 Lowest priority 3-104

Machine identification 3-4
 Machine recovery (MREC) utility
 K display 8-20
 Machine recovery table 13-2,10
 Magnetic tape
 Equipment EST entry 3-18
 Subsystem 3-95
 MAINLOG 18-44
 Maintenance logging transfer facility (MLTF) 14-1
 Maintenance mode 3-48
 MAP equipment EST entry 3-27
 MAP III or IV 3-95
 Mask
 Device 17-4
 Device examples 17-4,5
 Secondary 17-6
 Secondary examples 17-7
 Track 17-8

Mass storage
 Allocation of job files 3-50
 Archival subsystem 3-22
 Clear logical flaws 3-71
 Concepts 3-29; 17-3
 Configuration example 3-31
 Device functions 3-30
 Device recovery table 2-30
 Device types 3-43
 Equipment assignments 3-28
 Number of devices 3-42
 Set logical flaws 3-71
 Transport entry 3-21
 Validation 3-101
Mass storage adapter 12-1
Mass storage archival subsystem (MSAS) 3-95
 Description 11-1
 EST entry 3-22
 K display 8-27
Mass storage facility (MSF)
 Description 12-1
 EST entry 3-20
 Files on 12-9,20,27
 MSF catalog 12-5
Mass storage initialization status display 2-9
Mass storage status display 2-7
Mass storage subsystem (MSS) 3-95
 Description 12-1
 Files 12-1,10,27
 Initialization 12-45
 K display 8-28
 Utilities 12-2,10
Mass storage table
 Description 17-7
 Report, PFCAT 17-42
Mass storage transport 12-1,7
Master device 3-33,64
 Definition 17-3
 Number of catalog tracks 17-8
 User index relationship 17-3
Master file 15-3; 16-2
Master mainframe 11-10; 12-10; 13-2
Master mode
 For MSAS executive 3-100
 For MSS executive 3-100
Matrix algorithm processor (MAP) 3-27
Memory clearing 3-87
Memory control 3-77
Memory displays (DIS) 4-5
Memory entry commands 5-47
Memory entry commands (DIS) 4-16
Message control system 3-95
Messages, may appear under DIS 4-11
MINIT startup procedure 16-2,5
Mode of access 17-10
Modify
 APRDECK 3-69
 CMRDECK 3-1
 EQPDECK 3-8
 IPRDECK 3-73
Modifying
 APRDECK 2-15
 CMRDECK 2-2
 EQPDECK 2-3
 IPRDECK 2-16
MREC
 Commands 8-25
 Parameters 8-25
 Procedures 8-21
 Unit and controller reservations 8-26
 Utility 13-4,12
MRECOV startup procedure 16-2,5
MSAS commands 8-27
MSAS operational procedures 11-68
MSAS processing
 Backup and recovery 11-70
 Disk space management 11-10
 Error conditions and corrective actions 11-73
 Initialization 11-68
MSAS utilities 11-11
 List 11-1
MSF catalog
 Analysis 12-2,27,28,29
 Backup 12-9,18,49
 Chain of AUs 12-2
 Chain problems 12-53
 Creating 12-10,45,49
 Description 12-5
 FCT entry format 12-6
 Header format 12-5
 Ordinal 12-4
 Problems 12-52
 Subcatalog 12-5
 Subcatalog entry format 12-6
 Updating 12-9,14,15,18,27,42
MSF catalog/PFC mismatch 12-54
MSS commands 8-29
MSS processing
 Backup and recovery 12-47
 Disk space management 12-9,20,24,45
 Error conditions 12-3,50
 Initialization 12-2,45
 MSF space management 12-45,46
 Multiframe environment 12-11
 Pseudorelease 12-21,54
 Utilities 12-2,11
MSSEXEC 12-10,19,20,27,46
 In slave mode 3-20
MSSSLV 12-10
MULTI startup procedure 16-3,5

Multimainframe operation
 Deadstart 13-5
 Description 13-1
 Shared mass storage 13-1
 Multiplex line control blocks 16-13
 Multipliers for central processor 3-82
 Multispindle device 3-34

