UNISYS

Product Information Announcement

o New Release o Revision • Update o New Mail Code

Title

System 80 Models 3-6 and 8-20 OS/3 Operations Guide

This Product Information Announcement announces the release and availability of Update C to the *System 80 Models 3-6 and 8-20 OS/3 Operations Guide.*

This guide is a standard library item (SLI). It is part of the standard library provided automatically with the purchase of the product.

This update defines the following additional changes to this document for Release 13.0:

- Enhancement of DELETE command to allow deletion of some (or all) queued symbionts.
- Enhancement of MIX SQ command (display symbiont queue information) to provide a queue slot ID number for each symbiont in the queue. You can use these ID numbers to delete specific symbionts from the queue.
- Modification of system messages provided by the system utility services tape prep routine (INT).

All other changes in this document are corrections, deletions, or expanded descriptions applicable to items present in the hardware or software prior to this release.

You can order the update only, or the complete manual with all updates. To receive only the update, order UP-8859 Rev. 5-C. To receive the complete manual, order UP-8859 Rev. 5.

Additional copies of this document can be ordered through your branch representative or from Unisys Corporation, Corporate Software and Publication Operations, 13250 Haggerty Road North, Plymouth, Michigan 48170.

Announcement only:
MB00, SAB, and SAE

Announcement and attachments:
MB01, MBA9, and MBW

System: System 80
Release: 13
Date: May 1990
Part Number: UP-8859 Rev. 5-C

UNISYS

PUBLICATIONS UPDATE

System 80
Models 3-6 and 8-20
OS/3
Operations
Guide

UP-8859 Rev. 5-B

This Library Memo announces the release and availability of Update B to the *System 80 Models 3-6 and 8-20 OS/3 Operations Guide*, UP-8859 Rev. 5.

This guide is a standard library item (SLI). It is part of the standard library provided automatically with the purchase of the product.

Changes to this document for Release 13.0 include:

- Addition of new console commands to soft cancel a job (CJ) and control the console time-of-day clock (SET VC).
- Support of the M9720 disk subsystem for System 80 models 10, 15, and 20.
- Enhanced ICAM unsolicited down message that allows the console operator (or interactive services user with console privileges) to down an ICAM terminal that is attached to a DCP.
- Enhanced PAUSE and STOP commands that allow specific job steps to be specified.
- Enhanced SET AL command to allow warning beeper to sound only once.
- Additional disk or diskette VTOC information provided via the VTP command.
- Enhanced interactive services ST command to allow terminal information to be returned for a specified user-id.
- Changed default setting of the ONUERL program ELOG option.
- Revised ONUERL report format.

All other changes in this document are corrections, deletions, or expanded descriptions applicable to items present in the hardware or software prior to this release.

You can order the update only, or the complete manual with all updates. To receive only the update, order UP-8859 Rev. 5-B. To receive the complete manual, order UP-8859 Rev. 5.

LIBRARY MEMO ONLY	LIBRARY MEMO AND ATTACHMENTS	THIS SHEET IS
Mailing Lists MB00, SAB, and SAE	Mailing Lists MB01, MBA9, and MBW (97 pages plus Memo)	Library Memo for UP-8859 Rev. 5-B
		RELEASE DATE: January 1990

UNISYS

PUBLICATIONS UPDATE

System 80

OS/3 Operations Guide

UP-8859 Rev. 5-A

This Library Memo announces the release and availability of Update A to the *System 80 OS/3 Operations Guide*, UP-8859 Rev. 5.

This guide is a standard library item (SLI). It is part of the standard library provided automatically with the purchase of the product.

This update provides an index and a set of tab breakers. The content of the manual is unchanged.

Copies of Update A are now available. You can order the update only, or the complete manual with the update, through your local Unisys representative. To receive only the update, order UP-8859 Rev. 5-A. To receive the complete manual, order UP-8859 Rev. 5.

LIBRARY MEMO ONLY	LIBRARY MEMO AND ATTACHMENTS	THIS SHEET IS
Mailing Lists MBZ, MCZ, MMZ, M28U, and M29U	Mailing Lists MB00, MB01, and MBW (30 pages plus Memo)	Library Memo for UP-8859 Rev. 5-A RELEASE DATE: February 1989

UNISYS

PUBLICATIONS REVISION

System 80

OS/3
Operations
Guide

UP-8859 Rev. 5

This Library Memo announces the release and availability of the *System 80 OS/3 Operations Guide*, UP-8859 Rev. 5.

This guide is a standard library item (SLI). It is part of the standard library provided automatically with the purchase of the product.

Changes to this document for Release 12.0 include:

- Disk cache enhanced to support up to 8-megabyte cache buffer and 24K byte segment size
- Addition of new ICAM commands for monitoring status of line(s), terminal, and LOCAP
- VTOC verification routine enhanced to support a PACK facility, which can be used to reorganize your VTOC for more efficient searching
- Interactive services RECALL command enhanced for use from system console and to allow specification of message prefix
- For secure environments, enhancements to the LIMITS command to UP terminals deactivated after a specified number of invalid logon attempts and to display invalid logon retry status
- Addition of new console commands to discard system messages (FLUSH), display system buffer pool information (DI BI), and control the system alarm (SET AL)
- Support of model 15 processor, BT3200 tape subsystem, and 9246-14B printer

All other changes in this document are corrections, deletions, or expanded descriptions applicable to items present in the software prior to this release.

Additional copies may be ordered through your Unisys representative.

Destruction Notice: This revision supersedes and replaces the *OS/3 Operations Guide*, UP-8859 Rev. 4, released on Library Memo dated August 1987. Please destroy all copies of UP-8859 Rev. 4 and its Library Memo.

LIBRARY MEMO ONLY	LIBRARY MEMO AND ATTACHMENTS	THIS SHEET IS
Mailing Lists MBZ, MCZ, MMZ, M28U, and M29U	Mailing Lists MB00, MB01, and MBW (288 pages plus Memo)	Library Memo for UP-8859 Rev. 5 RELEASE DATE: October 1988

UNISYS

System 80
Models 3-6 and 8-20
OS/3
Operations Guide

Previous Title: System 80 OS/3 Operations Guide

October 1988

Printed in U S America
UP-8859 Rev. 5

Priced Item

UNISYS

System 80
Models 3-6 and 8-20
OS/3
Operations Guide

OS/3 Release 13

May 1990

Printed in U S America

Priced Item

UP-8859 Rev. 5 - Update C

Correspondence regarding this publication should be forwarded to Unisys Corporation either by using the Business Reply Mail form at the back of this manual or by addressing remarks directly to Unisys Corporation, OS/3 Systems Product Information Development, P.O. Box 500, Mail Station E5-114, Blue Bell, Pennsylvania, 19424, U.S.A.

PAGE STATUS SUMMARY
ISSUE: Update C - UP-8859 Rev. 5

Part/Section	Page Number	Update Level
Cover		B
Title Page/Disclaimer		C
PSS	iii	C
About This Guide	Tab Breaker	A
	v thru ix	Orig.
Contents	xi, xii	B
	xiii	C
	xiv	B
	xv thru xviii	Orig.
	xix thru xxi	B
1	Tab Breaker	A
	1 thru 8	Orig.
2	Tab Breaker	A
	1 thru 3	Orig.
3	Tab Breaker	A
	1 thru 16	Orig.
4	Tab Breaker	A
	1	Orig.
	2 thru 4	B
	5	Orig.
	6 thru 13	B
	14, 15	B
5	Tab Breaker	A
	1 thru 3	Orig.
6	Tab Breaker	A
	1 thru 19	Orig.
7	Tab Breaker	A
	1 thru 5	Orig.
8	Tab Breaker	A
	1 thru 10	Orig.
	10a	B
	11 thru 29	Orig.
	30	C
	30a	C
	31 thru 36	Orig.
	37 thru 46	B
	46a, 46b	B
	47	C
	48	B
	48a	B
	49	B
	50 thru 58	Orig.
	59	C
	60	Orig.
	61	B
	62 thru 70	Orig.
	71	C
	72 thru 77	Orig.
	78	B
	79, 80	Orig.
	81	B
	82	B
9	Tab Breaker	A
	1, 2	Orig.
	3	B
	4 thru 6	Orig.
	7	B
	8 thru 11	Orig.
10	Tab Breaker	A
	1 thru 6	Orig.
	7, 8	B
	9 thru 13	Orig.
11	Tab Breaker	A
	1 thru 10	Orig.
	11	B
	12 thru 19	Orig.
	20	C
	21 thru 29	Orig.
	30	B
	31 thru 40	Orig.
12	Tab Breaker	A
	1, 2	Orig.
	3	C
	4	B
	5 thru 9	Orig.
Appendix A	Tab Breaker	A
	1	B
	2, 3	Orig.
	4	B
	5 thru 12	Orig.
	13	B
	14	C
	14a	B
	15 thru 25	B
	26 thru 37	Orig.
Appendix B	Tab Breaker	A
	1 thru 4	Orig.
Appendix C	Tab Breaker	A
	1	B
	2, 3	Orig.
	4	C
	5	Orig.
Index	Tab Breaker	A
	1 thru 11	B
User Comments Form		
Back Cover		Orig.

* New page Technical changes are denoted by a change bar in the margin.

About This Guide

Purpose

This guide is designed to instruct the operator in the procedures required to operate the System 80 under the control of Operating System/3 (OS/3).

Audience

The intended audience for this guide is the operator with a basic knowledge of data processing operations but without experience on Unisys systems.

Organization

Your System 80 computer can be one of eight models: model 3, 4, 5, 6, 8, 10, 15, or 20. System 80 models 3, 4, 5, and 6 are identical in capability and in available features, with the exception that models 5 and 6 are configured with high-performance control storage. Model 8, on the other hand, is a completely different hardware system intended for larger interactive data processing system users. System 80 models 10, 15, and 20 are similar to model 8 with respect to hardware configurations, initialization procedures, and operator commands and messages.

Because the operating procedures for models 3, 4, 5, and 6 differ from those of models 8, 10, 15, and 20, this operations guide is organized as follows:

PART 1. SYSTEM 80 MODELS 3-6

Section 1. Models 3-6 System Definition

Briefly describes the minimum and maximum hardware configurations for models 3-6.

Section 2. Models 3-6 System Turn-On and Turn-Off Procedures

Provides the procedures for turning the model 3-6 system on and off.

Section 3. Models 3-6 System Initialization Procedures

Provides system procedures for initial microprogram load and initial program load for models 3-6.

PART 2. SYSTEM 80 MODELS 8-20

Section 4. Models 8-20 System Definition

Briefly describes minimum and maximum hardware configurations for models 8-20.

Section 5. Models 8-20 System Turn-On and Turn-Off Procedures

Provides the procedures for turning models 8-20 on and off.

Section 6. Models 8-20 System Initialization Procedures

Provides system procedures for initial microprogram load and initial program load for models 8-20.

PART 3. USING OS/3 ON ALL SYSTEM 80 MODELS

Section 7. OS/3 - The Operating System

Briefly describes the OS/3 software available for all the System 80 models.

Section 8. Job Processing Procedures

Provides job processing procedures performed by the system operator.

Section 9. Interactive Services

Describes the interactive services operator commands and messages used to control the interactive system environment.

Section 10. Integrated Communications Access Method (ICAM) Procedures

Describes the ICAM communications tasks, including how to load the ICAM symbiont, change the ICAM name established during SYSGEN, run and terminate the global user service task, and use the message instructions to facilitate communications.

Section 11. System Utility Services

Describes the system utility symbiont (SL$$SU) provided by OS/3 to perform card, tape, disk, and diskette functions.

Section 12. Disk Cache Facility (DCF)

Describes various aspects of using the disk cache facility.

PART 4. APPENDIXES

Appendix A. ONUERL Program

Describes the ONUERL error log editor program.

Appendix B. Supervisor Modification Procedure

Describes the output messages and operator responses required to modify the supervisor during the IPL procedure.

Appendix C. Disk Cache Facility Supportive Information

Provides background and illustrative information to help the operator use DCF optimally.

Notation Conventions

The conventions used to illustrate the messages and commands presented in this guide are:

- Commands, parameters, and input messages in uppercase letters must be keyed in literally. For example:

  ```
  SYSDUMP
  ```

  Output messages are shown in uppercase letters. For example:

  ```
  EARLY WARNING OVERTEMP. CONDITION EXISTS
  ```

- The delta symbol Δ indicates a space:

  ```
  SHUTDOWNΔDDP
  ```

- Lowercase letters represent variable information that is either displayed or keyed in. For example, the following command format implies that the command DELETE must be followed by the name of the job to be deleted.

  ```
  DELETEΔjobname
  ```

- Underlined letters in a command indicate that they are the only letters required to be keyed in to initiate processing of the command and its associated symbiont. For example, only the letters DE need be keyed in to initiate processing of the DELETE command and subsequent running of the delete symbiont. Its format is presented as:

  <u>DE</u>LETEΔjobname

- Braces { } illustrate alternate choices. For example, the format of the change command

```
CHANGEΔjobname, { PRE  }
                { HIGH }
                { NOR  }
                { LOW  }
```

indicates that PRE (P), HIGH (H), NOR (N) or Low (L) may be keyed in after the job name.

- Brackets [ ] denote optional entries.

For example, this portion of the FILE command format

```
FILE [ { (did)          } ]
     [ { ([did],label)  } ]
     [ { (RDR,label)    } ]
```

indicates that the FILE command can be keyed in by itself, or with a parameter as specified in the format.

- Default parameters are shaded. For example, the ENTER command format

```
ENTERΔ'filename',QUEUE=RDR [ ,HOLD= { N } ]
                           [        { Y } ]
```

specifies that the filename and QUEUE=RDR parameters must be included (parameters 1 and 2). If you omit parameter 3, HOLD=Y is used.

Note: *Not all optional parameters have a default specification.*

- An ellipsis (...) indicates the omission of a variable number of entries.

```
modifier-1,...,modifier-n
```

Related Product Information

Whereas the system operator uses the procedures and commands described in this guide to control the entire system, many of these commands are also available to a workstation user under certain limitations. For a description and explanation of workstation operation procedures and available functions, see the *Interactive Services Operating Guide* (UP-9972).

Another document relating to the operation of the System 80 under the control of OS/3 is the *System Messages Reference Manual* (UP-8076), which describes all the system messages you could encounter while operating the System 80 and the appropriate responses, when necessary.

Note: *Throughout this guide, when we refer you to another manual, use the current version that applies to the software level at your site.*

Contents

PART 2. SYSTEM 80 MODELS 8-20

Section 4. Models 8-20 System Definition

Section 5. Models 8-20 System Turn-On and Turn-Off Procedures

Section 6. Models 8-20 System Initialization Procedures

PART 3. USING OS/3 ON ALL SYSTEM 80 MODELS

Section 7. OS/3 Operating System

Section 8. Job Processing Procedures

Section 10. Integrated Communications Access Method (ICAM) Procedures

Section 11. System Utility Services

Section 12. Disk Cache Facility (DCF)

PART 4. APPENDIXES

Appendix A. ONUERL Program

Appendix B. Supervisor Modification Procedures

Appendix C. Disk Cache Facility Supportive Information

User Comments Form

Figures

Tables

Section 1
Models 3-6 System Definition

The System 80 is a general-purpose, disk-oriented computer designed to function in many different data processing environments with equal operating efficiency. This efficiency is achieved through the use of Operating System/3 (OS/3), a multiprogramming software system specifically designed to make maximum use of the capabilities of the system hardware. Different versions of the system are assembled from system configurations (see 1.1). Part 1 concerns models 3-6.

1.1. Models 3-6 System Configuration

Diagrams of the basic System 80 models 3-6 and expanded hardware options are presented in Figures 1-1 and 1-2, respectively.

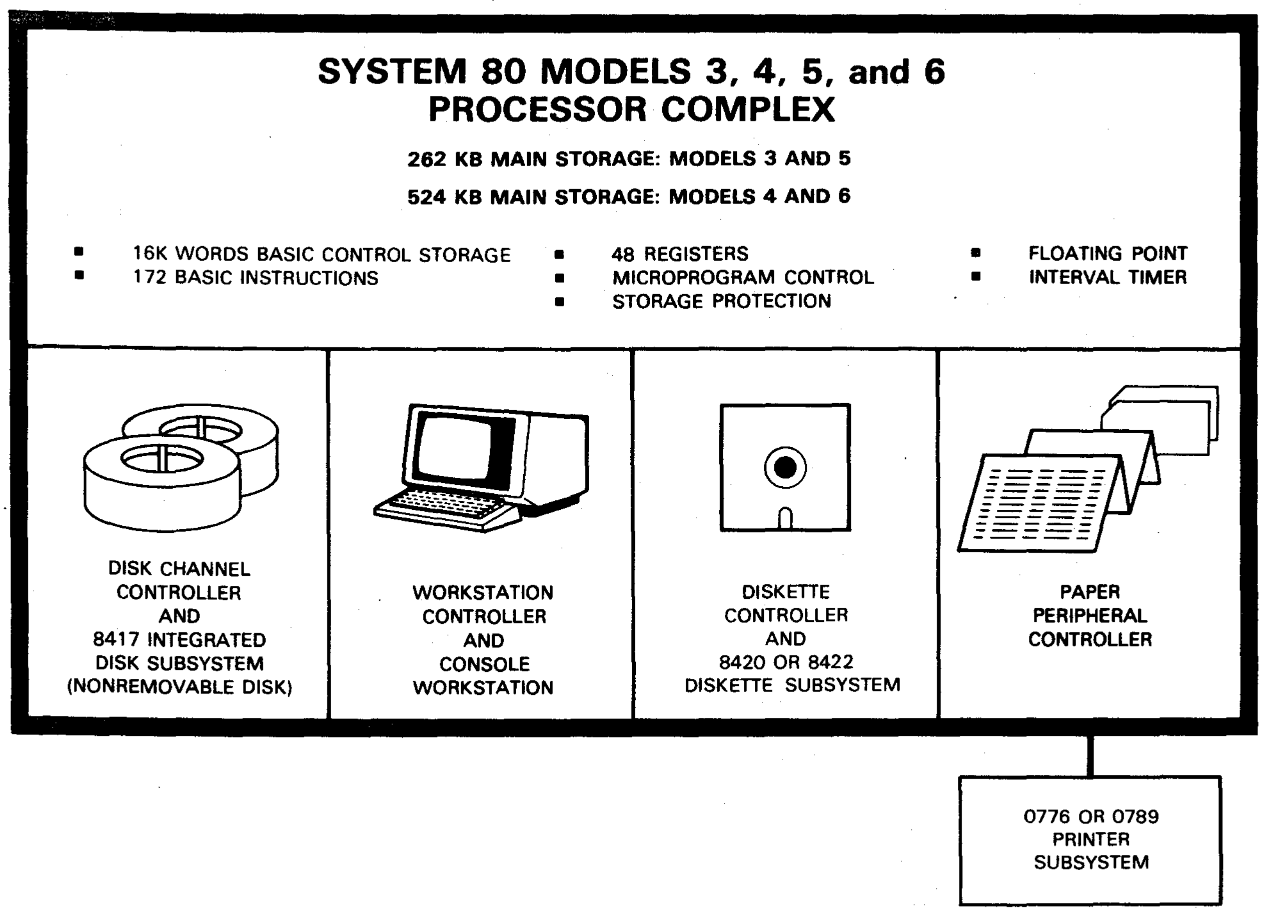

Figure 1-1. System 80 Models 3-6 Basic Configuration

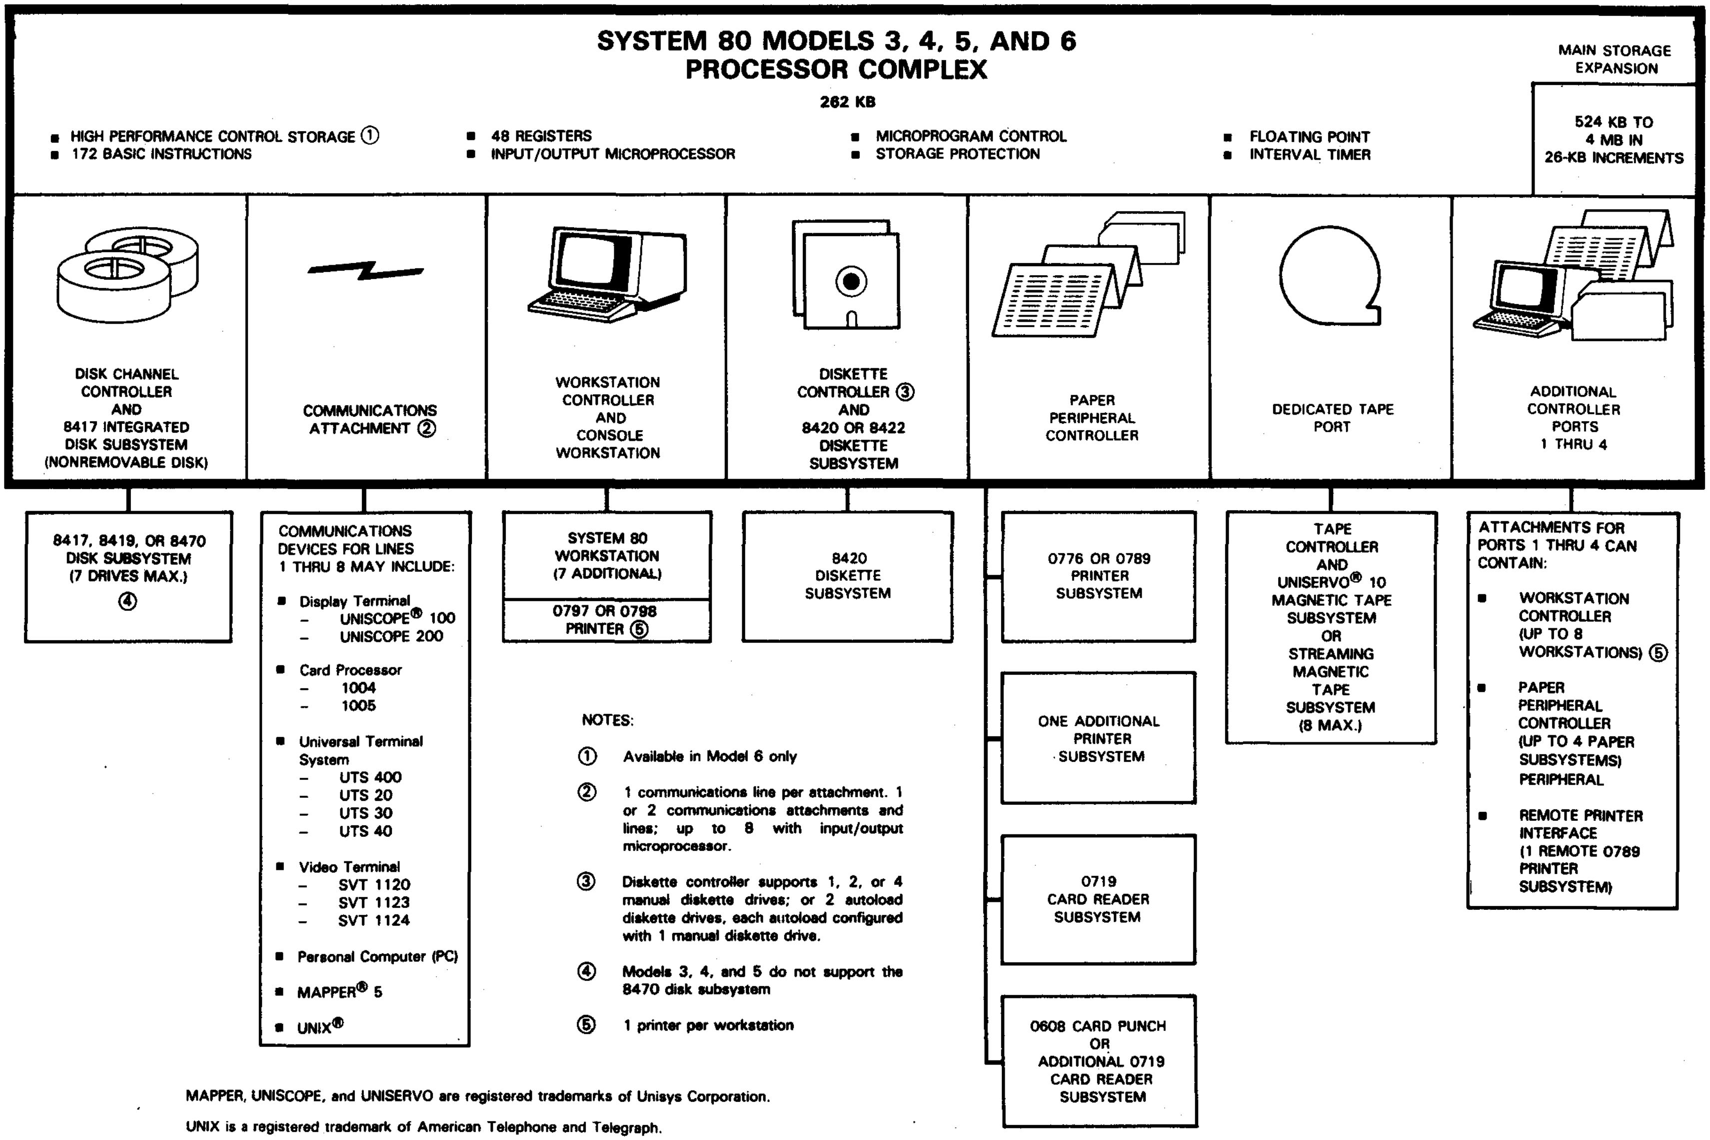

NOTES:

① Available in Model 6 only

② 1 communications line per attachment. 1 or 2 communications attachments and lines; up to 8 with input/output microprocessor.

③ Diskette controller supports 1, 2, or 4 manual diskette drives; or 2 autoload diskette drives, each autoload configured with 1 manual diskette drive.

④ Models 3, 4, and 5 do not support the 8470 disk subsystem

⑤ 1 printer per workstation

MAPPER, UNISCOPE, and UNISERVO are registered trademarks of Unisys Corporation.

UNIX is a registered trademark of American Telephone and Telegraph.

Figure 1-2. System 80 Models 3-6 Configuration with All Hardware Options Shown

1.2. Central Hardware

The central hardware of a basic System 80 models 3-6 is housed within a common cabinet called the processor complex that includes the console workstation, a disk and diskette subsystem, and the standard processor features and input/output controls (Figure 1-1). The central hardware for an expanded system includes an I/O expansion cabinet and additional integrated processor features and I/O controls (Figure 1-2). Enhanced system performance is obtained through high-performance control storage and the input/output microprocessor.

Within the processor complex, the central processor provides the data path, registers, and associated control logic for instruction executions, system control, and I/O channel support. The processor, along with the integrated I/O controllers and the communications attachment, provides microprogram control to the system. The functional components comprising the processor complex are summarized as follows:

- Basic control storage/high-performance control storage
- Expandable main storage
- 172 basic instructions, including instructions for floating-point and storage protection
- 48 registers for program and system use, plus floating-point and relocation registers
- Input/output microprocessor supports:
  - Paper peripheral controller
  - Workstation controller
  - Diskette controller
  - Remote printer interface
  - Tape controller (not integrated)
  - Communications attachment
- Disk channel controller
- Interval timer

A brief description of each of these components and associated integrated I/O devices follows.

1.2.1. Control Storage

Control storage consists of two modular units having microinstructions to control central processor operations. Each module includes resident microdiagnostics and capabilities for performing an initial microprogram load (IMPL) to main storage. The IMPL loads remaining control storage from an internal fixed-head disk. Control storage has a maximum capacity of 16K words. Each word is 32 bits long plus 4 parity bits.

The operating speed of the central processor is determined by the type of control storage (basic or high performance) configured in the system. High-performance control storage increases the central processor speed by 55 percent over that of basic control storage.

1.2.2. Main Storage

The main storage unit stores data in word format. Models 3 and 5 have 262K bytes of main storage and can be expanded to 4 MB in 262K byte increments. Models 4 and 6 have 524K bytes of main storage and can be expanded to 4 MB in 262K byte increments.

Main storage is of the semiconductor type with a 400-nanosecond, half-word read/write cycle time. Protection logic assures data and program protection in the system. Error correction, which also is provided for main storage data, includes double-digit detection and single-bit error correction on storage accesses.

Main storage is accessed directly by the processor (for programs), the disk channel controller (for disk), and the input/output microprocessor in expanded systems or the I/O control section in basic systems (for all I/O devices except disk).

1.2.3. System Instructions

The system instructions control the operation of the central processor. A repertoire of 172 basic instructions, 25 of which are privileged, is provided with the system. The privileged instructions are used for system control and input/output operations. Nonprivileged instructions are used for general control, data manipulation, decimal operations, list processing, and floating-point operation.

1.2.4. System Registers

The system registers are provided for accumulating, addressing, and controlling operations and data within the system. The 32 general registers may be used as accumulators in general arithmetic and logical operations. They are also used as base address and index registers in arithmetic and logical operations and in address generation. The 16 control registers are used to specify whether an operation can be performed in the system, or they may be used to provide special information that is required by a system facility. The four floating-point registers are available for floating-point operations when the floating-point instruction feature is included in the system. Main storage space is fully utilized by the system using dynamic storage relocation and the 16 relocation registers. Automatic relocation of all programs in main storage allows maximum loading of programs.

1.2.5. Input/Output Control Section

The input/output control section initiates, directs, and monitors the transfer of data between main storage and the integrated and peripheral I/O devices. I/O controls accept instructions from the processor (through user and system software), from the console workstation and workstations, and from operator controls on the devices. After an I/O instruction is initiated, the data is transferred independently of other processor functions; i.e., the I/O and the processor operate concurrently.

The I/O control section with integrated devices consists of the:

- Disk channel controller and 8417 disk subsystem
- Workstation controller and console workstation
- Diskette controller and 8420 and 8422 diskette subsystems
- Paper peripheral controller
- Remote printer interface
- Tape controller
- Communications attachment

In the expanded system, an input/output microprocessor provides the interface between main storage and the controls for all I/O devices except the disk. In the basic system, this interface forms part of the I/O control section.

Brief descriptions of these hardware components follow.

Disk Channel Controller and Integrated Disk Subsystem

Disk read, write, and search operations are accomplished through a disk channel controller that directly accesses main storage. The integrated 8417 disk subsystem and up to seven peripheral disk drives are controlled by the disk channel controller. Up to 1.2 MB/second of data flow can be transferred. (See Table 1-1 for a listing of the disk subsystem hardware manuals.)

Input/Output Microprocessor

In an expanded system, the controls for I/O devices other than disk drives communicate with the processor via a multiplexer-type input/output microprocessor (IOMP). The microprocessor supports devices having low or medium data transfer rates, including workstations, diskette drives, printers, card readers, a card punch, and a magnetic tape subsystem. Data throughput may be at the rate of up to 200K bytes/second.

Workstation Controller and Console Workstation

The integrated console workstation and up to seven peripheral workstations are attached to the system through the workstation controller. Four additional workstation controllers and 32 additional workstations may be added to expand system capability. These are added in increments of up to eight workstations per controller (to total 39 workstations maximum per system). An 0425, 0797, 0791, or 0798 printer can be directly connected to UTS 400 workstations to allow the user to print workstation screens. The controller contains dedicated buffers for each workstation to permit all workstations to transfer data concurrently.

The console workstation provides the facility for operator interaction with the processor. Special system functions, as well as all standard workstation functions, are included in the console workstation to provide total system control to the operator. The console workstation accepts data from the keyboard, displays the data, and transfers the data to the workstation controller. Data entered from the keyboard is displayed on the console workstation screen in an 80-character-per-line by 24-line format, providing a total display of 1,920 characters. An additional line indicates system status or operating mode. (See Table 1-1 for a listing of the console workstation and System 80 workstation hardware manuals.)

Diskette Controller and Integrated Diskette Subsystem

The integrated 8420 or 8422 diskette subsystem and peripheral diskette drives are attached to the system through the diskette controller. The controller supports one, two, or four manual diskette drives; or two autoload diskette drives, each autoload configured with one manual diskette drive. The autoload drive provides an automatic loading capability of up to 20 standard diskettes. (See Table 1-1 for a listing of the diskette subsystem hardware manuals.)

Paper Peripheral Controller

The paper peripheral controller provides several options for attaching peripheral printers and card processors to the system. One or two printers with a combined print capacity of up to 1,500 lines per minute, plus two card processors, may be attached. These card processors can be up to two card readers or one card reader and one card punch. Additional paper peripheral controllers, supporting duplicate printer and card processor options, may be added to the system.

Remote Printer Interface

The remote printer interface is included in the system by using one of the additional controller ports. The remote printer interface supports one remote printer subsystem in an expanded system. The remote printer may be located up to 5,000 feet from the processor complex. Used with the paper peripheral controller, the interface provides input/output control functions for the remote printer. Additional remote printer interfaces, each supporting one remote printer subsystem, may be included in the system via the additional controller ports.

Tape Controller

The tape controller supports up to eight magnetic tape subsystems in an expanded system. The tape controller is not integrated and is attached to the processor complex via a dedicated tape port.

Communications Attachment

The communications attachment connects communications devices, using a single communications line, to the system. Up to eight communications attachments and lines may be included in an expanded system. The communications attachment provides I/O control functions, such as interfacing data transfer to and from main storage and error checking, and includes special character recognition and other specialized services. Up to 9,600 bits per second can be transferred through lines.

1.2.6. Interval Timer

The interval timer is a register used by system software to perform various timing functions within the system and to keep track of the time of day.

1.3. Input/Output Subsystems

The I/O integrated and peripheral subsystems available for use with System 80 include the console workstation, disk and diskette subsystems, workstations, a magnetic tape subsystem, printer subsystems, card reader and punch subsystems, and remote communications subsystems.

1.4. Associated Hardware Manuals

The I/O subsystems available for models 3-6, along with the respective hardware references, are listed in Table 1-1. These references describe each subsystem from a hardware standpoint, giving the function, programming information, turn-on and turn-off procedures, use of the operating controls and indicators, recovery procedures, and directions for loading and unloading such subsystems as the printers, card readers, diskette drives, and the magnetic tape unit.

Table 1-1. System 80 Models 3-6 Hardware Documentation

Processor/Device	UP Document Number (current version) General Description	Subsystem/ Programmer Reference	Operator Reference	Programmer Reference
System 80 Processor	8800	–	8880	8881
Card Punch, 0608	8894	8742	8896	–
Card Readers, 0719	8619	8742	8617	–
Console Workstation	–	–	(See 8880)	–
Diskettes, 8420/8422	8699	8742	(See 8880)	–
Disks				
8417	8916	8742	8917	–
8419	8918	8742	8919	–
8470	10002		10004	
Magnetic Tape, UNISERVO 10 (Type 0871)	8891	8890	8609	–
Printers				
0776	8354	–	8250	–
0789	8897	9167	8908	–
0797	9159	–	9160	–
0798	8871	–	8882	–
System 80 Workstation	–	8742	(See 8880)	–
Card Processor, 1004/1005	4052	–	7839	–
UNISCOPE 100/200 Terminal	8155	–	7788	7807
UTS 400 Universal Terminal System	8357	–	8358	–
UTS 20	9134	–	9135	9136
UTS 30	10496	–	–	–
UTS 40	9141	–	9142	9143
SVT 1120	11535	–	11534	–
SVT 1122	12265	–	12263	–
SVT 1123	12267	–	–	–
SVT 1124	12267	–	–	–
PC /HT™	11393	–	12605	–
/IT™	11700	–	12605	–
/microIT™	12237	–	12605	–
UNIX	11760	–	11761	11762
MAPPER 5	9690	–	–	–
	10910	–	–	–

PC/HT, PC/IT, and PC/microIT are trademarks of Unisys Corporation.

Section 2
Models 3-6 System Turn-On and Turn-Off Procedures

2.1. Introduction

The System 80 provides simplified turn-on and turn-off procedures for models 3-6 that:

- Apply and remove power to and from the system
- Turn on and off the console workstation, integrated disk subsystem, and all diskette subsystems
- Turn on and off the peripheral subsystems
- Automatically perform the initial microprogram load (IMPL) procedure and start the initial program load (IPL) procedure

Normally, the circuit breakers on the processor cabinet and expansion cabinet (if configured), as well as the circuit breakers on all peripheral devices where the circuit breaker is separate from the ON/OFF switch, are left in the ON position and power is controlled through a main system circuit breaker. Power should be removed whenever the system is to remain unused for a period of time, such as over a weekend, overnight, or for the length of one work shift.

The console workstation, integrated disks, and all diskettes are turned on and off by using the SYS PWR ON/OFF switch; all other peripheral subsystems, such as printers and additional disk subsystems, are turned on and off at the device.

Caution

Before turning system power on or off, be sure that files are protected at each 8417 and 8419 disk drive, or portions of recorded data will be erased. Set the FILE PROT switch on the 8417 disk drive operator control panel to ON (up) position, and press the FILE PROTECT switch/indicator on the 8419 disk drive operator control panel.

Whenever power has been removed from the system, the IMPL and IPL initialization procedures must be performed each time power is reapplied. The system turn-on procedure, using the system-resident (SYSRES) integrated disk, automatically performs the IMPL procedure to load and initialize control storage; it then initiates the IPL procedure to load the resident portion of OS/3 (the supervisor). Operator action is not required until the interactive IPL (3.3.2) when you select how the supervisor is initialized. If you want to use another load device for the IMPL and IPL, you must initiate the manual IMPL when the automatic IMPL/IPL sequence completes. These procedures are explained further in Section 3.

2.2. System Turn-On Procedure

To turn on System 80 models 3-6 from a full power-off condition, proceed as follows:

1. Set the MAIN POWER circuit breaker at the rear of the processor cabinet to the ON position.

2. If an expansion cabinet is used, set the circuit breaker at the rear of that cabinet to the ON position.

Note: *Usually the MAIN POWER circuit breaker and the expansion cabinet circuit breaker are left in the ON position and power is controlled with the SYS PWR ON / OFF switch. Thus, steps 1 and 2 need not be performed after the first time power is applied to the system.*

3. Set all peripheral devices (except diskettes and workstations) to ON, according to the procedures described in the appropriate hardware reference manual. (See Table 1-1.) Workstations are turned on and off, as required, by the workstation user.

4. Set the SYS PWR ON/OFF switch, located on the console workstation table, to the ON position to turn on the console workstation, integrated disks, and all diskettes. After the SYS PWR ON/OFF switch flashes for a few moments during the power-on confidence test, the console workstation screen and the SYS PWR ON/OFF switch both light. System state messages are displayed, indicating that loading is in progress. Proceed with the interactive IPL (3.3.2).

2.3. System Turn-Off Procedure

The turn-off procedure for System 80 models 3-6 turns off the console workstation, the integrated disk, and all diskette subsystems; turns off the peripheral devices; then turns off the power to the system.

To turn off and remove power from System 80, proceed as follows:

1. Set the SYS PWR ON/OFF switch, located on the console workstation table, to the OFF position to turn off the console workstation, integrated disks, and all diskettes. The SYS PWR switch and console workstation screen lights go out.

2. Set all peripheral devices (except diskettes and workstations) to OFF, according to the procedures described in their respective hardware reference manual (see Table 1-1). Workstations are turned on and off, as required, by the workstation user.

Note: *Step 3 need not be performed unless complete power turn-off is desired due to an emergency or if the procedure is performed as a normal site requirement.*

3. Set the circuit breakers at the rear of the processor and expansion cabinets to the OFF position.

Caution

Do not open the processor cabinet doors to operate other circuit breakers. Air has been purged in the processor cabinet to maintain efficient operation of the internal disk drive.

Section 3
Models 3-6 System Initialization Procedures

3.1. Introduction

Initializing your system means preparing it for normal operation. There are two steps or procedures that do this. They are

- Initial microprogram load (IMPL) procedure. This procedure loads the microinstructions into control storage. These microinstructions (also called microcode) control the basic operation of your system.

- Initial program load (IPL) procedure. This procedure loads the resident portion of the OS/3 operating system (called the supervisor) into main storage.

There are three occasions when you initialize your system. First and foremost is when you turn the system on. (For details about the turn-on procedures, see Section 2.) During turn-on, the system does most of the work for you. It automatically performs the IMPL procedure and starts the IPL procedure. You don't do anything until the system reaches a stage called the interactive IPL.

The other two occasions that require some form of initialization occur during day-to-day operations. They are: 1) when a nonrecoverable error occurs while the system is running; and 2) when you want to use a different supervisor (possibly on another device).

In both of these situations, you must initialize the system yourself. It is not automatic.

When you perform both the IMPL and IPL procedures, you're completely re-initializing the system. This isn't always required. In fact, more often than not, the only procedure you need is the IPL procedure. For instance, when you need an alternate supervisor, all you're changing is the resident portion of OS/3. Since you're not changing the microcode, it isn't necessary to IMPL.

Similarly, many nonrecoverable errors only require you to re-IPL. When a nonrecoverable error does occur, the system displays an error message at the console. Refer to the message in the system messages manual. This manual instructs you what you should do. If it tells you to reboot the system, it still might only require the IPL procedure. However, it's a good rule of thumb to perform both the IMPL and IPL procedures in this case. The IMPL procedure is easy and fast and ensures that the system is properly loaded.

If an error occurs during either the IMPL or IPL procedure, then you should always start again with the IMPL procedure.

The IPL procedure is a two-phase operation. In the first phase you initiate the procedure by indicating which device you're loading your supervisor from. The second phase is the interactive phase. During this phase you supply the name of your supervisor and you select various load options. As previously mentioned, if you're powering your system up, the system itself does all the work up to this interactive phase.

3.2. Initial Microprogram Load Procedure

The IMPL procedure loads microcode into your system's control storage. Once it's loaded, it stays in control storage until you turn the system off or you choose to re-IMPL. When you're powering your system up, the IMPL is the first initialization step and the system does it for you.

It loads (or tries to load) the microcode from the system resident, integrated, fixed disk. This is the disk that's built into the System 80 cabinetry. In all probability, that's where your microcode resides, along with the rest of the OS/3 software. It's commonly referred to as the SYSRES volume.

Be aware, however, that it is possible to locate the OS/3 software and the microcode on a different disk. (That disk becomes your SYSRES volume.) If this is the case, you have to re-do the IMPL procedure after you turn the system on. When you initiate the procedure, you can tell the system where to look for the microcode.

Once the system is running, the occasions when you must perform the IMPL are rare. They are:

- When you want to reload the microcode using a different device from the last one you or the system used when the IMPL executed. For example, you might want to load the microcode from diskette or from another disk.
- When a previous IMPL is unsuccessful.
- When the IPL procedure is unsuccessful.
- When an unrecoverable error occurs and you want to completely reboot the system.

There are two ways to initiate the IMPL procedure. You can use the POCLR (power clear) button located under the console table or you can use the IMPL key on your console keyboard (see 3.2.1 and 3.2.2).

Notes:

1. *The POCLR button is available only on models 4 and 6.*

2. *The 8470 disk is not supported as an IMPL/IPL device for models 4 and 6; however, it can be used as a SYSRES device.*

3.2.1. Using the POCLR Button

The POCLR button simulates turning the power off and on again. It completely clears the contents of main storage. Then, it performs the IMPL procedure and initiates the IPL procedure. You don't do anything until the interactive IPL (3.3.2). When you use the POCLR button, the system loads the microcode from the fixed integrated disk (device 100). If you want to load from some other device, don't use the POCLR button.

3.2.2. Using the IMPL Key

To IMPL using the IMPL key on the console keyboard, follow these steps:

1. Get into system control mode.

   Press the FUNCTION key, hold it down, and press the D/SYS CONT key. This bumps the cursor below the band of reverse video (called the system state line) that is located two-thirds of the way down the screen.

2. Initiate the IMPL.

   Press the FUNCTION key, hold it down, and press the IMPL key.

   Note: *The IMPL key is part of a dual-function key. Depending on the model keyboard you are using, "IMPL" is lettered above one of the following keys: IN DISP INSERT IN LINE on model A keyboard. FCC ENABLE STATUS (shown in step 2) on model B, and F20/F11 on model C. For example, when instructed to press the IMPL key on a model B keyboard, press and hold the FUNCTION key, then press the FCC ENABLE STATUS key.*

The system displays the message IMPL=CDD? and positions the cursor at the beginning of the next line:

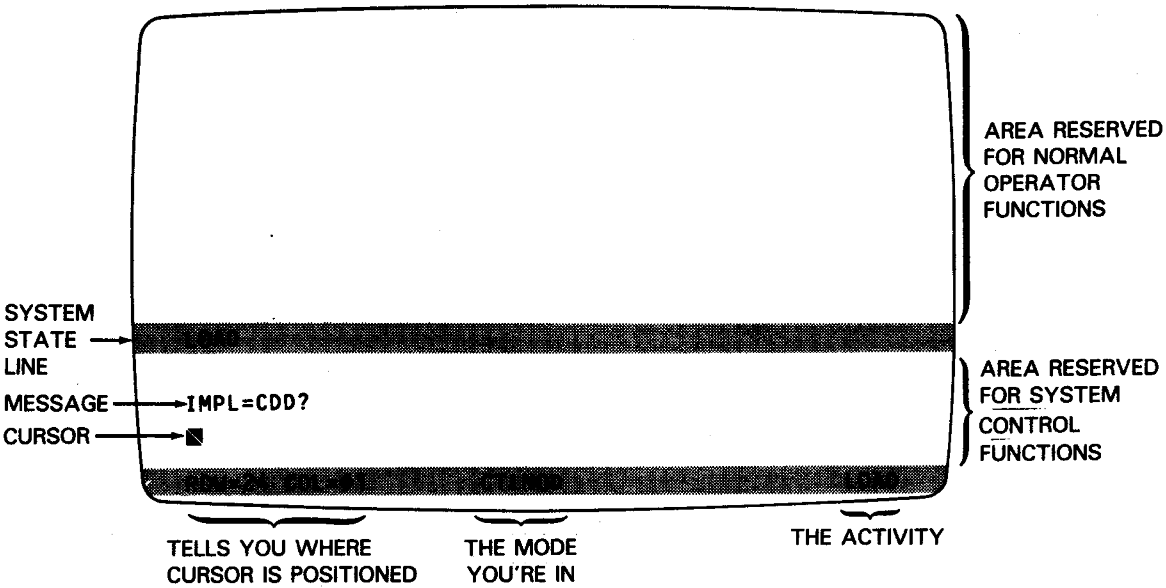

3. Identify where the microcode resides:

- Are you loading from the fixed integrated disk (device 100)?

  YES Press the XMIT key; do not key in a value. The system loads the microcode from device 100 (the system default), initiates the IPL procedure for you, and brings you up to the interactive IPL (3.3.2).

  NO Key in the device's 3-digit address and press XMIT. The system again displays the IMPL=CDD? message. Press XMIT. No other key-in is required. Go on to step 4.

- What is a device address?

  Each device in your configuration has a device address called the did. The did is a combination of the device's channel, subchannel, and selected device number. Normally, every device is labeled with its did (except for device 100). The only valid IMPL devices and their dids are:

Device	Channel	Sub-channel	Device No.	did
Integrated fixed disk	1	0	0	100
All other disks	1	0	1 - 7	100 - 107
Diskettes	3	2	0 - 3	320 - 323

Example

If you load the microcode from disk drive 105, you respond:

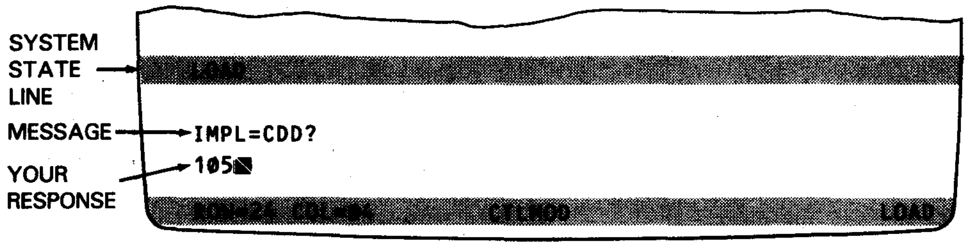

Press XMIT.

The IMPL=CDD? message appears again:

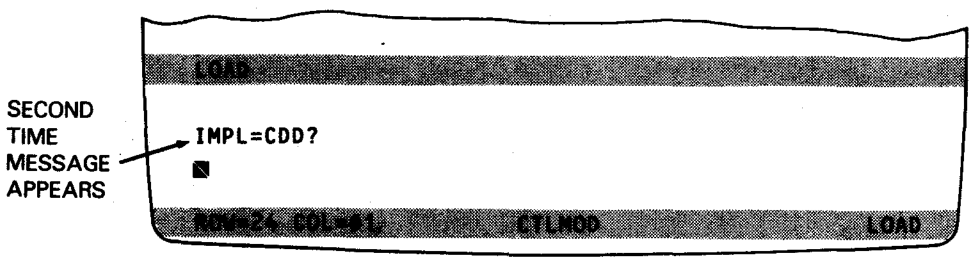

Don't enter anything. Just press XMIT.

4. The IMPL procedure completes.

   The system displays the first IPL message:

   ```
   IPL=CDD?
   ```

   You're ready to begin the IPL procedure. See 3.3.

What to Do If the IMPL Procedure Fails

If the IMPL procedure fails, the system displays a message on the system state line (the band of reverse video two-thirds of the way down the console screen).

These messages and what you should do are:

`LOAD*ERROR STOP*`
: See the *System Messages Reference Manual* (UP-8076); identify the error code displayed on the code displayed on. Follow any procedure the manual recommends. Retry the IMPL procedure.

If the error condition persists, record the error code displayed on the screen and see the *System 80 Operator Maintenance Guide* (UP-8915) for error definition and reporting procedures.

CHECK STOP

Press the POCLR button. The system retries the IMPL procedure.

3.3. Initial Program Load Procedure

The IPL procedure loads and initializes the part of the OS/3 operating system that stays in main storage while the system is running. This part of the operating system is called the supervisor.

When you IPL, you're telling the system where to find the supervisor and what options it should initialize it with. Since the procedure does both of these things, we split the discussion into two phases: the initiation phase (3.3.1) and the interactive phase (3.3.2).

3.3.1. How to Initiate the IPL Procedure

Have you just . . .

- Turned the system on?
- Pushed the POCLR button?
- Performed the IMPL procedure using device 100 as the load device?

If so, the system performs this step for you. Go on to the interactive IPL (3.3.2).

Caution

Make sure the system is idle (no jobs are running) before you IPL. Otherwise, the SYSRES volume table of contents (VTOC) may be damaged, requiring you to generate a new SYSRES volume.

To initiate the IPL procedure, follow these steps:

- If you have just completed the IMPL procedure using a device other than device 100 (3.2), go to step 3.
- If you have *not* just completed the IMPL procedure using a device other than device 100, begin with step 1.

1. Get into system control mode

   Press the FUNCTION key, hold it down and press the D/SYS CONT key. This bumps the cursor below the band of reverse video (called the system state line) that appears two-thirds of the way down the screen.

2. Initiate the IPL

   Press the FUNCTION key, hold it down and press the IPL key.

   Note: *The IPL key is part of a dual-function key. Depending on the model keyboard you use "IPL" is lettered above one of the following keys: EOD ERASE EOL on the model A keyboard, FCC GEN on model B (shown above) and F21/F22 on model C. For example, when instructed to press the IPL key on a model A keyboard, press and hold the FUNCTION key, then press the EOD ERASE EOL key.*

   The system displays the first IPL message and positions the cursor at the beginning of the next line.

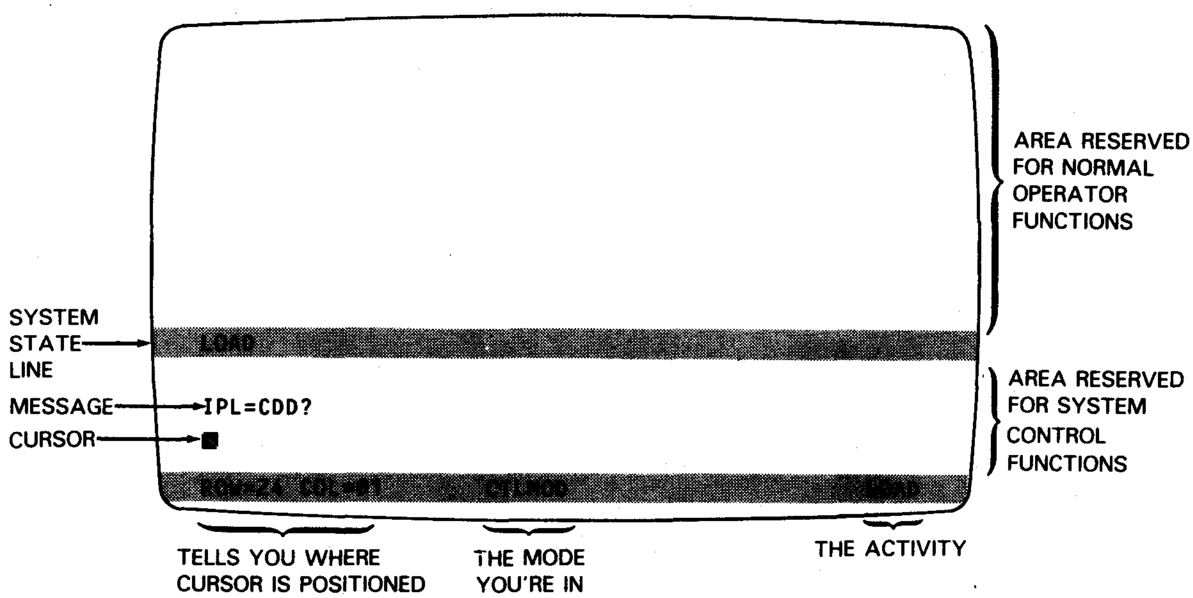

3. Identify the load device.

   - Is your supervisor on the integrated fixed disk (device 100)?

     YES Press the XMIT key; no key-in is required. The system displays the IPL01 screen; this is the beginning of the interactive IPL (see 3.3.2).

NO Is the device a diskette or a disk?

Diskette — Enter the diskette drive's 3-digit address and press XMIT. The system displays the IPL01 screen; this is the beginning of the interactive IPL (see 3.3.2).

Disk — Enter the disk drive's 3-digit address and press XMIT. The system displays another message:

```
CYL=CCC?
```

You don't need to enter a value; press XMIT. The system displays the IPL01 screen; this is the beginning of the interactive IPL (see 3.3.2).

- What is a device address?

  Each device in your configuration has a device address called the did. The did is a combination of the device's channel, subchannel, and selected device number. Normally, every device is labeled with its did (except for device 100). The only valid IPL devices and their dids are:

Device	Channel	Sub-channel	Device No.	did
Integrated fixed disk	1	0	0	100
All other disks	1	0	1 - 7	101 - 107
Diskettes	3	2	0 - 3	320 - 323

Example

If your supervisor is located on disk 104 you respond:

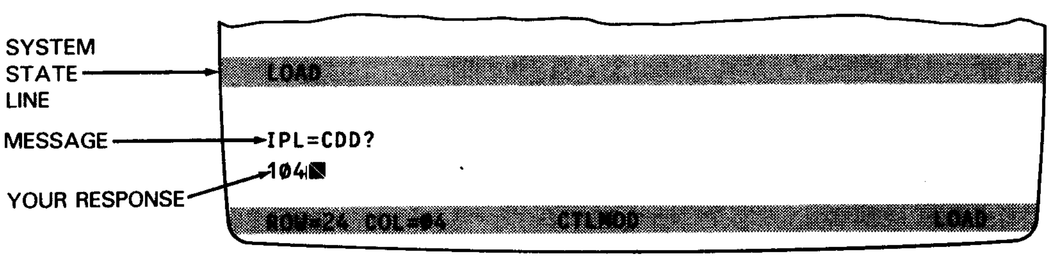

Press XMIT. The CYL=CCC message appears:

```
LOAD

IPL=CDD?
CYL=CCC?
1Ø4

ROW=24 COL=01          CTLMOD                    LOAD
```

You press XMIT.

3.3.2. Interactive IPL - How to Initialize the Supervisor

You're ready to initialize the supervisor. By the time you reach this phase, you've directly or indirectly completed the IPL procedure's initiation phase (3.3.1). Let's review how you could have arrived at this point:

- You turned the system on. The system performed both the IMPL procedure and the IPL procedure's initiation phase for you.
- You pressed the POCLR button. The system performed both the IMPL procedure and the IPL procedure's initiation phase for you.
- You manually initiated the IMPL procedure and loaded the microcode from device 100 (the default). The system initiated the IPL procedure for you.
- You just finished the IPL's initiation phase (the system didn't do it for you).

The display that appears at your console screen is the IPL01 screen. Note that it appears above the system state line. The system has taken you out of system control mode and into console mode.

```
Ø1 IPLØ1 KEY IN AND TRANSMIT
Ø2 SUPERVISOR NAME                                    (DEFAULT=SY$STD)
Ø3 OPTION S=STD, D=DEBUG, L=S/A LOAD, Q=QUICK (DEFAULT=S)
Ø4 LOAD DEVICE ADDRESS                                (DEFAULT=did)
Ø5?SY$STD,S,did
Ø5 ■
```

DEFAULTS

CURSOR POSITIONED FOR YOUR RESPONSE

SYSTEM STATE LINE

AREA RESERVED FOR NORMAL OPERATOR FUNCTIONS

Complete this phase by following these steps:

1. Select the supervisor, the load option, and the load device.

   On the IPL01 screen you supply the name of your supervisor, the load option, and the device where the supervisor resides. There are default values for each of these items. They're displayed on line 05 of the screen:

   ```
   05?SY$STD,S,did
   ```

   where:

   SY$STD
   : Is the default name for your supervisor. When you create a supervisor, you can assign any name to it that you want. If you don't name it when you create it, the system assigns the name SY$STD.

   S
   : Specifies the STANDARD default load option. Use this specification on the initial IPL after receiving a new release or after applying an SMC or SMP. This updates the table of symbiont sizes that is maintained on disk. This process is time-consuming and is necessary only when a symbiont size has changed.

     The other load options are:

     D
     : Specifies the supervisor DEBUG load option. (For details, see the *Supervisor Technical Overview* (UP-8831).)

     L
     : Specifies the special supervisor load for stand-alone programs.

     Q
     : Specifies the QUICK load option, which should normally be used, because your IPL will be significantly faster.

     Note: *This option cannot be used if an SMC has been applied or if a new supervisor or ICAM has been generated.*

   did
   : This is the device address where the supervisor resides. If you manually initiated the IPL procedure (3.3.1), the system takes the device address you entered and displays it here. For example, if you entered 102 when the system displayed the IPL=CDD? question, 102 appears as the did on line 05 of the IPL01 screen.

     If the system performed the IPL initiation step for you, then the device address 100 is displayed. Device 100 is the system's integrated fixed disk; it's the default load device for both the IMPL and IPL procedures.

Do you accept the defaults displayed?

YES Press the XMIT key; no key-in is required. The system displays another screen for you to fill in. See step 2.

NO Enter the values you wish to use. (The cursor is already positioned where you enter your selections.) Press XMIT. The system displays another screen for you to fill in. See step 2.

For example, if the name of your supervisor is SUP140 and it's located on device 103, you would complete the IPL01 screen as follows:

```
01 IPL01 KEY IN AND TRANSMIT
02 SUPERVISOR NAME                                  (DEFAULT=SY$STD)
03 OPTION S=STD, D=DEBUG, L=S/A LOAD, Q=QUICK       (DEFAULT=S)
04 LOAD DEVICE ADDRESS                              (DEFAULT=100)
05?SY$STD,S,100
05 SUP140,S,103 ■
```

YOUR RESPONSE

Press XMIT. (In this example, the system initiated the IPL procedure; therefore 100 appears as the default device address.)

2. Answer the next series of questions that are displayed on the console screen; press XMIT.

CURSOR POSITIONED AT FOR YOUR RESPONSE

```
                ***OS/3 VERSION xx.xx.xx***
DATE? (YY/MM/DD) █_/__/__
TIME? (HH:MM:SS) __:__:__
RUN LIBS DVC ADDR?  (DEFAULT=xxxxxx)___
DUMP FILE DVC ADDR? (DEFAULT=xxxxxx)___
RECOVER FILES?
 JOB QUEUE   (N,Y,H             DEFAULT=x)_
 ERROR LOG   (N,Y               DEFAULT=Y)_
 SPOOL FILES (N, A, C, L, H     DEFAULT=x)_
SPOOLING DVC ADDR?      (DEFAULTxxxxxx)___
MODIFY SUPERVISOR?      (N, Y     DEFAULT=N)_
```

Let's examine each field displayed:

```
*** OS/3 VERSION xx.xx.xx ***
```

Shows you what version of OS/3 you're working with.

```
DATE? (YY/MM/DD)  nn/  nn/  nn
```

Asks you for the current date. Enter the date in the displayed format. This screen tells you to enter the year (YY) followed by the month (MM) then the day (DD). But since this format is set at SYSGEN, it can be in a different order. The system itself uses the YY/MM/DD for all processing operations and output messages, regardless of the format you use on this screen.

After you complete this screen and you press XMIT, the system compares the date you enter with the date entered the last time a supervisor was loaded from this SYSRES volume. If the date you enter is six days less than or six days greater than the date of the last load, the system displays the message:

```
DATE QUESTIONABLE
```

You can either ignore the message by pressing XMIT or you can reenter the date.

```
TIME? (HH:MM:SS) __:__:__
```

Asks you for the time of day in hours (HH) minutes (MM) and seconds (SS) (seconds are optional). Use military notation (0100 to 2400 hours).

```
RUN LIBS DVC ADDR? (DEFAULT=xxxxxx) ___
```

Asks whether the system's job run library file (YRUN) is on the disk specified at SYSGEN and displayed as the default (xxxxxx) or on another disk volume. To accept the default, don't key anything in. You can go on to the next question (use the RETURN key or the TAB FORWARD key). To locate the run library on a different volume, key in the device address of the disk you want to use.

Remember, the disk you identify for YRUN must be online for the system to operate. If it isn't online, the system displays a message at the console, asking you to mount the volume.

```
DUMP FILE DVC ADDR? (DEFAULT=xxxxxx)
```

Asks whether the system dump file (YDUMP) is to be located on the SYSRES disk displayed as the default (xxxxxx) or on an alternate disk. To accept the default, don't key anything in. You can go on to the next question (use the RETURN key or the TAB FORWARD key). To specify an alternate dump file, key in the device address of the disk you want to use.

During system initialization, YDUMP will be allocated on this alternate disk when necessary. However, SYSRES will always contain a minimum YDUMP file for dumps taken during system initialization.

> **Note:** *When you are using an alternate dump file and YDUMP is extended on SYSRES, you can use a system jobstream, SG$SADMP, to scratch the original YDUMP and to reallocate and format a minimum YDUMP on SYSRES.*

`RECOVER FILES?`

Asks you if you want to recover the following three types of files.

`JOB QUEUE (N, Y, H  DEFAULT=x) _`

Whatever default was set for this option at SYSGEN appears as the default on this line. To delete jobs previously filed in the job queue, key in N. To retain these jobs for processing, key in Y. To place them in a hold status, key in H. To accept the default that's displayed, enter nothing. Move to the next question. If you're using the basic supervisor (SY@BAS), key in N the first time you perform a system IPL.

`ERROR LOG (N, Y  DEFAULT=Y) _`

To clear all accumulated error in the error log and start a new error log file, key in N. Otherwise, to preserve the contents of the error log and continue to list new errors, enter nothing. Move to the next question.

`SPOOL FILES (N,A,C,L,H  DEFAULT=x) _`

This message appears only if you have spooling configured in your system. It asks you for the level of recovery you want for the spool file in your system. The default is set as SYSGEN. Your choices are:

`N`

The spool file is empty; no recovery is required. All previously spooled input and output was processed before the system was brought down.

`A`

Recover (and then process) all spooled input and output files. This response recovers all spooled subfiles, whether they're complete, incomplete, or saved. You must specify this option or the H option if you want to recover the console log file. The system does not recover any console or workstation messages that were still in the main storage buffer when the system went down.

`C`

Recover only complete subfiles.

`L`

Recover only the user log directory.

`H`

Select this only if you took a dump of the previously loaded system. This option recovers all spooled subfiles as well as console messages that were in the main storage buffer when the system went down. (Workstation messages in the buffer are not recovered.)

Use this recovery method (sometimes referred to as a hot start) only if a system crash occurs. If you plan to shut the system down and recover spooled files at a later date, breakpoint the console log, then select the A, C, or L recovery option.

Note: To take a system dump, press and hold the FUNCTION key, then press the D/SYS CONT key to get into system control mode. Then press the FUNCTION key and the restart key. Refer to the Dump Analysis Programming Guide *(UP-9980) for system dump information.*

`SPOOLING DVC ADDR? (DEFAULT=xxxxx) ___`

This message only appears if your system has spooling configured. Further, if you have multivolume spooling, it appears once for each volume where the spool file can be located as specified at SYSGEN. Each of these messages asks you to identify the disk volumes that contain the spool file. To use the SYSGEN default (*xxxxxx*), press the spacebar three times. To locate the spool volume on a different disk, key in the disk's device address. Repeat this procedure for each message displayed.

Remember, all spool volumes must be online for the system to operate.

`MODIFY SUPERVISOR? (N, Y DEFAULT=N) _`

You can make certain modifications to the supervisor before you initialize it. Appendix B describes these modifications. To modify the supervisor, enter a Y. The screen clears, then displays various output messages. See Appendix B for details.

If you don't want to modify the supervisor, press XMIT.

Caution

Be sure to respond to all the preceding statements requiring keyin before pressing the XMIT key. Once the XMIT key is pressed, the questions and answers are lost and the entire procedure must be restarted if the procedure was not performed correctly.

3. If your system is *not* configured with the disk cache feature, go to step 4.

   Specify disk cache (see Section 12 for more information on disk cache). The system asks you how much memory it should use for disk cache. It displays the message:

   ```
   CM01 ENTER THE NUMBER OF 1024 BYTE BLOCKS OF MEMORY FOR DISK CACHE OR
   NONE.  VALID VALUES IN THE RANGE OF 64-1024?
   ```

   Enter NONE or the amount you require. We recommend you specify at least 100 blocks. For greater productivity, you should assign 256 more.

4. If your supervisor is *not* configured with the transient work area feature, go to step 5.

   Initialize the transient work area. The system should display the message

   ```
   TRANSIENT WORK AREA IS INITIALIZED
   ```

   as soon as the IPL procedure completes. If the message doesn't appear, initiate the transient work area yourself. To do so, issue the TW command (see 8.5.16).

5. System displays job slot header, job queue, and error log messages.

   When the supervisor is loaded and initialized, the system displays the job slot header and job queue and error log messages (if appropriate). The system is ready to process user jobs.

   You've completed the IPL procedure when the system displays this screen:

   ```
   JOB SLOTS { (1)   (2)   (3)   (4)   (5)   (6)   (7) yy/mm/dd hh:mm:ss
             { (8)   (9)   (1Ø)  (11)  (12)  (13)  (14)

   System 8Ø OS/3  version-no  supname      yy/mm/dd  hh:mm:ss
                               Name         Date      Time
                               of your
                               supervisor
   ```

   Each number represents a job slot. The number of slots that appear is the number of jobs that can run at the same time. This number is set at SYSGEN. When a user initiates a job, the jobname replaces the numbered slot on the screen.

   For example, if three jobs named PAYROL, MYJOB, and LIST are running, they each occupy a job slot:

   ```
   PAYROL   MYJOB   LIST   (4)    (5)    (6)    (7)
   (8)      (9)     (1Ø)   (11)   (12)   (13)   (14)
   ```

In addition to the job slots, the system tells you about the files it recovered. If you selected Y or H for the JOB QUEUE recovery option (see step 2) this message appears:

```
JOB QUEUE RECOVERED - x JOBS QUEUED
```

where x is the number of jobs queued.

If you selected N for the ERROR LOG recovery option (see step 2) the system displays this message:

```
ERROR LOG NOT RECOVERED
```

Note: *If a device (except the workstation) is not online when you perform the IPL procedure, the system sets the device to not available. Use the SET IO command to make the device available again.*

If you have configured an alternate console, the final message will be displayed on that device, and the hardware console will be cleared and deactivated.

What to Do If the IPL Procedure Fails

If the IPL procedure fails, the system displays an error message on the console screen. Messages preceded by an IPL id appear above the system state line (the band of reverse video two-thirds of the way down the screen). The other-three message types appear on the system state line.

For all messages with the IPL prefix, see the *System Messages Reference Manual* (UP-8076). For the other three message, follow these procedures:

LOAD*ERROR STOP*

Use the *System Messages Reference Manual* (UP-8076) to identify the error displayed on the screen. Retry the IPL procedure. If the error persists, record the error code and refer to the *System 80 Operator Maintenance Guide* (UP-8915).

STOPPED*HPR*

Use the *System Messages Reference Manual* (UP-8076) to identify the HPR code (included in the message text as INST=hpr-instruction-code). For a nonrecoverable error, repeat the IPL procedure. If the error persists, see the *System 80 Operator Maintenance Guide* (UP-8915) for more details.

CHECK STOP

Press the POCLR button.

Section 4
Models 8-20 System Definition

System 80 models 8-20 are interactive data processing systems designed for established users who want to upgrade their data processing capabilities. Models 8-20 offer increased main storage capacity, processor speed, and input/output rates. In addition, they support a significantly greater number of peripheral devices and as a result, offer faster processing, more disk storage space, and increased system size.

The System 80 models 8-20 operate on Operating System/3 (OS/3), the multiprogramming software system used by all System 80 models. Different versions of the system can be assembled from system configurations (see 4.2).

4.1. Models 8-20 System Configuration

The extensive capabilities of System 80 models 8-20 permit you an almost unlimited variety of system configurations from which you can tailor one to your particular needs. The system is configured from the basic processor complex to maximum expansion with one expansion cabinet.

Diagrams of basic System 80 models 8-20 and their expanded hardware options are presented in Figures 4-1 and 4-2.

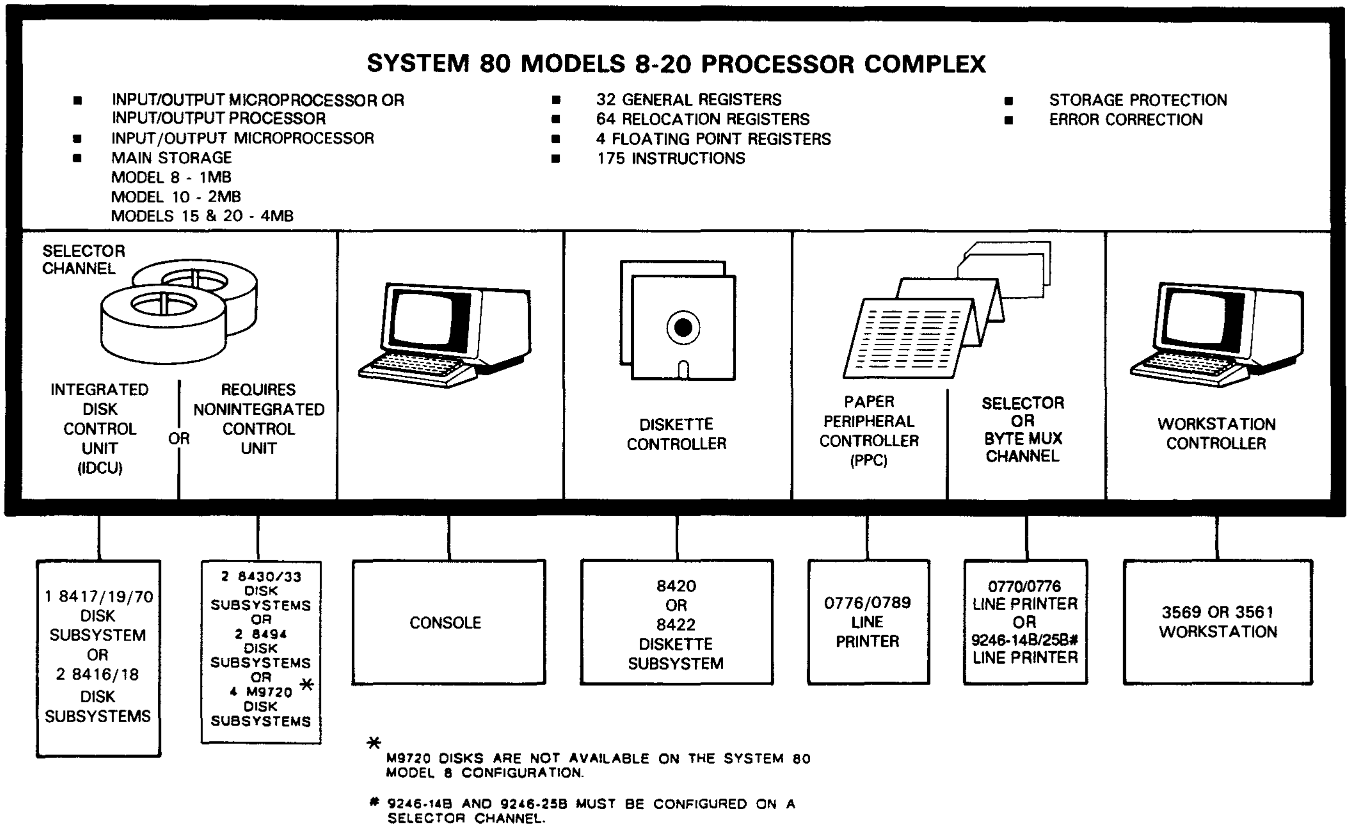

Figure 4-1. Minimum System 80 Configuration - Models 8-20

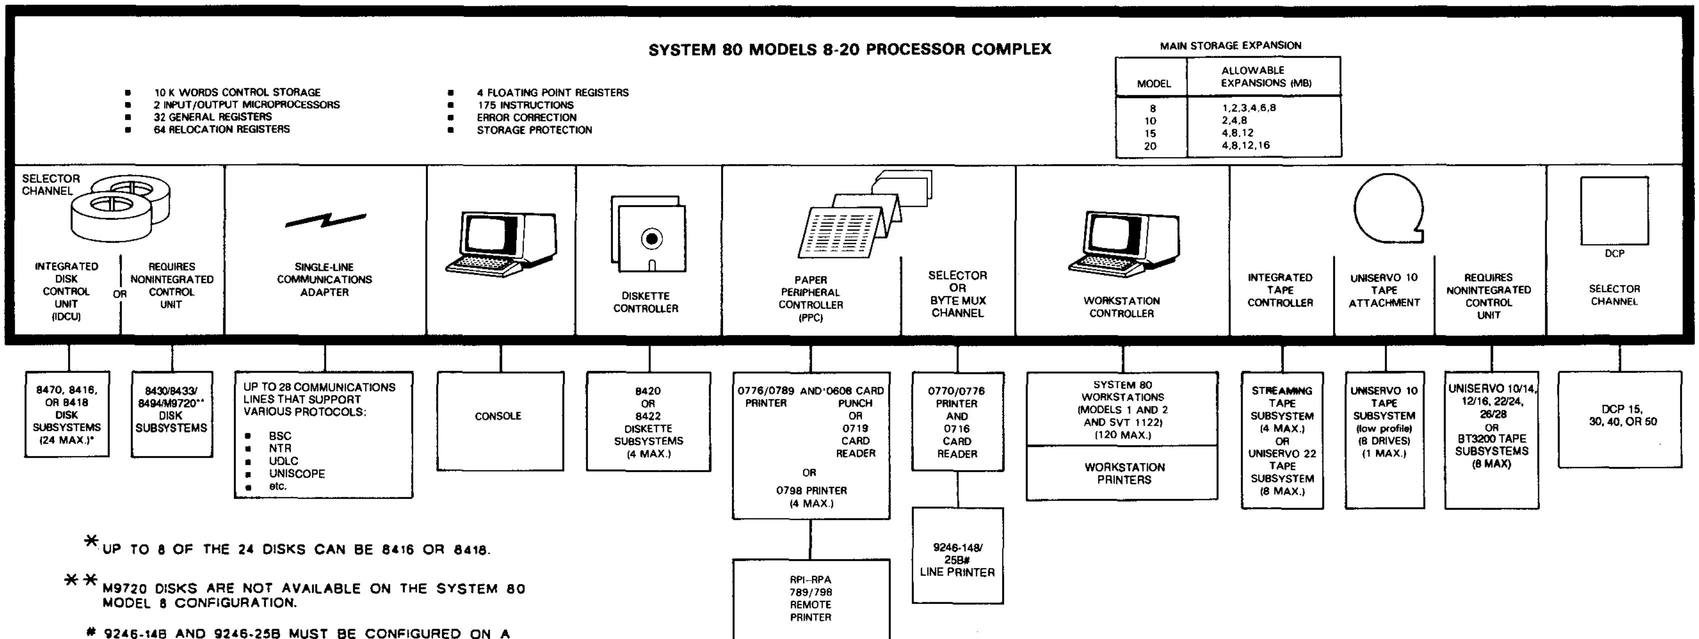

* UP TO 8 OF THE 24 DISKS CAN BE 8416 OR 8418.

** M9720 DISKS ARE NOT AVAILABLE ON THE SYSTEM 80 MODEL 8 CONFIGURATION.

9246-14B AND 9246-25B MUST BE CONFIGURED ON A SELECTOR CHANNEL.

Figure 4-2. Maximum System 80 Configuration - Models 8-20

4.2. Central Hardware

The basic System 80 models 8-20 hardware consists of the processor complex and the console, disk, and diskette subsystems, as well as a printer subsystem. It also includes the standard processor features and input/output controls as shown in Figure 4-1. The central hardware for a maximum system can consist of the processor complex plus an I/O expansion cabinet and additional integrated processor features, and I/O controls (Figure 4-2).

The central processor complex contains all processors, I/O channels, and integrated peripheral system controls. The following central processor components are provided:

- High-performance control storage
- Expandable main storage
- 175 basic instructions, including instructions for floating-point, decimal, and system control
- 32 registers for program and system use, plus 4 floating-point registers and 64 relocation registers
- System control processor
- For model 8 systems, an input/output microprocessor (maximum of 2) supporting:
  - Diskette controller
  - Workstation controller
  - Paper peripheral controller
  - Integrated tape controller
  - Single-line communications adapter (SLCA)
- For model 8 systems, a channel controller (maximum of 2) supporting :
  - One byte MUX channel (for console, printer, and card reader support)
  - Up to five integrated selector channels (for disk, printer, tape, and distributed communications processor (DCP) support)
- For model 10/15/20 systems, a selector channel (maximum of 6) for disk, printer, tape, and DCP support
- For model 10/15/20 systems, an input/output processor (IOP) (maximum of 2) that provides support equivalent to the IOMP on the model 8

  Brief descriptions of the major components follow.

4.2.1. Control Storage

Control storage contains the macroinstructions that control operations of the central processor. The CPU executes and controls instructions and processes I/O interrupts, interval timer activities, and general interrupts.

4.2.2. Main Storage

The main storage unit stores data and commands. The minimum main storage values are shown in Figure 4-1. The main storage can be expanded (by adding another main storage unit). The increments vary by processor as shown in Figure 4-2.

Main storage provides error correction for single bit errors and error detection for double bit errors. It also provides a protection scheme for data and programs using keys. Main storage is accessed directly by the central processor (for programs), the system control processor, the input/output microprocessor, and the channel controller.

4.2.3. System Instructions

The system instructions control the operation of the central processor. A repertoire of 175 basic instructions, 31 of which are privileged, is provided with the system. The privileged instructions are used for system control and input/output operations. Nonprivileged instructions are used for general control, data manipulation, decimal operations, list processing, and floating-point operation.

4.2.4. System Registers

The system registers are provided for accumulating, addressing, and controlling operations and data within the system. The 32 general purpose registers may be used as accummulators in general arithmetic and logical operations. They are also used as base address and index registers in arithmetic and logical operations and in address generation. The 16 control registers are used to specify whether or not an operation can be performed in the system, or they may be used to provide special information required by a system facility. The four floating-point registers are available for floating-point operations when the floating-point instruction feature is included in the system. Main storage space is fully used by the system through dynamic storage relocation and the 48 relocation registers. Automatic relocation of all programs in main storage allows maximum loading of programs.

4.2.5. System Control Processor (SCP)

The system control processor (SCP) functions independently of the central processor to control the console, control panels, system initialization, maintenance, and diagnostic functions, system recovery, and remote maintenance.

4.2.6. Input/Output Microprocessor (IOMP) or Input/Output Processor (IOP)

The IOMP (for model 8) or IOP (for model 8/15/20) initiates, directs, and monitors the transfer of data between main storage and some of the peripheral I/O devices. The IOMP/IOP accepts instructions from the central processor (through user and system software), from the workstation, and from operator controls on the devices. After an I/O instruction is initiated, the data is transferred independently of other processor functions; i.e., the I/O and the processor operate concurrently.

The IOMP/IOP is used to support low-speed peripheral subsystems, such as those attached to a diskette controller, workstation controller, paper peripheral controller, or integrated magnetic tape controller. Up to eight controllers and a UNISERVO® 10 tape subsystem (low profile) can be attached to an IOMP/IOP through the shared direct memory access interface (SDMA). Data transfer rate with SDMA is 200K bytes/second.

Each IOMP/IOP can also support a single-line communications adapter (SLCA) with up to 14 communications lines. The multiple line communications multiplexer (MLCM) is the interface between the SLCA and the IOMP/IOP. MLCM microcode and IOMP/IOP hardware control buffer pool management, data chaining, and command chaining. The MLCM transfers data at the rate of 128K bytes/second.

The IOMP or IOP supports:

- Diskette controller
- Workstation controller
- Paper peripheral controller
- Integrated tape controller
- Single-line communications adapter

The following subsections describe these hardware components. Manuals for these devices are identified in Table 4-1.

Diskette Controller

The 8420 or 8422 diskette subsystem is attached to the system through the diskette controller. The controller is capable of controlling up to four diskette drives - either two autoload drives and two manual drives or four manual drives. Your system must include at least one. The autoload type provides an automatic loading capacity of up to 20 standard diskettes. The controller uses microprocessors and interfaces the IOMP/IOP through shared direct memory access.

UNISERVO is a registered trademark of Unisys Corporation.

Workstation Controller

A workstation controller attaches up to eight workstations to the system. Additional workstation controllers can be attached to expand the number of workstations in the system to a maximum of 120.

The workstation controller accommodates three types of workstations on the System 80 models 8-20 - the model 1, the model 2, and the SVT 1122. All types may be located up to 5000 cable feet (1500 meters) from the processor cabinet. The controller uses dedicated buffers to allow data transfers at a rate of 19,200 bits/second. The model 2 workstation can directly connect to an 0791, 0797, 0798, or 0425 printer.

Paper Peripheral Controller

The paper peripheral controller attaches various peripheral printers and card processors to the system through the IOMP/IOP. The controller can support one or two printers and two card processors for a maximum of four devices operating concurrently. The controller accommodates the following types of devices: 0776, 0789, 0798 printers, 0608 card punch, and an 0719 card reader. The total print capacity of the paper peripheral controller is 1500 lines per minute. The System 80 models 8-20 support up to 12 paper peripheral controllers.

Integrated Tape Controller

The integrated tape controller connects up to four streaming tape subsystems or eight UNISERVO 22 tape drives to the system. The total number of drives the integrated controller can support is eight and only one controller is allowed per system.

Single-Line Communications Adapter (SLCA)

The SLCA provides the system with communications capabilities. Models 8-20 can support up to 14 SLCAs. You can have 14 low- or medium-speed half-duplex, 7 high-speed full-duplex, or 12 high-speed half-duplex SLCAs. Since your system can include two IOMPs or IOPs, each supporting 14 SLCAs, the total number of communications lines your system can have is 28. Each line can accept devices operating at speeds of up to 56 kilobits/second.

SLCAs connect remotely located workstations, terminals, and all other communications devices to the system. They perform special character recognition, integrity checks, data transfer control between main storage and devices, and other required control functions.

4.2.7. Channel Controller (Model 8 Only)

In model 8 systems, the channel controller initiates, directs, and monitors the transfer of data between main storage and those devices not connected via the IOMP. The system can include two channel controllers. The first channel controller supports the byte MUX channel and three integrated selector channels. A second channel controller can support up to two integrated selector channels.

Byte MUX Channel

The byte MUX channel connects the console to the system. The console includes two diskette drives for system initialization and maintenance. Via the byte adapter, the byte MUX channel also connects the 0770/0776 line printer and the 0716 card reader. The system can support only one byte MUX channel. In turn, the byte MUX channel can connect up to eight devices at a maximum data rate of 70 KB/second.

Integrated Selector Channel

An integrated selector channel controls high-speed I/O devices, such as disks, magnetic tapes, high-speed buffered printers, and distributed communications processors. It operates in burst mode with the subsystem for data transfers. There is a maximum of five integrated selector channels per model 8 system.

The 8416, 8417, 8418, 8419, and 8470 disk subsystems are connected to the integrated selector channel via an integrated disk control unit (IDCU). Each selector channel can support no more than three IDCUs (a maximum of six IDCUs per system) and 24 disks.

The 8430, 8433, and 8494 disk subsystems are connected to the integrated selector channel via nonintegrated (freestanding) disk control units. The system supports a maximum of 16 subsystems.

The integrated selector channel also connects, via nonintegrated controllers, UNISERVO 10, 12, 14, 16, 20, 22, 24, 26, and 28 tape units; 3200 series tape units; 0770 and 0776 printers; 9246-14B and 9246-25B printers; and DCP 15, 30, 40, and 50 units. The system supports a maximum of eight subsystems in any combination.

4.2.8. Selector Channel (Models 10/15/20 Only)

In model 10/15/20 systems, selector channels initiate, direct, and monitor the transfer of data between main storage and high-speed printer, disk, tape, and the DCP. A minimum system requires one selector channel; a maximum of six selector channels may be configured.

Each model 10/15/20 selector channel can be connected to I/O devices as defined for the model 8 integrated selector channel.

4.3. Input/Output Subsystems

The I/O peripheral subsystems available for use with System 80 include the console, disk and diskette subsystems, workstations, magnetic tape subsystems, printer subsystems, card reader/punch subsystems, and remote communications subsystems.

4.4. Associated Hardware Manuals

The I/O subsystems available for use with models 8-20, along with the respective hardware references, are listed in Table 4-1. These references describe each subsystem from the hardware standpoint, giving the function, programming information, turn-on and turn-off procedures, use of the operating controls and indicators, recovery procedures, and directions for loading and unloading such subsystems as the printers, card readers, diskette drives, and magnetic tape units.

Table 4-1. System 80 Models 8-20 Hardware Documentation

Document Number (current version)				
Processor/Device	General Description	Subsystem/ Programmer Reference	Operator Reference	Programmer Reference
System 80 Processor		UP-9692	UP-9608	
Card Punch 0608	UP-8894		UP-8896	
Card Readers 0716 0719	 UP-8619		 UP-7621 UP-8617	
Communications Adapter	UP-8723	UP-8247		
Disks 8416/8418 8417 8419 8430/8433 8470 8494 M9720	 UP-10005 UP-8916 UP-8918 UP-8325 UP-10002 UP-11625	 UP-8344 UP-11627	 UP-8361 UP-8362 UP-8917 UP-8919 UP-8343 UP-10004 UP-11626 98000011	

continued

Table 4-1. System 80 Models 8-20 Hardware Documentation (cont)

Processor/Device	General Description	Subsystem/ Programmer Reference	Operator Reference	Programmer Reference
	Document Number (current version)			
Diskettes				
8406	UP-8475		UP-8476	
8420/8422	UP-8699			
Printers				
0770		UP-8016	UP-7938	
0776	UP-8354	UP-8441	UP-8250	
0797	UP-9159		UP-9160	
0789	UP-8897	UP-9167	UP-8908	
0798	UP-8871		UP-8882	
9246-14B			1221139	
9246-25B			1220905	
UNISERVO Magnetic Tapes				
10	UP-8891	UP-8890	UP-8609	
10/14	UP-8206	UP-8205	UP-8607	
12/16			UP-7882	UP-7661
20			UP-7956	
22/24	UP-8904	UP-8902	UP-8903	
26/28	UP-10010		UP-10009	
3782 Streaming Tape	UP-9382		UP-9381	
B3200 Series Magnetic Tape	5016041			
Uniscope Display Terminal			UP-7788	
UTS400	UP-8357		UP-8358	UP-8359
Workstation			UP-8910	
UTS 20	UP-9134		UP-9135	UP-9136
UTS 30	UP-10496			
UTS 40	UP-9141		UP-9142	UP-9143

continued

Table 4-1. System 80 Models 8-20 Hardware Documentation (cont)

Document Number (current version)				
Processor/Device	General Description	Subsystem/ Programmer Reference	Operator Reference	Programmer Reference
SVT 1120	UP-11535		UP-11534	
SVT 1122	UP-12265		UP-12263	
SVT 1123	UP-12267			
SVT 1124	UP-12267			
PC /HT /IT /microIT	UP-11393 UP-11700 UP-12237		UP-12605 UP-12605 UP-12605	
UNIX	UP-11760		UP-11761	UP-11762
MAPPER 5	UP-9690 UP-10910			
DCP			UP-8502 UP-8743	

4.5. System 80 Models 8-20 Console and Control Panel

As the system operator, most of the functions and procedures you perform are accomplished at the system console and system control panel. Following is a brief general description of these items. Note that two different system console versions are used; one for the model 8 and another for models 10, 15, and 20. Refer to the *Processor and Central Peripherals Operating Guide* (UP-9608) for detailed information on switch and indicator functions for both system console versions.

4.5.1. System 80 Models 8-20 Console

The console consists of a cathode ray tube (CRT) display device and a keyboard. The model 8 United States national keyboard is shown in Figure 4-3. Seven additional national keyboard styles are available.

The keyboard is used to communicate with the system via input commands and to select display frames. When you press the ESC key, a command input mode is initiated. You can enter commands by pressing an appropriate key while pressing and holding the ESC key. On the model 8 system, the LOCK/UNLOCK switch must be set to the UNLOCK position to allow command inputs. Table 4-2 summarizes console operations that can be performed with the ESC key. Refer to the *Processor and Central Peripherals Operating Guide* (UP-9608) for more detailed information.

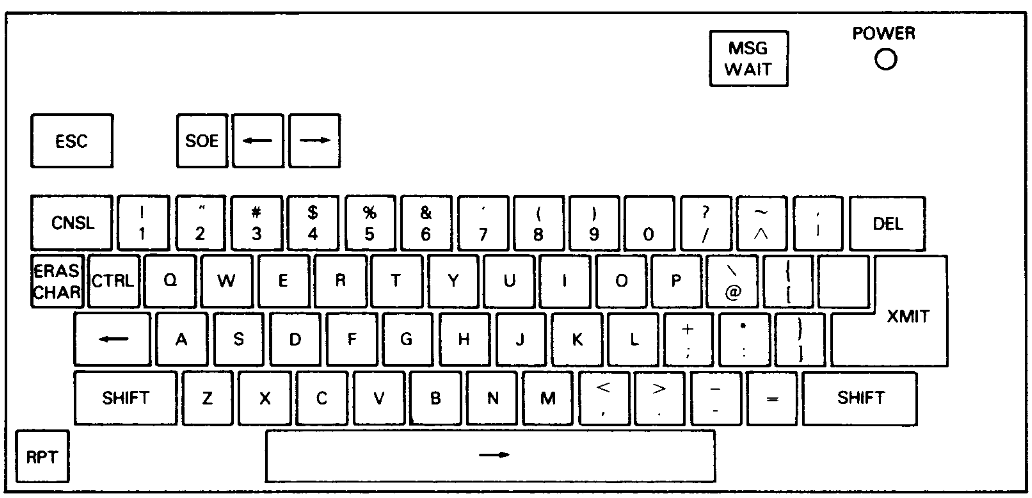

Figure 4-3. Model 8 Console Keyboard

Table 4-2. Console Control Selections

Key Input	Function
ESC and R	Switches processor from stop state to operational state so that software instructions can be executed.
ESC and S	Switches processor from operational state to stop state.
ESC and C	Puts console in system control mode and displays the configuration control frame. This frame is used to control auto IPL, autorecovery, and remote console modes. It also displays processor configuration data and the state of system control panel and console indicators
ESC and M	Puts console in system control mode and displays the manual frame. This frame is used to perform manual IPL and and IMPL functions, and to debug programs and operations.
ESC and A	Puts console in system control mode and displays the alter/display frame. This frame allows you to display and alter various registers, main storage locations, and the PSW.
ESC and P	Puts console in maintenance console mode and displays the service frame. This frame provides numerous error logging and maintenance-related functions.
ESC and T	Puts console in remote maintenance mode and displays the teleprocessing frame.
CNSL	Puts console in operator console mode and displays the console frame. This frame is used to exchange information between the operator and the operating system.