NAD maintenance utilities 14-1
 NAM
 Initialize 15-1
 K display 8-31; 15-3
 Shutdown 15-5
 Startup 15-3
 Status display 8-39
 NAM initialization
 Description 16-1
 Parameters 16-2
 Procedure calls 16-2
 NAM K display
 Control characters 8-31
 Fields 8-30
 Operation 8-32
 NAM mode
 Commands 8-32
 Maintenance commands 8-34
 DB command 8-35
 DE command 8-35
 DU command 8-35
 FL command 8-36
 LB command 8-36
 LE command 8-37
 LR command 8-37
 RS command 8-38
 ST command 8-39
 NAM procedure file 16-2
 NAMNOGO procedure file 16-2
 NCTF1 16-28
 NE equipment EST entry 3-15
 NETLOG 14-1
 NETOPS 16-28
 Network
 Command summary 15-42
 Commands
 Refer to tab for section 15; 15-1
 Invocation number (NIN) 15-3
 Organization 15-1
 Processing units (NPU's) 15-1
 Sample 15-2
 Supervisor (NS) 15-1
 Validation facility (NVF) 15-1
 Network access device (NAD)
 Controlware 1-2
 EST entry 3-26
 Network access method (NAM) 3-9
 K display 8-29
 Network configuration file 14-1; 16-28
 Network configuration statements
 Examples 10-15
 List of 10-8
 Order specified 10-9
 Network control
 DOP 15-30
 HOP 15-6
 NOP 15-16
 Network definition language processor 16-28
 Network description files 16-1,28
 Network dump collector
 Description 16-1,4
 Initialization 16-4
 Parameters 16-4
 Network element
 Channel 15-1
 Coupler 15-1
 Line 15-1
 Logical link 15-1
 NPU 15-1
 Terminal 15-1
 Trunk 15-1
 Network invocation number 16-2,4
 Network load file 16-6,8
 Network processing unit EST entry 3-24
 Network startup process 16-2
 Network supervisor 16-6
 Network utilities (refer also to individual utilities) 16-1
 NIN (refer also to Network invocation number) 15-3
 Nine-track tape units EST entry 3-18
 NLID statement 10-3
 Node/path linkage statements 11-63
 Nonincremental load 17-50
 Nonmass storage equipment assignments 3-14
 Nonremovable mass storage device 3-40
 Nonstandard equipment EST entry 3-14
 NOP
 Commands 15-20
 Network control 15-16
 NPU operator 15-3
 Status message formats 15-18
 Unsolicited status reports 15-20
 NPID statement 10-2
 NPU
 Diagnostic test commands 15-35
 Go command 15-26
 Load command 15-27
 Load/dump 15-8
 Load status command 15-8
 Memory dump command 15-28
 Memory dump indicator 15-9
 Operator (NOP) 15-3
 Saturation 15-20

NPU dump analyzer
 Description 16-1,9
 Dump interpretation 16-25
 Input directives 16-10
 A - LCB and TCB/CCB 16-12
 B - Port tables and MUXLCB 16-13
 0 - Comments 16-10
 1 - File 1 register 16-10
 2 - Macromemory word 16-10
 3 - Contiguous block data structure 16-10
 4 - Circular buffer data structure 16-11
 7 - Page registers 16-11
 8 - File 1 registers 16-11
 9 - Macromemory words 16-12
 NPU dump files 16-25
 Output file content 16-24
 Parameters 16-9
 Sample NPU dump file 16-14

NS
 Control commands 15-7
 Network supervisor 15-1
 Recent history command 15-11

NVF
 Control commands 15-11
 Network validation facility 15-1
 Recent history command 15-13

Obsolete file 12-9
 Off-cartridge link 12-8,36,37,39
 Off-cartridge link flag 11-9
 Off status for a device 3-57
 Operation under DIS 4-1
 Ordinal
 FCT 12-4,7
 MSF catalog 12-4
 Origin type 3-96
 List of 5-8
 Orphan file(s) 11-5,41,43; 12-10,27,30,31
 Output queue special handling level 3-101
 O26 keypunch mode 3-86
 O29 keypunch mode 3-86

Pack reformatting, 881/883
 Control directives F-4
 Data statements F-4
 Examples F-5
 FORMAT command F-2
 Input format F-3
 Output format F-5
 Special access procedures F-9
 Paper size for line printer 3-17
 Parameter record 15-3; 16-2
 RESTRT 15-3