4.5.2. System 80 Models 8-20 System Control Panel

The system control panel contains controls and indicators used for operator maintenance during system operation. On the model 8, this panel is located on the door of the console table. On models 10, 15, and 20, it is physically attached to the keyboard.

Figure 4-4 shows the model 8 system control panel. The model 10/15/20 system control panel is of a different design but provides similar functionality. Brief descriptions of the panel indicators and switches follow (refer to the *Processor and Central Peripherals Operating Guide* (UP-9608) for more detailed information on this panel).

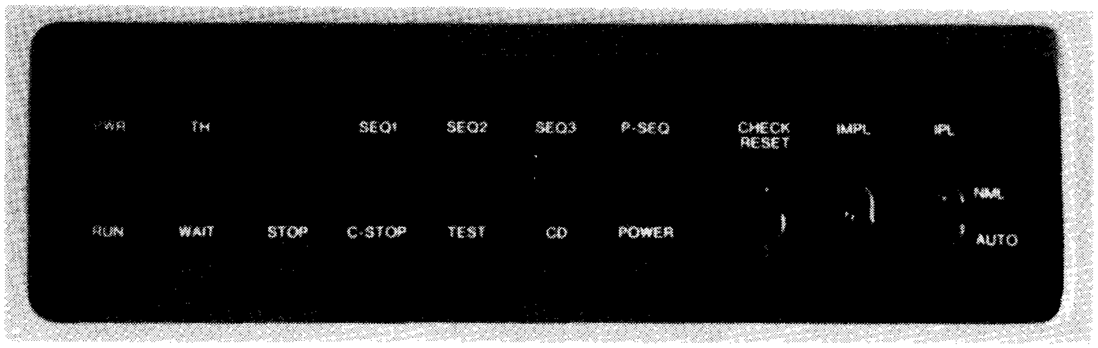

Figure 4-4. Model 8 System Control Panel

Indicator Lights

- PWR

  Indicates a fault in the dc power supply to the system. This light stays lit until the fault is corrected. When this light goes on, the system alarm buzzer also sounds.

- TH

  Indicates that the speed of the cooling fan in the system has dropped below a safe level. This light stays lit until the fan is fixed. When this light goes on, the system alarm buzzer also sounds.

- SEQ1, SEQ2, SEQ3

  Indicates that power has been applied to your system or that the initial microprogram load (IMPL) has begun. Once IMPL and initial program load (IPL) complete, these lights go off. Anytime an error occurs within the system control panel itself during IMPL, IPL, or system operation, these lights go on.

- P-SEQ

  Indicates that you are performing the power-on procedure. This light goes on as soon as you turn on the POWER switch behind the control panel door. It turns off when the P-OK signal is displayed on your console.

- RUN

  Indicates that your system is running jobs.

- WAIT

  Indicates that your system is in a WAIT state.

- STOP

  Indicates that your system is in a STOP state.

- C-STOP

  Indicates that your system is in a CHECK STOP state. This light stays lit until you perform system reset. When this light goes on, the system alarm buzzer also sounds.

- TEST

  Indicates that you are operating the system under other than normal conditions, such as processor STEP mode. This light stays lit until the abnormal condition is cleared.

- CD

  Indicates a remote console is connected. Your system receives the carry detect (CD) signal from the modem and turns this light on.

- POWER

  Indicates that your system is receiving power.

Switches

- CHECK RESET

  Stops the system alarm buzzer.

- IMPL

  Starts the service processor (SVP) and console initial microprogram load (IMPL) procedure. Note that additional IMPL functions must be performed using the manual frame display.

- IPL

  Specifies manual (MNL) or automatic (AUTO) initial program load.

Section 5
Models 8-20 System Turn-on and Turn-off Procedures

5.1. Introduction

The System 80 provides simplified turn-on and turn-off procedures for models 8-20 that:

- Apply and remove power to and from the system
- Turn on and off the console, integrated disk subsystems, and all diskette subsystems
- Turn on and off the peripheral subsystems
- Automatically perform the initial microprogram load (IMPL) procedure and start the initial program load (IPL) procedure (assuming the IPL switch is in the AUTO position)

Normally, the circuit breakers on the processor cabinet, I/O cabinet, and I/O expansion cabinet (if configured), as well as the circuit breakers on all peripheral devices where the circuit breaker is separate from the ON/OFF switch, are left in the ON position and power is controlled through the main system circuit breaker. You should perform power-off procedures whenever the system is unused for a period of time, such as over a weekend, overnight, or for the length of one work shift.

The console, integrated disks, and all diskettes are turned on and off by using the POWER switch on the control panel; all other peripheral subsystems, such as printers and additional disk subsystems, turn on and off at the device.

Whenever power has been removed from the system, the IMPL and IPL procedures must be performed each time power is reapplied. The system turn-on procedure automatically performs the IMPL procedure to load and initialize control storage. If the IPL switch on the control panel is in the AUTO position, the system turn-on procedure also initiates the IPL procedure to load the resident portion of OS/3 (the supervisor). In this case, operator action is not required until the interactive IPL (see 6.3.2) when you select how the supervisor is initialized. On the other hand, if you want to use another load device for IPL, other than the one identified in your configuration frame, you must perform the system turn-on procedure with the IPL switch in the MNL position. Then you perform the manual IPL procedure when the automatic IMPL procedure completes. Section 6 explains this procedure in more detail.

5.2. System Turn-On Procedure

To turn on System 80 models 8-20 from a full power-off condition, proceed as follows:

1. Set the MAIN POWER circuit breaker at the rear of the processor cabinet to the ON position.

2. Set the MAIN POWER circuit breaker at the rear of the I/O cabinet to the ON position.

3. If an I/O expansion cabinet is used, set the circuit breaker at the rear of that cabinet to the ON position.

Note: *Usually, the MAIN POWER circuit breakers are left in the ON position and power is controlled with the POWER ON / OFF switch on the control panel. Thus, steps 1, 2, and 3 need not be performed after the first time power is applied to the system.*

4. Set all peripheral devices (except diskettes and workstations) to ON, according to the procedures described in the appropriate hardware reference manual (see Table 4-1). Workstations are turned on and off, as required, by the workstation user. The console POWER ON/OFF switch should always be in the depressed, ON position.

5. Set the IPL, MNL/AUTO switch to the type of IPL procedure you require - either manual (MNL) or automatic (AUTO). Unless you want to use an alternate load device for IPL other than the one identified in your configuration frame, you should always perform automatic IPL.

6. Set the POWER ON/OFF switch, located behind the control panel door, to the ON position to turn on the console, integrated disks, and all diskettes. Your system first lights the SEQ1, SEQ2, SEQ3, and P-SEQ lights on the control panel. When power sequencing is complete, the P-SEQ light turns off and the PWR light turns on. Then your system displays the following messages on the console screen indicating that the IMPL procedure is in progress:

```
HEALTH CHECK RUNNING
BPU IMPL W08START
BPU IMPL W09 START
DMUX IOML START
```

When the IMPL procedure completes, the SEQ1, SEQ2, and SEQ3 lights go off. At this point, the IPL procedure must be performed either manually or automatically.

If the IPL switch on the control panel is in the AUTO position, the system turn-on procedure performs the IPL up until the interactive IPL procedure (6.3.2). It displays the IPL01 screen and turns on the RUN, WAIT, and POWER lights on the control panel. To complete the IPL procedure, see 6.3.2.

If the IPL switch on the control panel is in the MNL position, the system completes the IMPL procedure then displays the configuration frame. In this case, you must perform the manual IPL procedure yourself. The configuration display frame shows the ***SPECIAL MODE*** AUTO IPL as OFF and the STOP and POWER lights on the control panel light. To complete the IPL procedure, see 6.3.1.

Note: *If the IPL switch is in the AUTO position and your system encounters an error during the automatic IPL procedure, your system will display the configuration display frame. The frame will indicate what caused the error in the form of a message in the upper-right corner. To complete the IPL procedure, see 6.3.1.*

5.3. System Turn-Off Procedure

The turn-off procedure for System 80 models 8-20 turns off the console, the integrated disk, and all diskette subsystems; turns off the peripheral devices; then turns off the power to the system.

To turn off and remove power from System 80 models 8-20, proceed as follows:

1. Set the POWER ON/OFF switch, located behind the control panel door, to the OFF position to turn off the console, integrated disks, and all diskettes. The console screen clears and the POWER light on the console keyboard goes off. All lights on the control panel go off also.

2. Set all peripheral devices (except diskettes, workstations, and the console) to OFF, according to the procedures described in their respective hardware reference manual (see Table 4-1). Workstations are turned on and off, as required, by the workstation user. Be sure to leave the console POWER ON/OFF switch in the depressed, ON position.

Note: *Step 3 need not be performed unless complete power turn-off is desired due to an emergency or if the step is performed as a normal site requirement.*

3. Set the circuit breakers at the rear of the processor, the I/O cabinet, and the I/O expansion cabinet to the OFF position.

Section 6
Models 8-20 System Initialization Procedures

6.1. Introduction

Initializing your system means preparing it for normal operation. There are two steps or procedures that do this. They are:

- Initial microprogram load (IMPL) procedure. This procedure loads the microinstructions into control storage. These microinstructions (also called microcode) control the basic operation of your system.

- Initial program load (IPL) procedure. This procedure loads the resident portion of the OS/3 operating system (called the supervisor) into main storage.

There are three occasions when you initialize your system. First and foremost is when you turn the system on. (For details about the power-on procedures, see Section 5.) During turn-on, the system does most of the work for you. It automatically performs the IMPL procedure and, optionally, starts the IPL procedure. If you set the IPL switch on the processor cabinet to AUTO, you don't do anything until the system reaches a stage called the interactive IPL. (If the switch is set to MNL, you must initiate the IPL procedure yourself.)

The other two occasions that require some form of initialization occur during day-to-day operations. They are:

1. When a nonrecoverable error occurs while the system is running (such as a CHECK STOP, when the C-STOP indicator lights)

2. When you want to use a different supervisor (possibly on another device).

In both of these situations, you must initialize the system yourself. It is not automatic.

When you perform both the IMPL and IPL procedures, you're completely reinitializing the system. This isn't always required. In fact, more often than not, the only procedure you need is the IPL procedure. For instance, when you need an alternate supervisor, all you're changing is the resident portion of OS/3. Since you're not changing the microcode, it isn't necessary to IMPL.

Similarly, many nonrecoverable errors only require you to re-IPL. You rarely have to IMPL once you turn the system on. When a nonrecoverable error does occur, the system displays an error message on the manual frame, or lights an indicator on the processor cabinet. Refer to the error message in the system messages manual. This manual explains what you should do. If it tells you to reboot the system, it still might only require the IPL procedure. However, it's a good rule of thumb to perform both the IMPL and IPL procedures in this case. The IMPL procedure is easy and fast and ensures that the system is properly loaded.

If an error occurs during either the IMPL or IPL procedure, then you should always start again with the IMPL procedure.

The IPL procedure is a two-phase operation. In the first phase you initiate the procedure by indicating which device you're loading your supervisor from. The second phase is the interactive phase. During this phase you supply the name of your supervisor and you select various load options. As previously mentioned, if you're powering your system up, and the IPL switch is set to AUTO, the system itself does all the work up to this interactive phase.

6.2. Initial Microprogram Load Procedure

The IMPL procedure loads your system's control storage. Once it's loaded, it stays in control storage until you turn the system off or you choose to re-IMPL.

When you're powering your system up, the IMPL is the first initialization step and the system does it for you. It loads the microcode from the diskette mounted in the diskette drive behind the processor cabinet door. This drive is reserved for the microcode diskette; it should be mounted at all times.

To perform the IMPL procedure, follow these steps:

1. To display the manual frame, press the ESC key, then the M key. The manual frame appears:

```
  ENTER.

PROG LOAD(Ø-BFF)               CONTROL
N NORMAL                       U RUN
C CLEAR                        Q STOP
G NORMAL (NO RESET)
T IMPL                         RESET
  Ø ALL                        L SYSTEM
  1 BPU                        B BPU
  2 DMUX
                               COMPUTE CTRL
```

2. Enter the T option, followed by how much of the microcode you want to reload. Press XMIT.

   You have the option of loading all the microcode (0), just the microcode for the BPU (1), or the microcode for the input/output system (2).

   For example, to load all the microcode, complete the manual frame as follows:

```
  ENTER.TØ

PROG LOAD(Ø-BFF)               CONTROL
N NORMAL                       U RUN
C CLEAR                        Q STOP
G NORMAL (NO RESET)
T IMPL                         RESET
  Ø ALL                        L SYSTEM
  1 BPU                        B BPU
  2 DMUX
                               COMPUTER CTRL
```

   Press XMIT.

   While the IMPL procedure is running, the SEQ1, SEQ2, and SEQ3 indicators light briefly. When the procedure completes, the system redisplays the manual frame. It displays your previous IMPL selection on line 2:

```
ENTER.>TØ<

PROG LOAD(Ø-BFF)               CONTROL
N NORMAL                       U RUN
C CLEAR                        Q STOP
G NORMAL (NO RESET)
T IMPL                         RESET
  Ø ALL                        L SYSTEM
  1 BPU                        B BPU
  2 DMUX
                               COMPUTE CTRL
```

   You're ready to perform the initial program load (IPL) procedure.

6.3. Initial Program Load Procedure

The IPL procedure loads and initializes the part of the OS/3 operating system that stays in main storage while the system is running. This part of the operating system is called the supervisor.

When you IPL, you're telling the system where to find the supervisor and what options it should initialize it with. Since the procedure does both of these things, we split the discussion into two phases: the initiation phase (6.3.1) and the interactive (6.3.2) phase.

Note: *Because models 8-20 disk control units need loadable microcode, you must load IPL via the same control unit that contains the SYSRES devices.*

6.3.1. Initiation Phase

When you apply power to your system, you have the option of automatically initiating the IPL procedure. You simply set the IPL switch on the processor cabinet to AUTO. When the IMPL procedure completes, the system begins the IPL procedure, bringing it up to the interactive phase.

But powering up is the only time the procedure initiates automatically and the only time when the IPL switch is effective. All other times you must start and finish the IPL procedure yourself. The occasions when you do this include:

- STOPPED*HPR* message is displayed indicating an non-recoverable error occurred.
- LOAD*ERROR STOP* message appears during an IPL attempt.
- CHECK STOP error occurs (the C-STOP indicator on the processor cabinet lights).
- You need an alternate supervisor.

To initiate IPL manually, follow steps 1, 2, and 3:

1. Make sure no user jobs are active. The system must be idle before you initiate the IPL procedure. If the system just completed the IMPL procedure for you, then the configuration frame is displayed. If you've just completed the IMPL procedure, go on to step 3.

2. Press the ESC key, then the M key to display the manual frame. On the manual frame, you indicate that you're resetting the system by keying in L:

```
ENTER.L

PROG LOAD(Ø-BFF)              CONTROL
N NORMAL                      U RUN
C CLEAR                       Q STOP
G NORMAL (NO RESET)
T IMPL                        RESET
  Ø ALL                       L SYSTEM
  1 BPU                       B BPU
  2 DMUX
                              COMPUTE CTRL
```

Press XMIT. The manual frame reappears, echoing your reset choice on the second line.

3. Select a normal (N) program load (PROG LOAD) along with the address of the device where your supervisor resides.

   If you are loading from an autoload diskette, select the G option to prevent ejection of the diskette.

   For example, if your supervisor resides on a disk drive with the address of 180, you complete the manual frame as follows:

```
ENTER.N18Ø
                  >L<
PROG LOAD (Ø-BFF)             CONTROL
N NORMAL                      U RUN
C CLEAR                       Q STOP
G NORMAL (NO RESET)
T IMPL                        RESET
  Ø ALL                       L SYSTEM
  1 BPU                       B BPU
  2 DMUX
                              COMPUTE CTRL
```

   Press the XMIT key.

   This completes the IPL initiation phase. Go on to the interactive phase (6.3.2).

6.3.2. Interactive IPL - How to Initialize the Supervisor

You're ready to initialize the supervisor. By the time you reach this phase, you've directly or indirectly completed the IPL procedure's initiation phase (6.3.1). The display that appears at your console screen is the IPL01 screen:
Complete this phase by following these steps:

```
01  IPL01 KEY IN AND TRANSMIT
02  SUPERVISOR NAME                              (DEFAULT=SY$STD)
03  OPTION S=STD, D=DEBUG, L=S/A LOAD, Q=QUICK   (DEFAULT=S)
04  LOAD DEVICE ADDRESS                          (DEFAULT=did)
05  SY$STD,S,did
05  ┌
```

1. Select the Supervisor, Load Option, and Load Device

   On the IPL01 screen you supply the name of your supervisor, the load option, and the device where the supervisor resides. There are default values for each of these items. They're displayed on line 05 of the screen:

   ```
   05?SY$STD,S,did
   ```

   where:

   SY$STD
   : This is the default name for your supervisor. When you create a supervisor, you can assign any name to it that you want. If you don't name it, the system assigns the name SY$STD.

   S
   : Specifies the STANDARD, default load option. Use this specification on the initial IPL after receiving a new release or after applying an SMC or SMP. This updates the table of symbiont sizes that is maintained on disk. This update process is time-consuming and is only necessary when a symbiont size has changed.

     The other load options are:

     D
     : Specifies the supervisor DEBUG load option. (For details, see the *Supervisor Technical Overview* (UP-8831).)

     L
     : Specifies the special supervisor load for stand-alone programs.

Q

Specifies the QUICK load option, which should normally be used, because your IPL will be significantly faster.

Note: *This option cannot be used if an SMC has been applied or if a new supervisor or ICAM has been generated.*

did

This is the device address where the supervisor resides. If you manually initiated the IPL procedure (6.3.1), the system takes the device address you entered on the manual frame and displays it here. For example, if you entered 180 on the manual frame, 180 appears as the did on line 05 of the IPL01 screen.

If the system performed the IPL initiation step for you, it displays the address of the volume named as SYSRES on the configuration frame.

If you accept all the defaults, press XMIT; no other key-in is required. Go on to the next step.

If not, enter the values you wish to use. (The cursor is already positioned where you enter your selections.) Press XMIT. The system displays another screen for you to fill in. See the next step.

For example, if the name of your supervisor is SUP140 and it's located on device 180, you complete the IPL01 screen as follows:

```
Ø1  IPLØ1 KEY IN AND TRANSMIT
Ø2  SUPERVISOR NAME                                (DEFAULT=SY$STD)
Ø3  OPTION S=SUPERVISOR,D=DEBUG,L=S/A LOAD         (DEFAULT=S)
Ø4  LOAD DEVICE ADDRESS                            (DEFAULT=18Ø)
Ø5  SY$STD,S,18Ø
Ø5  SUP14Ø,Q
```

Press XMIT.

2. Complete the System Supervisor Initialization Screen

The system displays the OS/3 supervisor initialization screen. This screen displays the system date and time and some of the system generation (SYSGEN) selections. Each of these lines is identified by a letter (A through F) before the line. From this screen, you can either accept all of the default values shown, or you can change them. To accept all the defaults on the screen, press XMIT. Then go to step 3.

Here is what the screen looks like:

```
     *** OS/3 VERSION xx.xx.xx supname SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:      HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE  (N,Y,H       ) DEFAULT=N
     ERROR LOG  (N,Y         ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDR:  XXX (SYSRES)
F)  DUMP FILE DVC ADDR: XXX (SYSRES)








TRANSMIT TO ENTER DEFAULTS. OTHERWISE ENTER THE
LETTERS OF THE PARAMETERS(S) TO BE CHANGED
▶_____
```

If your system includes multivolume spooling files, line E is followed by these lines:

```
E)SPOOLING DVC ADDR:
    DEFAULT #1=XXX(SYSRES)  DEFAULT #2=___(      )
    DEFAULT #3=___(      )  DEFAULT #4=___(      )
    DEFAULT #5=___(      )  DEFAULT #6=___(      )
    DEFAULT #7=___(      )  DEFAULT #8=___(      )
```

If you're loading a supervisor configured without spooling, line E does not appear at all.

To change any of the default values shown, enter the letter of the line on which that default appears in the space after the cursor. You can enter the letters one at a time or all at once. Then press the XMIT key. The following paragraphs describe each of the five lines and how you specify different defaults for the values they give.

- OS/3 VERSION Line

  Indicates the release version (*xx.xx.x*) of OS/3 and the supervisor you specified on the IPL01 message screen as the supervisor being loaded. This line cannot be changed.

- Line A

  Shows the date and time supplied by the hardware clock. Usually, you'll have no need to change this information. However, if you do want to change it, enter letter A on the screen and press the XMIT key. Your system then displays the following screen:

```
    *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:    HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE  (N,Y,H      ) DEFAULT=N
     ERROR LOG  (N,Y        ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDR:  XXX (SYSRES)
F)  DUMP FILE DVC ADDR: XXX (SYSRES)






ENTER DATE, TIME sysgen date format, HHMMSS
-- -- --, -- -- --
```

  On this screen, enter the date in the format shown. The format is the one you specified at SYSGEN. YY is the year, MM the month, and DD for day. Then press the XMIT key.

  The date you enter is compared to the date keyed in at the last load from the same SYSRES. If the date entered is less than or six days greater than the date of the last load, the message DATE QUESTIONABLE appears and the cursor returns to the date field. If the date you entered is correct, press the XMIT key and the date is accepted. If the date is incorrect, key in the correct date and press the XMIT key.

  Also on this screen, enter the time in the format shown. Then press the XMIT key.

- Line B

  Shows SYSRES as the default address for the disk containing the system job run library YRUN. If you want to locate YRUN on a different volume, enter letter B on the screen and press the XMIT key. Your system then displays the following screen:

```
      *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:   HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE  (N,Y,H       ) DEFAULT=N
     ERROR LOG  (N,Y         ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDR:  XXX (SYSRES)
F)  DUMP FILE DVC ADDR: XXX (SYSRES)






ENTER RUN LIBS DVC ADDR
---
```

Enter the device address of the disk volume to contain YRUN. Remember, the volume you specify must be online for the system to be operational. If it is not, you'll receive two system messages on your console screen telling you to mount the volume.

- Line C

  Shows the default values for the SYSGEN options concerning file recovery of the job queue, error log, and spool file (if you configured your system with spooling). To change any of these values, enter C on the screen and press the XMIT key. Your system then displays the following screen:

```
     *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:    HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE  (N,Y,H      ) DEFAULT=N
     ERROR LOG  (N,Y        ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDR:  XXX (SYSRES)
F)  DUMP FILE DVC ADDR: XXX (SYSRES)









ENTER NEW VALUES FOR THE JOB QUEUE, ERROR LOG, SPOOL FILES: XXX
---
```

Enter the recovery option you want for each of the three files. The first blank contains your choice for the job queue file, the second blank your choice for the error log file, and the third blank your choice for the spool file. When choosing these options, consider the following information concerning these three files:

- Job Queue File

  To delete jobs previously filed in the job queue, accept the default value of N for no recovery by not entering a value in the first blank space. To retain these jobs for processing, key in Y in the first blank space. To place them into hold status, key in H. Then press the XMIT key.

- Error Log File

  To retain the present error log file and continue to list errors in the error log, accept the default value of Y for recovery by not entering a value in the second blank space. To clear all accumulated errors in the error log and start a new error log file, key in N for no recovery. Then press the XMIT key.

- Spool Files

  If previously spooled input and output was processed before the system was turned off and the spool file is empty, accept the default value of N for no recovery by not entering a value in the third blank space.

If you want to recover all spooled subfiles when the spool file is reinitialized because previously spooled input or output files in the spool file are to be processed, key in A for all files in the third blank space. With this response, all spooled subfiles, whether complete or incomplete, are saved. This response (or H) must be specified to recover the console log file and workstation log file if configured into the system). Console and workstation messages that were not copied from the main storage buffer to the spool file are not recovered; copying is done only when the buffer has been filled.

To recover only completed subfiles when the spool file is initialized, key in C for completed in the third blank space and press the XMIT key.

To recover only the user log directory when the spool file is initialized, key in L for log directory in the third blank space and press the XMIT key.

If you took a system dump of the previously loaded system, key in H for hot start in the third blank space and press the XMIT key. With this option, when the spool file is reinitialized, all spool subfiles (same as with A option), as well as console messages that have been accumulated in the main storage buffer, but not copied onto the spool file, are recovered. Copying is done only when the buffer has been filled. (Workstation messages that accumulated in the main storage buffer are not recovered, however.) Use this method of recovery, sometimes called a hot start, only if a system crash occurs. When there is a planned shutdown of the system and you plan to recover spool files later, breakpoint the console log, then select the A, C, or L spool recovery option when you reinitialize your system.

- Line D

This line asks whether you want to modify the supervisor in any way that you specified at SYSGEN. If you do, enter D on the OS/3 VERSION screen and press the XMIT key. Your system then displays the following screen:

```
     *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:    HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE   (N,Y,H       ) DEFAULT=N
     ERROR LOG   (N,Y         ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDR:  XXX (SYSRES)
F)  DUMP FILE DVC ADDR: XXX (SYSRES)




ENTER MODIFY SUPERVISOR OPTION (Y/N)
---
```

This screen actually tells your system whether or not you want to modify your supervisor. If not, key in N for no and press the XMIT key. If you do, key in Y for yes and press the XMIT key and the screen clears and your system displays output messages. Proceed with the appropriate operator action described in Appendix B. Then come back to this procedure.

- Line E (Assuming your supervisor is configured with spooling)

  No matter whether your system is configured with single-volume or multivolume spooling, line E shows your SYSRES (or the volume set at SYSGEN) as the volume containing your spooling volume. If you want to specify another device as your spooling volume, or if you have multivolume spooling and want to specify more than one disk volume as your spool volume, enter E on the OS/3 VERSION screen. Your system then displays one of the following two screens depending on whether you have single-volume or multivolume spooling:

Single-Volume Spooling

```
     *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:     HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE  (N,Y,H       ) DEFAULT=N
     ERROR LOG  (N,Y         ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDR:  XXX (SYSRES)
F)  DUMP FILE DVC ADDR: XXX (SYSRES)




 ENTER DUMP FILE DVC ADDR
 ---
```

On this screen, enter the device address of the single volume other than SYSRES that you want to contain your spooling volume and press the XMIT key.

Multivolume Spooling

```
     *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format   TIME:    HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE   (N,Y,H        ) DEFAULT=N
     ERROR LOG   (N,Y          ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H ) DEFAULT=N
D)  MODIFY SUPERVISOR?  (N,Y) DEFAULT=N
E)  SPOOLING DVC ADDRS:
      DEFAULT #1=XXX(SYSRES)     DEFAULT #2=___(       )
      DEFAULT #3=___             DEFAULT #4 ___(       )
      DEFAULT #5 ___             DEFAULT #6 ___(       )
      DEFAULT #7 ___             DEFAULT #8 ___(       )
F)  DUMP FILE DVC ADDR: XXX (SYSRES)




ENTER SPOOLING DVC ADDRS
#1  #2   #3  #4  #5  #6  #7  #8
___,___,___,___,___,___,___,___
```

Enter the device addresses of any of eight spooling devices in your system. Then press the XMIT key. Remember that all volumes identified as spool volumes at IPL must be online for your system to be operational.

- Line F

Shows SYSRES as the default address for the disk containing YDUMP. If you want to specify an alternate dump file on another disk, enter letter F on the screen and press the XMIT key. Your system then displays the following screen:

```
   *** OS/3 VERSION xx.xx.xx supnam SUPERVISOR INITIALIZATION ***
A)  DATE: sysgen date format  TIME:    HH:MM:SS
B)  RUN LIBS DVC ADDR:  XXX (SYSRES)
C)  FILE RECOVERY
     JOB QUEUE  (N,Y,H         ) DEFAULT=N
     ERROR LOG  (N,Y           ) DEFAULT=Y
     SPOOL FILES (N,A,C,L,H    ) DEFAULT=N
D)   MODIFY SUPERVISOR?  (N,Y)   DEFAULT=N
E)   SPOOLING DVC ADDR:  XXX (SYSRES)
F)   DUMP FILE DVC ADDR: XXX (SYSRES)









ENTER DUMP FILE DVC ADDR
---
```

Enter the device address of the disk to contain YDUMP. During system initialization, YDUMP will be allocated on this disk when necessary. However, SYSRES will always contain a minimum YDUMP file for dumps taken during system initialization.

Note: *When you are using an alternate dump file and YDUMP is extended on SYSRES, you can use a system job stream, SG$SADMP, to scratch the original YDUMP and reallocate and format a minimum YDUMP on SYSRES. The device containing the YDUMP file does not need to be allocated for any user.*

3. Specify Disk Cache

At this stage of the IPL procedure, your system displays the following message:

```
CM01 ENTER THE NUMBER OF 1024 BYTE BLOCKS OF MEMORY FOR DISK
CACHE OR NONE. VALID VALUES IN THE RANGE OF 160-8192?
```

Your response to this message must be one of the following:

$$\left\{\begin{array}{l}\texttt{NONE}\\ \texttt{xxxx}\end{array}\right\}$$

where:

NONE
: Disables disk cache operation.

xxxx
: Is a 4-digit value, in the range of 160 to 8,192, specifying the number of 1,024-byte blocks to be assigned.

Note: *See Section 12 and Appendix C for more information on the disk cache facility.*

4. System Ready Status

After responding to the system requests in the previous steps, your system loads the supervisor you selected and initializes your system. When it has completed the IPL process, it displays the following message, along with job queue and error log messages when appropriate. (If you have configured an alternate console, this message will appear on the designated device.) Your system is now ready to process user jobs.

```
SYSTEM 80 OS/3 version-no supnam yy/mm/dd  hh:mm:ss
```

where:

version-no
: Specifies the OS/3 software version loaded in the system

supnam
: Specifies the name of the supervisor loaded into the system

yy/mm/dd
: Specifies the year, month, and day used by the system and normally tracked by the hardware clock.

hh:mm:ss
: Specifies the time in hours, minutes, and seconds used by the system.

If you selected Y or H for the JOB QUEUE message in step 4, your system also displays the following message:

```
JOB QUEUE RECOVERED- n JOBS QUEUED
```

where:

n JOBS QUEUED
Specifies the number of jobs in the scheduling queues.

If you selected N for the ERROR LOG message in step 4, your system also displays the following message:

```
ERROR LOG NOT RECOVERED
```

If the IPL procedure is not successful, the preceding messages do not appear. Instead, you may receive an IPLnn error message, or the STOP or CHECK STOP indicators on the control panel light. Take the corrective action for the IPLnn error messages described in the *System Messages Reference Manual* (UP-8076).

If the STOP or CHECK STOP indicators on the control panel light, your system has stopped and the load procedure has halted. In this case, you must display the manual frame to determine why the machine stopped. Do this by pressing the ESC key, releasing it, then pressing M for the manual frame.

The console screen clears at this point, then displays the following manual frame. On the second line from the bottom of the screen, the status message (BPU STATUS=S) appears. Directly under this line, the manual frame displays INS and the number of the instruction that caused your machine to stop and the program status word (PSW) identification. Check the instruction number in the *System Messages Reference Manual* (UP-8076). This will explain why your machine stopped and gives the corrective action you should take to continue the IPL procedure.

The sample frame display shows the BPU STATUS as stopped (S) and shows that instruction 41D0C8C94820 caused the IPL procedure to halt and the system to stop.

```
ENTER.
PROG LOAD(0-BFF)                 CONTROL
N NORMAL                         U RUN
C CLEAR                          Q STOP
G (NO I/O RESET)
T IMPL                           RESET
  0 ALL                          L SYSTEM
  1 BPU                          B BPU
  2 DMUX
                                 COMPUTE
R OPER RECOVERY(0-BFF)           P NORMAL
                                 S STEP
ADR STOP(0-FFFFFFFF)
M NORMAL                         CHECK CTRL
I STOP IA                        H NORMAL
A STOP ANY                       K STOP
                                 V STORE STATUS
BPU STATUS=S           HPR=99080801     00000000
INS=41D0C8C94820   PSW=00000000  3003A2D8       RLR=00000000
```

Note: See the System 80 Model 8 Operator's Maintenance Guide *(UP-9710) for a complete description of the manual frame.*

If you configured your supervisor with the transient work area feature, your system should display the message TRANSIENT WORK AREA IS INITIALIZED as soon as the IPL procedure completes. If it doesn't, initiate the transient work area feature yourself at this time. To do so, issue the TW command (see 8.5.16).

During the IPL operation, the system automatically sets all devices or subsystems, except workstations, not online (not turned on) to not available. The devices or subsystems are not available for system use until you identify them as available via the SET IO command for a disk drive, until you mount a disk pack and initialize the drive. You can review the availability status for devices by using the MIX command. Jobs requiring more than the available devices are terminated with an R277 message.

Section 7
OS/3 Operating System

OS/3 is composed of a group of 11 major programs - supervisor, job control, data management, integrated communications access method (ICAM), utility programs, information management system (IMS), language processors, data base management system (DMS), application programs, diagnostic programs, and interactive services. (Figure 7-1). These programs are discussed in 7.1 through 7.11.

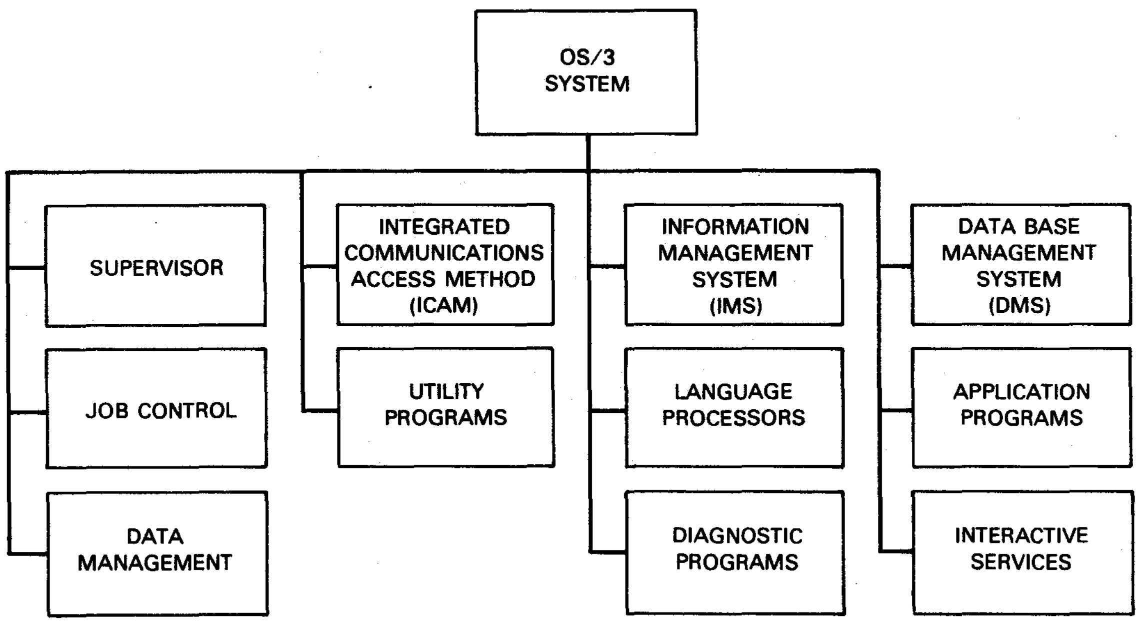

Figure 7-1. OS/3 Components

7.1. Supervisor

The supervisor is a resident program of the operating system; it interfaces with the user programs to provide the necessary control for the best use of the system hardware and software. It controls the physical I/O operations, system resource allocation on a dynamic basis, task switching to achieve multitasking, hardware interrupt servicing, communications with the system operator, and interface to user programs with the system hardware. To accomplish this, the supervisor is composed of these program elements:

- Interrupt control
- Priority control
- Transient control
- Physical I/O controls
- Resource allocation
- Task control
- Interrupt timer and day clock services
- Program management
- Console workstation management
- File services
- Program error handling
- Cooperative/symbiont operations
- Physical input/output control system (PIOCS)
- Debugging aids

7.2. Job Control

Job control is a nonresident program of the operating system responsible for controlling the orderly initiation and termination of jobs within a multiprogrammed environment. The job control services are performed prior to execution of the initial job step of a job, during the transition between job steps, and at the conclusion of a job. Some of the services of job control are:

- Volume label and file label storage
- Job control stream file maintenance

- Job scheduling by priority
- Main storage allocation and reallocation
- Peripheral device assignment
- Program restart

The functions of job control are implemented interactively or in batch mode by the programmer through the job control language. These sequenced control statements form the control stream that defines facility requirements of a job and directs the execution of the job. The job control statements, through the job control stream, function as an interface between the programmer and OS/3.

7.3. Data Management

Data management provides a consolidated data interface between the hardware-oriented I/O facility (card, tape, printer, and MIRAM disk files) and the user program. The data management facilities consist of logical input/output control stream (IOCS) modules, transient routines, declarative macroinstructions, and imperative macroinstructions.

7.4. Integrated Communications Access Method

The integrated communications access method (ICAM) may be used to control local workstations directly connected to the processor and other ICAM-supported terminals connected to the processor by means of communications lines. ICAM is loaded as a symbiont and acts as an extension to the System 80 supervisor.

Application programs communicate with ICAM by means of the following interfaces using imperative and declarative macroinstructions:

- CPI - communications physical interface
- DDI - direct data interface
- STDMCP - standard GET/PUT interface

7.5. Utility Programs

The OS/3 utility programs make available to the system the means for sorting data into a specified order, merging data to facilitate processing, maintaining files on magnetic disk storage, linking output modules of language processors into executable programs, and copying input cards, magnetic tapes, disk or diskette files to any other card, magnetic tape, disk or diskette, or printer device.

The major utility programs include:

- Data utilities
- Linkage editor
- System librarian
- Sort/merge
- Disk, diskette, and tape prep routines
- System utility

7.6. Information Management System

The information management system (IMS) is a transaction-oriented application package that facilitates access to information stored in data files. IMS provides a terminal-oriented data retrieval and update capability for managerial and clerical personnel and, thereby, relieves them of needing to learn complex methods employed by programming personnel.

7.7. Language Processors

Eight language processors are available with OS/3: basic assembly language (BAL), COBOL, FORTRAN IV, FORTRAN 77, Pascal, report program generator II (RPG II), BASIC, and ESCORT. Language processor input can be on punched cards, disk or diskette files, or magnetic tape; output can be recorded on disk or diskette files or on magnetic tape. All processor output is in a common system output format.

7.8. Diagnostic Programs

The diagnostic programs provided with OS/3 are hardware verification and maintenance routines that can be executed concurrently with user programs. These programs are intended to be run as confidence tests by the system operator and as diagnostic and maintenance tests by the customer engineer.

7.9. Data Base Management System

The data base management system (DMS) is a collection of system programs that support integrated data bases. These programs provide for the description, initialization, creation, accessing, maintenance, backup, and recovery of data base. The languages used in the description and manipulation of DMS data bases are derived from the CODASYL data base specifications. A data base may be accessed by batch application programs and communications application programs.

7.10. Application Programs

Application programs are those specialized programs that are available to a user but not provided as part of the standard software package. These programs are directed towards handling problems distinctive to a particular user and include program evaluation and review techniques/critical path method analysis (PERT/CPM) and linear programming (LP).

7.11. Interactive Services

Interactive services provides an extensive interactive command set that is available to a user as part of the standard software package. These commands enable you to control the interactive system environment, including all jobs within the system, all workstation users (local locations), and all terminal users (local and remote locations).

Interactive services also perform a variety of functions including program creation and file/data manipulation. The interactive facilities used to perform these functions are:

- Screen format services
- Interactive data utilities
- General editor
- Interactive job stream preparation
- Menu services
- Interactive dump/restore facility
- ESCORT
- BASIC

Section 8
Job Processing Procedures

After you successfully initialize your system, the system is ready to process user jobs. To begin processing, place the input media (cards, diskette, or disk) containing the job control stream (JCS) to be processed (or the input data required by a prefiled job to be processed) on the input device. Then proceed to the console workstation or console to initiate running the job, using the system commands described in this section. When card input is used, press the RUN switch on the card reader after you enter the command and the I/O stop state message is displayed.

Notes:

1. *To demonstrate that all installed products in the system are operational, refer to the* Installation Verification Procedures Operating Guide *(UP-8820).*

2. *Refer to About This Guide for the conventions used to illustrate the commands and messages appearing in this guide.*

8.1. Keyin Procedure

8.1.1. Models 3-6 Console Workstation

Before keying in any message or command on the console workstation for transmittal to the system, you must press the UNLOCK MSG WAIT key on the console workstation. Pressing this key:

- Notifies the system that you want to send a message or command to it

- Automatically opens a line on the console workstation screen for display of the input command or message

- Gives the start of entry (▷) signal

- Sets the cursor (◪) to the position where typing is to start

After keying in the message or command, press the XMIT key on the console workstation keyboard. This initiates transmittal of the message or command to the system.

After you enter a command or message, your keyboard locks and a *WAIT* message is displayed on the system state line (line 17) until the system accepts the command or message for processing. Once your entry is accepted, the *WAIT* message is cleared from line 17 and your keyboard unlocks to permit another entry.

The accepted command (or unsolicited message used instead of a command, as we describe in 8.3.3) is either acted upon immediately or placed on a queue for future processing; it is never ignored or lost. Commands are placed on queue when insufficient main storage exists or a required device is not available.

Queued commands are activated as soon as all requirements for execution are met. Therefore, commands keyed in twice are eventually executed twice. If a message or command is unacceptable, the system responds with a negative acknowledgment (NAK) error message indicating why. This error message appears in the last 12 character positions of the line of the unacceptable message or command and, if necessary, overwrites any message or command text that may be present in the last 12 character positions.

A pair of blinking marker symbols (◥and◥) bracket each error message. The message or command must then be retyped in accordance with the information furnished by the error message, so that the job involved with the unacceptable message or command can be executed. The error messages that can be displayed are described in the *System Messages Reference Manual* (UP-8076).

If console logging is configured in your system, the console log messages are recorded in a spool file for later printing.

You can perform all of the functions described in the *Interactive Services Operating Guide* (UP-9972) when you use the console workstation as a workstation. Press and hold the FUNCTION key, then press the SYS MODE or WS MODE key, to switch from console workstation use to workstation use. To switch from either SYS (system) MODE or WS (workstation) MODE back to a console function mode, press the FUNCTION key, and while holding it down, press the C/CSL (console) key or the D/SYS CONT (control) key for the appropriate console mode.

If you are using the console workstation as a workstation (operating in either SYS or WS MODE), and informational messages develop concerning a console function, the following is displayed on the console workstation indicator line (line 25):

```
CNSMSG
```

To view the messages, press the FUNCTION key, and while holding it down, press the C/CSL key. The messages are displayed on the second line of the screen, and CNSMSG is cleared from the indicator line.

If a message develops concerning a console function that requires a response, or if more than 14 informational messages accumulate, the following is displayed on the console workstation indicator line:

```
CNSREQ
```

The XMIT key is also locked, preventing you from performing any workstation functions until you respond to the console message. When you enter the console mode, the message is displayed, and CNSREQ is cleared from the indicator line. After informational messages are displayed, they roll up and off the screen.

If, prior to entering console mode, you had begun a job in which the console workstation is used as a workstation, informational messages concerning the workstation job could develop. When this occurs, the following is displayed on the console workstation indicator line:

```
WSMSG
```

To view the messages, press the FUNCTION key, and, while holding it down, press the SYS MODE key. The messages are displayed on the second line of the screen, and WSMSG is cleared from the indicator line. If a message develops concerning the workstation job that requires a response, the following is displayed on the console workstation indicator line:

```
WSREQ
```

Note that the console functions are *not* locked by the appearance of this indicator. To view the message and respond to it, press the FUNCTION key and, while holding it down, press the SYS MODE key. The message is displayed, and WSREQ is cleared from the indicator line. Since the console functions are not locked by this indicator, response messages may roll off the screen before you have had a chance to answer them. Should this happen, use the REBUILD command (see 8.5.3) to redisplay them.

8.1.2. Models 8-20 Console Keyin Procedure

Before keying in any message or command on the console for transmittal to the system, you must press the MSG WAIT key on the console. Pressing this key:

- Notifies the system that you want to send a message or command to it
- Automatically opens a line on the console screen for display of the input command or message
- Gives the start of entry (▷) signal
- Sets the cursor (▨) to the position where typing is to start

After keying in the message or command, press the XMIT key on the console keyboard. This initiates transmittal of the message or command to the system. The accepted command (or unsolicited message used instead of a command, as we describe in 8.3.3) is either acted upon immediately or placed on a queue for future processing; it is never ignored or lost. Commands are placed on queue when insufficient main storage exists or a required device is not available.

Queued commands are activated as soon as all requirements for execution are met. Therefore, commands keyed in twice are eventually executed twice. If a message or command is unacceptable, the system responds with a negative acknowledgment (NAK) error message indicating why. This error message appears in the last 12 character positions of the line of the unacceptable message or command and, if necessary, overwrites any message or command text that may be present in the last 12 character positions.

A pair of blinking marker symbols (◥and◥) bracket each error message. The message or command must then be retyped in accordance with the information furnished by the error message, so that the job involved with the unacceptable message or command may be executed. The error messages that may be displayed are described in the *System Messages Reference Manual* (UP-8076).

If console logging is configured in your system, the console log messages are recorded in a spool file for later printing.

8.1.3. Alternate Console Keyin Procedure for All Models

You can designate that a local workstation serve as the alternate console by system generation or the SET command. The local workstation takes the place of the hardware console and becomes the command entry and system control device.

Before keying in any message or command on the alternate console for transmittal to the system, you must press MSG WAIT on the keyboard. Pressing this key:

- Notifies the system that you want to send a message or command to it
- Automatically opens a line on the console screen for display of the input command message
- Positions the cursor at the beginning of the input line to allow you to type in your command

After keying in the command or message, press the XMIT key to enter it into the system. The system will exit the system line and repeat your input at the bottom of the screen. Input received is handled the same way as input from the regular console (see 8.1.1 or 8.1.2).

The spacebar is set as destructive so you need not clear the input line manually each time you use it. Just type in the new text on top of the previous command. The last command entered remains on the system command line so that you may edit and reenter it without having to rekey the entire command.

Although the device is physically a workstation, when it is functioning as the alternate console it cannot run programs or functions that allocate the entire screen (such as the editor or screen format services).

If the alternate console becomes unusable, the system will automatically reactivate the original console. Normally when the alternate console is active, input from the original console is rejected with a negative acknowledgement (NAK) error message.

If you specify a workstation at system generation to be the alternate console, the hardware console is still used to IPL the system. The system switches control to the alternate console when the system is successfully initialized. The hardware console screen is cleared and has the header line:

```
OS/3 INACTIVE CONSOLE - ALTERNATE CONSOLE IS DEVICE ddd
```

You can switch the alternate console to another device or back to the hardware console by using the command:

$$\texttt{SET IO,}\left\{\begin{array}{l}\texttt{did}\\\texttt{ORG}\\\texttt{ALT}\end{array}\right\}\texttt{,CON}$$

where:

`did`
: Is the device address.

`ORG`
: Is the original hardware console.

`ALT`
: Is the system generated.

`CON`
: Is a console change.

8.2. Command Characteristics

When a command is being typed in, there must be at least one space between the command and the first parameter, and commas between all parameters. The general format for these commands is:

$$\triangleright\texttt{command}\left[\left\{\begin{array}{l}\texttt{(did)}\\\texttt{([did],label)}\\\texttt{(RDR,label)}\end{array}\right\}\right]\Delta\texttt{[command-parameters]}$$

where:

$\triangleright$
: Is the start-of-entry (SOE) symbol that must precede all lines. This symbol is automatically generated by OS/3 when the UNLOCK MSG WAIT (models 3-6) or MSG WAIT (models 8-20) key is pressed.

`command`

Is two to eight alphabetic characters that identify the system command to be processed. At least two characters must be supplied.

`(did)`

Is a 3-digit device identification number that identifies the channel, subchannel, and selected device number to be used when carrying out the command. A did should be included when a particular peripheral device is to be specified or when no default option is provided.

If a did is not entered, the first appropriate device is used.

`([did],label)`

When a diskette is used to serve the function of a card reader or card punch, it must be formatted in data set label mode; a file identifier (label) is required to identify the specific file to be accessed on the diskette. The device address also must be included, unless the diskette is configured as the SYSRDR. The label, which may be a maximum of eight alphanumeric characters, is separated from the device address by a comma. The data set label diskette record size must be 128 bytes or less, and the records must be unblocked and unspanned.

`(RDR,label)`

Specifies that the device address to be used is the input spool file with the specified label. The label may be a maximum of eight alphanumeric characters and is separated from the RDR entry by a comma.

Examples

`RUN(010,DATANAM1) JOBNAM1`

Specifies that the diskette with a device address of 010 contains a file labeled DATANAM1 to be accessed by the command.

`RUN(,DATANAM1) JOBNAM1`

Specifies that the diskette configured as SYSRDR contains a file labeled DATANAM1 to be accessed by the command.

`RUN(RDR,DATANAM1) JOBNAM1`

Specifies that the input spool file contains a subfile labeled DATANAM1 to be accessed by the command.

`command-parameters`

Are optional positional parameters used to tailor the effect of the command being issued. Each specified parameter must appear in its own position. Commas are used to separate positions. For example, the following portion of the RUN command format indicates that two positional parameters are associated with the command; the comma separates the first from the second.

```
RUN [ {(did)          } ] [ {jobname [(new-name)]} ] [ , {PRE } ]
    [ {([did],label)  } ] [ {(new-name)          } ] [   {HIGH} ]
    [ {(RDR,label)    } ]                            [   {NOR } ]
```

If only the second parameter is to be specified, the command must be keyed in as follows:

```
RUN ,HIGH
```

As shown, the comma must be included to indicate the omission of the first parameter; otherwise, HIGH is considered the name of the job to be run.

Also, if a parameter takes more than one form (as is the case with the first parameter in the RUN command), the punctuation marks (in this case, parentheses) must be keyed in whenever shown in the format. For example, if a new-name is specified, the command could be keyed in either of the following ways:

```
RUN MYJOB(JOBA),PRE

RUN (TEMPNAME)
```

In either case, the parentheses must be included.

8.3. Message Characteristics

8.3.1. Output Messages

Output messages are displayed on the console workstation or console to provide you with information, to direct you to take some action, or to ask you a question that requires a response. The messages that can be output to the console workstation or console by the components of OS/3 are described in the *System Messages Reference Manual* (UP-8076), together with their associated operator responses, when appropriate. The format of an output message is:

```
nn { ? } { jobnamez }
   { Δ }.{ synnnn   } message-text
   { * } { command  }
```

where:

nn

Is a unique message number from 11 to 99 (numbers 1 through 10 are reserved for other system uses). This number serves as the message identification. Numbers 11 to 42 are used for questions.

?

Identifies an output message that must be responded to before the job that issued the message can continue. Output messages requiring replies are not rolled off the console workstation or console screen until they are answered.

Δ

Identifies an output message that requires no reply or operator action; it gives information only. Input messages, solicited and unsolicited, must include a space between the message-id and message text.

*

Identifies an output message that requires you take some action. The job that generated the message has placed itself in a *yield* state. A GO command is required from the operator to reactivate the job (see 8.4.3).

jobnamez

Is the name of the user job sending the message. This name is the 8-character name taken from the job preamble.

synnnn

Is the name of the symbiont sending the message. This name is the 2-character symbiont identification followed by a 4-digit binary job number inserted at task initiation. It is a 6-character name altogether.

command

Is the interactive command followed by the interactive command ID.

message-text

Comprises the actual message content and is a maximum of 60 characters. The message text includes the message prefix which you use to find the message in the OS/3 manual documenting it - either the *System Messages Reference Manual* (UP-8076) or the OS/3 manual describing the system program that displays the message.

8.3.2. Solicited Input Messages

Solicited messages are those messages you input in direct response to an output message that requires a reply (question mark immediately follows message-id). The format for all solicited input messages is:

```
nnΔmessage-text
```

where:

nn

Is the unique message number of the message you are responding to.

Note: In some cases, an apparent system halt is caused by the operator's failure to answer output messages that require a response. Before deciding that the system is in a halt condition, you should ensure that all output messages have been answered.

8.3.3. Unsolicited Input Messages

Unsolicited messages are those messages you input that are not in direct response to an output message that requires a reply. You may enter unsolicited messages from the console or from the workstation that initiated the job or symbiont. The format for all unsolicited input messages is:

```
⎧ 00Δsynnnn          ⎫
⎪ 00Δsymbname,       ⎪
⎪ 00Δsymid           ⎪  Δmessage text
⎪ 00Δsymid(did)      ⎪
⎨ UNSΔjobname,       ⎬
⎪ UNSΔsymid, S       ⎪
⎪ UNSΔsynnnn         ⎪
⎪ UNSΔsymbname,S     ⎪
⎩ UNSΔsymid(did),S   ⎭
```

where:

`00`
: Specifies that the unsolicited message is for a symbiont.

`UNS`
: Specifies that the unsolicited message is for a user job.

`symbname`
: Is the 8-character name of the symbiont receiving the message. For example, SL$$OW00 is the symbiont name of the output writer.

`symid`
: Is the 2-character symbiont identification used to invoke the symbiont receiving the message. For example, PR is the symbiont identification of the output writer.

`symid(did)`
: Is the 2-character symbiont identification currently used to invoke the symbiont receiving the message, along with the address of the device used or controlled by a specific copy of the symbiont in main storage. For example, PR(160) is the symbiont identification of the output writer using the printer that is at device address 160. The symbiont having this device allocation receives the unsolicited message. If the specified device is not assigned to the symbiont, the unsolicited message is not acknowledged. The did must be enclosed in parentheses.

`jobname`
: Is the name of the user job receiving the unsolicited message.

`s`
: Indicates a symbiont identification.

`synnnn`
: Is the name of the symbiont sending the message. This 6-character name is the 2-character symbiont identification followed by a 4-digit binary job number inserted at task initiation.

`message-text`
: Is the actual text of the message.

Note: *When you key in an unsolicited message to a symbiont, the system task control blocks (TCB) are searched to locate an active symbiont identified by the name specified in the message. When no address is specified, the unsolicited message is transferred to the buffer of the first symbiont active; only the first will get the message. Figure 8-1, line 11 is an example of an unsolicited keyin to a symbiont.*

Following in 8.4 and 8.5 are descriptions of the operator procedures for entering commands, unsolicited messages, and solicited messages according to the function required. Specific command and message formats are included, specifying the appropriate parameters and their order for that particular function.

Figure 8-1 shows typical messages, how OS/3 displays them, and how you respond to them. Note that the messages are not meant to be real; they simply show how you use the message characteristics described in this section.

```
11 PROGRAM7 THIS IS A COMMENT FROM PROGRAM 7
12?PROGRAM1 THIS IS QUESTION FROM PROGRAM 1
12 THIS IS THE ANSWER TO PROGRAM 1
13 PROGRM12 ACTIVATE USER ISLAND CODE FOR PROGRAM 12
14*PROGRAM4 MOUNT DEV=44Ø VSN=DSP614 LU=Ø5Ø DEV=441 VSN=DSP633 LU=Ø51
15*PROGRAM4 MOUNT DEV=442 VSN=DSP554 LU=Ø52 GO?
GO PROGRAM4
DISPLAY 1ØØ
16?DIØØ32 THIS IS A QUESTION FROM THE 'DISPLAY' SYMBIONT
16 THIS IS THE REPONSE TO DIØØ32'S QUESTION
ØØ PR ST
CANCEL PR,S
CANCEL PROGRAM12,D
17 PROGRM12  JCØ3  JOB PROGRM 12 TERMINATED ABNORMALLY. ERR CODE Ø49
```

Figure 8-1. Typical Messages

8.3A. Job Processing During the Daily Memory Refresh Routine

Model 10, 15, and 20 systems automatically run a memory refresh routine every day at midnight. When initiated, this routine accesses four megabytes of main storage and causes a hardware correction mechanism to correct any single-bit errors that are encountered. This operation increases overall system performance since it reduces the possibility that uncorrectable multibit errors will occur. Multibit errors require system recovery operations and may cause halts.

The memory refresh routine automatically starts at midnight and runs for approximately three minutes. During this time, you may notice a slight decrease in job processing performance.

Note: *Refer to 8.3.1. for output message format definition.*

When memory refresh is initiated, the following message is displayed at the system console:

```
nn synnnn MEMORY REFRESH INITIATED
```

When memory refresh is completed, the following message is displayed at the system console:

```
nn synnnn MEMORY REFRESH OF xxxxxxx-yyyyyyy COMPLETE
```

where:

xxxxxxx
: Is the starting main storage address.

yyyyyyy
: Is the ending main storage address.

8.4. Job Processing Commands

Job processing commands enable the operator to:

- Read job control streams into the system and assign scheduling priorities to them (job initialization)
- Control jobs awaiting execution within the scheduling priority queues (schedule jobs)
- Control jobs being executed (execute jobs)
- Stop jobs under execution (terminate jobs)

In the job processing commands that follow, scheduling priorities are defined as preemptive, high, normal, or low to specify in what order jobs begin execution. Jobs to be run are placed in one of the four scheduling priority queues:

1. PRE (preemptive) Queue

   Contains jobs to be executed first (i.e., before any jobs assigned HIGH, NOR, or LOW scheduling priority), even if the resources for any HIGH, NOR, or LOW priority jobs are available and the resources for any PRE jobs are not. PRE jobs are always executed first unless they're placed on hold via the HOLD command (see 8.4.2). If rollin/rollout is configured, a PRE job initiated for execution when sufficient main storage is not available may cause HIGH, NOR, or LOW jobs being processed to be rolled out to make main storage space available for the PRE job. Rolled out jobs are rolled in and continue processing when main storage is again available. When rollin/rollout is configured, a maximum of 14 jobs can be in the preemptive queue.

2. HIGH Queue

   Contains jobs to be executed before any jobs assigned a NOR or LOW scheduling priority. HIGH scheduling priority jobs are not executed unless the PRE queue is either empty or placed on hold. HIGH priority jobs are always executed before NOR or LOW priority jobs (unless they're placed on hold), even if the resources for the HIGH job are not available and the resources for the NOR or LOW jobs are. A maximum of 35 jobs can be in the high priority queue.

3. NOR (normal) Queue

   Contains jobs to be executed before any jobs assigned a LOW scheduling priority. NOR scheduling priority jobs are executed only when there are no jobs left in the PRE or HIGH queues or when the queues are placed on hold. NOR scheduling priority is the default for a job control stream and for some of the job scheduling commands (8.4.2). NOR scheduling priority jobs are always executed before the LOW priority jobs (unless they're placed on hold) even if the resources for the NOR job are not available and the resources for the LOW job are. A maximum of 70 jobs can be in the normal priority queue.

4. LOW Queue

   Contains jobs to be executed only when there are no jobs left in the PRE, HIGH, or NOR queues or when the queues are placed on hold. A maximum of 70 jobs can be in the low priority queue.

8.4.1. Job Initialization

Job control streams are read into the system by using one of the job initialization commands (FILE, RUN/RV, and SI/SC). These commands enable you to file the job for future use or to process the job immediately.

When a job is filed for future use, it is placed in the job control stream library (YJCS) file or in an alternate library file, as specified in the operator command.

A job may be initiated for processing from one of three places:

1. Input device (card reader, data set label diskette, or spool file)

2. YJCS or an alternate library where the job is filed

3. YSAVE file where it has been saved in its expanded *run* state via a statement included in the job control stream (// OPTION SAVE or // OPTION NOSCHED statement)

When a job is initiated, it is placed in a scheduling priority queue to await execution. This scheduling priority queue can be specified by the programmer submitting the job in the job control stream itself. You can override this specification by entering another scheduling priority in the job initialization command. The default is to use the normal priority queue.

Filing Job Control Streams (FILE)

The FILE command files jobs and JPROCs, read from an input device, into the permanent JCS library file (YJCS) or an alternate library file. (The alternate file may be a MIRAM or a SAT file, but keep in mind that the RUN processor cannot access data in a MIRAM file.) The input can be from a card reader, a diskette drive (data set label diskette only), or the input spool file.

Note: *The FILE command cannot be issued from an enter stream.*

If no device and label are identified, the first available card reader, as defined when the system was generated, is expected to contain the job control streams and/or JPROCs to be filed. If the job control stream is on a diskette, the label is required; if it is in the input spool file, RDR and label are required (see 8.2).

Jobs filed from the card reader must terminate with a // FIN job control statement. When jobs are filed from diskette or the spool file, the // FIN job control statement is not necessary. Jobs input from diskette to the spool file must be single volume and formatted in data set label mode.

Format

```
FILE [ { (did)          } ] Δ [ { :alt-filename                                            } ]
     [ { ([did],label)  } ]   [ { :(alt-filename, { RES })                                 } ]
     [ { (RDR,label)    } ]   [ {                 { RUN }                                  } ]
                              [ {                 { vsn }                                  } ]
                              [ { :(alt-filename, [ { RES } ] ,write-password)             } ]
                              [ {                 [ { RUN } ]                              } ]
                              [ {                 [ { vsn } ]                              } ]
```

Positional Parameter 1

Identifies an alternate library file where job control streams and/or JPROCs are to be filed. Omit this parameter when no alternate file is required; the job is filed into YJCS. You can use a disk or a format label diskette for the alternate library file.

`:alt-filename`

Specifies the name of the alternate library file, residing on SYSRES, to receive the job and/or JPROC. If the alternate file name is cataloged, the vsn of that file in the catalog is used. There must not be a write password for the alternate file in the catalog.

```
:(alt-filename, { RES })
                { RUN }
                { vsn }
```

Specifies the name of the alternate library file to receive the job and/or JPROC and identifies a volume serial number (RES, RUN, or vsn) for the file. You specify RES to identify SYSRES as the volume to contain the file or RUN to identify the system RUN pack as the volume to contain the file; or you may specify the volume serial number (vsn) of a disk pack or format label diskette to be used. If a file with the same file name is in the catalog, the volume serial number you include in the command makes the distinction between the files and overrides the catalog vsn. There must not be a write password for the alternate file in the catalog. The colon and enclosing parentheses are optional and included only for consistency with the format of the RUN/RV command.

```
: ( alt-filename, [ { RES } ] ,write-password )
                  [ { RUN } ]
                  [ { vsn } ]
```

Specifies the name of the alternate library file to receive the job and/or JPROC and includes the write password, identified in the catalog, required to write to that file. You specify RES to identify SYSRES as the volume to contain the file or RUN to identify the system RUN pack as the volume to contain the file; or you may specify the volume serial number (vsn) of a disk pack or format label diskette to be used. If you omit a volume serial number (RES, RUN, or vsn), your file is written to the volume associated with that file name in the catalog. You must specify RES, RUN, or a vsn if you want the file written to a different volume; the volume serial number you specify overrides the catalog vsn. The colon and enclosing parentheses are optional and included only for consistency with the format of the RUN/RV command.

Note: If you specify a disk or diskette vsn for the alternate library file in the command when the volume isn't already mounted, the system displays a mount message suggesting an available device. The volume can be mounted on the available device or any other suitable device. In rare instances when the system cannot determine whether the volume is a disk or diskette, the message may suggest an available disk drive for a diskette volume. Ignore this inconsistency and mount the diskette volume on any suitable device.

Example

Operator keyin:

```
FI :(ALTJCS,RUN)
```

Function requested:

The job control stream, residing on a card reader, is to be written to the alternate job control library file called ALTJCS on the system RUN pack.

Running Job Control Streams (RUN/RV)

Function

The RUN/RV commands read a job control stream from either an input device, or the YJCS or alternate job control library file. The input device can be a card reader, a diskette drive (data set label diskettes only), or the input spool file.

The commands cause the job control stream to be written to the job run library (YRUN) file and expanded to its *run* state (JPROCs are expanded), then scheduled for execution. YRUN is a temporary file; the job is deleted when it is run.

When an // OPTION SAVE job control statement is included in the job control stream, the job is scheduled to be run from YRUN and a copy of the expanded job is saved in the YSAVE file. When an // OPTION NOSCHED job control statement is included in the job control stream, a copy of the expanded job is saved in YSAVE; however, the job is not scheduled to be run.

The expanded job to be saved in YSAVE or an alternate save file can be run by using the SI/SC commands. (SI/SC commands are described next in this subsection.)

The RV command initiates the reading of a prefiled job control stream that does not contain a // CR statement indicating that there is input (cards, diskette, or spool file) to be read and inserted into the stream. The RUN command initiates the reading of a job control stream that requires an input device (i.e., card reader, diskette, or spool file). This means if the job is initiated from an input device or if the job contains a // CR statement to read input, you must use RUN. When the RUN command is issued, it is accepted only if an input device is available, whether or not one is needed by the job control stream being read. The RV command allows a job control stream to be initiated that does not require an input device. You must include a job name when you enter an RV command.

Remember that, when a system card reader is placed online, the RUN command to read a job control stream from cards in the hopper is initiated when the RUN switch on the card reader is pressed, or when the RUN command is keyed in at the console workstation or console. The RUN command can be initiated from either location, but not from both. If a duplicate RUN command is initiated for the same job, the supervisor queues the second command until the input device is available. Presuming that nothing is in the hopper when the second RUN command is executed, a hopper empty message results.

Format

```
{ RUN [ { (did)           } ] Δ [ { jobname[(new-name)] } ] }
{       { ([did],label)   }       { (new-name)          }     }
{       { (RDR,label)     }                                   }
{                                                             }
{ RVΔjobname;[(new-name)]                                     }

  [ { :alt-filename                                           } ]
  [ { :(alt-filename, { RES })                                } ]
  [ {                 { RUN }                                 } ]
  [ {                 { vsn }                                 } ]
  [ { :(alt-filename, [ { RES } ] ,read-password)             } ]
  [ {                 [ { RUN } ]                             } ]
  [ {                 [ { vsn } ]                             } ]

  [ ,{ PRE  } ] [time + n][,key-1=val-1,...,key-n=val-n]
  [  { HIGH } ]
  [  { NOR  } ]
  [  { LOW  } ]
```

Command Codes

RUN

Initiates the running of a job control stream that requires an input device. You must specify a job name when the job control stream is prefiled. If the job is prefiled and you don't specify an input device, the first available card reader is assigned to the job, whether or not it is required. If you omit an input device and job name, the first available card reader is expected to contain the job to be run.

You must specify a label when the job control stream is on a diskette. If it is in the input spool file, you must specify RDR and a label (see 8.2). For diskette and spool file input, the last // FIN job control statement is not needed because it is used only to terminate card reader operation. However, the // FIN statements that separate groups of card images read with // CR statements are still necessary. Jobs input from diskette to the spool file must be single volume and formatted in data set label mode.

RV

Initiates the running of a prefiled job control stream that does not require an input device; that is, it does not contain a // CR (read card reader, diskette, or spool file) statement. You must specify a job name.

Parameter 1

`jobname[(new-name)]`

Identifies the name of the job to be read from YJCS or an alternate job control library file and stored in a scheduling priority queue to await execution. The job name consists of one to eight alphanumeric characters. The job name is required with RV.

You include *new-name* to assign a new 1- to 8-character alphanumeric name to a job already stored in YJCS or an alternate job control library file. The job identified by the *jobname* parameter is read from YJCS or an alternate file, and stored in a scheduling priority queue under the name identified by the *new-name* parameter to await execution. The new name cannot contain blanks.

`(new-name)`

Used with the RUN command to assign a new 1- to 8-character alphanumeric name to a job input from the card reader. The job is read and stored in a scheduling priority queue under the new name to await execution. The new name cannot contain blanks.

If parameter 1 is omitted from the RUN command, the job is read and stored in a queue under the job name on the // JOB statement in the job control stream.

Parameter 2

Used when the job control stream resides in an alternate job control library file on disk or format label diskette, rather than in YJCS. When the job resides in an alternate library file, this parameter identifies the library file to be read. If omitted, the job is read from YJCS.

```
:alt-filename
```

Specifies the name of the alternate library file, residing on SYSRES, that contains the job. If the alternate file name is cataloged, the vsn of that file in the catalog is used. There must not be a read password for the alternate file in the catalog.

```
: (alt-filename, { RES })
                 { RUN }
                 { vsn }
```

Specifies the name of the alternate library file that contains the command which makes the distinction between the files and overrides the catalog vsn. There must not be a read password for the alternate file in the catalog.

```
: (alt-filename, [{ RES }] ,read-password)
                 [{ RUN }]
                 [{ vsn }]
```

Specifies the name of the alternate library file that contains the job stream and includes the read password, identified in the catalog, required to read from that file. You specify RES to identify SYSRES as the volume that contains the file or RUN to identify the system RUN pack as the volume that contains the file; or you may specify the volume serial number (vsn) of a disk pack or format label diskette to be read. If you omit a volume serial number (RES, RUN, or vsn), your file is read from the volume associated with that file name in the catalog. You must specify RES, RUN, or a vsn if you want the file read from a different volume; the volume serial number you specify overrides the catalog vsn.

Parameter 3

Used to specify scheduling priority and/or time (hour and day).

PRE

Places the job in the preemptive scheduling priority queue to await execution.

HIGH

Places the job in the high scheduling priority queue to await execution.

NOR

Places the job in the normal scheduling priority queue to await execution.

LOW

Places the job in the low scheduling priority queue to await execution.

time

Is a 4-digit number specifying a time of day in military format indicating when execution of the job is to begin.

n

Is an integer from 1 to 9 specifying the number of days after today to start execution of the job.

Parameter 3 can be specified with priority only, time only, or priority and time. If parameter 3 is omitted, the scheduling priority specified in the job control stream is used. If not specified in the job stream, the normal priority is used.

Parameters 4 through *n*

key-1=val-1,...,key-n=val-n

Are the keywords and their values, which may be used by the job being run. The keywords and values must be supplied by the user requesting the job.

Notes:

1. *The total length of all the parameters specified in this command, from the first character of parameter 1 to the last character of the last keyword value specified, is limited to 60 characters.*

2. *If you specify a disk or diskette vsn for the alternate library file in the command when the volume isn't mounted, the system displays a mount message suggesting an available device. The volume can be mounted on the available device or any other suitable device. In rare instances when the system cannot determine whether the volume is a disk or diskette, the message may suggest an available disk drive for a diskette volume. Ignore this inconsistency and mount the diskette volume on any suitable device.*

Examples

1. Operator keyin:

```
RU MYJOB:(ALTJCS,RUN)
```

Function requested:

The job named MYJOB, filed in the alternate job control library file ALTJCS on the system RUN pack, is run under the priority specified in the job control stream. The first available card reader is expected to contain some input for MYJOB, which contains a // CR statement.

2. Operator keyin:

```
RV MYJOBA(NETPAY)
```

Function requested:

The job named MYJOBA, filed in YJCS, is run under the new name NETPAY according to the priority specified in the job control stream.

3. Operator keyin:

```
RV JOBABC:(CPYLIB,PUBDSK)
```

Function requested:

The job named JOBABC, filed in the alternate library file CPYLIB on volume serial number PUBDSK, is placed in the scheduling queue specified in the job control stream.

4. Operator keyin:

```
RV MYJOB,2230
```

Function requested:

Execution of the job, MYJOB, will begin as soon as resources are available after 2230 hours (10:30 p.m.).

5. Operator keyin:

```
RV MYJOB,P2230+1
```

Function requested:

The job, MYJOB, will be placed in the preemptive scheduling priority queue at 10:30 p.m. tomorrow.

Models 3-6 Console Workstation Screen Display

```
JOBABC     (2)     (3)     (4)     (5)     (6)     (7)
  JC06  USING   DEV=104 VSN=PUBDSK
  JC01  JOB JOBABC   EXECUTING JOB STEP LIBS0000 #001 11:52:19
  (1)      (2)     (3)     (4)     (5)     (6)     (7)
  JC02  JOB JOBABC   TERMINATED NORMALLY              11:52:38
```

Explanation:

The job JOBABC is running in job slot (1). The system messages provide device information and the job step currently executing. The job terminated normally and job slot (1) is now available.

Models 8-20 Console Screen Display

```
JC06  USING   DEV=104 VSN=PUBDSK
JC01  JOB JOBABC   EXECUTING JOB STEP LIBS0000 #001 11:52:19
JC02  JOB JOBABC   TERMINATED NORMALLY              11:52:38
```

Explanation:

The system messages provide device information and the job step currently executing. However, the models 8-20 console does not display job slots, so you do not see your job occupying a job slot on the screen.

Running Saved Job Control Streams (SI/SC)

The SI/SC commands initiate the running of a job control stream from the YSAVE MIRAM library file or from an alternate library file, then schedule the job for execution. In either case, the control streams have been saved in their expanded *run* state. A control stream is expanded in the YRUN file when the RUN/RV command is issued for the job. When an // OPTION SAVE or // OPTION NOSCHED job control statement is included in the job control stream, a copy of the expanded control stream is stored in the YSAVE or alternate file for subsequent runs using SI or SC.

The SC command is used only to initiate the reading of a job control stream that does not require an input device to replace embedded data. The SI command initiates the reading of a job control stream that requires an input device (i.e., card reader, data set label diskette, or spool file) to replace embedded data. When the SI command is issued, it is accepted only if an input device is available. The SC command allows a job control stream to be initiated that does not require an input device.

Format

```
{ SI [ { (did            } ] } Δjobname[(new-name)]
{    [ { ([did],label)   } ] }
{    [ { (RDR,label)     } ] }
{ SC                         }

      [ { :alt-filename                                              } ] [ ,{ PRE  } ] [time + n]
      [ { : (alt-filename, { RES })                                  } ] [  { HIGH } ]
      [ {                  { RUN }                                   } ] [  { NOR  } ]
      [ {                  { vsn }                                   } ] [  { LOW  } ]
      [ { : (alt-filename, [ { RES } ] ,read-password)               } ]
      [ {                  [ { RUN } ]                               } ]
      [ {                  [ { vsn } ]                               } ]
```

Command Codes

SI

Initiates the running of a job control stream that requires an input device. If no device and label are identified, the first available card reader, as defined when the system was generated, is expected to contain the replacement data required by the job. If the data is on a diskette, the label is required; if it is in the input spool file, RDR and label are required. (See 8.2.) For diskette and spool file input, the last // FIN job control statement is not necessary.

SC

Initiates the running of a job control stream that does not require an input device to replace embedded data.

Parameter 1

jobname[(new-name)]

Identifies the name of the job to be read from YSAVE or an alternate library file and stored in a scheduling priority queue to await execution. The name consists of one to eight alphanumeric characters.

You include *new-name* to assign a new 1- to 8-character alphanumeric name to a job stored in YSAVE or an alternate file. The job identified by the *jobname* parameter is read from YSAVE or an alternate file and stored in a scheduling priority queue under the name identified by the *new-name* parameter to await execution. The new name cannot contain blanks.

Parameter 2

Used when the job control stream resides in an alternate job control library file on disk or format label diskette, rather than in YSAVE. When the job resides in an alternate library file, this parameter identifies the library file to be read. If omitted, the job is read from YSAVE.

`:alt-filename`

Specifies the name of the alternate library file, residing on SYSRES, that contains the job. If the alternate file name is cataloged, the vsn of that file in the catalog is used. There must not be a read password for the alternate file in the catalog.

$$:\left(\texttt{alt-filename,}\begin{Bmatrix}\texttt{RES}\\ \texttt{RUN}\\ \texttt{vsn}\end{Bmatrix}\right)$$

Specifies the name of the alternate library file that contains the job or JPROC and identifies a volume serial number (RES, RUN, or vsn) for the file. You specify RES to identify SYSRES as the volume that contains the file or RUN to identify the system RUN pack as the volume that contains the file; or you may specify the volume serial number (vsn) of a disk pack or format label diskette to be read. If a file with the same file name is in the catalog, the volume serial number you include in the command makes the distinction between the files and overrides the catalog vsn. There must not be a read password for the alternate file in the catalog.

$$:\left(\texttt{alt-filename,}\left[\begin{Bmatrix}\texttt{RES}\\ \texttt{RUN}\\ \texttt{vsn}\end{Bmatrix}\right]\texttt{,read-password}\right)$$

Specifies the name of the alternate library file that contains the job stream and includes the read password, identified in the catalog, required to read from that file. You specify RES to identify SYSRES as the volume that contains the file or RUN to identify the system RUN pack as the volume that contains the file; or you may specify the volume serial number (vsn) of a disk pack or format label diskette to be read. If you omit a volume serial number (RES, RUN, or vsn), your file is read from the volume associated with that file name in the catalog. You must specify RES, RUN, or a vsn if you want the file read from a different volume; the volume serial number you specify overrides the catalog vsn.

Parameter 3

Used to specify scheduling priority and/or time (hour and day).

`PRE`

Places the job in the preemptive scheduling priority queue to await execution.

HIGH

Places the job in the high scheduling priority queue to await execution.

NOR

Places the job in the normal scheduling priority queue to await execution.

LOW

Places the job in the low scheduling priority queue to await execution.

time

Is a 4-digit number specifying a time of day in military format indicating when execution of the job is to begin.

n

Is an integer from 1 to 9 specifying the number of days after today to start execution of the job.

Parameter 3 can be specified with priority only, time only, or priority and time. If parameter 3 is omitted, the scheduling priority assigned to the job via the job control stream is in effect.

Example

Operator keyin:

```
SI MYJOB(AVGYTD),H2230+1
```

Function requested:

The job called MYJOB, expanded and filed in YSAVE, is run under the new name AVGYTD. The job is to be placed in the high scheduling priority queue to await execution. Replacement embedded data for AVGYTD is expected to be found on the first available card reader after 10:30 p.m. tomorrow.

8.4.2. Job Scheduling

A job is placed in a scheduling priority queue to await the availability of system resources (e.g., main storage, disk drive, printer) to execute that job. While waiting for these resources, you can exercise control over any specific job in a queue, all jobs in a specific queue, and all jobs in all queues, by using job scheduling commands. You may also control, in a similar manner, jobs initiated by a specific workstation user. These commands allow you to:

- Defer jobs from being executed
- Permit jobs to be executed
- Delete jobs from a queue
- Display contents of a queue
- Change a job scheduling priority

Deferring Jobs Scheduled for Execution (HOLD)

The HOLD command permits you to defer the scheduling of jobs according to the command parameters specified. You can defer scheduling for all jobs in all queues or in a specific queue; for a specific job within a queue; for a specific workstation user's jobs in all queues or a specific queue; or for a specific host's jobs in all queues or a specific queue. Scheduling remains deferred until the jobs are reactivated via the BEGIN command.

Format

```
HOLDΔ { JBQ [ , { PRE  } ] } { [ , { OLD } ] [ , { DDP   } ] }
      {         { HIGH }   } { [   { NEW } ] [   { LOCAL } ] }
      {         { NOR  }   } {               [   { RBP   } ] }
      {         { LOW  }   } {               [   { WKSTN } ] }
      {         { ALL  }   } {                               }
      {                    } { [,UID=user-id] [,HOST=host-id] }
      { jobname            }
```

Positional Parameter 1

JBQ

Specifies the command applies to the jobs in a job scheduling priority queue, as further defined by PRE, HIGH, NOR, LOW or ALL, where:

PRE

Defers the jobs in the preemptive scheduling priority queue.

HIGH

Defers the jobs in the high scheduling priority queue.

NOR

Defers the jobs in the normal scheduling priority queue.

LOW

Defers the jobs in the low scheduling priority queue.

ALL

Defers the jobs in all scheduling priority queues.

jobname

Defers a particular job from being scheduled for execution. No further parameters are permitted.

Keyword Parameters

OLD

Defers only jobs already in the scheduling priority queue defined in parameter 1. Jobs subsequently entered in this queue are not deferred.

NEW

Defers only jobs subsequently placed in the scheduling priority queue defined in parameter 1. All existing jobs are still available for execution.

Once a new job is placed in queue, it becomes an old job for any subsequent commands to defer or permit execution. Therefore, the NEW parameter is used with the HOLD command to defer new jobs entering a queue, while the old jobs already residing in the queue remain unchanged and are still able to be scheduled for execution. Likewise, the OLD parameter is used with the HOLD command to defer old jobs already residing in the queue while the new jobs entering the queue are still able to be scheduled for execution. Once a new job enters a queue that is under the influence of a HOLD NEW command, a subsequent command to permit the old jobs in that queue to be scheduled for execution will also release the new job from its deferred status.

Note: *When issued from a workstation, the HOLD console command will hold all OLD jobs on the job queue that were initiated by the same workstation, but will not hold NEW jobs entering the job queue. There is no mechanism at a workstation to hold NEW jobs.*

DDP

Defers only distributed data processing jobs initiated from a remote host.

LOCAL

Defers only locally entered jobs.

RBP

Defers only jobs entered remotely (i.e., from a remote batch terminal).

WKSTN

Defers only jobs either initiated from a workstation or using the // OPTION ORI= or // OPTION MAS= job control statement.

Note: All of the preceding keyword parameters may be interchanged.

UID=user-id

Defers all old jobs associated with a particular workstation user-id in the scheduling priority queue defined in parameter 1. *Jobs subsequently entered with this user-id are not deferred.* The user-id is a 1- to 6-character alphanumeric identification that does not start with Y.

HOST=host-id

Defers all old jobs associated with a particular host in the scheduling priority queue defined in parameter 1. *Jobs subsequently entered with this host-id are not deferred.* The host-id is one to four alphanumeric characters.

Notes:

1. *The UID and HOST parameters may be interchanged. No other keyword parameters are permitted.*
2. *If a command is entered from the system console (local or remote), the HOST parameter should not be used; that site's host-id is used automatically.*
3. *The special user-id OPERATOR indicates that the command applies to console-initiated jobs.*
4. *Any command from the system console (local or remote) pertains only to jobs with a matching host-id regardless of user-id.*
5. *DDP users submitting the HOLD command via a parameterized enter stream may use HOST=$HOST to specify that the host-id of the command's submitter be used.*

Examples

1. Operator keyin:

```
HO JBQ,A,NE,WKSTN
```

Function requested:

All jobs subsequently initiated from workstations and placed in all scheduling priority queues are to be deferred.

2. Operator keyin:

```
HO JBQ,H,UID=WKSTA1
```

Function requested:

All jobs in the high scheduling priority queue associated with the workstation user WKSTA1 are to be deferred.

3. Operator keyin:

```
HO JBQ,N,HOST=B
```

Function requested:

All jobs in the normal priority queue and associated with host B are to be deferred. New jobs entered into this queue are not to be deferred.

Scheduling Deferred Jobs (BEGIN)

The BEGIN command reinstitutes the scheduling for execution of currently deferred jobs, according to the command parameters specified. You can reinstitute scheduling for all jobs in all queues or in a specific queue; for a specific job within a queue; for a specific workstation user's jobs in all queues or a specific queue; or for a specific host's jobs in all queues or a specific queue. Jobs remain deferred by a HOLD command until you enter a BEGIN command to permit their rescheduling for execution.

Format

```
BEGINΔ { JBQ [ , { PRE  } ] { , { OLD }  [ , { DDP   } ] }
       {          { HIGH }   {    { NEW }      { LOCAL } ] }
       {          { NOR  }   {                 { RBP   } ] }
       {          { LOW  }   {                 { WKSTN } ] }
       {          { ALL  }   {                             }
       {                     { [,UID=user-id];[,HOST=host-id] }
       { jobname }
```

Positional Parameter 1

JBQ

Specifies that the command applies to the jobs in a job scheduling priority queue, as further defined by PRE, HIGH, NOR, LOW or ALL, where:

PRE

Permits the jobs in the preemptive scheduling priority queue to be scheduled for execution.

HIGH

Permits the jobs in the high scheduling priority queue to be scheduled for execution.

NOR
: Permits the jobs in the normal scheduling priority queue to be scheduled for execution.

LOW
: Permits the jobs in the low scheduling priority queue to be scheduled for execution.

ALL
: Permits the jobs in all scheduling priority queues to be scheduled for execution.

jobname
: Permits the particular job to be scheduled for execution. No further parameters are permitted.

Keyword Parameters

OLD
: Permits only jobs already in the scheduling priority queue defined in parameter 1 to be scheduled for execution.

NEW
: Permits newly entered jobs to be placed in the scheduling priority queue defined in parameter 1 to be scheduled for execution.

If omitted, both old and new jobs are permitted to be scheduled for execution. (See the HOLD command earlier in this subsection.)

DDP
: Permits only distributed data processing jobs to be scheduled for execution.

LOCAL
: Permits only locally entered jobs to be scheduled for execution.

RBP
: Permits only jobs entered remotely (i.e., from a remote batch terminal) to be scheduled for execution.

WKSTN
: Permits only jobs either initiated from a workstation or using the // OPTION ORI= or // OPTION MAS= job control statement to be scheduled for execution.

Note: *All of the preceding keyword parameters may be interchanged.*

UID=user-id
: Permits all jobs associated with a particular workstation user-id to be scheduled in the priority queue defined in parameter 1. The user-id is a 1- to 6-character alphanumeric identification that does not start with Y.

`HOST=host-id`

Permits all jobs associated with a particular host to be scheduled in the priority queue defined in parameter 1. The host-id is one to four alphanumeric characters.

Notes:

1. *The UID and HOST parameters may be interchanged, but no other parameters are permitted.*
2. *If a command is entered from the system console (local or remote), the HOST parameter should not be used; that site's host-id is used automatically.*
3. *The special user-id OPERATOR indicates that the command applies to console-initiated jobs.*
4. *Any command from the system console (local or remote) pertains only to jobs with a matching host-id regardless of user-id.*

Examples

1. Operator keyin:

   ```
   BE MYJOB
   ```

   Function requested:

   The currently deferred job named MYJOB is permitted to be scheduled for execution.

2. Operator keyin:

   ```
   BE JBQ,UID=WKSTA2
   ```

   Function requested:

   All currently deferred jobs in all scheduling priority queues associated with the workstation user WKSTA2 are permitted to be scheduled for execution.

3. Operator keyin:

   ```
   BE JBQ,H,HOST=ABC
   ```

   Function requested:

   All currently deferred jobs in the high priority queue associated with the host ABC may be scheduled for execution.

The BE command produces the same screen display as the RN/RV command (see 8.4.1).

Deleting Jobs or Symbionts from Scheduling Priority Queues (DELETE)

The DELETE command permits you to delete jobs or symbionts according to the command parameters specified.

You can delete all jobs in all queues or in a specific queue; a specific job within a queue; a specific workstation user's jobs in all queues or a specific queue; or a specific host's jobs in all queues or a specific queue. Only those jobs residing in a scheduling priority queue, and thus waiting to begin execution, can be deleted.

You can delete all queued symbionts or specific symbionts by queue slot ID number. These ID numbers are provided in the MI SQ command display.

Note: The DELETE command cannot be issued from an enter stream.

Format

```
DELETEΔ⎡JBQ, ⎧PRE ⎫⎡⎡ ,⎧DDP  ⎫⎤                  ⎤⎤[,LOG]
       ⎢     ⎪HIGH⎪⎢⎢  ⎪LOCAL⎪⎥                  ⎥⎥
       ⎢     ⎨NOR ⎬⎢⎢  ⎨RBP  ⎬⎥                  ⎥⎥
       ⎢     ⎪LOW ⎪⎢⎣  ⎩WKSTN⎭⎦                  ⎥⎥
       ⎨     ⎩ALL ⎭⎢                             ⎥⎬
       ⎢           ⎣[,UID=user-id][,HOST=host-id]⎦⎥
       ⎢SQ,⎧id[id,...,id]⎫                        ⎥
       ⎢   ⎩ALL          ⎭                        ⎥
       ⎣jobname                                   ⎦
```

Positional Parameter 1

JBQ,

Specifies the command applies to the jobs in a job scheduling priority queue, as further defined by PRE, HIGH, NOR, LOW, or ALL, where:

PRE
: Deletes the jobs in the preemptive scheduling priority queue.

HIGH
: Deletes the jobs in the high scheduling priority queue.

NOR
: Deletes the jobs in the normal scheduling priority queue.

LOW
: Deletes the jobs in the low scheduling priority queue.

ALL
: Deletes the jobs in all scheduling priority queues. When UID=*user-id* is specified in the command, ALL is the default condition.

SQ,

Specifies the command applies to symbionts in the outstanding symbiont queue, as further defined by id or ALL, where:

`id`

Is the queue slot ID number for the symbiont to be deleted. You obtain this ID number by executing the MI SQ command.

`ALL`

Deletes all symbionts from the outstanding symbiont queue. The MI SQ command lists all symbionts in the queue.

`jobname`

Specifies that a particular job is to be deleted from being scheduled for execution.

Keyword Parameters

DDP
: Deletes only distributed data processing jobs.

LOCAL
: Deletes only locally entered jobs.

RBP
: Deletes only jobs entered remotely (i.e., from a remote batch terminal).

WKSTN
: Deletes only jobs either initiated from a workstation or using the // OPTION ORI= or // OPTION MAS= job control statement.

Note: *All of the preceding keyword parameters may be interchanged.*

`UID=user-id`
: Deletes all old jobs associated with a particular workstation user-id in the scheduling priority queue defined in parameter 1. Jobs subsequently entered with this user-id are not deleted. The user-id is a 1- to 6-character alphanumeric identification that does not start with Y.

`HOST=host-id`
: Deletes all old jobs associated with a particular host in the scheduling priority queue defined in parameter 1. Jobs subsequently entered with this host-id are not deleted. The host-id is one to four alphanumeric characters.

Notes:

1. *The UID and HOST parameters may be interchanged, but no further parameters are permitted.*

2. *If a command is entered from the system console (local or remote), the HOST parameter should not be used; that site's host-id is used automatically.*

3. *The special user-id OPERATOR indicates that the command applies to console-initiated jobs.*

4. *Any command from the system console (local or remote) pertains only to jobs with a matching host-id regardless of user-id.*

`LOG`
: Specifies that the job log is printed for all jobs deleted. If omitted, the log is not printed.

Examples

1. Operator keyin:

```
DE JBQ,H,HOST=CDE
```

Function requested:

All jobs residing in the high scheduling priority queue and originating from host CDE are to be deleted.

2. Operator keyin:

```
DE JBQ,A,UID=WKSTA1,LOG
```

Function requested:

All jobs in all scheduling priority queues associated with the workstation user WKSTA1 are deleted and their logs are printed. A system message tells you the number of jobs deleted.

Displaying Jobs in Scheduling Priority Queues (DISPLAY)

The DISPLAY command permits you to display the contents of any or all job scheduling queues on the console workstation or console screen as specified by command parameters. You can display all jobs in all queues or in a specific queue; a specific job within a queue; a specific workstation user's jobs in all queues or a specific queue; or a specific host's jobs in all queues or a specific queue. All requested jobs within the specified queue are displayed. Jobs in a deferred status (HOLD command) are displayed with parentheses around the job name. When all queues are requested, PRE is displayed first, followed by HIGH, then NOR, then LOW. If no jobs are found in the queue you request, a system output message is displayed stating that condition.

A system output message is displayed prior to the list of jobs. The message specifies:

- Whether the request was for LOCAL, RBP, DDP, WKSTN, or for all jobs (QUEUED)
- The user-id and host-id
- The priority of the queue display to follow
- Whether a hold local (HL), hold RBP (HR), hold workstation (HW), or hold DDP (HD) status is in effect for that queue

Jobs are displayed in three different formats, depending on the parameters you specify on the DISPLAY command. In all cases, jobs in hold status are displayed in parentheses.

- Two jobs per line

  If the DDP parameter is specified, the host-id is included in the following format:

  ```
  jobname: host-id/user-id        jobname: host-id/user-id
  ```

- Three jobs per line

  This format is the most common and is used when you specify the HOST parameter or the WKSTN parameter, or the command is from the system console (local or remote) with no other parameters. The format is:

  ```
  jobname: user-id        jobname: user-id        jobname: user-id
  ```

- Five jobs per line

  If there can be no user-id (LOCAL or RBP) or if there is only one user-id to be selected (UID= or DI JBQ from workstation), the jobs are displayed in the following format:

  ```
  jobname    jobname    jobname    jobname    jobname
  ```

Note: The DISPLAY command cannot be issued from an enter stream.

Format

```
DISPLAY∆JBQ, { PRE  }  { [ ,{ DDP   } ]                        }
             { HIGH }  { [  { LOCAL } ]                        }
             { NOR  }  { [  { RBP   } ]                        }
             { LOW  }  { [  { WKSTN } ]                        }
             { ALL  }  {                                       }
                       { [,UID=user-id][,HOST=host-id]         }
```

Positional Parameter 1

JBQ

Specifies the command applies to the jobs in a job scheduling priority queue, as further defined by PRE, HIGH, NOR, LOW, or ALL, where:

PRE

Displays the jobs in the preemptive scheduling priority queue.

HIGH

Displays the jobs in the high scheduling priority queue.

NOR

Displays the jobs in the normal scheduling priority queue.

LOW
: Displays the jobs in the low scheduling priority queue.