Parameter record name
 Feature description 15-5
 INIT 15-5
 MINIT 15-5
 MRECOV 15-5
 MULTI 15-5
 RECOVR 15-5
 RESTRT 15-5
 Partial dump 17-46
 Peripheral equipment control commands 5-18
 Peripheral interface package 16-25
 Permanent file catalog
 Analysis 11-41; 12-29
 Entry 17-9
 Format 17-9
 Problems 11-76; 12-52
 Permanent file device 17-19
 Permanent file family, default 3-31
 Permanent file recovery 11-74; 12-51
 Permanent file supervisor 17-1
 Permanent file utilities
 Calling
 Batch 17-28
 Console 17-23
 DIS 17-28
 Procedure files 17-29
 Terminal 17-29
 Example 17-52
 List of 17-1
 Overview 17-1
 Parameters
 Descriptions 17-31
 Table 17-30
 Routines (refer also to individual routine name) 17-38
 Permanent files
 Maximum number 3-108
 Size 3-109,110
 Validation 3-102
 Permanent files device assignment
 entry 3-51
 Permission mode 17-38
 Permit information 3-33
 PFATC (refer also to Permanent file utilities) 17-41
 PFCAT (refer also to Permanent file utilities) 17-41
 PFCOPY (refer also to Permanent file utilities) 17-45
 PFDUMP (refer also to Permanent file utilities)
 Description 17-45
 Example 17-52
 PFLoad (refer also to Permanent file utilities) 17-47
 Example 17-52
 Selection process 17-49
 Physical identifier (PIP) 10-1

Request/return codes 2-6; 3-2
 Requests, system 1-2
 RERUN macro F-6
 Residence, file 5-14
 Residence, preferred 5-13
 Retention cycle 2-11
 RETURN
 Function 3-20
 Macro 3-45
 Versus EVICT macros 3-51
 Versus UNLOAD macros 3-45
 Return code 2-6
 RETURN macro F-3
 Returning a tape file 3-45
 REWIND function 3-20
 REWIND macro 3-43; F-3
 Rewinding a file 3-43
 REWRITE macro 3-38; F-3
 REWRITEF macro 3-39; F-3
 REWRITER macro 3-39; F-3
 RFILEB macro 2-12; F-3
 RFILEC macro 2-15; F-3
 ROFT file type E-5,8
 ROJS symbol E-6
 Rollout
 Defined C-12
 File C-12
 File type E-5
 Rollout file type E-8
 ROLLOUT macro F-6
 Rotating mass storage C-12
 ROUTE macro F-3
 RPHR macro 3-25; F-3
 RPHRLS macro 3-30; F-3
 RPVBLK macro F-7
 RTIME macro F-4
 Running field length 6-14

 S format tape J-3
 Reading 3-31
 Writing 3-41
 SAVE macro F-4
 Sector C-12
 Secure system memory 1-17
 Secured system 1-17; C-12
 CIO processing 3-12
 Permanent file access 5-3
 Security
 Access categories 1-17,20; 5-2
 Access level 4-3
 Access levels 1-17
 Conflict 1-21
 Considerations 1-17
 Count C-12
 Security administrator C-12
 Security conflict 3-33

 Security unlock status C-12
 Semiprivate files 5-12
 Sense switches C-12
 Sequential
 Access C-12
 File C-12
 Service class E-9
 Set identification 2-12; 4-16
 Set, magnetic tape 3-46
 SETASL macro F-7
 SETFAL macro 4-3; F-6
 SETFET macro F-7
 SETFET parameter processing 2-19
 SETFS macro 4-3; F-4
 SETGLS macro F-6
 SETJAL macro F-6
 SETJCI macro F-7
 SETJCR macro F-6
 SETJOB macro F-6
 SETJSL macro F-7
 SETLC macro 6-13; F-6
 SETLOF macro F-8
 SETMFL macro F-8
 SETPAGE macro F-6
 SETPFAC macro F-4
 SETPFAL macro F-4
 SETPR macro F-6
 SETRFL macro 6-14; F-8
 SETSC macro F-6
 SETSIM macro F-7
 SETSS macro 6-16; F-7
 SETSSM macro F-7
 SETTTL macro F-8
 SHELL macro F-8
 Short PRU C-12
 SI format tape J-4
 SIJS symbol E-6
 SKIPB macro 3-53; F-4
 SKIPEI macro 3-53; F-4
 SKIPF macro 3-52; F-4
 SKIPFB macro 3-54; F-4
 SKIPFF macro 3-52; F-4
 SOJS symbol E-6
 Special
 Entry point C-12
 Error flags E-12
 File treatment 12-3
 Handling 12-5
 Specifying subsystem 6-16
 SSSC service class E-9
 SST pseudoinstruction 1-6
 Stack purging E-4
 Standard labeled tape C-12
 Status bit of FET 1-3
 STATUS macro F-4
 STATUS macro (012) 4-4
 STATUS macro (013) 4-5
 STIME macro F-4