ALL
: Displays the jobs in all scheduling priority queues.

Keyword Parameters

DDP
: Displays only distributed data processing jobs.

LOCAL
: Displays only locally entered jobs.

RBP
: Displays only jobs entered remotely (i.e., from a remote batch terminal).

WKSTN
: Displays only jobs either initiated from a workstation or using the // OPTION ORI= or // OPTION MAS= job control statement.

Note: *All of the preceding keyword parameters may be interchanged.*

UID=user-id
: Displays all old jobs associated with a particular workstation user-id in the scheduling priority queue defined in parameter 1. Jobs subsequently entered with this user-id are not displayed. The user-id is a 1- to 6-character alphanumeric identification that does not start with Y.

HOST=host-id
: Displays all jobs associated with a particular host and residing in the scheduling priority queue defined in parameter 1. The host-id is one to four alphanumeric characters.

Notes:

1. *The UID and HOST parameters may be interchanged, but no further parameters are permitted.*

2. *If a command is entered from the system console (local or remote), the HOST parameter should not be used; that site's host-id is used automatically.*

3. *The special user-id OPERATOR indicates that the command applies to console-initiated jobs.*

4. *Any command from the system console (local or remote) pertains only to jobs with a matching host-id regardless of user-id.*

Examples

1. Operator keyin:

```
DI JBQ
```

Function requested:

Display all jobs in all scheduling queues.

Screen display:

```
    JC23 NO     QUEUED     JOBS WITH PREEMPTIME PRIORITY
    JC23 NO     QUEUED     JOBS WITH HIGH       PRIORITY
    *******     QUEUED     JOBS WITH NORMAL     PRIORITY
(JOBABC)
```

Explanation:

The job JOBABC is the only job presently in the system. It is in the normal priority queue and the parentheses indicate that the job is being held.

2. Operator keyin:

```
DI JBQ,N,LO,RBP
```

Function requested:

Display all locally and remotely entered jobs in the normal scheduling priority queue.

3. Operator keyin:

```
DI JBQ,UID=WKSTA2
```

Function requested:

Display all jobs in all scheduling priority queues that are associated with the workstation user WKSTA2.

Changing a Job Scheduling Priority (CHANGE)

The CHANGE command changes the scheduling priority of a specific job. If you place a deferred job into a new (changed) scheduling priority queue, the job retains its deferred status. Likewise, if you place a job into a queue that is under the influence of a HOLD NEW command, it too will become deferred. The job is put on the end of the new queue; it is the last examined for scheduling for execution in that queue.

Format

```
CHANGEΔjobname, { PRE  }
                { HIGH }
                { NOR  }
                { LOW  }
```

Positional Parameter 1

jobname
: Specifies the particular job to have its scheduling priority changed.

Positional Parameter 2

PRE
: Moves the specified job into the PRE scheduling priority queue.

HIGH
: Moves the specified job into the HIGH scheduling priority queue.

NOR
: Moves the specified job into the NOR scheduling priority queue.

LOW
: Moves the specified job into the LOW scheduling priority queue.

Example

Operator keyin:

```
CH JOBABC,H
```

Function requested:

Moves JOBABC from a previously assigned (via the job control stream, command entry for running the job, or NOR by default) scheduling priority queue into the HIGH scheduling priority queue. The following message appears on your screen:

```
CH OF JOB JOBABC
```

8.4.3. Job Execution

When a job is being executed, you can control the processing of that job through job execution commands. These commands allow you to:

- Suspend a job under execution
- Restart a job that has been suspended
- Raise or lower the switching priority level of a job being executed

Suspending a Job in Progress (PAUSE)

The PAUSE command suspends processing of a job. You may enter the command at any time, at a workstation or at the system console, and job processing suspends immediately or at the completion of a specified job step. If the job is between job steps, PAUSE (without a job-step-number parameter) takes effect at the beginning of the next job step. The PAUSE command permits you to mount a new volume on a tape, disk, or diskette drive, replace paper on the printer, or place more cards in the card reader. The suspended job is reactivated by the GO command. Furthermore, a suspended job cannot be cancelled until it is reactivated with the GO command.

Note: A job paused by the console operator cannot *be reactivated with a GO command from the workstation user. It can only be reactivated by the console operator.*

Format

```
PAUSEΔjobname[,job-step-number]
```

Positional Parameter 1

`jobname`
Specifies the name of the job whose processing is suspended.

Positional Parameter 2

`job-step-number`
Specifies the number of a job step that must complete before processing is suspended. This is an optional entry.

Examples

1. Operator keyin:

```
PA JOBA
```

Function requested:

Suspends JOBA processing immediately.

2. Operator keyin:

```
PA JOBA,02
```

Function requested:

Suspends JOBA following completion of the second job step.

Reactivating a Suspended Job (GO)

The GO command reactivates a job suspended by the PAUSE command or by job control operations. Job control suspends processing of a job when it issues instructions to mount a new volume on a tape, disk, or diskette drive. The GO command also is required as a response to a message from the system preceded by an asterisk (*).

Jobs started at a workstation and paused at the system console can be restarted only with a GO command at the system console. A GO command at a workstation on a job that was paused at the system console results in an NAK (negative acknowledgment) error message.

Format

```
GOΔjobname[,nn]
```

Positional Parameter

jobname
: Specifies the job to be reactivated after execution was temporarily suspended.

nn
: Is the number of the message requiring the GO command. This specification allows multitask jobs to suspend and start tasks other than the primary ones.

Example

Operator keyin:

```
GO JOBA
```

Function requested:

Reactivate job processing for the suspended job named JOBA.

Changing a Job Switching Priority (SWITCH)

The SWITCH command changes the switching priority level for a job under execution. The switching priority level is changed for either the currently executing job step or for the current job step and all subsequent job steps. If the priority level is changed for the current job step only, any subsequent job step executes under the priority established for it (via // EXEC job control statement or default to the lowest level established at SYSGEN) unless changed by another SWITCH command. A job assigned a higher switching priority level has priority over lower switching priority level jobs for control of the central processor. If a job is changed to a higher switching priority level than another job currently being executed, the lower switching level job will often be processed slower than the higher switching level job.

The number of switching priority levels a job can be raised or lowered is governed by the number of switching priority levels established at system generation time (maximum 60 levels). Switching priority levels are from 1 to *n*, where 1 is the highest priority level and *n* is the lowest.

When you change any job's switching priority, all tasks of that job retain the same switching priority relative to each other; therefore, if a job's task exceeds the upper or lower switching priority limit, all the job's tasks move only by the number of priority levels that the highest or lowest priority task can be moved within the switching priority limits.

Format

```
SWITCHΔjobname, {+number-of-priority-levels} [,ALL]
                {-number-of-priority-levels}
```

Positional Parameter 1

`jobname`
Specifies the name of the job whose task switching priority is changed.

Positional Parameter 2

`+number-of-priority-levels`
Specifies the number of switching priority levels a job is raised.

`-number-of-priority-levels`
Specifies the number of switching priority levels a job is lowered.

Positional Parameter 3

`ALL`
Indicates that the priority to which the job is switched is in effect for all subsequent job steps. If omitted, the priority is only for the current job step.

Example

Operator keyin:

```
SW PAYJOB,-5
```

Function requested:

Lowers the switching priority level of the currently executing PAYJOB job step by five queues. The switching priority remains changed for the duration of the job step.

8.4.4. Job Termination

Job termination commands permit you to terminate the processing of a job or symbiont, as defined by the command parameters.

Canceling a Job in Progress (CANCEL)

The CANCEL command immediately halts all processing of a job or symbiont. The CANCEL command may be issued at any time during processing of a job and results in the immediate termination of the job step being executed at the time the CANCEL command is given, plus any subsequent job steps scheduled for the job. The job run library file for the job also is deleted.

Format

```
CANCEL∆ ⎧ jobname[, {D} ] ⎫
        ⎪          {N}    ⎪
        ⎨ symbiont,S[,N]   ⎬
        ⎪ synnnn,S         ⎪
        ⎪ sy(did),S        ⎪
        ⎩ sy               ⎭
```

Positional Parameter 1

jobname
: Specifies the name of the job whose processing is immediately terminated and whose job run library file is deleted.

synnnn
: Is the name of the symbiont sending the message. This 6-character name is the 2-character symbiont identification followed by a 4-digit binary job number inserted at task initiation.

symbiont
: Specifies the 8-character symbiont name.

sy

Specifies the 2-character ID used to call the symbiont to be terminated.

did

Is the device address.

Positional Parameter 2

D

Specifies that a dump is to be taken when the job terminates, regardless of the dump option specified in its job control stream.

N

Specifies that no dump is taken when the job terminates, regardless of the dump option specified in its job control stream.

S

Specifies that the name specified in positional parameter 1 is the name of a symbiont.

If omitted, the job control dump options remain in effect. Positional parameter 2 must be specified when a symbiont is cancelled.

Positional Parameter 3

N

Specifies that no dump is taken when the symbiont terminates.

If omitted, a symbiont dump is taken.

Examples

1. Operator keyin:

   ```
   CA JOBA,D
   ```

   Function requested:

   Terminates JOBA immediately, deletes the JOBA run library file, and provides a JOBA dump.

2. Operator keyin:

   ```
   CA SU,S,N
   ```

   Function requested:

   Terminates the system utility symbiont immediately without a dump.

Soft Cancel of a User's Job (CJ)

The CJ command delays the standard cancel procedure until the job is executing under "normal" job conditions. This allows the operator to cancel a job without compromising a MIRAM file. Normal job execution requires:

- Execution under the job's key
- No wait flags set
- Not executing in external code

If this command does not terminate a job, the CANCEL command can be used. When CJ is issued, an unknown period of time may pass before the job actually cancels (for instance, the job may have been paused or yielded when the command was executed). A dump requested at the time of the command may not reflect the status of the command at that time; instead, it will reflect the job status when the cancel becomes effective.

Format

$$\underline{\text{CJ}}\Delta\text{jobname}\left[,\begin{Bmatrix}\text{D}\\\text{N}\end{Bmatrix}\right]$$

Positional Parameter 1

jobname
: Specifies the name of the job that will be terminated the next time it is executing under normal job conditions.

Positional Parameter 2

D
: Specifies that a dump is to be taken at the time when the job is effectively canceled, regardless of the dump option specified in its job control stream.

N
: Specifies that no dump is taken when the job terminates, regardless of the dump option specified in its job control stream.

Examples

1. Operator keyin:

```
CJ JOBA,D
```

Function requested:

Terminates JOBA as soon as JOBA is executing under normal conditions, deletes the JOBA run library file, and provides a JOBA dump.

2. Operator keyin:

```
CJ JOBA,N
```

Function requested:

Terminates JOBA as soon as JOBA is executing under normal conditions, deletes the JOBA run library file, and does not provide a JOBA dump.

Stopping Execution of a Dump (END)

The END command terminates execution of a cancel or end-of-job dump for a particular job.

Format

```
END∆DUMP, ⎧jobname     ⎫
          ⎨symbiont,S  ⎬
          ⎪sy,S        ⎪
          ⎩sy(did),S   ⎭
```

Positional Parameter 1

DUMP
: Specifies that the execution of a dump is to stop.

Positional Parameter 2

jobname
: Specifies the name of the job whose cancel or end-of-job dump you want stopped.

symbiont
: Specifies the 8-character symbiont name.

sy
: Specifies the 2-character ID used to call the symbiont to be terminated.

did
: Is the device address.

Positional Parameter 3

S
: Specifies that the name specified in positional parameter 1 is the name of a symbiont.

Example

Operator keyin:

```
EN DUMP,JOBA
```

Function requested:

Terminates the execution of the dump being taken for JOBA.

Terminating a Job (STOP)

The STOP command terminates a specific job at the end of the current or specified job step. This command provides for orderly termination of the job.

Format

```
STOPΔjobname[,job-step-number]
```

Positional Parameter 1

jobname
: Specifies the job whose processing is terminated in an orderly sequence.

Positional Parameter 2

job-step-number
: Specifies the number of a job step that must complete before processing is suspended. This is an optional entry.

Examples

1. Operator keyin:

   ```
   ST JOBC
   ```

   Function requested:

   Terminates JOBC at the end of the currently executing job step.

2. Operator keyin:

   ```
   ST JOBC,02
   ```

   Function requested:

   Terminates JOBC at the end of job step 02.

8.5. Selected Occasion Operator Commands

During the course of processing a job, you may be required to enter system-oriented commands to obtain information or make changes not involved with the execution of a particular job. These commands enable you to:

- Display an area of main storage
- Display information on active jobs and symbionts, current system I/O device status, and outstanding requests and commands
- Clear the console workstation of all but outstanding output messages
- Change the system time or date
- Control software-detected hardware error logging
- Set an I/O device down or up as required for normal maintenance or device malfunction
- Read the volume serial number of a mounted disk or tape volume
- Display the status of jobs in the system
- Dump contents of main storage
- Set the YDUMP file to an unlocked condition
- Set one or more blocks of main storage up or down
- Terminate system activity
- Specify resource management
- Verify the volume table of contents (VTOC)

The commands enabling you to perform these functions are described in the following subsections. Commands required to control interactive services, data communications, and system utility functions are described in Sections 9, 10, and 11.

8.5.1. Displaying Portions of Main Storage (DISPLAY)

The DISPLAY command displays selected areas of main storage at the console workstation or console. You usually enter the command when your system administrator needs a job address displayed for program debugging purposes.

```
DISPLAYΔaddr[,jobname]
```

Positional Parameter 1

addr
: Is a hexadecimal number used for a specific (absolute) main storage address or a job-relative main storage address. A job-relative address is identified by the job name; otherwise, an absolute address is displayed.

Positional Parameter 2

jobname
: Identifies the job name of the job in main storage whose relative address is displayed.

When no parameter 2 is entered, the address entered in positional parameter 1 is an absolute address.

After you enter DISPLAY to load the display symbiont (DI), the following output message appears on the console workstation or console:

```
addr[JOBNAME] contents-of-selected-addr Y,N,NXT?
```

where:

addr
: Is the address of the main storage location being displayed, in hexadecimal.

JOBNAME
: Identifies the address being displayed as a job-relative address and identifies the job region by the job's name. If JOBNAME is not displayed, the address being displayed is an absolute address.

contents-of-selected-addr
: Is the hexadecimal representation of the contents of the selected main storage address.

Y
: Is a message response to display the next sequential main storage location.

N
: Is a message response to terminate the display symbiont.

NXT

Is a reminder that you can display another nonconsecutive main storage location without recalling the display symbiont by responding to this output message with the solicited input message:

```
addr[,jobname]
```

where:

addr

Is the address of the main storage location to be displayed, in hexadecimal.

jobname

Identifies the job name of the job whose relative address is to be displayed.

When there are no more addresses to display, terminate the display symbiont by responding to the last display message with an N for none.

8.5.2. Displaying System Information (MIX)

The MIX command displays tables of different aspects of system information.

Format

```
      ┌                                  ┐
MIXΔ  │DA  ┌ , ┌jobname       ┐ ┐        │
      │    └   └symbiont-name ┘ ┘        │
      │                                  │
      │VI  ┌ , ┌did     ┐ ┐              │
      │    └   └c/ca/caa┘ ┘              │
      │SQ                                │
     {│SI                                │}
      │DS[,did]                          │
      │SC                                │
      │FR                                │
      │MM                                │
      │EN                                │
      │EL                                │
      └HF (models 10/15/20 only)         ┘
```

Positional Parameter 1

DA

Displays the following information on the console workstation or console screen for the job or symbiont you specify (positional parameter 2) or for all active jobs and symbionts when you omit positional parameter 2:

- Job name or symbiont name
- Job slot number; 0 for symbiont
- Priority
- Allocated device numbers
- Allocated device types
- Volume serial number or user identification

VI

Displays the following information for the devices you specify in positional parameter 2 that have a mounted volume, or for all devices having a mounted volume, if you omit positional parameter 2:

- Device address of each device
- Volume serial number
- Mode setting of tape devices (U26 and U28 only)

SQ

Displays a list of outstanding symbiont requests, including unprocessed queued operator commands. (These are console commands that could not be processed immediately.) The displayed list includes a queue slot ID number that can be used to delete items from the oustanding symbiont queue (see DELETE command).

SI

Displays system information, including supervisor name, release-id, date, time, RES device address, RUN device address, system reader (RDR) device address, spool (SPL) device address, and total available main storage in decimal.

DS

Displays the following information on the console workstation or console screen for the device you specify in positional parameter 2, or for all devices in the system if you omit positional parameter 2 (use MI DS to determine if a diskette drive is available when the device requires maintenance; when the drive is in use, an I/O error occurs if its door is opened):

- Device address of each device
- Whether the device is up or not (Y or N)
- Whether the device is available or not (Y or N)
- Whether the device is sharable or not (Y or N)
- Job numbers of all the jobs to which the device is allocated

SC

Displays the name, address, and decimal size of each shared code module currently residing in main storage.

FR

Displays the address and decimal size of all regions currently unused in main storage.

MM

Displays the type, name, address, and decimal size of every main storage region. The following types are displayed:

- FREE Free region
- JOB Job region
- JOB SCHD Job scheduler
- SYMBIONT Symbiont (RUN processor, ICAM, etc.)
- RESERVED Region reserved for a job just scheduled

- BUF POOL — Dynamic buffer pool. These storage pools are internally subdivided to provide dynamically allocated main storage for system software.

  Note: See Section 12 regarding disk cache storage allocations.

- SHR CODE — Shared code. When shared code occupies its own main storage region, you display it by using either the SC or the MM parameters. When shared code is part of a dynamic buffer pool, you display it by using the SC parameter only.

- DISABLED — Not available. The region is either temporarily or permanently disabled. A region is temporarily disabled when it is in the process of being changed from one type to another (e.g., from SYM to FREE). This temporary condition exists for no more than a few seconds. A region is permanently disabled when a system error, such as a parity error, prevents the termination and clearing out of a region.

- DOWN — Down. The region is set down as the result of either an operator command or too many recoverable main storage errors occurring in that region.

EN

Terminates the MIX function currently in progress.

EL

Displays error log information, including error logging functions that are in effect, percentage of error records used, and number of records remaining.

HF

Displays the enabled or not enabled status of various model 10/15/20 hardware features (this parameter is not supported on other models). The displayed response indicates the model number (10, 15, or 20) and whether the following features are enabled (Y) or not enabled (N): bypass cache (BYP), floating point processor (FPP), instruction cache (ICM), and operand cache sets 1 through 4 (OC1 - OC4).

Positional Parameter 2

jobname

Used with the DA parameter to identify a job name for which information is desired. If you omit this parameter, information for all active jobs and symbionts is displayed.

`symbiont-name`

Used with the DA parameter to identify a symbiont name for which information is desired. If omitted, information for all active jobs and symbionts is displayed.

`did`

Used with the VI and DS parameters to obtain status information for a device or group of devices. One, two, or three characters may be used. Key in three characters to specify the exact device address of the device to be displayed. If you key in two characters, the group of devices whose device addresses start with those two characters is displayed. If you key in one character, information is displayed about all devices whose addresses begin with that character. If you omit this parameter, the status of all the devices in the system is displayed.

caa

Used with the VI parameter to indicate a device or group of devices. C specifies the channel; a, the address. One, two, or three characters may be used, as with did.

Example 1

Operator keyin:

```
MI DA
```

Function requested:

Display information pertaining to all active jobs and symbionts.

Screen display:

```
JOB NAME    JOB NO    PRI    DEV      TYPE          VSN/UID
RC$$IS00      0         0    312    3560 WRKSTN     LKT
```

The symbiont RC$$IS00 is recognized as a symbiont by a job slot number of 0. Its priority is 0, and it is allocated to device number 312, which is a 3560 workstation. LKT is the user-id.

Example 2

Operator keyin:

```
MI SI
```

Function requested:

Display system information.

Screen display:

```
SUP204  07. 1.0   81/05/19 13:34:04  USABLE MEMORY=    971,520
        RDR=332   RES=103   RUN=103   SPL=103
```

Supervisor SUP204 is running under release 12.0. The date and time is followed by the available main storage, which is 971,520 bytes. The reader device address is 332, the RES device address is 103, the RUN device address is 103, and the spool device address is 103.

Example 3

Operator keyin:

```
MI DS
```

Function requested:

Display information pertaining to all devices in the system.

Screen display:

DVC	UP	AVL	SHR	USERS
31Ø	Y	Y	N	Ø
33Ø	Y	Y	N	
331	Y	Y	N	
332	Y	Y	N	
333	Y	Y	N	
32Ø	Y	Y	N	
321	Y	Y	N	
341	Y	Y	N	
34Ø	Y	Y	N	
1ØØ	Y	Y	Y	
1Ø1	Y	Y	Y	
1Ø2	Y	Y	Y	
1Ø3	Y	Y	Y	Ø,1,2,3,4,5,6,7,8,9,1Ø,11,12,13,14
1Ø4	Y	N	Y	
1Ø5	Y	N	Y	
1Ø6	Y	N	Y	
311	Y	Y	N	
312	Y	Y	N	Ø
313	Y	Y	N	
314	Y	Y	N	
315	Y	N	N	
316	Y	N	N	
317	Y	N	N	

The DVC column contains the device addresses of all the devices in the system. The next three columns tell you whether the device is up, available, and sharable. Also provided are the number of users of each device and the job slot numbers of all the jobs to which the device is allocated.

8.5.3. Reconstructing Console Display (REBUILD)

The REBUILD command clears all information from the console workstation, then restores the job number header lines and rewrites all outstanding question and action request output messages on the console workstation screen. All displays other than unanswered questions and action requests are lost.

Format

```
REBUILD
```

Positional Parameters

No positional parameters are required for the REBUILD command.

Note: *If you're working on a System 80 models 8-20, the REBUILD command does not restore the job number header lines on the console screen.*

8.5.4. Setting Simulated Day Clock (SET CLOCK)

The SET CLOCK command resets the time of day in the system-simulated day clock; it updates the models 8-20 hardware clock. The date and time are changed automatically at midnight of each day.

Format

```
SET△CLOCK,hh:mm:ss
```

Positional Parameter 1

CLOCK
: Sets the simulated day clock and hardware clock to the time specified in positional parameter 2.

Positional Parameter 2

hh:mm:ss
: Specifies the time to set the simulated day clock and hardware clock as follows:

  hh
  : Specifies hours (00 through 99).

mm

Specifies minutes (00 through 59).

ss

Specifies seconds (00 through 59).

8.5.5. Setting Date Field (SET DATE)

The SET DATE command resets the calendar date in the system information block date field, updates the models 8-20 hardware clock, and resets the job date for every job currently in main storage (except those jobs containing a // SET DATE job control statement). The date and time are changed automatically at midnight of each day.

Format

```
SET DATE,yy/mm/dd[,yyddd]
```

Positional Parameter 1

DATE

Changes the calendar date in the system information block date field, the hardware clock, and in the preambles for current jobs to the date specified in positional parameter 2.

Positional Parameter 2

yy/mm/dd

Specifies the date to be used for the calendar date in the system information block date field, the hardware clock, and job preambles, as follows:

yy

Specifies year (00 through 99).

mm

Specifies month (01 through 12).

dd

Specifies day (01 through 31).

Positional Parameter 3

yyddd

Specifies the ordinal date, where yy is the year (00 through 99) and ddd is the day of year (001 through 366). This date is maintained in a separate part of the system information block and is used by data management routines that check the label fields.

If omitted, this field in the system information block is set to the ordinal date corresponding to the yy/mm/dd date specified in positional parameter 2.

8.5.6. Setting Error Log (SET ELOG)

The SET ELOG command should not be used without first consulting your customer engineer. Communications, I/O device, machine check, and I/O termination record error logging can be turned on or off when you enter specified combinations of the command and parameters. The IPL procedure automatically turns on error logging and all error logging functions. Any changes to this all-on condition that you enter are lost when the system is reloaded.

Format

```
SET∆ELOG, { {ON } [ ,{ALL } ] }
          { {OFF}    { COMM }   }
          {          { MCHK }   }
          {          { IO   }   }
          {          { TERM }   }
          {          { CLOG }   }
          {          { RSTK }   }
          {          { MSE  }   }
          {          { did  }   }
          { factor,{RSTK}       }
          {        {MSE }       }
```

Positional Parameter 1

ELOG

Specifies that the PIOCS error logging processor is to be set to some condition.

Positional Parameter 2

ON

Turns on the function of logging errors. When parameter 3 is included, error logging is turned on for that parameter specification only.

OFF

Turns off the logging of errors for the function specified by parameter 3.

factor

Specifies a 1- to 3-character decimal value that is used to redefine main storage error (MSE) or retry stack error (RSTK) limits. The error is suppressed when its established limit is exceeded. The factor value entered is multiplied by 32 to redefine the millisecond time factor (F) in the error limit formula:

```
F*S:E
```

where:

F
: Is the millisecond time factor to be redefined.

S
: Is the size or number of MSE or RSTK errors.

E
: Is the elapsed time (in milliseconds) since the last interrupt of the same type.

When F times S is greater than E, the reported type of suppressible machine error is disabled and a message to that effect is displayed on the console workstation or console screen.

Positional Parameter 3

ALL
: Is valid only when ON is specified in parameter 2. ALL specifies that all loggable errors (i.e., communications, machine check, I/O device, and I/O termination record error logging) are to be logged for all devices in the system.

COMM
: Turns on or off communications error logging as directed by parameter 2.

MCHK
: Turns on or off machine check error logging as directed by parameter 2.

IO
: Turns on or off all I/O device error logging as directed by parameter 2.

TERM
: Turns on or off all I/O termination record error logging as directed by parameter 2.

CLOG
: Turns on or off channel logout error logging as directed by parameter 2.

RSTK
: Is valid for all parameter 2 specifications and has the following functions:

ON
: Turns on retry stack error logging.

OFF
: Turns off retry stack error logging.

`factor`

Specifies that the factor specified in parameter 2 is to be used in the algorithm for establishing retry stack error limits.

`MSE`

Is valid for all parameter 2 specifications and has the following functions:

`ON`

Turns on main storage error logging.

`OFF`

Turns off main storage error logging.

`factor`

Specifies that the factor specified in parameter 2 is to be used in the algorithm for establishing main storage error limits.

`did`

Represents a device address as one, two, or three hexadecimal characters specifying a channel, subchannel, and device, respectively. A did can be used with any parameter 2 specification.

- If *did* is specified as one character, all devices on that channel are directed by the parameter 2 specified.
- If *did* is specified as two characters, all devices on that channel and subchannel are directed by the parameter 2 specified.
- If *did* is specified as three characters, a specific device on that channel and subchannel is directed by the parameter 2 specified.

If omitted, the error logging condition for each parameter 3 remains unchanged, (i.e., turned on or off as previously set). Thus, if you omit parameter 3 and enter the command SET ELOG,ON, the error logging function is turned on, but only for those parameter 3 specifications previously turned on.

Operator Considerations

When the ELOG file (YELOG) is almost full, the system displays the message:

```
LOG FILE IS NEARLY FULL
```

At this time you should run the system-supplied ONUERL program (Appendix A).

The ELOG file can become completely full only if the ONUERL program was not run when the nearly full message was displayed. If the ELOG file fills up, ELOG asks you whether you want to turn off error logging or wrap around to the beginning of the log file.

When the message

```
LOG FILE IS FULL W(RAP) OR O(FF)
```

appears, respond using the solicited message format:

```
{O}
{W}
```

where:

`O`

Turns off the PIOCS error logging processor.

`W`

Informs the PIOCS error logging processor to wrap around the log file and continue logging.

If ONUERL has already been run, respond with the letter W. If ONUERL has not been run, enter O to turn off logging, then initiate ONUERL.

Note: *It is recommended that you run ONUERL when the message LOG FILE IS NEARLY FULL is displayed. Allowing the ELOG file to wrap around causes loss of potentially valuable data.*

8.5.7. Discarding System Messages (FLUSH)

The FLUSH command allows you to discard system messages queued for delivery at your workstation. You can discard all queued messages or only those with a specified message prefix.

Format

```
FLUSH {*ALL      }
      {msg-prefix}
```

Parameters

`*ALL`
Discards all queued messages.

`msg-prefix`
Discards only messages with this prefix.

8.5.8. Setting Physical Unit Blocks (SET IO)

The SET IO command allows you to set specific bits in the physical unit blocks (PUB), which define operational characteristics and assignments of I/O devices. There is one physical unit block, comprising a 3- or 4-character did (device address), for each physical device in your system. You must set all devices or subsystems DOWN before attempting operation on the device, such as forms loading or changing ribbon; or in case of malfunction, before turning the device off. This is required if the processor is to continue operation with other peripheral devices while the subject device undergoes isolated operations. Before performing a procedure or turning power off for the device, key in SET IO,*did*,DOWN. After the offline procedure completes or after turning on power independently of the processor, key in SET IO,*did*,UP.

Note: *The system automatically sets all devices or subsystems not online during IPL time to not available (NA). After IPL, any attention interrupt from a device causes that device to be set available (AV).*

Format

```
SETΔIO,did,⎧ AV            ⎫ ⎡ ⎧ O ⎫ ⎤
           ⎪ CON           ⎪ ⎢,⎨   ⎬ ⎥
           ⎪ DOWN          ⎪ ⎣ ⎩ H ⎭ ⎦
           ⎪ EON           ⎪
           ⎪ EOF           ⎪
           ⎪ FEA,type-code ⎪
           ⎨ HOME          ⎬
           ⎪ NA            ⎪
           ⎪ NOSHARE       ⎪
           ⎪ RDR           ⎪
           ⎪ SHARE         ⎪
           ⎪ TYPE,type-code⎪
           ⎩ UP            ⎭
```

Positional Parameter 1

`IO`

Specifies the change is made in the device address specified by positional parameter 2.

Positional Parameter 2

`did`

Specifies a 1- to 4-character numeric field identifying the device address to be changed (device addresses are usually attached at a visible location on the device cabinet) as follows:

`ALT`	Alternate console specified in SYSGEN
`caa`	Alternate console specified by channel (*c*) and address (*aa*)

ORG	Original (hardware) console
100	SYSRES integrated disk
DDD	Indirect printer
10n	All other disks
31n	Workstation
32n	Diskette
33n	First printer
34n	All other printers
35n	Card reader
36n	Card punch

where:

n
Is the 1- or 2-digit actual physical unit number assigned to the selected device.

Note: *Positional parameter 2 DDD can be associated only with positional parameter 3 UP or DOWN.*

Positional Parameter 3

AV

Specifies the device identified by positional parameter 2 is recognized by the system and available for assignment to user jobs.

CON

Designates the specified device as the current console.

DOWN

Specifies the device identified in positional parameter 2 is not to be considered available for assignment to user jobs, although the device is recognized by the system.

The DOWN[,*n*] format is applicable only if the device identified in positional parameter 2 is an indirect printer. The *n* is a single-digit number indicating the indirect printer PUB on which the SET IO command is to be performed. If *n* is omitted, the first indirect PUB is assumed.

EON

Turns on error logging for the specified device.

EOF

Turns off error logging for the specified device.

FEA, type-code

Modifies the feature bytes of the device specified in positional parameter 2, where type-code is a 1- to 4-character field specifying the device, its options, and features desired. See the *Job Control Programming Reference Manual* (UP-9984) for the definition of type and feature codes.

HOME

Synchronizes the operating system with the physical paper position of an 0768, 9200, or 9300 printer during a home operation.

NA

Specifies the device identified by positional parameter 2 is not recognized by the system; it is not available for assignment to user jobs.

NOSHARE

Forbids allocation of the device specified in positional parameter 2 to more than one program simultaneously.

RDR

Assigns the new did specified in positional parameter 2 as the system card reader.

SHARE

Permits the device specified in positional parameter 2 to be shared by more than one program simultaneously.

TYPE, type-code

Modifies the type bytes of the device specified in positional parameter 2, where type-code is a 1- to 4-character field specifying the device, its options, and features desired. See the *Job Control Programming Reference Manual* (UP-9984) for the definition of type and feature codes.

UP

Specifies the device identified in positional parameter 2 is considered available for assignment to user jobs. The device remains recognized by the system.

The UP[,*n*] format is applicable only if the device identified in positional parameter 2 is an indirect printer. The *n* is a single-digit number indicating the indirect printer PUB on which the SET IO command is to be performed. If *n* is omitted, the first indirect PUB is assumed.

Notes:

1. *For more information on the indirect printer, see the* Installation Guide *(UP-8839).*
2. *Positional parameter 4 is required in order to set a device to available, not available, down, or up when your system is configured with stability monitoring.*

Positional Parameter 4

O

Indicates that the device status change specified in positional parameter 3 (AV, DOWN, NA, or UP) is required for operational reasons (for example, to set a device down to change forms or ribbons).

H

Indicates that the device status change specified in positional parameter 3 (AV, DOWN, NA, or UP) is required for hardware reasons (for example, to set a device to not available because the hardware is malfunctioning).

Example

Operator keyin:

```
SET IO,321,DOWN
```

Function requested:

Set the device identified by device number 321 down. Device 321 is a diskette. No screen display results after this command is issued. To determine if the device is now unavailable to user jobs, issue a MI DS command (see 8.5.2).

8.5.9. Reading a Mounted Volume Serial Number (AVR)

The AVR command initiates reading the volume serial number of a premounted prepped disk pack, diskette, or magnetic tape volume and storing it in the device physical unit block.

Format

```
AVR△did[,did][,did]
```

Positional Parameters 1 through 3

did

Specifies the device addresses for the volumes to be recognized.

8.5.10. Displaying Job Status (DISPLAY JS/SY)

The DISPLAY JS command displays the status of jobs in the system at the console workstation or console. You can display the status of a specific job or all jobs in main storage, a specific job in a scheduling priority queue, a job being processed by the RUN or OCL processor, or all tasks attached to symbionts. The display includes the job name along with the CPU time used when the job is under execution, the reason why the job is not executing (such as waiting for IO, waiting for mount message, under a pause), or the scheduling priority queue in which the job resides. It also includes the status of all subtasks attached to the job.

Format

```
DISPLAYΔ{JS,[jobname][,ALL]          }
        {SY[,symbiont-name][,ALL]    }
```

Positional Parameter 1

JS

Displays job status at the console workstation or console.

SY

Displays at the console the status of all tasks attached to symbionts.

Positional Parameter 2

jobname

Displays the status of the job name specified. If omitted, the status of all jobs in main storage is displayed except for those on the job queues.

symbiont-name

Displays the tasks attached to the symbiont named. This name is the name displayed by the MIX DA command. If omitted, the tasks attached to all symbionts are displayed.

ALL

Displays the status of all tasks attached to a job or symbiont, instead of just the primary task.

Example

Operator keyin:

```
DI JS,MYJOB
```

Function requested:

Produce a display of status information about a job named MYJOB.

Typical informational messages:

The following are examples of the job status informational messages produced by the DISPLAY JS command. They show what data DI JS might display if it were entered against MYJOB as the job proceeds through the job processing steps.

For the message:

```
MYJOB IN STEP 01(LNKEDT00)-PRI=10 CPU-TIME=00:01:43.874
```

MYJOB is active in its first step, performing linkage editing. The CPU TIME portion of the display indicates the linkage editor had control of the CPU for 1 minute, 43 seconds, and 874 milliseconds. If the job is proceeding, you can reenter DI JS for MYJOB and see an increase in the CPU TIME figure.

For the message:

```
MYJOB IN STEP 02(LIBS0000)-WAITING FOR I/O #00005736
```

MYJOB is in its second step, executing the librarian. Currently, the 5736th I/O operation of this step is being performed. If you reenter DI JS, you may see the I/O number increase. If MYJOB remains at #00005736, it might be stuck, requiring your intervention.

For the message:

```
MYJOB IN STEP 03 -IN STEP PROCESSOR
```

MYJOB is between job steps. Step 03 has just completed or is about to start.

For the message:

```
MYJOB NOT YET SCHEDULED-INSUFFICIENT MAIN STORAGE
```

MYJOB is not executing; it is placed on a job queue but is not scheduled for execution because not enough main storage is available.

8.5.11. Dumping the Contents of Main Storage (SYSDUMP)

The SYSDUMP command is used to dump the entire contents of main storage to the YDUMP file on SYSRES. Use this command whenever a system dump is required without supervisor reloading (no manual IPL is required). After the contents of main storage are dumped, the job SYSDMP*nn* (where *nn* is a unique number assigned by the system) is automatically initiated to print the YDUMP file. The YDUMP file locks until the SYSDMP*nn* job completes, to prevent other system functions (caused, for example, by an // OPTION SYSDUMP statement or by system errors) from also using the file (see the *Dump Analysis Programming Guide* (UP-9980).)

If you don't want a printout of the dump, enter NONE when the SYSDMP*nn* job asks what type of dump should be printed. If you delete the SYSDMP*nn* job from the job queue or cancel it before it sends you the SD01 DUMP OPTION message, you must enter the SET SY command to unlock the YDUMP file.

You can't use the console workstation or console keyboard until the main storage contents are completely written (only a few seconds). After the keyboard unlocks, you can resume system activity without impairing the integrity of your dump.

Format

```
SYSDUMP
```

Notes (Models 3-6 Users):

1. *If a nonrecoverable error occurs and you want a dump of everything in main storage at that time, press the FUNCTION and RESTART keys simultaneously. This writes the contents of main storage to the YDUMP file and automatically initiates IPL. (You must reinitiate IPL when a nonrecoverable error occurs.) The IPL procedure, in turn, automatically schedules SYSDUMPO to print the dump.*

   If you get an HPR of 999F, it means the YDUMP file contains data from a previous main storage write operation that hasn't been processed by SYSDUMP or SYSDUMPO. You can overwrite the contents of the file by pressing the FUNCTION and START keys simultaneously, or you can follow the next procedure.

2. *If there's information in the YDUMP file you want to preserve and a nonrecoverable error occurs, reinitiate IPL by pressing the FUNCTION and IPL keys simultaneously. This reinitializes the system without affecting the YDUMP file. To print the contents of YDUMP file, run the SYSDUMP or SYSDUMPO job stream when the IPL procedure is completed.*

Notes (Models 8-20 Users):

1. *If a nonrecoverable error occurs and you want a dump of everything in main storage at that time, press the ESC key, then the M key to display the manual frame. Then key in L for system reset and press the XMIT key. Next, key in U for run and press the XMIT key.*

   If you get an HPR of 999F, it means the YDUMP file contains data from a previous main storage write operation that hasn't been processed by SYSDUMP or SYSDUMPO. You can overwrite the contents of the file by keying in U for run on the manual frame and pressing XMIT, or you can follow the next procedure.

2. *If there's information in the YDUMP file that you want to preserve and a nonrecoverable error occurs, perform the IPL procedure according to 6.3.1. This reinitializes the system without affecting the YDUMP file. To print the contens of YDUMP, run the SYSDUMP or SYSDUMPO job stream when the IPL procedure is completed.*

8.5.12. Setting the YDUMP File to Unlocked Condition (SET SY)

The SET SY command unlocks the YDUMP file after an SY command (8.5.10) or system error has locked it. You use the SET SY command to unlock YDUMP if you previously entered the SYSDUMP command and deleted the SYSDMPnn job from the job queue or cancelled it before it displayed the SD01 output message.

Format

SETΔSY,LOFF

Positional Parameter 1

SY
: Specifies the YDUMP file condition is set.

Positional Parameter 2

LOFF
: Specifies the YDUMP file is unlocked.

Note: *No dump is provided in stand-alone IPL.*

8.5.13. Setting Main Storage Condition (SET MEM)

The SET MEM command sets one or more blocks of main storage up or down. The length of one block is determined by the replacement element size (RES), that is, the size of the replacement element that contains the main storage block determined unusable. Normally, main storage is in an up (usable) condition. The system automatically determines when a main storage block becomes unusable. The system sets the block down and displays a listing on the console workstation or console screen of all main storage blocks currently set to down. The down block is unavailable for use by new jobs or symbionts until it is set to up via the SET MEM command. However, any jobs or symbionts currently running in the down block continue to run until completed.

The map of usable storage is kept in the error log. If you add main storage to your system or take a system-resident pack from a smaller system to a larger system and the error log is not initialized, the additional main storage will not be usable. To make the main storage usable, initialize the error log when you perform an IPL or use the SET MEM command to set the additional main storage up.

You enter the SET MEM command at the request of your system administrator to set a main storage block up or down as required. Normally, SET MEM is used to set a block to up after your Unisys customer engineer corrects the condition that made it unusable.

Format

```
SET△MEM, { DOWN } ,address [ , { no-of-blocks } ]
         { UP   }              {  1           }
```

Positional Parameter 1

MEM
: Specifies main storage is to be set up or down.

Positional Parameter 2

DOWN
: Sets one or more blocks of main storage down.

UP
: Sets one or more blocks of main storage up.

Positional Parameter 3

address
: Specifies the beginning address, in hexadecimal, of a block to be set up or down. This address must be on a RES boundary (a multiple of the replacement element size).

Positional Parameter 4

no-of-blocks
: Specifies the number of contiguous blocks, in decimal, set up or down beginning with the address specified in positional parameter 3. If omitted, one block is assumed.

8.5.14. Setting the Unattended Console Feature (SET UNCON)

The unattended console feature is a SYSGEN option that allows messages requiring operator intervention to be answered by the operating system if the operator has not responded in 15 minutes. These messages are informational only and pertain to device I/O errors. The unattended console feature decreases the need for operator intervention by requiring operator responses only in extreme cases. The SET UNCON command allows you to remove the previously SYSGENed option, reactivate it once it has been removed, and change the unattended console auto-answer time of 15 minutes.

Format

```
SET△UNCON { OFF  }
          { ON   }
          { time }
```

Positional Parameter 1

UNCON
: Specifies that the unattended console feature enabled during SYSGEN is to be disabled, reactivated, or changed.

Positional Parameter 2

OFF
: Specifies that the previously enabled (during SYSGEN) unattended console feature is to be removed.

ON
: Specifies that the unattended console feature enabled at SYSGEN and removed by a previous SET UNCON command is to be reactivated.

time
: Specifies the time in minutes for the unattended auto-answer. If this parameter is not specified, the time is 15 minutes.

8.5.15. Terminating System Activity (SHUTDOWN)

The SHUTDOWN command terminates system activity or distributed data processing (DDP) in an orderly manner.

Format

```
SHUTDOWN∆ { SYSTEM }
          { DDP    }
```

Positional Parameter 1

SYSTEM
: Specifies that system activity will be terminated. The spooler and job scheduler will not start any new files or jobs. Interactive services will not start any new functions and will terminate when its current activity ceases. This command *cannot* be abbreviated; the entire word SYSTEM is required to protect the user against accidental keyin and shutdown.

DDP
: Specifies that only distributed data processing will be terminated when its activity ceases. During DDP shutdown processing, all DDP activity is logged into the interactive services log file. The activity information remains in the log file and will be printed only when interactive services is shut down by using the IS SHUTDOWN message (see 9.2.4).

8.5.16. Initiating Transient Work Area Feature (TW)

The TW command initiates the transient work area feature. This feature can improve system performance. The most recently used transients are stored in the transient work area of main storage. Therefore, an I/O delay can be avoided by satisfying the transient load request directly from the work area rather than from SYSRES.

Format

```
TWΔ[SZ=nnn]
```

Positional Parameter 1

SZ=nnn

Specifies the number (*nnn*) of 1,024-byte blocks for the work area, where *nnn* must be in the range of 32 to 250. Any value outside of this range causes your system to terminate the transient work area and display the message INVALID PARAMETER.

If omitted, 65K is allocated for the work area by default.

You can initiate the transient work area whenever sufficient free main storage is available to satisfy the work area space requirements. However, you should initiate TW immediately after the IPL procedure to avoid possible fragmentation of main storage.

Once the transient work area is activated, the message TRANSIENT WORK AREA IS INITIALIZED appears. If you receive this message during the IPL procedure, it indicates that the work area in your supervisor generation has already been allocated to the specified size during system initialization.

To change the work area size, you must terminate the transient work area and then key in the TW command with a new value. To terminate the transient work area, key in:

```
00ΔTWΔEOJ
```

The message TRANSIENT WORK AREA TERMINATED then appears.

8.5.17. Specifying Resource Management (LIMITS)

The LIMITS command lets you control the percentage of your system's main storage space used for symbionts, interactive processing, and batch processing. This enables you to maximize system turnaround time and workstation response. Resource management lets you control these activities:

- Number of batch jobs run from both the console and workstations
- Number of batch jobs run from the workstations only

- Number of logged-on workstations (including workstations logged on to both interactive services and OS/3 interactive)
- Number of ENTER processing tasks
- Number of RUN symbionts

You can issue the LIMITS command at any time during the session.

If your system is primarily a batch or interactive system, you can specify resource management parameters during system generation (SYSGEN). See the *Installation Guide* (UP-8839) for a description of the resource management SYSGEN parameters.

If your system operations change throughout the day or week from batch to interactive, or you use a mixture of both, then use the LIMITS command to change your resource management controls as your processing needs change. When you reuse the LIMITS command, any parameters not specified remain unchanged.

Before using the LIMITS JOBMEM command to reduce the percentage of main storage available for batch jobs, be sure there is sufficient main storage for the jobs in the current running system. The MAXJOBS limit can be temporarily lowered to insure that no new jobs are initiated. After some jobs have terminated normally and the percentage of main storage in use falls below the new target level, then you can issue the LIMITS JOBMEM command and reset the MAXJOBS limit. This precaution helps avoid jobs from aborting due to insufficient memory.

Format

```
LIMITSΔ[SYMBMEM={5-100}][,INTMEM={5-100}][,JOBMEM={0-100}][,MAXJOBS=1-n]
                {NLMT }          {NLMT }          {NLMT }

      [,MAXWSJOBS=0-n][,MAXSWSJOBS=0-n][,MAXINTUSERS=0-255][,MAXENTERS=0-255]

      [,MAXRUNSYMBS=1-10][,MAXLOGONS={1-99}][,UPTERMINAL={tttt}]
                                     {5   }             {*ALL}
```

Parameters

```
SYMBMEM={5-100}
        {NLMT }
```

Specifies the percentage of available main storage your system uses for symbiont use. Either specify a percentage value between 5 and 100, or specify NLMT. NLMT specifies that resource management does not control that area of main storage allocated to symbionts. The value you specify can be greater than or less than the SYSGEN value.

INTMEM= {5-100 | NLMT}

Specifies the percentage of available main storage your system uses for interactive use. Either specify a percentage value between 5 and 100, or specify NLMT. NLMT specifies that resource management does not control that area of main storage allocated for interactive use. The value you specify can be greater than or less than the SYSGEN value.

JOBMEM= {0-100 | NLMT}

Specifies the percentage of available main storage your system uses for batch use. Either specify a percentage value between 0 and 100, or specify NLMT. NLMT specifies that resource management does not control that area of main storage allocated for batch use. The value you specify can be greater than or less than the SYSGEN value.

MAXJOBS=1-n

Specifies the number of jobs that can be executed concurrently. The value you choose cannot exceed the value you have specified for the SYSGEN JOBSLOTS parameter. The maximum number of jobs on models 3-6 is 14; the maximum number of jobs on models 8-20 is 48.

MAXWSJOBS=0-n

Specifies the maximum number of jobs that can be executed concurrently from workstations only. The value you choose cannot exceed the value you have specified for the SYSGEN MAXJOBS parameter, or the last LIMITS command MAXJOBS parameter.

MAXSWSJOBS=0-n

Specifies the maximum number of jobs that can be initiated concurrently from any single workstation. The value you choose cannot exceed the value you have specified for the SYSGEN MAXWSJOBS parameter or the last LIMITS command MAXWSJOBS parameter.

MAXINTUSERS=0-255

Specifies the maximum number of interactive users that can be logged on at one time. This parameter changes the value of the ISINTLMT parameter defined at SYSGEN or set with a previous LIMITS command.

MAXENTERS=0-255

Specifies the maximum number of batch tasks initiated with the ENTER command that can execute concurrently in the system. This parameter changes the value of the the ISBATCHLMT parameter defined at SYSGEN or set with a previous LIMITS command.

MAXRUNSYMBS=1-10

Specifies the maximum number of RUN symbionts that can be concurrently executed. This parameter changes the value of the MAXRUNSYMBS parameter defined at SYSGEN or set with a previous LIMITS command.

```
          { 1-99 }
MAXLOGONS={      }
          { 5    }
```

Specifies the maximum number of invalid logon attempts before a terminal is deactivated. The default value is 5.

```
           { tttt }
UPTERMINAL={      }
           { *ALL }
```

Specifies the terminal (*tttt*) to be reactivated. *ALL reactivates all terminals, whether they have been deactivated or not.

Examples

1. Let's say your system is used primarily for interactive jobs in the morning, for both batch and interactive jobs in the afternoon, and then only batch jobs in the evening.

   Upon user requests for a better balance between system turnaround time and workstation response time during each of these shifts, the operator issues the LIMITS command.

   First, before the morning shift, he issues:

   ```
   LIMITS JOBMEM=20,INTMEM=80,MAXRUNSYMBS=1
   ```

   This example shows a small value for the JOBMEM parameter, the parameter that specifies the percentage of main storage to be used for batch jobs. It shows a high value for INTMEM, the parameter that specifies the percentage of main storage to be used for interactive jobs. And it shows a low value for the number of RUN symbionts that can execute concurrently.

   Next, before the afternoon shift, the operator issues:

   ```
   LIMITS JOBMEM=70,INTMEM=70,MAXRUNSYMBS=2
   ```

   This example shows a balance of the percentage of main storage to be used for batch and interactive use. (Both JOBMEM and INTMEM are specified as 70). This time, the operator increases the number of RUN symbionts that can execute concurrently.

   Finally, before the evening shift, the operator issues:

   ```
   LIMITS JOBMEM=80,INTMEN=10,MAXRUNSYMBS=3
   ```

   This example shows a much higher percentage of main storage to be used for batch use (JOBMEM is specified as 80; but, INTMEM is specified as only 10.) The operator increases the number of RUN symbionts that can execute concurrently to 3.

2. The following example reactivates workstation C12 and changes the invalid logon limit to 3:

   ```
   LI  UP=C12,LOG=3
   ```

8.5.18. Verifying and Correcting the VTOC (VV)

The preventive maintenance procedures, VTOC verification and VTOC correction, are called using the VV system command. The VTOC verification routine determines whether volume table of contents (VTOC) corruption has occurred on an OS/3 disk. It performs a thorough analysis of a VTOC, looking for any inconsistencies between the labels in the VTOC. When it detects an inconsistency, it displays a message providing details on the error.

Use this routine under the following circumstances to detect VTOC problems as soon as possible:

- Initially, to verify that all volumes pass verification.
- After a system stop, to verify that all volumes in use at the time of the stop pass verification.
- When one of the error codes described under "Symptoms of a Compromised VTOC" is experienced.

The VTOC correction routine can correct volume label errors (VV08) only. See the VV08 message description in the *System Messages Reference Manual* (UP-8076) for a list of the specific error codes that are correctable.

It is not necessary to prevent other jobs from accessing the device specified when using the VTOC correction routine; however, you should use this facility only during periods of minimal activity in the system.

Before using the VTOC correction routine, you should make an alternate copy of the VTOC being corrected. (See Note 5, Figure 8-2.)

Format

```
VVΔ[did[,{END|FIX|ECnn|COPY|PACK}][,ALT]]
```

Positional Parameter 1

did
: Specifies a 1- to 4- character numeric field identifying the device address of the device whose VTOC is to be verified. If you omit this parameter, the system prompts you for the device address.

The VV command automatically locks the specified VTOC to prevent it from being accessed while it is being read or updated.

Positional Parameter 2

END

Terminates the routine when it completes the verification. If you omit this optional parameter:

- The system prompts you for another device address when the VTOC sucessfully passes verification.
- The routine terminates if verification fails.

FIX

Calls the VTOC correction routine and performs the following steps:

1. Complete VTOC verification
2. VTOC correction
3. Reverification

If the error encountered is not correctable, the system informs you that the error cannot be automatically corrected (VV22).

If you omit the FIX parameter, only verification is performed.

Note: *It is recommended that you specify the FIX parameter.*

ECnn

Causes a dump to be taken when the error code *nn* is encountered.

For example,

```
VV 190,EC18
```

causes a dump to be taken when VV08 EC18 is encountered. (See the *System Messages Reference Manual* (UP-8076).)

The symbiont produces an error code 33A, calls SYSDUMP to print the dump, and continues processing. (Select the MINI SYSDUMP option.)

Use this debug facility to provide Unisys with additional documentation for an unexplained error condition.

COPY

Copies the VTOC to a file named ALT$VTOC. You must create this file before using the COPY parameter.

PACK

Indicates that the VTOC is to be packed (i.e., reorganized). Format 1 labels are searched on all ALLOCATE, SCRATCH, EXTEND, OPEN, and CLOSE file management operations. The pack process consists of moving the Format 1 labels to the front of the VTOC to reduce search times when performing these common operations.

First a verify is performed to guarantee that the VTOC passes verification. Then the VTOC is copied to the alternate VTOC file and the VTOC is reorganized. Should a problem occur during the pack process, the VTOC is restored to its original contents. Do *not* cancel VV while a pack function is in progress.

The alternate VTOC file (ALT$VTOC must be allocated prior to performing the pack.

You should use PACK only when the device in question is *not* being used by any other jobs (i.e., VV can get exclusive use of the device). If you pack the VTOC on a system disk (i.e., RES, RUN, or SPOOL), you must guarantee that the system is *completely* idle and, immediately upon completion of the pack, you must re-IPL the system because there may be internal tables that contain invalid (i.e., old) format label addresses. You can pack multiple system disks in succession without a re-IPL after each one, but you must pack the RES device last and then re-IPL immediately.

Positional Parameter 3

ALT

Performs the specified routine on the alternate VTOC (ALT$VTOC). The VTOC must be copied to the alternate file, using the COPY parameter, before the ALT parameter is used.

This parameter is optional. If omitted, the specified routine is performed on the original VTOC.

When you specify ALT, all VTOC verify messages (information and error) include the label ALT$VTOC at the end of the message.

Note: *It is suggested that you correct an alternate VTOC (VV did,FIX,ALT) as a first trial before correcting your VTOC file.*

The volume table of contents (VTOC) is a file that resides on every disk and serves as a directory for all files on that disk. The address of the VTOC is contained in the VOL1 label, which resides in a fixed location on each disk. Each physical record (i.e., sector) in the VTOC file is one of six types of format labels.

The following summarizes the various format label types, presented according to their logical progression.

- FORMAT 4 (F4) - Contains device-dependent information about the volume and a pointer to the F0 chain. There is only one F4 label.
- FORMAT 5 (F5) - Contains information regarding available space on the disk. There can be multiple F5 labels.
- FORMAT 0 (F0) - An unreserved label type available for label usage by the system as necessary. There can be multiple F0 labels that are logically linked together by pointers.
- FORMAT 1 (F1) - Contains information regarding the physical space allocated to a specific file (i.e., physical extents).
- FORMAT 2 (F2) - Contains supplementary F1 information that specifies how the physical space allocated to the file is subdivided (i.e., logical extents).
- FORMAT 3 (F3) - Contains additional F1 logical extent information.

Cause of Compromised VTOC

The VTOC is updated when a new file is added or deleted, or when an existing file requires that two or more labels in the VTOC be updated. If the update does not complete due to a system stop (HPR) or unrecoverable I/O error, the VTOC may be left in a compromised state, which can lead to files being destroyed.

Symptoms of Compromised VTOC

- Error codes 30, 31, or 32

  Caused by an unrecoverable I/O error during a VTOC update.

- Error code 3A

  A definite indication of a compromised VTOC. Disk space management (DSM) reports this error when it detects a label inconsistency.

- Error codes 35 or 36 or DM41

  Generally occurs when there are no available (F0) labels or no available space. In this case, these errors are generally not a symptom of a compromised VTOC. However, if they occur when there should be labels or space available, then it could be a sign of a compromised VTOC.

- Invalid data

  If invalid data (e.g., data from a different file) is being returned to a program, this could be sign of a compromised VTOC.

Error Conditions during Verification

Several conditions prevent the verification routine from performing. When any of these is encountered, it is identified by a VV07 message followed by a VV02 message prompting you to enter a device address. Correct the condition specified by the VV07 message and reenter the device address.

Volume and file type errors are identified by VV08 and VV09 messages, respectively. It is possible to get several VV08 or VV09 messages or a combination of both.

During verification, the routine attempts to identify all volume and file type error conditions. However, certain error conditions may be too serious for the routine to continue. Correct the conditions listed, then perform the VTOC verification routine again to identify any remaining VV08 and VV09 errors.

Refer to the *System Messages Reference Manual* (UP-8076) for an explanation of VTOC verification messages and their error codes.

Corrective Action

The flow chart in Figure 8-2 shows the procedure to follow whenever a disk fails the VTOC verification routine.

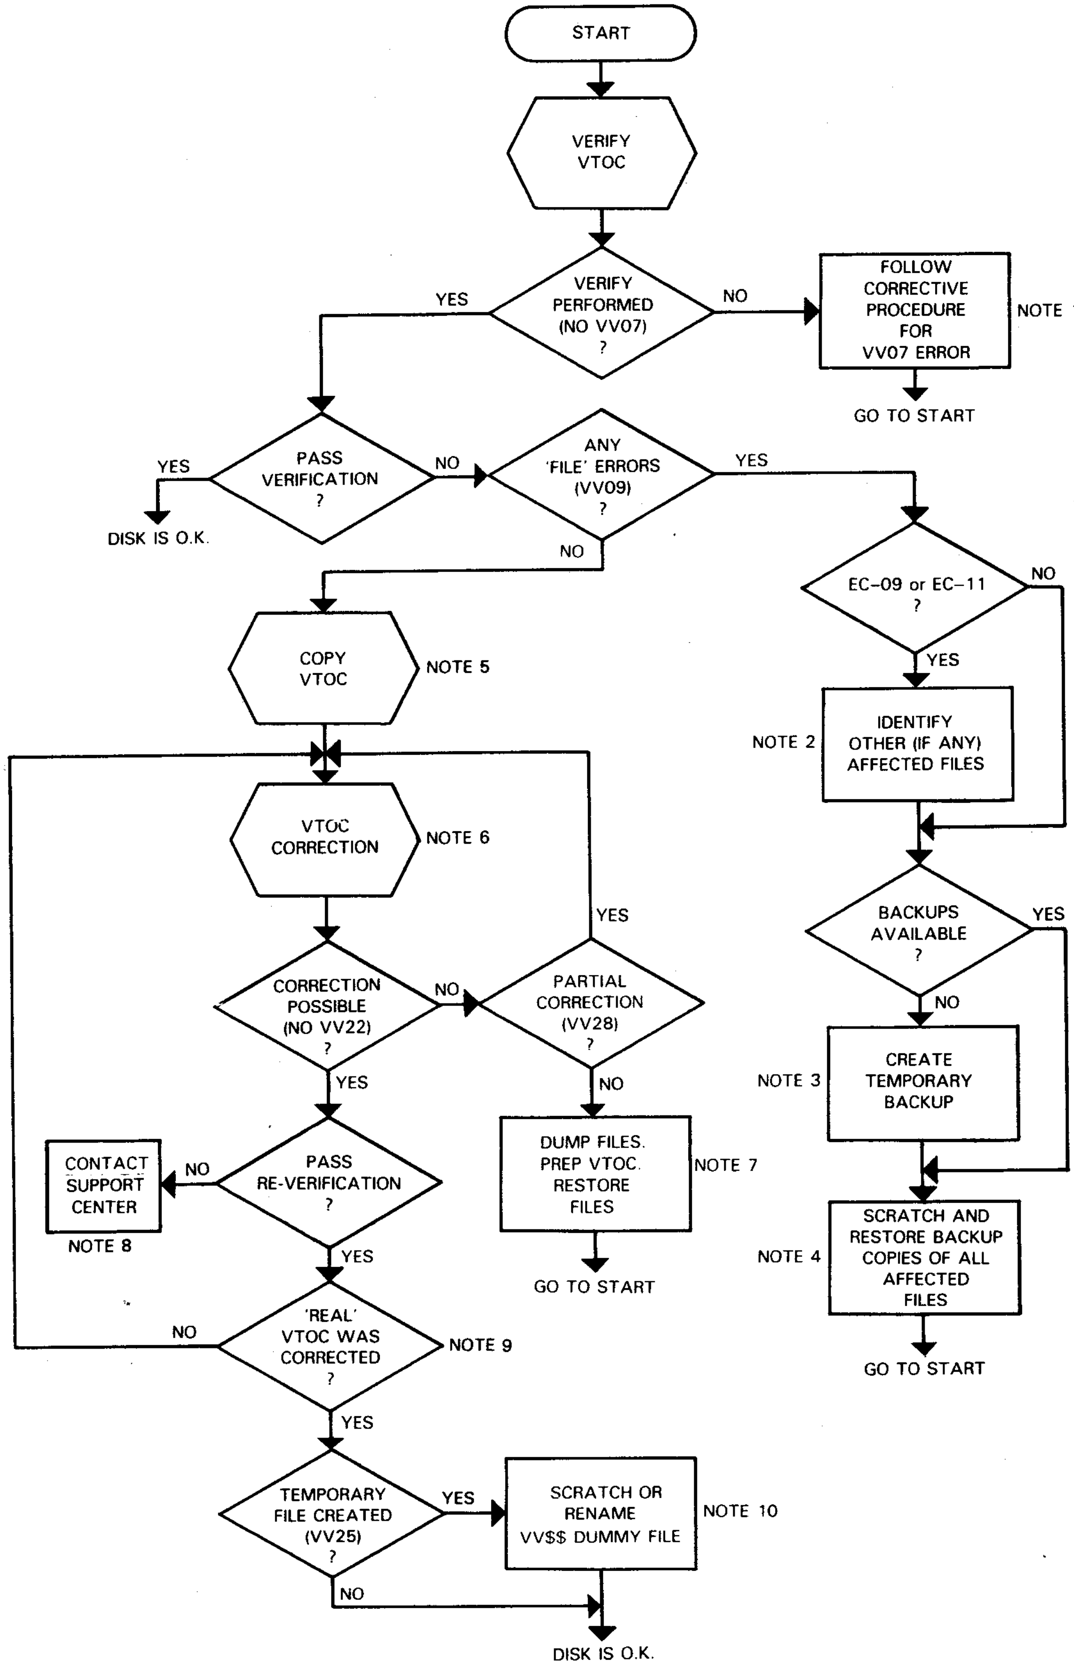

Figure 8-2. VTOC Verification/Correction Procedure (Part 1 of 2)

NOTES:

1. Some VV07 errors are user errors, such as wrong device address or device not on-line.

2. If you get a VV09 message with error code 09 or 11, other files besides the ones identified in the message may be affected.

   For example, when a cylinder is assigned to two different files, a VV09 message is displayed for the second file. The verification routine cannot determine which file is corrupted. You must use the SU VTP listing to determine whether there are other files whose disk space allocations (physical or logical) overlap the allocations of the file named. If so, these files might have been affected also.

3. If you do not have a previously created backup copy of the affected files, create a temporary backup by copying the files onto another disk. Remember, however, that the integrity of this copy may be affected by the incorrect labels.

4. Restore all potentially corrupted files (identified by VV09 messages or as in note 2 above) from a previously created backup copy using a copy program such as DATA utility or DUMP/RESTORE in file mode.

5. Using the COPY parameter, make an alternate VTOC for a trial correction.

6. Make two passes through the correction routine: the first, a trial correction on the alternate VTOC and the second, the correction of your VTOC file. If the trial correction is unsuccessful, your VTOC file is not affected.

7. Dump each file on the disk (except those identified by VV09 messages or identified as in note 2) to another medium using the DUMP/RESTORE utility in file mode. Then prep the VTOC of the bad disk and restore all of the files just dumped.

8. Should this situation ever occur, verify the disk using the ECnn parameter. Send the resulting dump and a copy of the old VTOC (on tape or diskette) to the Unisys Support Center.

9. If the correction was performed against the alternate VTOC, correct your VTOC file.

10. Space that was not accounted for in the previous step is assigned to a temporary file (VV$$DUMMY), giving you the opportunity to access that data. If you want to keep the data, rename the file. If not, you can scratch the file and return the space to the free space list.

Figure 8-2. VTOC Verification/Correction Procedure (Part 2 of 2)

Verifying the VTOC during AVR (SET VV)

The verify routine can be invoked explicitly via the system command VV or implicitly whenever an automatic volume recognition (AVR) is performed. AVR is performed under the following circumstances:

- During IPL
- Whenever a disk is set online
- Explicitly, via the AVR command

Automatic VTOC verification AVR time can be requested in two ways:

1. Via the VVAVR SYSGEN parameter. See the *Installation Guide* (UP-8839) for details on the VVAVR parameter.

2. Via the SET VV command.

Format

```
SET VV,{ON }
       {OFF}
```

Positional Parameter 1

ON

Performs automatic VTOC verification at AVR time.

OFF

Does not perform automatic VTOC verification at AVR time. This overrides the VVAVR SYSGEN specification, described in the *Installation Guide* (UP-8839).

Submitting a User Communication Form

If a volume that previously passed verification suddenly fails verification when there have been no system stops or unrecoverable I/O errors to cause the VTOC to be compromised, report this occurrence to Unisys in a User Communication Form (UCF).

You should enclose with the form as much information as possible about the jobs that accessed the volume that was compromised. If possible, copy the disk and repeat the steps taken, verifying the disk after each program access to it. Then enclose with the UCF:

- A formatted VTOC print (SU VTP) and a disk print of the VTOC (SU DD) taken both before and after the program is run
- The JCL for the job(s) in question
- A console log
- An error log
- Any additional information useful in duplicating the problem (for example, operational information, copies of programs, or copies of disks)

8.5.19. Setting Memory Consolidation Conditions (SET MC)

The SET MC command enables or disables the memory consolidation (MEMCON) feature. MEMCON is enabled to allow movable jobs to be consolidated at the highest possible addresses, thus creating the largest possible free regions. In order to move each job, all the jobs must be inactive (no processing occurs within that job) prior to moving the job. In a system that is not fully utilizing its memory resources, it may not be desirable for this to occur. In this case, MEMCON should be disabled.

Ongoing MEMCONs at the time of the keyin are not affected by the command.

Format

```
SET MC, { ON  }
        { OFF }
```

Positional Parameter 1

MC
: Specifies that the MEMCON condition is to be altered.

Positional Parameter 2

ON
: Enables memory consolidation.

OFF
: Disables memory consolidation.

8.5.20. Setting Job Immovability Conditions (SET IM)

The SET IM command enables or disables the job immovability (IMVJOB) feature. Some jobs present a problem to MEMCON because they become immovable. Normally IMVJOB is enabled to move those jobs to the high addressed portion of memory. This reduces the possibility that that particular job will become a roadblock to an ongoing memory consolidation. To perform this move, several jobs may have to be idled simultaneously, thus temporarily preventing any processing by those jobs involved. If this condition is not desirable to the operator, IMVJOB can be disabled by this command.

An ongoing move at the time of the keyin is not affected by the command.

ROLLOUT must have been configured for the current supervisor in order for IMVJOB to be enabled by this command.

Format

```
SET IM, { ON  }
        { OFF }
```

Positional Parameter 1

`IM`
Specifies that the IMVJOB condition is to be altered.

Positional Parameter 2

`ON`
Enables job immovability moves.

`OFF`
Disables job immovability moves.

8.5.21. Displaying Buffer Pool Information (DI BI)

The DI BI command displays system buffer pool information at the system console.

Format

```
DI BI
```

Example

The following is a sample system console display of buffer pool information for the DI BI command.

```
33 BR0022    BREAKPOINT TAKEN FOR CONSOLE LOG 88/09/08
DI BI
34 DI0024     EXPANSION REGION SIZE (BYTES)...................          4,096
35 DI0024     RESERVE REGION SIZE (BYTES).....................            500
36 DI0024     DISK CACHE BUFFER SIZE (BYTES)..................      2,050,048
37 DI0024     TRANSIENT CACHE BUFFER SIZE (BYTES).............              0
38 DI0024     TOTAL JOB/SYMBIONT BUFFER SIZE (BYTES)..........         59,272
39 DI0024     TOTAL ALLOCATED BUFFERS (BYTES).................      2,109,320
40 DI0024     CURRENT NUMBER OF BUFFERS ALLOCATED.............             43
41 DI0024     LARGEST NUMBER OF BUFFERS ALLOCATED.............             49
42 DI0024     LARGEST NO. OF JOB/SYMB. BUFFERS USED (BYTES)...         77,812
43 DI0024    DISPLAY END
```

where:

`EXPANSION REGION SIZE (BYTES)`
Specified in a SYSGEN/IPL parameter.

`RESERVE REGION SIZE (BYTES)`
Specified in a SYSGEN parameter.

`DISK CACHE BUFFER SIZE (BYTES)`
Specified in a SYSGEN/IPL parameter/console command.

`TRANSIENT CACHE BUFFER SIZE (BYTES)`
: Specified in a SYSGEN parameter/console command.

`TOTAL JOB/SYMBIONT BUFFER SIZE (BYTES)`
: Is the total size of buffers currently allocated to jobs or symbionts.

`TOTAL ALLOCATED BUFFERS (BYTES)`
: Is the total size of disk cache, transient cache, and job/symbiont buffers.

`CURRENT NUMBER OF BUFFERS ALLOCATED`
: Is the number of buffer regions currently allocated in memory.

`LARGEST NUMBER OF BUFFERS ALLOCATED`
: Is the largest number of buffer regions allocated in memory since the last IPL.

`LARGEST NO. OF JOB/SYMB. BUFFERS USED (BYTES)`
: Is the largest number of buffer allocated to jobs or symbionts since the last IPL.

8.5.22. Enabling/Disabling the Console/Workstation Alarm Beeper (SET AL)

This command turns alarm beeper feature on or off at the console or workstation.

Format

```
SET ALARM, {ONE}
           {ON }
           {OFF}
```

Parameters

`ONE`
: Turns the alarm beeper feature on. Only a single warning beep is sounded.

`ON`
: Turns the alarm beeper feature on. Continuous warning beeps are sounded.

`OFF`
: Turns the alarm beeper feature off.

8.5.23. Enabling/Disabling the Console Video Clock (SET VC)

This command turns on or off the console time-of-day clock. This clock is located at the upper right corner of the console and is normally incremented every 10 seconds.

Format

```
SET VC,{ON }
       {OFF}
```

Parameters

ON

Turns on the console time-of-day clock.

OFF

Turns off the console time-of-day clock. The clock display goes blank and it is no longer incremented.

Section 9

Interactive Services

You use interactive services with an extended set of commands and messages to control the interactive system environment. (For a description of all interactive services commands, see the *Interactive Services Operating Guide* (UP-9972.) These commands and messages enable you to exercise control over the interactive OS/3 operating system, all jobs within the system, all workstation users (local locations), and all terminal users (local and remote locations).

9.1. Interactive Services Commands

The interactive services components are loaded automatically whenever required for the system operator and workstation users. Provided ICAM and the global user service task (GUST) are ready, a workstation user's interactive entry or your command loads interactive services for terminal users. Interactive services commands are used to:

- Send messages to a workstation
- Ask questions of workstation and terminal users
- Display the volumes in use and the status of active workstations, terminals, jobs, and functions
- Display status of system resources in use and available
- Display the log file
- Modify bulletin/log logon default setting and override values
- Load interactive services
- Terminate workstation and terminal user tasks or sessions

9.1.1. Sending Messages to Users (TELL)

The TELL command sends a message not requiring a response to a specific workstation or terminal user or to all users.

Format

```
TELLΔ { ALL     } ,'text'
      { user-id }
```

Positional Parameter 1

ALL
Displays the message on all active user screens.

user-id
Specifies the 1- to 6-character alphanumeric identification of the user to receive the message display.

Positional Parameter 2

text
Specifies the text of the message to be sent. The text may be a maximum of 88 characters long. Text must be preceded and followed by apostrophes.

9.1.2. Asking Questions of Users (ASK)

The ASK command sends a message that requires a response to a specific workstation or terminal user. The command displays your question to the specified user, accepts the reply, and returns the reply to the console workstation screen.

Format

```
ASKΔuser-id,'text'
```

Positional Parameter 1

user-id
Specifies the 1- to 6-character alphanumeric identification of the local or remote user to receive the message display.

Positional Parameter 2

text
Specifies the text of the message to be sent. The text may be a maximum of 88 characters long. Text must be preceded and followed by apostrophes.

9.1.3. Displaying System Status (STATUS)

The STATUS command displays the volumes currently in use and the status of active workstations, terminals, jobs, and functions. It also displays the status of system resources in use and their availability and limitations.

Format

```
STATUSΔ[ { JOBS              } ]
       [ { FUNCTIONS         } ]
       [ { RESOURCES         } ]
       [ { TERMINALS         } ]
       [ { TERMINALS[,uid]   } ]
       [ { VOLUMES           } ]
       [ { LIMITS            } ]
```

Positional Parameter 1

JOBS
: Displays a listing of the jobs and symbionts currently active in the system. The list includes the job or symbiont name and the amount of storage and CPU time used by each, information on what program each job is executing and its job step number, the job slot number for the job, and the master user-id for the job or symbiont. A summary line displays the amount of free main storage in the system and the largest contiguous region.

FUNCTIONS
: Displays a listing of all active commands and tasks initiated by users.

RESOURCES
: Displays a listing of the amount of storage being used by all users; the number of interactive, batch, and batch with interactive commands (ENTER) tasks; the number of active jobs and configured job slots; and the amount of total system storage.

TERMINALS
: Displays a listing of all currently active (logged on) workstations and terminals and the associated user-id.

TERMINALS,[uid]
: Displays a listing of all currently active (logged-on) terminals and associated user-ids that match the user-id you specify. The *uid* parameter must be between 1 and 6 alphanumeric characters.

VOLUMES
: Displays a listing of the tape, disk, and diskette volumes currently mounted on the system.

LIMITS
: Displays the system resource management limits values and their current status.

Example 1

Operator keyin:

```
STA V
```

Function requested:

Display a listing of the tape, disk, and diskette volumes currently mounted on the system.

Screen display:

```
D-REL110 D-REL12X D-OS3WRK D-REL120 D-REL11X
IS90 STATUS  COMMAND TERMINATED NORMALLY
```

Disk volumes REL110, REL12X, OS3WRK, REL120, and REL11X are currently mounted on the system.

Example 2

Operator keyin:

```
STA T
```

Function requested:

Display a listing of all the currently active workstations and terminals and their associated user-ids. In a secure system, display all terminals with at least one invalid log-on attempt and number of log-on attempts remaining for them.

Screen display:

```
TRM USERID     TRM USERID     TRM USERID     TRM USERID
313  PUBSI     314 PUBS
```

Workstations 313 and 314 are currently logged on under user-ids PUBSI and PUBS, respectively.

Additional screen display if security is active:

```
TRM            LOGON TRIES LEFT

     C12                 3
     (MAXLOGONS)         5
IS90 STATUS   COMMAND TERMINATED NORMALLY
```

Example 3

Operator keyin:

```
STA J
```

Function requested:

Display a listing of the jobs and symbionts currently active in the system.

Screen display:

```
JOB NAME   SIZE    CPU TIME  STEP   EXEC    JOB NO MASTER
RC$$IS00  006912       7.3    00   RC$$IS   00000
SL$$VT00  001536        .0    00   SL$$VT   00000  PUBS
SFGEN03   029696       2.9    01   SF$GEN   00003  PUBSI
I@DATA02  036864       4.8    01   DATA00   00002  PUBS
UNUSED MEMORY 0781312
IS90 STATUS   COMMAND TERMINATED NORMALLY
```

The jobs and symbionts currently active are listed in the first column. The SIZE column indicates the amount of storage each job is using. The CPU time is given next, followed by the job step number and the name of the program each job is executing. The JOB NO column is the job slot number of each job; MASTER provides the master user-id for the job or symbiont. The amount of unused memory is also shown.

Example 4

Operator keyin:

```
STA L
```

Function requested:

Display a listing of the current values and available limits for the resource management controls. These controls were set by the last LIMITS command or by resource management parameters specified during system generation.

Screen display:

```
JOBS     WSJOBS     SWSJOBS     INTUSER     ENTERS     RUNSYMBS
02/10    01/07        01          12/15      10/30       01/01
  SYMBMEM                   INTMEM                 JOBMEM
05/10% (030K/060K)     55/80% (330K/480K)     08/10% (048K/060K)
```

The values on this screen show that this user system will perform as follows:

JOBS 02/10
: Can execute 10 jobs concurrently, but is now executing only 2.

WSJOBS 01/07
: Can execute seven jobs concurrently from workstations only, but is now executing only one.

SWSJOBS 01
: Can execute only one job from a single workstation at a time.

INTUSER 12/15
: Can support up to 15 interactive users logged on at any one time, but only 12 users are logged on at this moment.

ENTERS 10/30
: Can support 30 batch sessions executing concurrently, but is now supporting only 10.

RUNSYMBS 01/01
: Can support one run symbiont executing at a time and is now supporting it.

SYMBMEM 05/10% (030K/060K)
: Can use up to 10 percent (or 60K) of available main storage for symbiont use, but is now using only 5 percent (or 30K) for symbiont use.

INTMEM 55/80% (330K/480K)
: Can use up to 80 percent (or 480K) of available main storage for interactive use, but is now using 55 percent (or 330K) for interactive use.

JOBMEM 08/10% (048K/060K)
: Can use up to 10 percent (or 60K) of available main storage for batch use, but is now using only 8 percent (or 48K) for batch use.

Note: *If you specified NLMT for the SYMBMEM, INTMEM, or JOBMEM parameters, then the STATUS command listing would display NL as your upper limit for those three values.*

9.1.4. Displaying Log File (RECALL)

The RECALL command lets you display all or part of your log file at your workstation screen or at the system console. RECALL only uses the lines that are available on the screen for the recall.

Format

```
RECALLΔ⎧LASTΔnn            ⎫[,prefix]
       ⎩hh:mm:ss-hh:mm:ss⎭
```

Parameters

LASTΔnn

Indicates the number of messages (*nn*) in the log file that you wish to see.

hh:mm:ss-hh:mm:ss

Indicates that you wish to see the contents of the log file for a particular time period. All time must be specified in military time. The minutes (*mm*) and seconds (*ss*) are optional. You must, however, separate the beginning time from the ending time with a dash (-). To display all messages from a specific time until the end of the log file, specify only a beginning time. To display all messages from the beginning of the log file until a specific time, specify only an ending time. When keying in an ending time only, you must prefix it with a dash.

prefix

Indicates that you only wish to view those messages that begin with this prefix. The prefix can be from one to eight characters.

9.1.5. Modifying Bulletin and/or Log Values (SET IS)

The SET IS console command allows you to modify the bulletin/log logon default setting and override values during normal system processing.

Format

```
SETΔIS,⎧BULLDEF⎫,⎧YES⎫
       ⎪BULLOVR⎪ ⎩NO ⎭
       ⎪WLOGDEF⎪
       ⎩WLOGOVR⎭
```

Parameters

BULLDEF and WLOGDEF
: Specify whether YES or NO appears as the default on the logon screen. These specifications are also valid for a user who logs on batch or does a smart logon (from system mode).

BULLOVR and WLOGOVR
: Specify whether the interactive services user logging on can change the values that appear as defaults. These specifications also apply to batch and smart logons. If NO is specified and the user attempts to change the value, the following message appears, but the user will remain logged on providing there are no other logon errors:

```
IS136  USER NOT ALLOWED TO CHANGE BULLETIN AND/OR LOG VALUES
```

9.1.6. Starting Interactive Services (IS)

Normally, the interactive services facility is started automatically for the system operator and all workstation users, whenever needed. However, there are occasions when you must manually load interactive services via the IS command. Enter the IS command with the REMOTE START parameter to manually start interactive services when ICAM and GUST are ready and interactive services are needed for terminals; but no workstation entry has automatically loaded them. When ICAM and GUST are ready and the interactive services facility is active (already loaded) before being needed for terminals, use the IS REMOTE unsolicited message (see 9.2.3) to start it. If, however, ICAM and GUST are not ready, terminal users cannot initiate interactive sessions. (See Section 10 for the operator procedures on loading ICAM and running the global user service task.)

You also use the IS command to manually start interactive services after a completed shutdown (termination). When you terminate interactive services via the SHUTDOWN message (9.2.4), the interactive services facility is placed in a closed condition and is unable to start automatically when needed. After a completed shutdown, you enter the IS command to remove the closed condition and restart interactive services.

Format

```
ISΔ[REMOTEΔSTART]
```

Positional Parameter

REMOTE START
: Loads interactive services for terminals when no workstation entry has automatically loaded them.

9.1.7. Terminating User Tasks or Sessions (REMOVE)

The REMOVE command terminates a single command for a specific task, a specific workstation or terminal user session, or all user sessions. Following the termination, a cancellation message is displayed on the terminated user's screen.

Format

$$\texttt{REMOVE}\Delta \left\{ \begin{array}{l} \texttt{task-id} \\ \texttt{user-id} \\ \texttt{ALL} \end{array} \right\}$$

where:

`task-id`
: Specifies the particular task under which a command is to terminate. Task-ids can be determined by using the STATUS FUNCTION command.

`user-id`
: Specifies the particular user whose session is to terminate. User-ids can be determined by using the STATUS TERMINAL command. All interactive functions for the user-id specified are terminated and the user is logged off. If the user is running an interactive session as a batch job (via ENTER), the user is logged off; however, the session is not affected. A message is displayed on the console workstation or console screen to indicate when the user cannot be logged off.

`ALL`
: Specifies all user sessions and batch runs are to terminate for each user-id that can be controlled by this user. All interactive functions are terminated and the users are logged off. Users running interactive sessions as batch jobs (via ENTER) are logged off. A message indicating which users cannot be logged off is displayed on the console workstation or console screen.

9.2. Interactive Services Messages

You use unsolicited messages provided by interactive services for additional control over the interactive environment. These messages are used to:

- Restrict and release new workstation and terminal user sessions
- Control interactive services for terminals
- Terminate interactive services

9.2.1. Restricting New User Sessions (CLOSE Message)

The CLOSE unsolicited message restricts any new workstation or terminal user sessions from starting. Currently active sessions are not affected.

Format

```
00ΔISΔCLOSE
```

9.2.2. Releasing New User Session Restrictions (OPEN Message)

The OPEN unsolicited message removes a previously entered CLOSE or reverses an incomplete SHUTDOWN to permit new workstation and terminal user sessions to start.

Format

```
00ΔISΔOPEN
```

9.2.3. Controlling Interactive Services for Terminals (REMOTE Message)

The REMOTE unsolicited message provides additional control over interactive services for terminals. You must load ICAM and run the global user service task (GUST) before you start interactive services for terminals. (See Section 10 for a description of the ICAM and GUST operator procedures.)

Format

$$\texttt{00}\Delta\texttt{IS}\Delta\texttt{REMOTE}\Delta\left\{\begin{array}{l}\texttt{START}\\ \texttt{SHUTDOWN}\\ \texttt{CANCEL}\end{array}\right\}$$

Parameters

START
: Starts interactive services for terminals after ICAM and GUST are ready. You use this parameter when interactive services are already active in the system (via a prior IS command or an automatic start by a workstation user). If you enter the REMOTE START message with ICAM and GUST not ready, an error message is displayed on the console workstation or console screen.

SHUTDOWN
: Terminates the interactive services facility for terminals as soon as the last terminal session completes. (Workstation users are not affected.)

CANCEL
: Immediately terminates the interactive services facility for terminals. Terminal users are not logged off.

9.2.4. Terminating Interactive Services (SHUTDOWN Message)

The SHUTDOWN unsolicited message terminates interactive services after all sessions have completed. If no sessions are active when you enter the command, it takes effect immediately. Otherwise, interactive services do not terminate until all active sessions complete. An OPEN message can be used to stop a shutdown in progress (incomplete). After the interactive services facility is completely shut down, you can restart it only via the IS command (9.1.6); it is not automatically loaded for workstation users via an interactive services command entry.

Format

```
00ΔISΔSHUTDOWN
```

Section 10
Integrated Communications Access Method (ICAM) Procedures

The integrated communications access method (ICAM) is an extension of the supervisor (a symbiont) that handles data communications tasks. Each symbiont can contain multiple network definitions (CCAs); each CCA can handle one or more communications lines.

At system generation time, the ICAM symbionts are tailored to each user's requirements. One or more ICAM symbionts can be configured during SYSGEN, each symbiont satisfying specific communications network requirements; or a single ICAM symbiont can be configured to satisfy all communications network requirements.

You must load the appropriate ICAM symbiont before the programs requiring it can execute or before interactive services can start for terminals. In addition, when interactive services or global networks are required, you must initiate the running of the global user service task (GUST) after loading ICAM.

The ICAM symbiont remains in main storage until GUST is shut down and all other CCAs are released by their associated CUPs. Then ICAM shuts itself down unless the system operator loaded ICAM with a KEEP operand. In this case, ICAM must be terminated with a CANCEL command.

The following subsections describe how to load ICAM, change the ICAM name established during SYSGEN, run and terminate the global user service task, and use messages to control the active communications environment.

10.1. ICAM Operator Commands

10.1.1. Loading the ICAM Symbiont (Cn/Mn)

The Cn/Mn command loads the ICAM symbiont. The symbiont name you specify in the command (C1-C9 or M1-M9) is normally assigned during SYSGEN.

Format

$$\left\{\begin{matrix}\texttt{Cn}\\\texttt{Mn}\end{matrix}\right\}\Delta\texttt{[KEEP]}$$

where:

```
{Cn}
{Mn}
```

Specifies the name of the required ICAM symbiont, where n is 1 to 9. It is the name specified on the MCPNAME parameter in the COMMCT phase of system generation.

```
[KEEP]
```

Keeps the ICAM symbiont loaded until cancelled by the system operator or if ICAM suffers an unrecoverable error.

When ICAM is successfully loaded, the following output message is displayed on the console workstation or console screen (where *xx.x* is the ICAM release level; for example, 12.0):

```
ICAM xx.x READY
```

Note: *Users can also load ICAM from a job control stream with the / / CC job control statement. This statement allows users to load ICAM as part of their IMS execution job or as part of a job that executes the global user service task.*

10.1.2. Changing the ICAM Name (SET IC)

The SET IC command changes the name of the ICAM symbiont that is loaded if remote batch output has output ready and ICAM is not loaded. The ICAM name entered at SYSGEN or by the operator using SET IC must be C1-C9 or M1-M9. By specifying C?, you are asked to supply the symbiont name to be used the next time output is ready and ICAM is not loaded.

Format

```
SETΔIC,{Cn}
       {Mn}
       {C?}
```

where:

```
IC
```

Specifies the SPOOLICAM SYSGEN parameter is to change.

```
{Cn}
{Mn}
{C?}
```

Specifies the ICAM symbiont name to use (C1-C9 or M1-M9) or specifies that you be asked for the symbiont name to be used when ICAM is not currently loaded to send remote batch output (C?).

10.2. Initializing and Terminating the Global User Service Task (GUST)

You must initialize the global user service task before starting interactive services for terminals or before executing user programs requiring global networks. You initiate the running of the job that executes the global user service task program ML$$GI through a console workstation or console command entry. (See your system administrator for the name of the GUST job to initiate.) When global network processing is no longer required for interactive services at terminals or for user programs, you enter an unsolicited message to shut GUST down.

10.2.1. Running the Global User Service Task Job

The RUN command is entered with the GUST job name to initiate ML$$GI.

Format

```
RUNΔjobname
```

where:

`jobname`
: Is the name of the prefiled job control stream run to execute the GUST ML$$GI program.

10.2.2. ML$$GI Program Operator Messages

When you execute ML$$GI, it sends the following messages to the console workstation or console to obtain the information it requires to initialize the global network. Respond to each message as indicated. (Only those messages pertaining to GUST initialization are presented here. See the *System Messages Reference Manual* (UP-8076) for all global user service task messages.)

For the message:

```
MC#420 ENTER NETREQ: CCA, PASS, RESTART, LINEREQ
```

Enter the name, password, restart, and line information associated with the global network to be initialized. When multinode global networks are used, ML$$GI automatically requests all virtual lines.

where:

`CCA`
: Is the name of the global network to be initialized. This name must be the same as the label of the CCA macroinstruction that begins the network definition for the global network.

PASS

Is a 1- to 8-character password.

RESTART

Is Y for yes or N for no. This determines the type of restart GUST requests for the network. If Y, GUST sets a restart flag to initiate a warm restart, i.e., messages are recovered from existing disk files. If you respond N, normal initialization is begun. That is, new disk files are created if required.

LINEREQ

Specifies the lines for which line requests are to be issued. Response may be one of the following:

ALL,Y

Issue line requests for all lines in the network. Any errors are treated as non-fatal and the network is brought up.

line-1,...,line-n

Specifies the lines for which line requests are to be issued. Lines specified must be identical to those specified in label fields of related LINE macroinstructions.

*

Issue line requests for the lines specified in forthcoming message MC#421.

For the message:

```
MC#421 ENTER LNEREQS: LINE-1, LINE-2,...* OR BLANK
```

Enter the lines for which line requests are issued. This message is displayed only when you respond to the MC#420 message with an asterisk for the line request.

where:

LINE-1, LINE-2,...line-n

Are the names of lines to be activated by means of line requests. GUST issues a LNEREQ statement for each line-name you specify. The line-name must be identical to that specified in the label field of the related LINE macroinstruction.

BLANK

Signals the GUST program to issue no specified line requests at this time. (VLINES and RWS lines are automatically requested during GUST processing.)

For the message:

```
MC#430 GUST ACTIVE FOR CCA nnnn
```

No response is required. This message specifies the GUST initialization is complete for the network named by nnnn.

10.2.3. Terminating the Global User Service Task

The following unsolicited message ends global network processing by canceling the GUST job.

Format

```
00Δ { Cn } ΔGUΔS,network-name
    { Mn }
```

where:

```
{ Cn }
{ Mn }
```

Specifies the currently loaded ICAM symbiont name (C1-C9 or M1-M9).

`GU`

Specifies this message is for GUST.

`S`

Specifies a shutdown is required.

`network-name`

Is the 1- to 4-character name of the active global network. This name must be the same as the label of the CCA macroinstruction in the global network definition.

If the shutdown request is accepted, the following message is displayed when the global user service task job is canceled:

```
MC#401 GUST SHUTDOWN COMPLETE
```

Note: *To resume communications processing with global networks or to start interactive services for terminals, the operator procedures for loading ICAM and initializing the GUST must be repeated. You should never cancel the GUST job the CANCEL operator command. If you do so, the CANCEL command and the requested GUST shutdown command will be rejected and the following message will be displayed:*

```
MC#402 CANCEL REJECTED, ISSUE SHUTDOWN COMMAND
```

10.3. ICAM Operator Messages

Occasionally, the system console operator must communicate with ICAM to change the communications environment or to advise ICAM of some external event. The console operator uses unsolicited messages to do this. We have divided our descriptions of the ICAM console messages into three types:

- Standard messages - those not related to public data networks or remote batch processing
- PDN unsolicited messages - those related to public data networks
- RBP unsolicited messages - those having to do with remote batch processing

10.3.1. Standard Unsolicited Messages

Standard unsolicited messages to ICAM have the following format:

$$00\Delta\left\{\begin{matrix}\texttt{Cn}\\\texttt{Mn}\end{matrix}\right\}\Delta cc\Delta f,\left\{\begin{matrix}\texttt{xxxx}\\\texttt{ccpp}\end{matrix}\right\},nnnn$$

where:

`Cn`
`Mn`
: Specifies the name of the required ICAM symbiont.

`cc`
: Is the command code. Some command codes are 2 characters long; some are 4 characters.

`f`
: Is a 1-character facility type (L=line, P=port, T=terminal).

`xxxx`
: Is a 1- to 4-character name of a line or terminal as defined in the label field of a LINE or TERM macroinstruction.

`ccpp`
: Specifies the input/output microprocessor (IOMP) channel number (*cc*) and the single line communications adapter id (*pp*) on the specified channel. Select your parameters from the following list:

Processor Model	Channel No. (cc)	SLCA No. (pp)
Models 3-6	02	08 to 15
Models 8-20 with 1 IOMP	13	01 to 14
Models 8-20 with 2 IOMPs	13 or 15	01 to 14

nnnn

Is the 1- to 4-character name of the network as specified in the label field of the CCA macroinstruction.

Table 10-1 lists the standard unsolicited messages ICAM provides for you.

A typical ICAM response to an unsolicited message is:

```
NET1 TERM TRM1 MARKED UP
```

This response tells the user that the terminal whose name is TRM1 on the network whose name is NET1 is now available.