Stranger format tape J-3
 SUBR macro F-7
 Subsystem
 Ordinal 6-16; 11-30
 Service class E-9
 SUJS symbol E-7
 SYFT file type E-5,8
 Symbol C-12
 SYOT origin type E-8
 SYS= entry point F-11
 SYSC service class E-9
 SYSCOM macro 1-8; F-4
 System
 File C-13
 File name table C-13
 File type E-5
 Internal format tape J-4
 Labeled tape C-13
 Library C-13
 Origin type E-8
 Request 1-2; C-13
 Resource unit C-13
 Service class E-9
 Text C-13
 System access category set C-13
 System access levels C-13
 System file type E-8
 System library E-2
 SYSTEM macro F-4
 Systems texts 1-6
 SYSTEXT 1-6

Tape

Error messages 4-14
 Format C-13
 Mark C-13

Temporary file C-13

Terminal

Class 11-31
 Classes 12-10,11
 Control keys 12-12,13
 Display mode 11-29
 Input 12-4,5
 Model 11-29
 Prompt 12-8
 Redefinition 12-10
 Status bits 11-31
 Status 11-30

Text

Format G-4
 Record format G-9

Time

Slice C-13
 TIME macro F-4
 Timed/event rollout C-13
 TOJS symbol E-7
 Transmission code 11-31

TSSC service class E-9
 TSTATUS macro 11-30; F-7
 TXOT origin type E-8

UCW parameter 2-18
 UHL label 3-18
 ULIB format G-7
 Underline symbol A-3
 Unified extended memory C-13
 Unlabeled tape C-14
 UNLOAD function 3-23
 UNLOAD macro 3-44; F-4
 UNLOCK macro 4-4; F-4
 Unprintable characters A-3
 Unsecured system 1-17; C-14
 Unused bit count 2-10; 3-41
 USECPU macro F-4

User

Header label 3-18
 Index C-14
 Library group format G-7
 Processing bit 2-6
 Trailer label 3-22
 Volume label 3-18