Table 10-1. Standard ICAM Unsolicited Messages

Message	Description
00Δ{Cn / Mn}ΔUPΔL,xxxx,nnnn	Marks line specified as available (up). Activates a communications line. Equivalent to a line request issued by a program. ICAM can poll, send output, and prepare to receive input. No traffic can flow over a nonactivated communications line.
00Δ{Cn / Mn}ΔDOΔL,xxxx,nnnn	Marks line specified as unavailable (down). Stops all activity on a communications line. ICAM does not poll, send, or receive traffic. A switched line is disconnected. ICAM does not answer calls or dial out.
00Δ{Cn / Mn}ΔUPΔT,xxxx,nnnn	Marks terminal (UNISCOPE or DCP) or remote workstation specified as available (UP). Enables ICAM to send output to a terminal previously marked down.
00Δ{Cn / Mn}ΔDOΔT,xxxx,nnnn	Marks terminal (UNISCOPE or DCP) or remote workstation specified as unavailable (down). Stops ICAM from sending output to the terminal specified and aborts any current session at the terminal. If the last terminal in a poll group (UNISCOPE or remote workstation) is marked down, the entire poll group is placed on slow poll (the group is polled about once every minute). If input is entered at one of these terminals, a response is delayed until the next time the terminal is polled and the terminal marked up. When a remote workstation is marked down, its session as well as its partner's session is aborted. (Partners are defined by the primary/secondary specifications on adjacent TERM macroinstructions.)
00Δ{Cn / Mn}ΔUPΔP,ccpp,nnnn	Marks port specified as available (UP). Useful when an ICAM symbiont contains more than one network and communications lines are shared between them. To reassign a line: - Down the port where the line is connected (specify the network currently using the line). - Using this command (UP port), reactive the port. (Specify the network to receive the line). - Issue an UP command to active the communications line. (Specify the network now using the line.)

continued

Table 10-1. Standard ICAM Unsolicited Messages (cont.)

Message	Description
00Δ{Cn}ΔDOΔP,ccpp,nnnn {Mn}	Marks port specified as unavailable (down). See the preceding message description.
00Δ{Cn}ΔCNΔL,xxxx,nnnn {Mn}	Notifies ICAM that dialing is completed on switched line specified. This is a response to a dialing message from ICAM to dial a manually dialed line. ICAM does not pass traffic on the switched line until it receives a response.
00Δ{Cn}ΔCNΔL,ALL,nnnn {Mn}	Notifies ICAM that dialing is completed on all switched lines. This is a response to one or more ICAM messages to manually dial switched lines.
00Δ{Cn}ΔSTΔL,llll,cccc {Mn}	Displays status of a communications line (llll) in a CCA (cccc).
00Δ{Cn}ΔSTΔL,*ALL,cccc {Mn}	Displays status of all communications lines in a CCA (cccc).
00Δ{Cn}ΔSTΔT,tttt,cccc {Mn}	Displays status of a terminal/workstation (tttt) within a CCA (cccc).
00Δ{Cn}ΔSTΔU,uuuu,cccc {Mn}	Displays status of a LOCAP (uuuu) within a CCA (cccc).

10.3.2. Public Data Network (PDN) Unsolicited Messages

There are two kinds of public data networks: packet-switched and circuit-switched. The formats for entering unsolicited messages to ICAM vary with the type of public data network involved. The following paragraphs describe these formats. (See the *ICAM Operations Guide* (UP-9745) for information about public data networks.)

Note: *In addition to the commands described in the following paragraphs, ICAM lines connected to public data networks use the standard ICAM UP and DOWN commands to activate and deactivate lines, terminals, or ports.*

How to Enter Packet-Switched PDN Unsolicited Messages

Packet-switched public data network unsolicited messages have the following format:

```
00Δ {Cn} ΔccccΔnnnn,llll[,dddd]
    {Mn}
```

where:

00

Means this message is destined for the supervisor.

Cn or Mn

Is the name of the ICAM symbiont to which this type-in is destined (C1-C9 or M1-M9).

cccc

Is the 4-character command code: STAT, CONN, DISC, or TEST.

where:

STAT

Causes the status of the level 2 link, the virtual circuits, or the subscribers to be displayed on the operator console. The llll parameter specifies the type of status requested.

CONN

Reestablishes a connection to the level 2 link after it is disconnected.

DISC

Drops (disestablishes) a connection to the level 2 link.

TEST

For DDX only, initiates level 2 self-testing procedures.

nnnn

Is the 1- to 4-character network name as specified in the label field of the CCA macroinstruction.

llll

For CONN, DISC, and TEST, this is a 1- to 4-character label of the VLINE macroinstruction that defines the link. For the STAT command code, this parameter indicates the type of status to be displayed:

- Enter link name llll as instructed for the other command codes for the status of the link.
- Enter VCT to display the status of the virtual circuit.
- Enter SUBS to display the status of the subscriber.

dddd

Is the 1- to 4-character label of the PDN macroinstruction that identifies the packet-switched public data network affected by this type-in. This parameter is optional and, if omitted, defaults to the label of the first PDN macroinstruction in the network definition.

How to Enter Circuit-Switched PDN Unsolicited Messages

The following formats are available for entering circuit-switched public data network unsolicited commands to ICAM.

In general, when the PDN keyword is omitted, the default is the label on the first PDN macroinstruction in the ICAM network definition. When the LDTE keyword is omitted, the default is the label on the first LDTE macroinstruction in the ICAM network definition.

Format 1

```
00Δ {Cn} ΔSTATΔnnn, {LDTE=ldte-name       } [,PDN=pdn-name]
    {Mn}            {RDTE=rdte-name       }
                    {TERM=term-name       }
                    {CPS= {cc }           }
                    {     {ALL}           }
```

Displays the status of the specified LDTE, RDTE, terminal (TERM), or call progress signal (CPS), where:

{Cn | Mn}
: Is the name of the ICAM symbiont for which this command is destined (C1-C9 or M1-M9).

nnnn
: Is the 1- to 4-character network name as specified in the label field of the CCA macroinstruction.

ldte-name
: Refers to the label on an LDTE macroinstruction in the ICAM network definition.

rdte-name
: Refers to the label on an RDTE macroinstruction in the ICAM network definition.

term-name
: Refers to the label on a TERM macroinstruction in the ICAM network definition.

cc
: Refers to a valid call progress signal number.

pdn-name
: Refers to the label on a PDN macroinstruction in the ICAM network definition.

Format 2

```
00Δ { Cn } ΔNETSΔnnn, { LDTE=ldte-name          }  [,PDN=pdn-name]
    { Mn }            { RDTE=rdte-name          }
                      { CPS= { cc  }            }
                      {      { ALL }            }
```

Displays the network statistics for the specified LDTE, RDTE, or call progress signal (CPS).

Format 3

```
00Δ { Cn } ΔCONNΔnnn,RDTE=rdte-name[,LDTE=ldte-name][,PDN=pdn-name]
    { Mn }
```

Establishes a connection to the specified RDTE.

Format 4

```
00Δ { Cn } ΔDISCΔnnn,RDTE=rdte-name[,LDTE=ldte-name][,PDN=pdn-name]
    { Mn }
```

Disconnects from the specified RDTE.

Format 5

```
00Δ { Cn } ΔCHNGΔnnn,RDTE=rdte-name { ,ADDR=address } [LDTE=ldte-name][,PDN=pdn-name]
    { Mn }                          {  DSCO=t       }
                                    {  DSCI=t       }
                                    {  TIME=t       }
```

For the specified RDTE, changes:

a. Network address (ADDR=address) (maximum 16 digits)

b. For terminal connections

   - Time before disconnect following output (DSCO) (range 0-32767)

   - Time before disconnect following input (DSCI) (range 0-32767)

c. For host to host connections, the activity timer (TIME) (range 0-32767)

Format 6

```
00Δ { Cn } ΔCHNGΔnnn,CPS=cc { ,NIR=n } [LDTE=ldte-name][,PDN=pdn-name]
    { Mn }                  {  TIR=n }
                            {  NSR=n }
                            {  TSR=n }
                            {  TBD=n }
```

For the specified call progress signal (CPS), changes:

a. ICAM retry parameters where

NIR = number of ICAM retries (range 0-32767)

TIR = time between ICAM retries (range 0-32767)

b. SLCA retry parameters where

NSR = number of SLCA retries (range 0-63)

TSR = time between SLCA retries (range 0-255)

c. Time before disconnect (TBD) (range 0-254)

For further descriptions of these parameters and their usage, refer to the *ICAM Operations Guide* (UP-9745).

10.3.3. Remote Batch Processing (RBP) Unsolicited Messages

The OS/3 console operator communicates with the remote batch processing ICAM symbiont by means of console commands in the form of unsolicited messages. These commands enable the operator to specify functions he would like performed. The format of a remote batch processing unsolicited message is:

```
00Δ { Cn } Δcc[,D]
    { Mn }
```

where:

{Cn / Mn}
Is a 2-character remote batch processing ICAM symbiont identifier (C1-C9) or (M1-M9).

cc
Is a 2-character command code.

D
Is an optional dump operand associated with the SH command.

The following command codes are accepted by remote batch processing. They are explained in detail in the *ICAM Utilities Programming Guide* (UP-9748).

- SH

  Specifies shutdown and removes the ICAM symbiont from the system when remote batch processing is the only active ICAM interface or idles remote batch processing when other ICAM interfaces are active in the same ICAM symbiont. If you specify the D option on the SH command, your system produces a dump during termination of the ICAM symbiont.

- RB

  Means remote batch and reactivates remote batch processing after it has been idled by an SH command.

- RD

  Starts the reading of batch jobs or data files from the central processor card reader under control of remote batch processing.

Section 11
System Utility Services

The system utility symbiont (SL$$SU) is a multipurpose utility that allows you to perform many card, tape, disk, and diskette functions at the console workstation or console. To initiate a function, you include its function code as a parameter with the SU/TU command or enter it as a solicited message. The functions you can perform and their function codes are listed in Table 11-1. The SU/TU command and the functions for card, tape, disk, and diskete are described in 11.1 through 11.5.

Table 11-1. System Utility Functions

Function Code	Function Performed
Card Functions	
CC	Reproducing cards punched in Hollerith code
CCB	Reproducing cards punched in binary and Hollerith code
CCS	Reproducing and resequencing source programs
CT	Writing card to tape in unblocked format
CTR	Writing card to tape in blocked format
CP	Listing cards
CH	Listing cards containing compressed mode
JCP	Punching cards from the system console
Tape Functions	
TT	Copying a tape to another tape
TH	Printing a tape in character and hexadecimal format
THR	Printing a tape in character, hexadecimal, deblocked format
TP	Printing a tape containing only standard characters
TPR	Printing a tape in character and deblocked format
TRS	Locating a specific record on tape
TC	Punching cards from tape

continued

Table 11-1. System Utility Functions (cont.)

Function Code	Function Performed
Tape Functions (cont)	
INT	Prepping a tape
FSF	Forward space to a specific file
BSF	Backward space to a specific file
FSR	Forward space to a specific record
BSR	Backward space to a specific record
WTM	Writing tape marks
REW	Rewinding a tape
RUN	Rewinding a tape with interlock
ERG	Erasing a portion of a tape
Disk and Format Label Diskette Functions	
DD	Printing a disk in unblocked format
DDR	Printing a disk in deblocked format
VTP	Printing the volume table of contents of a disk
SVT	Printing short format VTOC file
AVX	Displaying available disk extents on console screen
Data Set Label Diskette Functions	
DD	Printing a diskette in unblocked format
DDR	Printing a diskette in deblocked format
VTP	Printing the data set labels of a diskette

NOTE:

If XXX is entered in place of the function code, all function codes are displayed.

11.1. System Utility Commands and Messages (SU/TU)

The SU/TU command loads the system utility symbiont. SU and TU can be used interchangeably for all functions. However, we recommend that you use TU for tape operations because it increases the buffer size for tapes from 8,189 to 32,767 bytes. You can include the required function as a parameter with the SU/TU command.

A spooling parameter also can be entered with the command, if spooling is configured in your system. (For the description of all spooling commands, see the *Spooling and Job Accounting Operating Guide* (UP-9975).)

When you enter the command alone to load the symbiont, you enter the function as a solicited message. After the symbiont is loaded, you control it by responding with solicited messages. You use unsolicited messages only to terminate the symbiont or current symbiont functions on certain occasions.

Note: *The SU/TU command cannot be issued from an enter stream.*

Format

$$\left\{\begin{matrix}\texttt{SU}\\\texttt{TU}\end{matrix}\right\}\Delta\left[\begin{matrix}\texttt{function-code}\\\texttt{XXX}\end{matrix}\left[,\left\{\begin{matrix}\underline{\texttt{Y}}\\\texttt{N}\\\texttt{R}\\\texttt{H}\end{matrix}\right\}\right]\right]$$

Positional Parameter 1

`function-code`

Specifies the appropriate 2- or 3-character function code. (Function codes are listed in Table 11-1.) If omitted, the symbiont displays a message requesting that you enter a function code. All possible function codes that are recognized by either symbiont can be displayed on the console or console workstation by entering XXX in place of the function code. Following this display, the symbiont requests that you enter the required function code.

Positional Parameter 2

`Y`

Spools the system utility output.

`N`

Specifies that the system utility output is not spooled.

`R`

Retains system utility output in the HOLD condition in the spool file after it is printed or punched. The retained output is unavailable for additional processing until released via the BEGIN SPL command.

H

Places system utility output in the HOLD condition in the spool file before it is printed or punched. The output must be released via the BEGIN SPL command to permit processing.

This parameter is entered only if spooling is configured in your system. When Y is specified or taken as the default condition, the output writer automatically prints any spooled output at the end of every SU function.

Operator Considerations

- When the symbiont is loaded, the following message is displayed:

```
nn∆SUNNNN SYSTEM UTILITY SYMBIONT LOADED
```

where:

nn

Is the 2-digit message number

SUNNNN

Is the name of the symbiont sending the message. This name is the 2-character symbiont identification concatenated with a 4-digit binary jobnumber inserted at task initiation.

- If the command entry includes a function code, the symbiont completes the requested function, then displays an ENTER REQUIRED FUNCTION message to allow you to either terminate the symbiont or request another function.

If you omit the function code, the symbiont displays the ENTER REQUIRED FUNCTION message to allow you to enter a function.

The ENTER REQUIRED FUNCTION message is displayed as follows:

a. With spooling:

```
nn?SUNNNN ENTER REQUIRED FUNCTION AND SPOOL OPTION [,Y,N,R,H];DEFAULT=Y
```

b. Without spooling:

```
nn?SUNNNN ENTER REQUIRED FUNCTION
```

Initiate the required function with the keyin:

a. With spooling:

```
nn∆function-code, ⎧Y⎫
                  ⎪N⎪
                  ⎨R⎬
                  ⎩H⎭
```

b. Without spooling:

```
nnΔfunction-code
```

Each time the symbiont completes a requested function, it transmits the ENTER REQUIRED FUNCTION message to allow you to either terminate the symbiont or request another function.

- If your system supports spooling and the spool option is incorrectly entered (i.e., a character other than Y, N, R, or H is entered), the following message is displayed:

```
nn?SUNNNN IS {PRINTED} OUTPUT TO BE SPOOLED FROM SU  Y,N,R,H
             {PUNCHED}
```

Enter the required spooling option with the keyin:

```
nnΔ { Y }
    { N }
    { R }
    { H }
```

To terminate the symbiont, you reply to the ENTER REQUIRED FUNCTION with the end-of-job keyin:

```
nnΔEOJ
```

The symbiont terminates and displays the message:

```
nnΔSUNNNN SYSTEM UTILITY SYMBIONT ENDED
```

- You can also end the symbiont by pressing the MSG WAIT key on the console workstation or console and keying in the applicable unsolicited message:

```
00ΔSUΔEOJ
00ΔTUΔEOJ
```

This permits you to terminate the symbiont before it completes a function.

- To terminate only the current function of the symbiont, key in the applicable unsolicited message:

```
00ΔSUΔEND
00ΔTUΔEND
```

Note: *When message replies are keyed in incorrectly or the reply cannot be honored, the symbiont requests the information to be keyed in again. If no determination can be made on why the keyed input is not accepted, use the previous unsolicited message to terminate the current function or to terminate the symbiont.*

11.2. Card Functions

All the card functions you can perform are described in this subsection. All input card files must be terminated by a card with the words END OF DATA punched in columns 1 through 11.

Proceed as follows to perform the card functions.

1. As described in 11.1:

   a. Enter the SU symbiont command.

   b. Enter the appropriate function code either as a command parameter or as a solicited message response to the ENTER REQUIRED FUNCTION message.

   c. Enter the spooling option; otherwise, default is Y (applicable only if spooling is configured).

2. If a card file is being read, place it in the card reader designated the system reader (SYSRDR). If this card reader is unavailable, the first available card reader is assigned to the symbiont, causing the following message to be displayed on the console workstation or console:

   ```
   nnΔSUNNNNN USE READER did
   ```

   where:

   did
   : Is the device address of the card reader assigned to read the input file.

   If no card readers are available, the function aborts and the following message is displayed on the console workstation or console:

   ```
   nnΔSUNNNNN NO READER AVAILABLE
   ```

   Likewise, if the required output device is not available, the function aborts and the following message is displayed on the console workstation or console:

   ```
   nnΔSUNNNNN NO { PUNCH   } AVAILABLE
                 { TAPE    }
                 { PRINTER }
   ```

If the required devices are available, the operation continues for each function code as described in 11.2.1 through 11.2.8.

11.2.1. Reproducing Cards Punched in Hollerith Code (CC)

You use the CC function code to reproduce cards in 80 x 80 format containing the Hollerith code. All job control cards, even the /*, can be reproduced by using this function. You must submit an END OF DATA card with your input deck indicating the end of file to the symbiont.

11.2.2. Reproducing Cards Containing Binary Data (CCB)

You use the CCB function to reproduce cards containing binary data in addition to the Hollerith code. Again, you must submit an END OF DATA card as the last card in your input deck. When punching column binary, the output must not be spooled.

11.2.3. Reproducing and Resequencing Source Programs (CCS)

You use the CCS function code to reproduce and resequence an assembler (BAL), COBOL, or RPG II source language program. For a BAL program, the program name can be up to three characters in length; COBOL can be up to eight characters in length; and RPG II can be up to six characters in length. If you supply a name having fewer characters than the number permitted, the name is left-justified and space-filled. You must submit an END OF DATA card as the last card of your source program. Tables 11-2 through 11-4 show the formats of the source programs being reproduced and resequenced.

Table 11-2. Assembler Format

Column	Description
1–72	Source statement
73–75	3-character program name
76–80	5-character sequence number; first number is 00010. Succeeding numbers are incremented by 10.

Table 11-3. COBOL Format

Column	Description
1–3	Page sequence number; starts at 001 and is incremented by every 20 lines (cards)
4–6	Line sequence number; starts at 010 and is incremented by 10 for every line up to 200
7–72	Source statement
73–80	8-character program name

Table 11-4. RPG Format

Column	Description
1-2	Page number sequence; starts at 01
3-5	Line number sequence; starts at 010 and is incremented by 10 for every line up to 200
6-74	Source statement
75-80	6-character program name

11.2.4. Writing Cards to Tape in Unblocked Format (CT)

You use the CT function code to write cards to an unblocked tape in 80 x 80 format. You must submit an END OF DATA card in your input deck indicating the end of file whereby two tape marks are then written on your tape.

Procedure

Copies a deck of cards onto an unlabeled output tape in 80 x 80 format. No leading tape marks are written, but two tape marks are written when the END OF DATA card is read. Each data block is 80 bytes long. You specify the tape unit device address, and the density and mode in which the tape is written at execution time. The tape is not rewound at either the beginning or the end of this function.

For the message:

```
nn?SUNNNN CUUMMB OUTPUT TAPE B=BLK CNT
```

You respond with the tape unit device address, mode setting, and block count characteristics of the output tape (as described in 11.3.1).

If the output tape reaches EOV before all the card input is processed, SU writes two tapemarks, rewinds, and unloads the tape.

For the message:

```
nn?SUNNNN MOUNT NEXT OUTPUT TAPE. REPLY R OR END
```

Key in the following if the CT function is to be terminated:

```
nn△END
```

Otherwise, dismount the unloaded tape, mount the next output tape; then

Key in:

```
nnΔR
```

For the message:

```
nn?SUNNNN DOES USER WISH TO PREP THIS TAPE? REPLY YES OR NO
```

Key in:

`nnΔN`
: If no tape prep is required. The CT function resumes, using the new output tape.

`nnΔY`
: To prep the tape. The SU INT tape function is activated (see "Prepping a Tape (INT)" in 11.3.3). The CT function resumes when INT function is completed.

11.2.5. Writing Cards to Tape in Blocked Format (CTR)

You use the CTR function code to write cards to a tape in blocked format. The blocking factor is in the range of 1 to 100 for SU and 1 to 400 for TU.

Procedure

Same as the CT function, except that tape blocks written are 80 x *f* long, where *f* is the blocking factor. If the last block is not full, the length of the last block is 80 x *r*, where *r* is the number of records in the last block.

For the message:

```
nn?SUNNNN BLK FACTOR
```

You respond with blocking factor 1 to 100 (SU) or 1 to 400 (TU).

For the message:

```
nn?SUNNNN CUUMM OUTPUT TAPE B=BLK CNT
```

You respond with the tape unit device address, and the mode setting of the output tape (see 11.3.1).

For the message:

```
nn?SUNNNN OUTPUT EXCEEDS ALLOCATED BUFFER
```

You must reenter the function and blocking factor because the factor supplied was greater than the maximum allowed and the function was terminated.

Multivolume output tape is described in "Tape Addressing" in 11.3.3.

11.2.6. Listing Cards (CP)

You use the CP function code to list cards in 80 x 80 format on the printer. You must submit an END OF DATA card as the last card in your input. Only printable punch configurations are printed; however, any standard punch configuration is accepted. The function is useful when listing your job control cards.

11.2.7. Listing Cards Containing Compressed Mode (CH)

You use the CH function code to list cards containing the compress mode (hexadecimal characters) and the standard characters. Again, you must have an END OF DATA card as the last card in your input deck.

11.2.8. Punching Cards from the System Console (JCP)

The JCP function punches job control or data cards entered through the console workstation or console. Up to 60 columns may be entered in reply to the scale message. If 60 or more columns are needed, position the cursor under the 0 of 60 in the scale message and transmit. Another message requesting 20 more characters will appear. If a card having a blank in column 1 or column 61 is required, key in a right parenthesis instead of a blank in that column. The symbiont replaces a right parenthesis in column 1 or 61 with a blank. The right parenthesis is needed because all messages received by SU are returned left-justified. To terminate the JCP function, key in /END and transmit immediately.

11.3. Tape Functions

Use the SU command if your tape block size does not exceed 8189 bytes; TU supports tape blocks up to 32,767 bytes.

11.3.1. Tape Addressing

The tapes to be used for a TU function must be identified using the tape unit ID. Communication by you is initiated by the following message:

```
cuummb
```

where:

cuu
: Is device ID (channel and unit).

mm
: Is mode setting of the tape. If *mm* is blank, 00, or not entered (e.g., response is 102), the SYSGEN mode settings are assumed. Commonly used mode settings are C0 (for 1600 BPI PE) and D0 (for 6250 BPI GCR). Refer to the *System 80 OS/3 Job Control Programming Guide* (UP-9986) for information on mode settings.

b
: Is block count characteristics of the tape. If blank or omitted, the tape is assumed not to have a block count. If b is entered, the tape is assumed to have a block count.

Notes:

1. *The block count specification is not needed for tape functions used to position a tape (e.g., FSF).*

2. *The record number (REC) printed by the system utility for the TRS function is relative to the beginning of scan. For other tape functions, it is relative to where printing begins. The REC number does not correspond to the 3-byte block number on block numbered tapes.*

3. *The block number (TBLK) printed for the TRS function is a display of the 3-byte block number prefixed to block-numbered tapes.*

11.3.2. Tape Error Processing

If an error is encountered on an input tape, control is turned over to a tape error correction routine, where communication is established with you to determine whether the error should be ignored or bypassed or the function terminated.

If an error is ignored, the record is processed as is and the tape is forward spaced to the next block. If an error is bypassed, no information from the bad block is processed and the tape is forward spaced to the next block.

The messages:

```
nnΔSUNNNN TAPE ERROR ON INPUT TAPE
nn?SUNNNN B-BYPASS, I-IGNORE, OR E-END FUNCTION
```

indicates a tape error has occurred and one of the following options may be replied:

```
nnΔBYPASS     (This reply is not applicable during a TRL function.)
nnΔIGNORE
nnΔEND
```

The message:

```
nnΔSUNNNN WARNING: TAPE BLK FILLS AVAILABLE BUFFER OR MAY BE TRUNCATED
```

warns that a tape input or output record completely fills the allocated buffer. SU reads/writes a maximum of 8189 bytes and TU reads/writes a maximum of 32,767 bytes; any tape block exceeding this is truncated.

11.3.3. Tape Operating Instructions

To perform a tape function, proceed as follows:

1. As described in 11.2:

   a. Enter the TU symbiont command (or enter SU if the tape block size does not exceed 8189 bytes).

   b. Enter the appropriate function code either as a command parameter or as a solicited message response to the ENTER REQUIRED FUNCTION message.

   c. Enter the spooling option; otherwise, default is Y (applicable if spooling is configured).

2. Place the input tape volume on an available tape unit, and identify the tape unit to the TU (or SU) symbiont by keying in its unit address as described in 11.4.1.

If the required devices are available, the operation continues for each function code, as described in this subsection.

Copying a Tape to Another Tape (TT)

You use the TT function code to copy from one tape to another tape of the same or different device type. You can copy either labeled or unlabeled tapes in blocked or unblocked format. If you are copying standard label tapes, the file marks are used as controlling devices. The first file mark on the tape includes the header record; the second file mark includes all of your data; and the third file mark includes the trailer record. Therefore, three file marks make up one complete file.

Procedure

For the messages:

```
nn?SUNNNN CUUMMB-INPUT TAPE B=BLK CNT
```

Respond with the tape unit device address, and the mode setting and block count characteristics of the input tape (11.3.1).

```
nn?SUNNNN CUUMMB-OUTPUT TAPE B=BLK CNT
```

Respond with the tape unit device address, and the mode setting and block count characteristics of the output tape.

```
nn?SUNNNN #FILES 1-99
```

Respond with the number of files or file marks to be copied. The number must be from 1 through 99. See the previous description of TT for the number of files on a standard labeled tape.

If a high density tape is copied to a lower density, or if you copy a longer tape to a shorter tape, all of the input tape data may not fit on the output tape. SU/TU provides the option of mounting another output tape to continue. When the output tape reaches EOV before all the tape input is processed, SU/TU:

- Writes two tape marks
- Writes EOV labels, when appropriate, to correspond to a previously copied HDR1 HDR2 label set
- Rewinds and unloads the tape

For the message:

```
nn?SUNNNN MOUNT NEXT OUTPUT TAPE. REPLY R OR END
```

Key in:

```
nnΔEND
```

If the TT function is to be terminated.

Otherwise, dismount the unloaded tape, mount the next output tape; then

Key in:

```
nnΔR
```

If the input tape and the newly mounted output tape are prepped, the TT function resumes. If the input tape is prepped but the new output tape is not, SU/TU queries the operator.

For the messages:

```
nnΔSUNNNN SUPPLY VOLUME SERIAL NUMBER FOR NEW OUTPUT TAPE
nn?SUNNNN XXXXXX
```

Key in:

```
nnΔvsn
```

where:

vsn

Is the 6-character volume serial number for the new output tape. If the vsn is less than six characters, it is left-justified by SU/TU. The TT function resumes, using the new output tape.

For the message:

```
USER LABEL RESTRICTIONS APPLY.
```

This message cautions you that SU/TU does not generate user trailer or header labels for multivolume output in addition to those actually encountered on the input tape.

For the message:

```
nn?SUNNNN END OF VOLUME? Y OR N
```

This message is issued after all files specified are copied. Key in Y to write a second tape mark after the file; otherwise N if it is not to be written. After processing this response, the function is completed.

An example of tape copy with block count is as follows:

System message:

```
nn?SUNNNN CUUMMB - INPUT TAPE
```

Operator response:

```
nnΔ100ΔΔB
```

System message:

```
nn?SUNNNN  CUUMMB - OUTPUT TAPE
```

Operator response:

```
nnΔ101ΔΔB
```

System message:

```
nn?SUNNNN #FILES 1-99
```

Operator response:

```
nnΔ3
```

System message:

```
nn?SUNNNN END OF VOLUME? Y/N
```

Operator response:

```
nnΔY
```

Printing a Tape in Character or Hexadecimal Format (TH)

You use the TH function code to print a tape containing either standard or compressed mode (hexadecimal) characters. The tape error correction routine is enabled to allow you to bypass or ignore tape errors with this function. Tape positioning does not occur at either the beginning or the end of this function. The block number printed is relative to the location where printing begins.

Procedure

An option allows you to print the entire field or only a specified number of blocks:

```
nn?SUNNNN CUUMMB - INPUT TAPE B=BLK CNT
```

Respond by entering the tape unit device address and the mode setting and block count characteristics of the input tape (see 11.3.1). If you enter a block count, data is considered to begin in position 3, relative to position 0, for a length of blocksize minus 3 bytes.

```
nn?SUNNNN #BLKS OR END
```

Respond END to print the entire tape, or respond with the number of blocks to be printed. If the end-of-file code (that is, two tape marks side by side) is detected before the specified number of blocks is printed, the function terminates.

Printing a Tape in Character or Hexadecimal Deblocked Format (THR)

You use the THR function code to print records from a tape individually rather than in a contiguous string. Prints logical tape records in character and vertical hexadecimal formats. The tape error correction routine is enabled to allow you to correct, bypass, or ignore tape errors for this function.

Procedure

An option allows you to print the entire file or only a specified number of blocks.

```
nn?SUNNNN CUUMMB - INPUT TAPE B=BLK CNT
```

Enter the tape unit device address, and the mode setting and block count characteristic of the input tape (see 11.3.1).

```
nn?SUNNNN #BLKS OR END
```

If the entire tape is to be printed, enter END or the number of blocks to be printed. If the end of file (two tapemarks side by side) is detected before the specified number of blocks are read, the function is terminated.

```
nn?SUNNNN LOGICAL REC LNTH
```

Respond with the logical record length.

Printing a Tape Containing Only Standard Characters (TP)

You use the TP function code to print a tape in blocked format containing only standard characters. This function is identical to the TH function, except that your output is only in character format rather than character and hexadecimal format.

Printing a Tape in Character and Deblocked Format (TPR)

You use the TPR function code to print records from a tape individually in character format only. This function is identical to the THR function, except that your output is in character format rather than character and hexadecimal format.

Locating a Specific Record (TRS)

You use the TRS function code to search for one or more specific logical records in your file. The following information is needed for the symbiont to locate the record:

- Logical record length
- Length scan argument (length of the data field)
- Starting data position in your record
- Whether the data is in character or hexadecimal format
- Actual data needed

Procedure

The scan field is identified by the operator with the length and location within the logical record. The logical record length is also required for deblocking purposes. When you enter the scan argument in either hexadecimal or character format, it is compared to the scan field of each logical tape record and printed on the printer for future reference. When an equal is located, the block having the record is printed in character and hexadecimal format.

After each find, you are given the option of continuing the scan for more finds or terminating the function. Tape positioning does not occur at either the beginning or the end of this function. The tape error correction route is enabled, so you can ignore or bypass tape errors (see 11.3.2). The function is terminated either at the end-of-file (two tape marks side by side) or when a find is made and you terminate the function. The scale is printed to improve readability. You also enter the scan argument for use in the search.

For the messages:

```
nn△SUNNNN INVALID HEX CHAR
```

An error message is printed when the scan argument is entered in hexadecimal and an invalid hexadecimal character was detected. You reenter the sequence to obtain the scan argument.

```
nn△SUNNNN NO MATCH FND
```

This message is printed when an end-of-file mark (two tape marks back to back) is encountered before any finds were made. The function is terminated.

```
nn△SUNNNN CONTINUE SCAN? Y/N
```

This message is printed each time a find is made. You enter Y to continue the scan or N to terminate the scan.

The displayed find record is preceded by a header that includes:

`REC`
: Specifies the sequential number of the record, which is relative to the beginning of the search.

`DATA`
: Specifies the logical record length.

`TBLK`
: Specifies the block number, which is maintained by the system in a 3-byte prefix to each tape block.

An example of a tape with block count that is to be searched for a tape record containing a key field of "4637275467" in the first 10 bytes of the record is as follows:

System message:

```
nn?SUNNNN CUUMMB - INPUT TAPE
```

Operator response:

```
nnΔ10000B
```

System message:

```
nn?SUNNNN LOGICAL REC LENGTH
```

Operator response:

```
nnΔ125
```

System message:

```
nn?SUNNNN LENGTH ARGUMENT (1-30|0/)
```

Operator response:

```
nnΔ10
```

System message:

```
nn?SUNNNN STARTING DATA POSITION IN REC
```

Operator response:

```
nnΔ1
```

System message:

```
nnΔSUNNNN ENTER IN HEX-H, CHAR-C
```

Operator response:

```
nnΔC
```

System messages:

```
nnΔSUNNNN ENTER 10 BYTES, 1 CHAR PER BYTE
nn?SUNNNN 1...5...10...
```

Operator response:

```
nnΔ4637275467
```

System message:

```
nn?SUNNNN CONTINUE SCAN? Y/N
```

Operator response:

```
nnΔN
```

Punching Cards from a Tape (TC)

You use the TC function code to punch cards from either a blocked or unblocked tape. The data portion of the tape is transferred 80 bytes at a time to the card. If the data portion of the block is not a multiple of 80, the remaining data is left-justified and space-filled. When a tape mark is detected, a card containing all asterisks (*) is punched. When two tape marks are detected back to back, the function ends.

Procedure

When a tape mark is encountered, a card with all * is punched. The function is terminated when two tape marks back to back are encountered. Tape error processing is described in 11.3.2.

For the message:

```
nn?SUNNNN CUUMMB - INPUT TAPE B=BLK CNT
```

Respond with the tape unit device address, and the mode setting and block characteristics of the input tape (as described in 11.3.1).

Prepping a Tape (INT)

You use the INT function code to prep a tape by writing a standard label on a specified tape. When you enter the volume serial number and file label, the tape is positioned to the load point prior to the prep and the VOL1, HDR1, and HDR2 labels are written, along with two tape marks. When the prep is completed, the tape is positioned following the first tape mark. No AVR is performed; the volume serial number is not entered in the pub. It is, therefore, possible to prep multiple tapes with the same vsn.

Procedure

For the messages:

```
nn?SUNNNN CUUMMB - OUTPUT TAPE B=BLK CNT
```

Respond with the tape unit device address and the mode setting and block count characteristic of the output tape (refer to 11.3.1).

```
nnΔSUNNNN ENTER NEW FILE LABEL
nn?SUNNNN XXXXXXXXXXXXXXXXX
```

Enter up to six digits for VOL1 record. If you enter fewer than six digits, the new volume serial number is left-justified and space is filled on the right.

```
nnΔSUNNNN ENTER OWNER I.D. OR SCRATCH
nn?SUNNNN XXXXXXXXXXXXXXXXX
```

Enter up to 17 characters for a new owner identifier.

Example

System message:

```
nn?SUNNNN CUUMMB - OUTPUT TAPE
```

Operator response:

```
nnΔ10100B
```

System messages:

```
nnΔSUNNNN ENTER NEW VOL#
nn?SUNNNN XXXXXX
```

Operator response:

```
nnΔ001036
```

System messages:

```
nnΔSUNNNN ENTER NEW FILE LABEL
nn?SUNNNN XXXXXXXXXXXXXXXXX
```

Operator response:

```
nnΔTAPE1
```

System messages:

```
nnΔSUNNNN ENTER OWNER I.D. OR SCRATCH
nn?SUNNNN XXXXXXXXXXXXXXXXX
```

Operator response:

```
nnΔMr. Smith
```

Forward Space to a Specific File (FSF)

You use the FSF function code to advance the tape to the next tape mark.

Procedure

Advances the specified tape to the next tape mark.

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be spaced forward (refer to 11.3.1).

Backward Space to a Specific File (BSF)

You use the BSF function code to backspace the tape until a tape mark is encountered.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be backspaced (refer to 11.3.1).

Forward Space to a Specific Record (FSR)

You use the FSR function code to forward space a specific number of blocks on your tape.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be forward spaced (refer to 11.3.1).

```
nn?SUNNNN #BLKS
```

Enter the number of blocks to be advanced (maximum number is 9999).

Backward Space to a Specific Record (BSR)

You use the BSR function code to backspace a specific number of blocks on your tape.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be backspaced (refer to 11.3.1).

```
nn?SUNNNN #BLKS
```

Enter the number of blocks to be backspaced (maximum number is 9999).

Writing Tape Marks (WTM)

You use the WTM function code to write tapemarks on your file.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the output tape (refer to 11.3.1).

Rewind Tape (REW)

You use the REW function code to rewind your tape to load point.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be rewound (refer to 11.3.1).

Rewind and Unload Tape (RUN)

You use the RUN function to rewind either UNISERVO 16 or UNISERVO 20 tapes to load point with interlock.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be rewound (refer to 11.3.1).

Erasing Tape Record Gap (ERG)

You use the ERG function code to erase a portion of your tape. This function is useful to erase known defective areas on your tape. Erases the specific tape for approximately 3.5 inches.

Procedure

For the message:

```
nn?SUNNNN CUUMM - OUTPUT TAPE
```

Enter the tape unit device address and the mode setting of the tape to be erased (refer to 11.3.1).

11.4. Disk and Format Label Diskette Functions

The disk and format label diskette functions that can be performed are described in the following subsections.

11.4.1. Operating Considerations

Consider the following when operating with the system utility on disks or format label diskettes:

- An end-of-file record is a disk record on which data length is in binary zeros (not applicable on IDA disks).
- The system utility should be run only when no other job will be starting up or performing multiple steps that will allocate the same volumes. This could cause the system utility to terminate abnormally.

11.4.2. Disk and Format Label Diskette Operating Instructions

Proceed as follows to perform a disk or format label diskette function:

1. Place the subject disk/diskette volume on an available unit.
2. As described in 11.1:
   a. Enter the SU symbiont command.
   b. Enter the function code as a command parameter or as a solicited message response to the ENTER REQUIRED FUNCTION message.
   c. Enter the spooling option; otherwise, default is Y (applicable if spooling is configured).

11.4.3. Printing a Disk or Format Label Diskette in Unblocked Format (DD)

You use the DD function code to print your disk pack or format label diskette in character and hexadecimal format. The DD function does not deblock your logical records.

Procedure

For the messages:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the disk or diskette unit device address of the disk pack or format label diskette to be displayed.

```
nn?SUNNNN CCCHH - BEGIN OR FILE-ID
```

Enter in decimal the beginning cylinder (CCC) and head (HH) to be displayed, or the file identifier (up to 44 characters) as used on the // LBL job control statement when the file was created. If you enter less than 44 characters, the file ID is padded with blanks on the right. Printing begins at the start of the file when a file is entered.

Note: *System utility permits embedded blanks in file names. However, you must not place apostrophes or quotes around file names as delimiters, or they will be taken by system utility as part of the file names.*

```
nn?SUNNNN CCCHH - END
```

Enter the last cylinder (CCC) and head (HH) in decimal to be displayed.

```
nn?SUNNNN NO OF TRACKS TO PRINT UP TO 9
```

If you entered the file ID, now enter the number of tracks to be printed. Note that printing begins with the low cylinder and head numbers of the first extent. If the file is not laced, printing continues for the requested number of tracks or until end of file (EOF) is encountered, whichever comes first. If the file is laced, printing continues for the requested number of tracks. Printing is not confined to the end extent specified in the format label. For DD, a format label diskette is considered to have 74 cylinders, with 2 heads (0 and 1) in each cylinder. (Laced files are discussed in the *System Service Programs (SSP) Operating Guide* (UP-8841) in the section describing diskette prep.)

Example 1

The data on cylinder 3, head 4 through cylinder 6, head 6 on device 300 is to be printed.

For the message:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Key in:

```
nnΔ300
```

For the message:

```
nn?SUNNNN CCCHH - BEGIN OR FILE-ID
```

Key in:

```
nnΔ00304
```

For the message:

```
nn?SUNNNN CCCHH END
```

Key in:

```
nnΔ00606
```

Example 2

The first three tracks from a file called SEQUENTIAL DISC on device 440 are to be printed.

For the message:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Key in:

```
nnΔ440
```

For the message:

```
nn?SUNNNN CCCHH - BEGIN OR FILE-ID
```

Key in:

```
nnΔSEQUENTIAL DISC
```

For the message:

```
nn?SUNNNN NO OF TRACKS TO PRINT UP TO 9
```

Key in:

```
nnΔ3
```

11.4.4. Printing a Disk or Format Label Diskette in Deblocked Format (DDR)

You use the DDR function code to print your disk pack or format label diskette in both character and hexadecimal formats. This function is similar to the DD function, except that the printing is reformatted according to the logical record length that is entered.

Procedure

For the message:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the disk/diskette unit device address.

For the message:

```
nn?SUNNNN CCCHH-BEGIN OR FILE-ID
```

Enter in decimal the first cylinder (CCC) and head (HH) to be printed or the file identifier (up to 44 characters) as on the // LBL job control statement when the file was created. If you enter less than 44 characters, the file-id is padded with blanks on the right. Printing begins at the start of the file when a file-id is entered.

If cylinder and head are entered, SU displays these messages:

```
nn?SUNNNN CCCHH-END
```

Enter in decimal the last cylinder (CCC) and head (HH) to be displayed.

```
nn?SUNNNN ENTER LOGICAL RECORD LENGTH
```

Enter in decimal the logical record size to determine the print format.

```
nn?SUNNNN ENTER BLOCKSIZE
```

Enter in decimal the logical block size. However, if the block size is not an exact multiple of the record size, the following message will be displayed:

```
BLOCKSIZE NOT MULTIPLE OF RECORD SIZE
```

Reenter the record size/block size sequence.

If a file-id is entered, SU displays these messages:

```
nn?SUNNNN NUMBER OF TRACKS TO BE PRINTED UP TO 9
```

Enter the number of tracks to be printed. Note that printing begins with the low cylinder and head numbers of the first extent and continues for the requested number of tracks, or until end of file (EOF) is encountered, whichever comes first. Printing is not confined to the extents specified in the format label.

```
nn?SUNNNN ENTER LOGICAL RECORD LENGTH
```

Enter in decimal the logical record size to determine the print format. The block size will be taken from the VOL1 label.

11.4.5. Printing the Disk or Format Label Diskette Volume Table of Contents (VTP)

You use the VTP function code to get a copy of your VTOC. You can print:

- A full VTOC listing, giving you all the device information plus extents and other information for all your allocated files
- Device information only, giving you the available space left on your volume and other information regarding your volume
- File information only, giving you the extent and other information regarding the file

VTP edits and prints the volume table of contents (VTOC) information for the requested volume. It will not process other than OS/3-created VTOCs. Use the DD function to print a VTOC from a disk or diskette not created by OS/3.

Procedure

Three list options are available, as follows:

- VSN, DI - Device Information Only

  Lists the available free extents and other information on the requested volume.

- VSN, FILE-ID - File Information

  Lists device information plus the extent and other information on the requested file.

- VSN, ALL - Full VTOC Listing

  Lists device information plus the extent and other information for all files allocated on the requested file.

For the message:

```
nn?SUNNNN ENTER DVC/VSN, DI, ALL, FILE-ID, END OR EOJ
```

Key in one of the following, where vsn is the volume serial number of the disk pack or format label diskette whose VTOC is to be printed.

For full VTOC listing:

```
nnΔvsn,ALL
```

For a listing of only the device information:

```
nnΔvsn,DI
```

For a listing of up to 44 characters as used on the // LBL card when the file was created:

```
nnΔvsn,FILE-ID
```

To terminate the VTP function:

```
nnΔEND
```

To terminate SU:

```
nnΔEOJ
```

Notes:

1. *The device address of the disk pack or format label diskette may replace the vsn in the preceding messages.*

2. *When a VTP function terminates, SU produces the following message:*

   ```
   nn?SUNNNN ENTER DVC/VSN, DI, ALL, FILE-ID, END OR EOJ
   ```

   You may specify another volume for VTP display, END to end VTP functions, or EOJ to terminate the SU.

3. *System utility permits embedded blanks in file names. However, you must not place apostrophes or quotes around file names as delimiters, or they will be taken by system utility as part of the file names.*

Disk and Format Label Diskette VTP Listing Summary

The information listed by VTP is either taken directly from the disk or diskette labels or calculated from data contained in the labels. A summary of the VTOC information listed by the VTP function is provided in Table 11-5. (Refer to the *Consolidated Data Management Macroinstructions Programming Guide* (UP-9979) for details.)

Table 11-5. Summary of Disk/Diskette VTP Information

Field Heading	Field Label
VOLUME INFORMATION	
VOLUME SERIAL NUMBER	DL$VSN
VTOC ADDRESS (CCC HHRR)	DL$VTC
VOLUME SECURITY	DL$VSB
OWNER NAME/ADDR CODE	DL$ONR
DEVICE INFORMATION	
ADDRESS LAST FORMAT 1	DL$LF4
ADDRESS HIGHEST ALT TRACK	DL$HA4
NO CYL/TRK THIS DISK	DL$DS4
TOLERANCE	DL$TO4
NO UNUSED VTOC RECORDS	DL$AF4
NO OF ALT TRACKS	DL$AT4
NO BYTES PER TRACK	DL$TL4
NO VTOC LABELS PER TRACK	DL$LT4
AVAILABLE EXTENT DATA	
START ADDRESS	DL$XT5
END ADDRESS	DL$XC5 (calculated)
EXTENT SIZE	(calculated)

continued

Table 11-5. Summary of Disk/Diskette VTP Information (cont.)

Field Heading	Field Label
FILE INFORMATION	
FILE NAME	DL$ID1
FILE SERIAL NUMBER	DL$FS1
VOLUME SEQ NO	DL$VS1
CREATION DATE	DL$CD1
EXPIRATION DATE	DL$ED1
EXTENT COUNT	DL$XC1
OPTION CODES	DL$OC1
PCA COUNT	DL$PC1
TRACKS/CYLINDERS	DL$TPC2
LOW HEAD	DL$FLH2
FILE TYPE	DL$FT1
LAST ACCESS DATE	DL$ACDTE
LAST ACCESS TIME	DL$ACTME (milliseconds)
LAST UPDATE DATE	DL$UPDTE
LAST UPDATE TIME	DL$UPTME (milliseconds)
UPDATE JOBNAME	DL$UPID
UPDATE ACCESS CODE	DL$UPACC
PRIOR UPDATE DATE	DL$PUPD
PRIOR UPDATE TIME	DL$PUPT (milliseconds)
PRIOR UPDATE JOBNAME	DL$PUPID
PRIOR UPDATE ACCESS CODE	DL$PUPAC
DATE FILE LAST DUMPED	DL$DRDTE
TIME FILE LAST DUMPED	DL$DRTME

continued

Table 11-5. Summary of Disk/Diskette VTP Information (cont.)