User library E-2
 USERNUM macro F-4
 USN parameter 2-18
 UVL label 3-18

Validation file C-14

VERSION macro F-8

Volume

Defined C-14
 Serial number C-14

VOL1 label 3-18

WAIT macro F-4

WEC= entry point F-10

WNB= entry point F-11

Word C-14

WPHR macro 3-34; F-4

Write

Lockout bit 4-4

Ring C-14

WRITE macro 3-34; F-4

WRITEC macro 3-59; F-4

WRITECW macro 3-36; F-4

WRITEF macro 3-35; F-4

WRITEH macro 3-59; F-4

WRITEN macro 3-41; F-4

WRITEO macro 3-60; F-4

WRITER macro 3-35; F-5

WRITES macro 3-41,61; F-5

Shared mass storage
 Initialization 13-7,14
 Multimainframe complex 13-6,12
 Recovery tables 13-7,13
 Unloading 13-8,14
 Site debugging modifications 3-14
 Slave mainframe 11-10; 12-10; 13-2
 Slave mode
 For MSAS executive 3-100
 For MSS executive 3-100
 SM/CU linkage statements 11-64
 SM identifier 11-2
 SM map
 Analysis 11-40
 Backup 11-72
 Creating 11-12
 Description 11-2
 Problems 11-76
 Spindown of 834 disk units 3-110
 SSALTER
 Console input 11-65
 CU K display 11-67
 Description 11-2,65
 SM K display 11-66
 SSBLD
 Description 11-2,62
 Example 11-64
 Processing sequence 11-62
 Statements 11-62
 SSDEBUG
 Description 11-2,53
 Directives 11-54
 Parameters 11-55
 Restrictions 11-58
 SSDEF
 Description 11-1
 Utility 11-12
 SSEXEC 11-10
 In slave mode 3-22
 SSF
 Enabling and disabling file
 transfers 8-65
 File transfer limit commands 8-67
 Operator interface 8-65
 Spooled file transfer commands 8-67
 Staged file transfer commands 8-67
 Station disconnection and logout 8-66
 Station login 8-65
 Station recovery 8-66
 SSLABEL
 Description 11-1,13
 Directive file 11-13
 Directives 11-14
 Example 11-22
 Parameters 11-17
 Restrictions 11-21
 Update sequence 11-21
 SSMOVE
 Description 11-1,22
 Directives 11-24
 Example 11-31
 Restrictions 11-31
 Selection algorithms 11-29
 Specific file directives 11-25
 Value specifier directives 11-25
 SSSLV 11-10
 SSUSE
 Basic usage report 11-46,48
 Description 11-2,46
 Optional report A 11-50
 Optional report B 11-51
 Optional report C 11-52
 Optional report D 11-53
 Optional reports 11-46
 SSVL
 Description 11-2,37
 Error detection and
 classification 11-40
 Examples 11-44
 Problem fixing 11-42
 Problem reporting 11-38
 Release processing 11-37,42
 Restrictions 11-46
 Validation report 11-42
 Stage 11-10; 12-10
 Stage to disk 3-99
 Start of fragment flag 11-9
 Status/control register simulation 3-89
 Status/control register simulator (refer
 also to SCRSIM) D-1
 Status host element command 15-14
 Status message formats
 DOP 15-31
 HOP 15-6
 NOP 15-18
 Status network element command 15-39
 STIMULA 3-95; 16-29
 Stimulator
 Description 16-1,29
 Multiple sessions 16-42
 Output evaluation 16-46
 Parameters
 Mixed mode 16-39,41
 Output format 16-42
 Session file 16-37
 Stimulator 16-38
 Postprocessing 16-42
 Procedure files 16-45
 Session file construction 16-30
 Session files
 Description 16-30
 Dynamic login 16-34
 Optional think time specification
 16-35
 Processing 16-37
 Task definitions 16-36

Stimulator equipment EST entry 3-23
 STMffff 16-37
 Stream
 Chain control 12-8
 Conflict flag 12-8,10
 Description 12-2
 Detail information 12-9,35
 Free/busy flag 12-8
 Label verification procedure 12-3
 Selected flag 12-8
 Subcatalog 11-4; 12-5
 Subcontrol point processing 3-91
 Subfamily 11-4,11; 12-5,10; 17-3
 SUBSYST L display 9-10
 Subsystem code 17-11
 Subsystem control commands 5-33
 Suspended job 3-110
 SYSEEDIT
 Command 19-1
 Description 19-1
 Directives 3-112
 Input directives 19-2
 Record types 3-113; 19-8
 System access level limits 3-106
 System checkpoint file device assignment
 entry 3-55
 System console lock 3-99
 System console unlock 3-111
 System control commands 5-37
 System control point facility 3-88
 System date line 3-5
 System dayfile 3-11
 System deadstart file 7-1
 System delay parameters 3-97
 System device 3-40
 System display console equipment EST
 entry 3-15
 System file initialization 20-1
 System library 3-54
 System library device assignment entry 3-54
 System tape density 3-92
 System version name 3-7

 TAF
 Commands 8-72
 CRM data base status K display 8-78
 CRM file status K display 8-78
 CRM status K display 8-77
 CRMTASK commands 8-76
 CRMTASK K display 8-76
 Initiating TAF K display 8-68
 K display 8-71
 Restarting TAF K display 8-71
 Tape
 Conversion mode 3-92
 Density 3-92
 Track type 3-93
 Unit EST entry 3-18

 Task name 16-36
 Task usage parameter 16-36
 TE equipment EST entry 3-15
 Temporary file device 3-41
 Terminal control blocks 16-12
 Terminal definition directive 16-28
 Terminal I/O EST entry 3-15
 Terminal verification facility (TVF) 15-1
 Test mode 12-21
 Threshold count 3-51
 Time-out delay 3-110
 Time slice 3-74,75,77,108
 Total initialization 3-66
 TRACER
 Commands 21-4
 Conditions monitored 21-1
 Data items reported 21-18
 Description 21-1
 Example 21-8
 Fast loop items 21-20
 Medium loop items 21-25
 Output file format 21-7
 Programs 21-3
 Slow loop items 21-29
 Statistical summary 21-30
 Summary file format 21-18
 TRACIAF 16-27
 Track flawing options table 8-18
 Track mask 17-8
 Track, mass storage
 First-track bits 17-16,17
 Flawed 17-17; F-1
 Interlock bits, TRT 17-16,17
 Reservation bits, TRT 17-16,17
 Track reservation table (TRT) 3-33,63,69;
 8-12; 17-16; F-1
 Track reservations 3-69
 Track type 3-93
 Transaction facility (TAF) 3-95
 Transaction facility (TAF) K display 8-68
 TRT (refer to Track reservation table) 8-12
 Trunk control units 1-3
 TST
 Drop option 15-35
 Message option 15-35
 Request option 15-36
 TT equipment EST entry 3-15
 TVF (refer to Terminal verification
 facility) 15-1
 Two-port multiplexer EST entry 3-26

 Unit device table 11-2,62
 Unit record equipment EST entry 3-15
 Unsolicited status reports
 DOP 15-33
 HOP 15-7
 NOP 15-20
 Unusable disk areas 3-69

USER command, multiple entries 3-105
 User control word 17-15
 User extended memory 3-111
 User index
 Permanent file device assignment 17-3
 Using the simulator D-1
 Utilities
 Controlware loading 1-1
 Dayfile dumping 18-1,43
 MSAS 11-11
 MSS 12-2,11
 Network 16-1
 Permanent file 17-1,38
 Queue file 18-1,20
 System data gathering 21-1
 Utility map, disk pack F-1

 Validation
 Mass storage 3-101
 Permanent file 3-102
 Validation report 11-44; 12-31,32
 Value specifier directive format 11-26
 Variable-length data structures 16-11
 Verify file 17-36
 Volume 11-11
 Volume serial number and file name
 association 3-15

 Weighting factor for queue 3-74,104

 x,y coordinates 12-3
 X.25 packet switching network (PSN) 15-1

 y,z relationship with SM map ordinal 11-58

 6250 cpi 3-19
 63-character set 3-84
 64-character set 3-84
 6683 satellite coupler EST entry 3-27
 7021 controller block identification 3-20
 7990 block header and trailer 11-61
 7990 catalog
 Analysis 11-40
 Backup 11-72
 Creating 11-12
 Description 11-4
 FCT entry format 11-6
 Header format 11-5
 Problems 11-76
 Subcatalog 11-4
 Subcatalog entry format 11-5
 7990 catalog chain problems 11-77
 7990 catalog/PFC mismatch 11-77
 7990 catalog/SM map mismatch 11-76
 7990 control unit 11-1
 7990 controller 3-22
 7990 equipment configuration
 constraints 11-64
 7990 files 11-10
 7990 hardware components 11-1
 7990 space management 11-69
 7990 stripe header and trailer 11-61
 7991 storage module 3-22; 11-1
 819 disks
 Buffer in LCME 3-41
 844 disk drives 3-34
 844 expander 3-34
 881/883 pack formatting 3-98
 885 disk drives 3-34
 885-42 disk drive 3-41

COMMENT SHEET

MANUAL TITLE: CDC NOS Version 2 Analysis Handbook

PUBLICATION NO.: 60459300

REVISION: D

NAME: _____

COMPANY: _____

STREET ADDRESS: _____

CITY: _____ STATE: _____ ZIP CODE: _____

This form is not intended to be used as an order blank. Control Data Corporation welcomes your evaluation of this manual. Please indicate any errors, suggested additions or deletions, or general comments below (please include page number references).

Please Reply


No Reply Necessary

CUT ALONG LINE

NO POSTAGE STAMP NECESSARY IF MAILED IN U.S.A.

FOLD

FOLD



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 8241 MINNEAPOLIS, MINN.

POSTAGE WILL BE PAID BY
CONTROL DATA CORPORATION

Publications and Graphics Division
ARH219
4201 North Lexington Avenue
Saint Paul, Minnesota 55112



CUT ALONG LINE

FOLD

FOLD