Field Heading	Field Label
FILE PARTITION (PCA) INFORMATION	
PCA NUMBER	(1 through 7)
BLOCK SIZE	DL$BL1
RECORD SIZE*	DL$RL1
RECORD FORMAT*	DL$RF1
BLOCK FACTOR*	(calculated: records per block for fixed-length blocked record format.)
BLOCKS PER TRACK	DL$SLA2
NO. OF RECS IN LAST BLOCK	DL$SPC2
KEY LENGTH OR LACE FACTOR	DL$SLF2
END OF DATA BLOCK NO.	DL$SEP2
TOTAL # RECORDS	(calculated: total number of records for sequential FIXBLK and FIXUNB files, and for VARUNB files.)
LOGICAL EXTENT INFORMATION	
PCA NO	(0 through 7)
START CCC' HH	DL$SXAR2 (calculated from relative track address)
END ADDR CCC HH	DL$SXAR2 (calculated from number of tracks in extent)
ALLOCATION SUMMARY	
EXT SIZE CCC HH	DL$SXAR2 (calculated: END ADDR minus START)
CUM/PCA CCC HH	(calculated: cumulative CCC HH for this PCA)
CUM/FIL CCC HH	(calculated: cumulative CCC HH for this file)
EXTENT UTILIZATION SUMMARY	
OCCUPIED/PCA CCC HH BLOCK	(calculated: allocated space used for this PCA)
AVAILABLE/PCA CCC HH BLOCK	(calculated: unused space for this PCA)
RECORD CAPAC REMAINING	(calculated: approximate number of logical records that fit in available space; not calculated for variable-length blocked record format)
END OF DATA CCC HH RRR	DL$SEP2 (calculated)

*Not applicable to MIRAM files on PCA basis.

continued

Table 11-5. Summary of Disk/Diskette VTP Information (cont.)

Field Heading	Field Label
PHYSICAL EXTENT INFORMATION	
EXTENT TYPE	DL$XT1
EXTENT SEQ NO.	DL$SX1
EXTENT START CCC HH	DL$XL1
EXTENT END CCC HH	DL$XU1
CUM CCC HH	(calculated)
AUTO EXTENT INCREMENT	DL$SA1
SAT FILE INFORMATION	
DIRECTORY PCA LACE FACTOR	DL$DIRL2
DIRECTORY PCA LACE ADJUST	DL$DIRF2
TEXT PCA LACE FACTOR	DL$TXTL2
TEXT PCA LACE ADJ	DL$TXTF2
(M)IRAM FILE INFORMATION	
KEY LOCATION	DL$XILOC
KEY LENGTH	DL$XILOC +2
DUPLICATES ALLOWED	DL$XILOC+3, X'80'
CHANGES ALLOWED	DL$XILOC+3, X'40'
DATA RECORD COUNT	DL$COUTR
RECORD SIZE	DL$DREC
INDEX BUFF SIZE	DL$CSIZ

11.4.6. Printing a Short Format VTOC for a Disk or Format Label Diskette (SVT)

You use the SVT function to obtain an abbreviated VTOC listing, consisting of a single print line for each physical extent for each file.

Procedure

For the message:

```
nn?SUNNNN ENTER DVC OR VSN OR END
```

Enter the volume serial number (up to six characters) of the disk unit or diskette device address of the disk pack or diskette whose VTOC is to be printed; either entry prints the abbreviated VTOC. Enter END to terminate the SVT function.

If SVT is attempted on a data set label diskette, the following message is displayed.

```
SVT VALID ON FORMAT LABEL VOLUMES ONLY
```

SVT Listing Summary for Disk and Format Label Diskette

The various information listed by SVT is either taken directly from the disk or diskette labels or calculated from data contained in the labels. A summary of the short format VTOC information listed by the SVT function is provided in Table 11-6. (Refer to the *Consolidated Data Management Macroinstructions Programming Guide* (UP-9979) for details.)

Table 11-6. Summary of SVT Information for Disk and Format Label Diskette

Field Heading	Field Label
VSN	DL$VSN
SECUR	DL$VSB
OWNER	DL$ONR
FILE-NAME	DL$ID1
FILE TYPE	DL$FT1
FILE SER. NO	DL$FS1
SEQ NO	DL$VS1
CREATION DATE	DL$CD1
EXPIRATN DATE	DL$ED1
OPT COD	DL$OC1
PCA CT	DL$PC1
START CCC HH	DL$XL1
END CCC HH	DL$XU1
CUM CCC HH	(calculated: cumulative)
AI	DL$SA1

11.4.7. Displaying the Available Disk and Format Label Diskette Extents (AVX)

You use the AVX function to display a list of available disk or diskette extents on the console workstation or console screen. The display is similar to the listing printed by the VTP function with the DI option; however, output is displayed at the console workstation console and printed at the console output printer rather than at the customary line printer.

Procedure

For the message:

```
nn?SUNNNN ENTER DEVICE ADDRESS
```

Enter the did of the disk or diskette drive containing the disk or diskette whose available extents are to be displayed.

If AVX is attempted on a data set label diskette, the following message is displayed:

```
AVX VALID ON FORMAT LABEL VOLUMES ONLY
```

AVX Listing Summary

The various information listed by AVX is either taken directly from the disk or diskette labels or calculated from data contained in the labels. A summary of the available disk or diskette extent information listed by the AVX function is provided in Table 11-7. (Refer to the *Consolidated Data Management Programming Guide* (UP-9978) for details.)

Table 11-7. Summary of AVX Information for Disk and Format Label Diskette

Field Heading	Field Label
VSN	DL$VSN
SECUR	DL$VSB
VTOC ADRS	DL$VTC
OWNR	DL$ONR
START CCC HH	DL$XT5
END CCC HH	DL$XC5 (calculated)
SIZE CCC HH	(calculated)

11.5. Data Set Label Diskette Functions

The data set label diskette functions that can be performed are described in the following subsections.

11.5.1. Data Set Label Diskette Operating Instructions

Proceed as follows to request a data set label diskette function:

1. Place the subject data set label diskette volume on an available diskette unit.
2. As described in 11.1:
   a. Enter the SU symbiont command.
   b. Enter the DD or DDR function code either as a command parameter or as a solicited message response to the ENTER REQUIRED FUNCTION message.
   c. Enter the spooling option; otherwise, default is Y (applicable if spooling is configured).

11.5.2. Printing the Data Set Label Diskette in Unblocked Format

You use the DD function to print a single-sided or double-sided data set label diskette in unblocked format.

Single-Sided Diskette in Unblocked Format (DD)

The DD function code prints your data set label diskette in character and hexadecimal format. The DD function does not deblock your logical records.

Procedure

For the messages:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the diskette unit device address of the data set label diskette to be displayed.

```
nn?SUNNNN TTRR-BEGIN
```

Enter in decimal the beginning track (TT) and sector (RR) to be displayed.

```
nn?SUNNNN TTRR-END
```

Enter the last track (TT) and sector (RR) in decimal to be displayed.

Double-Sided Diskette in Unblocked Format (DD)

The DD function code prints your data set label diskette in character and hexadecimal format. The DD function does not deblock your logical records. Your diskette is considered to have 74 cylinders with 2 heads (0 and 1) in each cylinder.

Procedure

For the messages:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the diskette unit device address of the data set label diskette to be displayed.

```
nn?SUNNNN CCCHH - BEGIN
```

Enter in decimal the beginning cylinder (CCC) and head (HH) to be displayed.

```
nn?SUNNNN CCCHH - END
```

Enter the last cylinder (CCC) and head (HH) in decimal to be displayed.

Example

The data on cylinder 3, head 4 through cylinder 6, head 6 on device 300 is to be printed.

For the message:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Key in:

```
nnΔ300
```

For the message:

```
nn?SUNNNN CCCHH - BEGIN
```

Key in:

```
nnΔ00300
```

For the message:

```
nn?SUNNNN CCCHH END
```

Key in:

```
nnΔ00601
```

11.5.3. Printing the Data Set Label Diskette in Deblocked Format

You use the DDR function to print a single-sided or double-sided data set label diskette in deblocked format.

Single-Sided Diskette in Deblocked Format (DDR)

The DDR function code prints your data set label diskette in character and hexadecimal format. The DDR function deblocks your logical records.

Procedure

For the messages:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the diskette unit device address of the data set label to be displayed.

```
nn?SUNNNN TTRR-BEGIN
```

Enter in decimal the first record (RR) and track (TT) to be printed.

```
nn?SUNNNN ENTER LOGICAL RECORD LENGTH
```

Enter in decimal the logical record size to determine the print format.

Double-Sided Diskette in Deblocked Format (DDR)

The DDR function code prints your data set label diskette in character and hexadecimal format. The DDR function deblocks your logical records. Your diskette is considered to have 74 cylinders with 2 heads (0 and 1) in each cylinder.

Procedure

For the messages:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the diskette unit device address of the data set label diskette to be displayed.

```
nn?SUNNNN CCCHH-BEGIN
```

Enter in decimal the first cylinder (CCC) and head (HH) to be printed. Double-sided diskettes have 72 cylinders and 2 heads, 00 and 01.

```
nn?SUNNNN CCCHH-END
```

Enter in decimal the last cylinder (CCC) and head (HH) to be printed.

```
nn?SUNNNN ENTER LOGICAL RECORD LENGTH
```

Enter in decimal the logical record size to determine the print format.

11.5.4. Printing the Data Set Label Diskette Volume Table of Contents (VTOC)

You use the DD function to print the diskette volume table of contents (that is, the data set labels).

Single-Sided Diskette VTOC (DD)

The DD function code prints your single-sided data set label diskette VTOC.

Procedure

For the message:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the diskette unit device address of the data set label diskette whose VTOC is to be printed.

For the message:

```
nn?SUNNNN TTRR-BEGIN
```

Key in:

```
nnΔ0008
```

For the message:

```
nn?SUNNNN TTRR-END
```

Key in:

```
nnΔ0026
```

Double-Sided Diskette VTOC (DD)

The DD function code prints your double-sided data set label diskette VTOC.

Procedure

For the message:

```
nn?SUNNNN ENTER DVC ADDRESS
```

Enter the diskette unit device address of the data set label diskette whose VTOC is to be printed.

For the message:

```
nn?SUNNNN CCCHH-BEGIN
```

Key in:

```
nnΔ00000
```

For the message:

```
nn?SUNNNN CCCHH-END
```

Key in:

```
nnΔ00001
```

Section 12
Disk Cache Facility (DCF)

The OS/3 disk cache facility (DCF) increases system performance by reducing the I/O bottleneck to disk devices. I/Os are slow relative to processor speed. The I/O bottleneck is reduced simply by reducing the number of physical I/Os.

When a read is done of a particular area on a disk, the disk cache facility (DCF) reads into a reserved cache storage area all or a portion of the track being read (depending on the segment size when DCF was initialized). Consequently, in subsequent reads of that same disk area, a physical I/O is unnecessary, since that data is already present in storage.

The advantage of DCF is that a reduction in the number of physical I/Os required results in:

- Reduced queuing delays (as I/O requests must wait to be issued in turn)
- Reduced disk seek time (as the disk accessor moves to the proper cylinder)
- Reduced latency (as the disk rotates so that the disk read/write heads reach the proper record)
- Reduced data transfer time (from disk to storage)
- Reduced process time overall

Main storage disk cache is supported via software. However, the software is included in the system microcode file (YMIC) and is called via the system definition file (YSDF) to maintain consistency between all System 80 models. Refer to the *Installation Guide* (UP-8839) for information on the system microcode and system definition files.

You can selectively activate or remove individual disk drives to DCF (12.3 and 12.4) or turn it off completely (12.5). Also, you can bypass the caching of specific files using the CACHE parameter of the DD job control statement and the CACHE IOGEN parameter. (See the *Consolidated Data Management Programming Guide* (UP-9978) and the *Installation Guide* (UP-8839).) Appendix C provides information to help you optimize DCF performance.

To use DCF, your system must be configured with the disk cache feature. (See 3.3.2 for models 4 and 6.) Following are descriptions of DCF features. Refer to the *System Messages Reference Manual* (UP-8076) for an explanation of messages and error codes.

12.1. Initializing Disk Cache (CM)

DCF is normally initialized during IPL. However, if it was not initialized during IPL or was shut down, you can initialize it manually by entering the CACHE MEMORY command at the system console. When DCF is initialized, all devices are activated or removed from DCF according to their CACHE IOGEN parameter specifications. Refer to the *Installation Guide* (UP-8839) for details on this parameter.

Format

```
CM
```

12.1.1. Specifying Cache Buffer Size (CM SIZ)

When cache is initialized at IPL time, the size of the cache buffer is determined from the YSDF file information or the CM01 message prompt, which asks for a buffer size. You can override this cache buffer size when you initialize DCF with the CM SIZ command.

Format

```
CM SIZ=n
```

where:

n

Specifies the number of 1,024-byte blocks allocated to the cache buffer. Valid specifications are 160 to 8,192 for models 8-20; and 64 to 1,024 for models 4 and 6. Note that larger cache buffers generally result in increased performance.

See Appendix C for more information on cache buffer and segment size.

12.1.2. Specifying Cache Buffer Segment Size (CM SEG)

The cache buffer is divided into fixed size segments. The amount of data that DCF reads from a disk is equal to the cache buffer segment size. You can specify the cache buffer segment size by entering the CM SEG command.

Format

```
CM SEG=n
```

where:

n

Specifies the segment size in 1,024-byte blocks. Valid specifications are 2 to 24 for models 8-20, and 2 to 15 for models 4 and 6. This command overrides the cache segment size previously defined through the CACHESEGSIZE SUPGEN parameter, as described in the *Installation Guide* (UP-8839).

If you don't specify the CM SEG command, the size specified in the CACHESEGSIZE SUPGEN parameter is used. If you don't specify either parameter, the default values used are as follows:

Cache Buffer Size	Default Cache Buffer Segment Size (in 1,024-byte blocks) Models 8-20	Models 4 & 6
64 - 255K	-	2.5
160 - 255K	3	-
256 - 511K	6	7.5
512 - 1,024K	12	15.0
1,025 - 8,192K	12	-

Note that using a smaller segment size causes a greater number of segments in your cache buffer. Generally, sequential processing performs better with large segments, and random processing performs better with a larger number of small segments. You can experiment to find the combination that is best suited to your environment. Use the statistics that DCF provides (see 12.6) to help tune these specifications.

The SIZ and SEG options can be separated by a comma and specified together on the CM command.

Format

```
CM SEG=15, SIZ=1000
```

12.2. Resegmenting the Existing Cache Buffer (CM RESEG)

With the unsolicited CM RESEG command, you can resegment the cache buffer that you allocated during DCF initialization with a different segment size. You can use this command without shutting down or reinitializing DCF. However, if you want to change the buffer size, then you must reinitialize DCF. The RESEG command resets the cache statistics to zero.

Format

```
00 CM RESEG=n
```

where:

n
: Specifies segment size in 1,024-byte blocks. Valid specifications are 2 to 24.

Notes:

1. *This command is available only on models 8-20.*
2. *Resegmenting the cache buffer resets all cache statistics but does not affect the activated or removed status of the disks on the system.*

12.3. Selective Caching

Normally, all drives and files are cached. The performance of cache memory may be optimized by using cache memory only with the most frequently accessed drives and files. Use selective caching to remove the less frequently used drives and files.

12.3.1. Removing Drives from DCF (CM REM)

You can remove a disk drive from DCF via the CACHE IOGEN parameter (also see the *Installation Guide* (UP-8839)) or via the unsolicited CM REMOVE command. Note that when a disk drive is shared by two processors, the drive must be removed from DCF in any processor that writes to it.

Format

```
00 CM REMOVE dvc# [,dvc,...,dvc#]
```

where:

dvc#
: Specifies the device address (drive number) of the drive to be removed from the DCF. You can specify more than one address.

Note: *This command is available only on models 8-20.*

12.3.2. Removing Files from DCF

You can bypass the caching of specific files using the CACHE parameter of the DD job control statement and the CACHE IOGEN parameter. For more information, see the *Consolidated Data Management Programming Guide* (UP-9978) and the *Installation Guide* (UP-8839).

12.4. Activating Drives to DCF (CM ACT)

You can activate a drive to DCF via the CACHE IOGEN parameter, as previously mentioned (also see the *Installation Guide* (UP-8839)), or via the unsolicited CM ACTIVATE command.

Format

```
00 CM ACTIVATE dvc# [,dvc#,...,dvc#]
```

where:

dvc#
: Specifies the device address (drive number) of the drive to be activated to DCF. You can specify more than one address.

Note: *This command is available only on models 8-20.*

12.5. Removing DCF from the System (CM SHUT)

Remove DCF from the system at any time using the unsolicited CM SHUTDOWN command:

Format

```
00 CM SHUTDOWN
```

12.6. Displaying DCF Statistics (CM STA, STARES)

You can monitor DCF operations by displaying DCF statistics. You can display DCF statistics on demand or at periodic, timed intervals. You can also display DCF statistics for a specific drive.

To display DCF statistics immediately, use the unsolicited cache commands CM STA and CM STARES.

Formats

`00 CM STA`
: Displays statistics on the system console.

`00 CM STARES`
: Displays statistics on the system console and resets all of the counters including disk device statistics maintained by DCF. It can be used to monitor DCF operations during a specific period of time.

12.6.1. Displaying and Turning Off Timer Statistics (CM STATIME, STARESTIME, TIMER)

Formats

`00 CM STATIME=n`

Displays statistics on the system console repeatedly.

where:

`n`

Specifies the time interval between displays. Valid values are 1 to 999 minutes.

`00 CM STARESTIME=n`

Repeatedly displays statistics and resets all counters maintained by DCF. This command can be used to monitor DCF operations for a specified time interval.

where:

`n`

Specifies the time interval between displays. Valid values are 1 to 999 minutes.

`00 CM TIMER=OFF`

Turns off the timer you set using the STATIME or STARESTIME command.

12.6.2. Displaying Statistics on a Specified Drive (CM STADISK, STARESDISK)

Formats

`00 CM STADISK=did`

Displays statistics for the specified disk on the system console.

where:

`did`

Specifies the device address of the disk.

`00 CM STARESDISK=did`

Displays statistics and resets the counters for the specified disk.

where:

`did`

Is the device address of the disk.

Notes:

1. *You can use the STADISK and STARESDISK commands on models 8-20 only.*
2. *If you generate an attention interrupt, AVR a disk, or mount another disk, the counters for that disk's device statistics are reset.*

12.6.3. Interpreting DCF Statistics

Figure 12-1 shows samples of the statistics displayed at the system console for models 8-20 and for models 4 and 6. Statistics that you can display include:

- Hit rates

  A hit occurs when a read or search/read operation does not have to access the disk because the data is in the cache buffer. Hit rates are the most significant indication of cache performance. The search hit rate is the number of search hits divided by the number of search HI/EQs. (Search equal commands are not cached.)

- Read errors

  Read errors are I/O errors encountered by DCF. When a read error occurs, the I/O is reissued directly into the program's I/O buffer, rather than into the cache buffer. The reissued I/O may be successful.

- Write-throughs

  A write-through occurs when the data being written resides in the cache buffer. The buffer is updated and the data is written to the disk.

- Unreferenced I/Os

  Unreferenced I/Os are a count of the number of segments that were reused without ever having been referenced by read or search/read commands. Segments are reused when space is required for an I/O. The first segments reused are those that were never referenced through a read, search/read, or write (LRU).

- Unreferenced I/O rate

  The unreferenced I/O rate is the number of unreferenced I/Os divided by the number of reads.

- Reads over half segment

  This count is the number of reads that are greater than half the segment size and, therefore, are not cached. This count is *not* included in the hit rate calculations.

- Reads not cached

  DCF automatically bypasses read or search/read commands that are multitrack or without data transfer. This count is included in the hit rate calculations.

- (CACHE=NO) I/Os

  This count records the number of read or search/reads that bypass DCF because of the selective noncaching by file option (see 12.3.2). This count is not included in the hit rate calculations.

Note: *If a read is larger than half of the segment size, it is counted as a read over half segment even when CACHE=NO has been specified.*

```
                        MODELS 8-20

00 CM STA
31 CM0005     CM16 READ HIT RATE                 73.0%
32 CM0005     CM17 SEARCH READ HIT RATE          99.5%
33 CM0005     CM18 UNREFERENCED I/O RATE          0.3%
34 CM0005     CM19 CACHE BUFFER SIZE             1024k
35 CM0005     CM20 SEGMENT SIZE (x 1024)            12
36 CM0005     CM21 TOTAL NUMBER DISK I/O         46170
37 CM0005     CM22 NUMBER READ HITS              30320
38 CM0005     CM23 NUMBER WRITES                  1846
39 CM0005     CM24 NUMBER SEARCH EQUALS            248
40 CM0005     CM25 NUMBER SEARCH HI/EQ            1121
41 CM0005     CM26 NUMBER READ ERRORS                1
42 CM0005     CM27 NUMBER READS NOT CACHED        8945
43 CM0005     CM28 NUMBER SEARCH HITS             1116
44 CM0005     CM29 NUMBER WRITE THROUGHS          1407
45 CM0005     CM30 NUMBER UNREFERENCED I/OS        140
46 CM0005     CM31 NUMBER (CACHE=NO) I/OS           20
47 CM0005     CM32 READS OVER HALF SEGMENT          50
48 CM0005     CM33 NUMBER READS                  41508

                      MODELS 4 AND 6

00 CM STA
68 CM0006     CM21 TOTAL NUMBER DISK I/Os         2574
69 CM0006     CM22 NUMBER READ HITS               1056
70 CM0006     CM23 NUMBER WRITES                   451
71 CM0006     CM24 NUMBER SEARCH EQUALS             59
72 CM0006     CM25 NUMBER SEARCH HI/EQ               3
73 CM0006     CM27 NUMBER READS NOT CACHED           2
```

Figure 12-1. DCF Statistics Displayed at the System Console

Appendix A
ONUERL Program

ONUERL is an error log editor program that is resident on the system while a maintenance agreement is in effect. ONUERL formats and prints the contents of the ELOG file. Models 3-6 users can optionally initiate the transmission of summary error data to a Unisys TRACE center. Models 8-20 users can optionally spool summary error data to a diskette for eventual transmission to a Unisys TRACE center.

You have several options when running ONUERL. You can:

- Choose to hold the output in a spool file for later printing
- Select where processing will start in the ELOG file
- Suspend error logging
- Transmit data to TRACE (models 3-6 users)
- Spool error summary data to a diskette (models 8-20 users)
- Obtain a printout that documents what ONUERL does
- Request graphic and tabular displays of accumulated error data (ONELAN program, described later in this appendix)

Note: *ONUERL should be run on a daily basis or other planned schedule to check system performance.*

A.1. Executing ONUERL

To execute ONUERL with preset options, enter the following command:

```
RV ONUERL
```

The program will execute with the following preset options:

```
OPEN = current
TRACE/SUMMARY = threshold
ELOG = yes
DOC = no
THRESHOLDS = no
DUMP = no
ESUMFIL = u
DMPSUM = no
```

To execute ONUERL and modify options, enter the following command:

```
RV ONUERL,,OPT=Y
```

A.2. ONUERL Program Options

The first message ONUERL displays after loading is:

```
ENTER SPOOL OPTION (YES,NO)
```

To hold the output in a spool file for later printing, enter the message-id followed by YES. Otherwise, enter only the message-id and press the XMIT key. The output from ONUERL will then go to the first available printer.

After you have indicated a spool option, ONUERL displays the following message:

```
MODIFY PROGRAM OPTIONS (Y,N,?) ?
```

Respond with the message-id and N to run ONUERL using the preset options (defaults). To specify alternate options, respond with the message-id followed by Y. If you're not sure what the options are, enter the message-id followed by a question mark. In the ONUERL program, a question mark prompts the system to display a brief explanation of the option or message.

If you have elected to modify program options (you've entered a Y), then ONUERL displays this message:

```
ENTER OPTIONS, SEPARATED BY COMMAS (1,2,3,4,5,6,7,8?)
```

The numbers displayed correspond to program options. To find out which options go with what numbers, either key in the message-id followed by a question mark, or refer to Table A-1. It lists the program options, what they mean, and what choices and defaults are available. Since ONUERL does not automatically transmit error summary data to a Unisys TRACE center when you run the program on a System 80 models 8-20, Table A-1 shows different choices for option 2, depending on whether you're running the program on a model 4/6 or 8-20.

After you've entered the numbers that correspond to the options you want to modify, the program lists the choices for each option you've selected. Figure A-1 is a step-by-step example of the modification procedure. Note in the example that the program is questioned for help at each step. Also note that the example is for model 4 and 6 users.

Table A-1. ONUERL Program Options

1 OPEN= $\begin{Bmatrix} C \\ P \\ B \end{Bmatrix}$

C

Current. ONUERL opens the error log file at the record where the last ONUERL program ended.

P

Prior. The error log file is opened at the same record where it was opened the last time ONUERL was run. The data displayed the last time is displayed again, along with any new data.

B

Begin. The entire error log file is opened, beginning with the oldest record and continuing to the most recent record.

2 TRACE= $\begin{Bmatrix} Y \\ T \\ N \\ R \\ D \end{Bmatrix}$

Models 3 – 6	Models 8 – 20
Y Yes. TRACE summary gathered during execution of ONUERL is to be transmitted to the TRACE center. Transmission occurs after error data has been printed locally and after the physical connection to the TRACE center (via an acoustic coupler) has been made. T Threshold. Summary data gathered during the execution of ONUERL is examined for over-threshold occurrence. If, for any device, the total error count exceeds 5 percent of the good I/O count, the program automatically requests connection to TRACE. If the connection is made, summary data is transmitted and the program terminates. If TRACE transmission is denied, the program terminates. However, future execution of ONUERL will be denied until TRACE data is transmitted. If over-threshold condition did not occur, summary data gathered is merged into the summary data file YESUM for future transmission.	Y Yes. TRACE summary gathered during execution of ONUERL is to be spooled to a system journal diskette. T Threshold. Summary data gathered during the execution of ONUERL is examined for over-threshold occurrence. If, for any device, the total error count exceeds 5 percent of the good I/O count, the program automatically displays a mount message for you to mount the system journal diskette. The program waits for the operator to type in GO ONUERL, then spools the error summary data to the diskette. ONUERL terminates. If over-threshold condition did not occur, summary data gathered is merged into the summary data file YESUM for future transmission.

continued

Table A-1. ONUERL Program Options (cont.)

Models 3 – 6	Models 8 – 20
N No. ONUERL operates in local mode and no TRACE transmission is made, even if over-threshold condition occurs. However, TRACE summary data is still saved in the summary data file (YESUM) for future transmission. (See the REPEAT option, following.) Note that any data present in YESUM before ONUERL is run is destroyed unless TRACE=REPEAT.	**N** No. ONUERL operates in local mode and does not spool error summary data to the system journal diskette even if over-threshold condition occurs. However, TRACE summary data is still saved in the summary data file (YESUM) for future transmission. (See the REPEAT option, following.) Note that any data present in YESUM before ONUERL is run is destroyed unless SUMMARY=REPEAT.
R Repeat. Indicates that TRACE summary previously saved in the summary data file (YESUM) is to be sent to a TRACE center. In this case, only YESUM is opened, not the error log file (YELOG). If there are no records in the YESUM file, an error message is displayed and the program is terminated.	**R** Repeat. Indicates that TRACE summary previously saved in the summary data file is to be spooled out to the system journal diskette. In this case, only YESUM is opened, not the error log file (YELOG). If there are no records in the YESUM file, an error message is displayed and the program is terminated.
D Disable. This option permits the permanent disabling of TRACE transmission. It is intended for those installations where, for security or other reasons, transmission of summary data is disallowed. Once TRACE is disabled, selecting TRACE=YES or THRESHOLD prompts a console workstation message informing you that TRACE has been disabled and asking if TRACE is to be restored. Also, since no summary data has previously been collected, selecting TRACE=REPEAT prompts a console workstation message informing you that TRACE request cannot be honored and ONUERL will be cancelled.	**D** Disable. This option permits the permanent disabling of spoolout of error summary data to the system journal diskette. It is intended for those installations where, for security or other reasons, transmission of summary data is not allowed. Once spoolout is disabled, selecting SUMMARY=YES or THRESHOLD prompts a console message informing you that TRACE has been disabled and asking if TRACE is to be restored. Also, since no summary data has previously been collected, selecting SUMMARY=REPEAT prompts a console message informing you that TRACE request cannot be honored and ONUERL will be cancelled.
3 ELOG= $\left\{ \begin{matrix} Y \\ N \end{matrix} \right\}$ **Y** Yes. Error logging is suspended during ONUERL execution. **N** No. ONUERL has no effect on error logging.	

continued

Table A-1. ONUERL Program Options (cont.)

```
4 DOC= { Y }
       { N }
```

Y

Yes. Operating instructions are printed at the local printer.

N

No. Printing of operating instructions is inhibited.

```
5 THRESHOLDS= { (Y) }
              { (N) }
              { (D) }
```

Y

Yes. Modify recoverable error thresholds.

N

No. Run with preset recoverable error thresholds.

D

Display. Display current recoverable error thresholds. This option will broadcast the following message:

```
ONUERL - DO YOU WISH TO ALTER THRESHOLDS (Y,N)
```

```
6 DUMP= { (Y) }
        { (N) }
```

Y

Yes. Causes a hexadecimal dump of the entire contents of YELOG. (This is a debug option.)

N

No. Do not dump YELOG.

```
7 ESUMFIL= { (I) }
           { (U) }
```

I

Initialize YESUM file.

U

Update YESUM file.

```
8 DMPSUM= { (Y) }
          { (N) }
```

Y

Yes. Causes a hexadecimal dump of the entire contents of YESUM. (This is a debug option.)

N

No. Do not dump YESUM.

```
RV ONUERL,,OPT=YES                                                              C 13:28:25
21 RV0030   ENTER SPOOL OPTION (YES,NO)                                         C 13:28:30
22?RV0030   JOB=ONUERL   SYMBOL=SPOOL    VALUE=NO     *ENTER VALUE              C 13:28:30
22 YES                                                                          C 13:28:36
 ONUERL     (2)      (3)      (4)       (5)       (6)    (7)                    C 13:28:48
23 ONUERL   JC01  JOB ONUERL   EXECUTING JOB STEP ONUERL00 #001 13:28:50        C 13:28:50
24 ONUERL    - OS/3 ERROR LOG EDITOR. VER. 48/20.                               C 13:28:53
25?ONUERL    - MODIFY PROGRAM OPTIONS (Y,N,?)    ?                              C 13:28:53
25 ?                                                                            C 13:28:56
26 ONUERL    - PRESET OPTIONS :                                                 C 13:28:57
27 ONUERL   OPEN=CURRENT, TRACE=THRESHOLD, ELOG=YES, DOC=NO,                    C 13:28:57
28 ONUERL   THRESHOLDS=NO, DUMP=NO, ESUMFIL=I, DMPSUM=NO                        C 13:28:57
29?ONUERL    - MODIFY PROGRAM OPTIONS (Y,N,?) ?                                 C 13:28:57
29 Y                                                                            C 13:29:05
30?ONUERL    - ENTER OPTIONS, SEPARATED BY COMMAS (1,2,3,4,5,6,7,8,?)           C 13:29:05
30 ?                                                                            C 13:29:07
31 ONUERL    - 1=OPEN, 2=TRACE†, 3=ELOG, 4=DOC, 5=THRESHOLDS,                   C 13:29:08
32 ONUERL   6=DUMP, 7=ESUMFIL, 8=DMPSUM                                         C 13:29:08
33?ONUERL    - ENTER OPTIONS, SEPARATED BY COMMAS (1,2,3,4,5,6,7,8,?)           C 13:29:08
33 1,2,3,4,5,6,7,8                                                              C 13:29:12
34?ONUERL    - 1: OPEN= (C,P,B,?)                                               C 13:29:12
34 ?                                                                            C 13:29:14
35 ONUERL    - C=CURRENT.(ERRORS LOGGED SINCE LAST EDIT).                       C 13:29:14
36 ONUERL      P=PRIOR.(REPEAT LAST EDIT).                                      C 13:29:15
37 ONUERL      B=BEGIN.(ENTIRE CONTENTS OF ERROR LOG).                          C 13:29:15
38?ONUERL    - 1: OPEN= (C,P,B,?)                                               C 13:29:15
38 P                                                                            C 13:29:26
39?ONUERL    - 2: TRACE†= (Y,T,N,R,D,?)                                         C 13:29:26
39 ?                                                                            C 13:29:30
40 ONUERL    - Y=ERROR SUMMARIES TO TRACE CENTER.                               C 13:29:30
41 ONUERL      T=SUMMARY TO TRACE IF THRESHOLD EXCEEDED.                        C 13:29:30
42 ONUERL      N=NO TRACE TRANSMISSION THIS RUN.                                C 13:29:30
43 ONUERL      R=RETRANSMIT LAST SUMMARY TO TRACE CENTER.                       C 13:29:31
44 ONUERL      D=INHIBIT TRACE PERMANENTLY.                                     C 13:29:31
45?ONUERL    - 2: TRACE†= (Y,T,N,R,D,?)                                         C 13:29:31
45 T                                                                            C 13:29:45
46?ONUERL    - 3: ELOG=(N,Y,?)                                                  C 13:29:45
46 ?                                                                            C 13:29:48
47 ONUERL    - N=SUSPEND ERROR LOGGING DURING EDIT.                             C 13:29:49
48 ONUERL      Y=WILL HAVE NO EFFECT ON LOGGING.                                C 13:29:49
49?ONUERL    - 3: ELOG=(N,Y,?)                                                  C 13:29:49
49 Y                                                                            C 13:29:58
50?ONUERL    - 4: DOC= (Y,N,?)                                                  C 13:30:58
50 ?                                                                            C 13:30:02
51 ONUERL    - Y=DISPLAY OPERATING INSTRUCTIONS ON PRINTER.                     C 13:30:02
52 ONUERL      N=INHIBIT DISPLAY OF INSTRUCTIONS.                               C 13:30:02
53?ONUERL    - 4: DOC= (Y,N,?)                                                  C 13:30:02
```

† Models 8, 10, 15, and 20 printouts print SUMMARY instead of TRACE at these three locations.

Figure A-1. Responding to the ONUERL Program (Part 1 of 3)

```
53 N                                                                    C 13:30:05
54?ONUERL    - 5: THRESHOLDS= (Y,N,D,?)                                 C 13:30:05
54 ?                                                                    C 13:30:07
55 ONUERL    - Y=MODIFY RECOVERABLE ERROR THRESHOLDS                    C 13:30:07
56 ONUERL      N=RUN WITH CURRENT THRESHOLDS                            C 13:30:07
57 ONUERL      D=DISPLAY CURRENT THRESHOLDS                             C 13:30:07
59?ONUERL    - 5: THRESHOLDS= (Y,N,D,?)                                 C 13:30:08
59 Y                                                                    C 13:30:30
60?ONUERL    - 6: DUMP= (Y,N,?)                                         C 13:30:31
60 ?                                                                    C 13:30:54
61 ONUERL    - Y=HEX. DUMP OF ENTIRE $Y$ELOG.                           C 13:30:54
62 ONUERL      N=INHIBIT DUMP.                                          C 13:30:55
63 ONUERL   NOTE: THIS OPTION WILL NOT BE ACCEPTED IF TRACE ON.         C 13:30:55
64?ONUERL    - 6: DUMP= (Y,N,?)                                         C 13:30:55
54 N                                                                    C 13:31:00
                                                  JOB NOT FND           C 13:31:00
64 N                                                                    C 13:31:03
65?ONUERL    - 7: ESUMFIL= (I,U,?)                                      C 13:31:03
65 ?                                                                    C 13:31:07
66 ONUERL    - I=INITIALIZE ESUM.                                       C 13:31:07
67 ONUERL      U=UPDATE ESUM                                            C 13:31:08
68?ONUERL    - 7: ESUMFIL= (I,U,?)                                      C 13:31:08
68 I                                                                    C 13:31:13
69?ONUERL    - 8: DMPSUM= (Y,N,?)                                       C 13:31:13
69 ?                                                                    C 13:31:18
70 ONUERL    - Y=HEX. DUMP OF $Y$ESUM FILE.                             C 13:31:19
71 ONUERL    - N=INHIBIT DUMP                                           C 13:31:19
72?ONUERL    - 8: DMPSUM= (Y,N,?)                                       C 13:31:19
72 N                                                                    C 13:31:27
73?ONUERL    - ENTER CUSTOMER ID. (----------,?)                        C 13:31:32
73 C42722                                                               C 13:31:39
11?ONUERL    - CUSTOMER-ID IS: C42722   . CONFIRM (Y OR N)              C 13:31:39
11 Y                                                                    C 13:31:41
12?ONUERL    - ENTER ONE (WSTA,DISC,DSKT,TAPE,URCD,?)                   C 13:31:51
12 ?                                                                    C 13:32:02
13 ONUERL    - WSTA : WORK-STATION SUB-SYSTEM                           C 13:32:03
14 ONUERL      DISC : DISC SUB-SYSTEM                                   C 13:32:05
15 ONUERL      DSKT : DISKETTE SUB-SYSTEM                               C 13:32:06
16 ONUERL      TAPE : TAPE SUB-SYSTEM                                   C 13:32:06
17 ONUERL      URCD : UNIT-RECORD (PRINTER, READER, PUNCH)              C 13:32:06
18?ONUERL    - ENTER ONE (WSTA,DISC,DSKT,TAPE,URCD,?)                   C 13:32:06
18 WSTA                                                                 C 13:32:16
19?ONUERL    - WSTA - ENTER ADDRESS/ES (XXX(,XXX),ALL,?)                C 13:32:16
19 ?                                                                    C 13:32:21
20 ONUERL   EXAMPLE: 100,101,103                                        C 13:32:21
21?ONUERL    - WSTA - ENTER ADDRESS/ES (XXX(,XXX),ALL,?)                C 13:32:22
21 315                                                                  C 13:32:35
```

Figure A-1. Responding to the ONUERL Program (Part 2 of 3)

```
22?ONUERL     - WSTA - ENTER VALUE (Ø.XXX,?)                              C 13:32:35
22 ?                                                                       C 13:32:4Ø
23 ONUERL   Ø.ØØØ = VALUE = 1.ØØØ                                         C 13:32:4Ø
24?ONUERL     - WSTA -ENTER VALUE (Ø.XXX,?)                               C 13:32:41
24 Ø.ØØØ1                                                                  C 13:32:46
25?ONUERL     - ANY FURTHER THRESHOLD MODIFICATION ?(Y,N)                 C 13:32:46
25 Y                                                                       C 13:32:49
26?ONUERL     - ENTER ONE (WSTA,DISC,DSKT,TAPE,URCD,?)                    C 13:32:49
26 DSKT                                                                    C 13:32:56
27?ONUERL     - DSKT - ENTER ADDRESS/ES (XXX(,XXX),ALL,?)                 C 13:32:56
27 ALL                                                                     C 13:33:ØØ
28?ONUERL     - DSKT - ENTER VALUE (C.XXX,?)                              C 13:33:ØØ
28 Ø.ØØ1                                                                   C 13:33:1Ø
29?ONUERL     - ANY FURTHER THRESHOLD MODIFICATION ?(Y,N)                 C 13:33:1Ø
29 N                                                                       C 13:33:29
3Ø ONUERL     SORT MIØØ END OF SORT                                       C 13:33:43
31 ONUERL   SORT AI86 RECORD COUNT          9Ø IN,         Ø DELETED     C 13:33:43
32 ONUERL     - CONNECTION TO REMOTE SUPPORT CENTER REQUIRED              C 13:33:47
33 ONUERL   ** CONNECT OR CANCEL                                          C 13:33:47
CA ONUERL,NO                                                               C 13:34:Ø1
   (1)       (2)      (3)      (4)      (5)       (6)      (7)            C 13:34:Ø5
34 ONUERL   JCØ2  JOB ONUERL   TERMINATED NORMALLY          13:34:Ø6     C 13:34:Ø6
```

Figure A-1. Responding to the ONUERL Program (Part 3 of 3)

Note: If this example were of an ONUERL session on the System 80 models 8-20, some of the messages would be different. First, all messages containing the word "TRACE" would display the word "SUMMARY" instead. Also, the following messages would replace messages 28, 29, 30, and 45 in the example:

```
29 ONUERL - Y=SPOOL SUMMARY TO SYSTEM JOURNAL DISKETTE
            T=SPOOL SUMMARY IF THRESHOLD EXCEEDED
29          N=NO SUMMARY SPOOLING THIS RUN
30          R=REPEAT LAST SUMMARY SPOOLING
            D=INHIBIT SUMMARY SPOOLING PERMANENTLY
                        .
                        .
                        .
45 ONUERL - LOAD SYSTEM JOURNAL DISKETTE IN FDD1 AND GO
```

A.3. Making the TRACE Connection

- Models 3-6 Users

If you've taken the default for the TRACE option (TRACE=THRESHOLD), a request for the TRACE connection may or may not be made, depending on whether error threshold was exceeded. The next message the program displays is:

```
TRACE CONNECTION REQUIRED ** CONNECT OR CANCEL **
```

Use an acoustic coupler to connect the system to the TRACE center. The *System 80 Operator Maintenance Guide* (UP-8915) gives detailed procedures for making the TRACE connection. Refer to these procedures to avoid common errors that cause the TRACE connection to abort. If the TRACE connection is made, the local screen goes completely blank and "RMI" should appear on the system state (indicator) line in reverse video. If "RMI" does not appear, it means the connection wasn't successful; cancel ONUERL using the CANCEL command, and TRACE data will be saved. At the next execution of ONUERL, the program will automatically request the TRACE connection and display the following message:

```
SUMMARY NOT SENT TO TRACE. SEND? (Y,N)
```

Respond N to terminate ONUERL, which prevents loss of data.

- Models 8-20 Users

For models 8-20 users, ONUERL does not automatically transmit error summary data to a Unisys TRACE center. Instead, if error threshold has been exceeded or you've chosen the TRACE=YES option, ONUERL requests that you mount the system journal diskette in flexible disk drive 1 (FDD1). Then, it spools the error summary data to the diskette and terminates. The following message instructs the operator to transmit summary data to a Unisys TRACE center:

```
- CONTACT REMOTE SUPPORT CENTER FOR TRANSMISSION -
```

The system operator is responsible for transmitting the error summary data to a Unisys TRACE center. To do so, you display the configuration frame (press ESC, then C) to enable remote maintenance. Next, mount the system journal diskette in FDD1 and display the manual frame (press ESC, then M). Through the manual frame, you send the error summary data on the system journal diskette to the TRACE center. See the *System 80 Model 8 Operator's Maintenance Guide* (UP-9710) for complete instructions on sending the data to the TRACE center when working on a model 8.

A.4. SDF$DSP System Program

If you run ONUERL on a System 80 models 3-6 and TRACE is active, the system automatically schedules the SDF$DSP system program for execution when ONUERL terminates. SDF$DSP prints the device-ids and microcode from the system definition file (SDF). Its purpose is to relay information about your system to the TRACE center.

If the TRACE connection is made, the information goes directly to the TRACE center; it is not printed or displayed locally. If you don't make the physical TRACE connection you should cancel ONUERL using the CANCEL command.

If you run ONUERL on a System 80 models 8-20, the content of the system definition file spools to the system journal diskette after all summary data is spooled. It also spools simultaneously to the local printer.

A.5. ONUERL Console Messages

Table A-2 lists the console messages that may be displayed during execution of the ONUERL program. An explanation and operation action, where required, also is provided.

Table A-2. ONUERL Console Messages

Message/Explanation/Operator Action
`ERROR READING PARAM, RØ=e-code. TERMINATE? (Y,N)` An error has been detected reading a // PARAM statement from the job control stream, where *e-code* specifies the error encountered. Respond Y to terminate program. Respond N to ignore.
`INVALID ENTRY nnnnn. IGNORE? (Y,N)` A // PARAM statement has the wrong format, where nnnnn is a display of the PARAM statement. Respond Y to ignore statement and continue. Respond N to terminate program.
`$Y$ELOG IS EMPTY` YELOG contains no new error data since the last time ONUERL was run. ONUERL is terminated.
`ERROR ACCESSING $Y$ELOG. RØ=e-code` An error has been detected attempting to access YELOG, where *e-code* is the error encountered. The program makes one attempt at recovering. If unsuccessful, ONUERL terminates.
`INVALID DATA IN $Y$ELOG. JOB CANCELLED` The data in YELOG is meaningless. ONUERL terminates.
`FIRST RECORD NOT HISTORICAL RECORD. JOB CANCELLED` The beginning of YELOG cannot be found. ONUERL terminates.
`COTABL FULL. JOB CANCELLED` The capacity of the internal device list has been exceeded. ONUERL terminates.
`MEDIA TABLE FULL` The capacity of the media error table has been exceeded. Further media errors are not saved.
`NO TRACE DATA AVAILABLE. JOB CANCELLED` In running ONUERL in the TRACE, REPEAT mode, no trace data is found in YESUM. ONUERL terminates.

continued

Table A-2. ONUERL Console Messages (cont.)

Message/Explanation/Operator Action
YESUM nnnn error. CD$FNMC: sense-data An error has been detected in accessing YESUM, where nnnn is OPEN, CLOSE, INPUT, or OUTPUT. *Sense-data* is the contents of the four sense bytes returned by the system. ONUERL terminates.
YESUM ERROR: EOF. FILE CAPACITY EXCEEDED An attempt to create a new record has failed for lack of space. ONUERL terminates. YESUM must be erased and reallocated, after which ONUERL can be re-executed. File YESUM will be reinitialized with the latest device list.
NEW SUMRCD CREATED FOR DEVICE nnn This channel/device number has been assigned to a new or different device type. A new record must be created and stored in YESUM.
INVALID DATE/TIME: xxx. JOB CANCELLED. A test of system date and time shows invalid values that are displayed. Error records cannot be summarized reliably. ONUERL is cancelled. Rerun ONUERL when date and time have been properly entered.
YESUM NOT INITIALIZED. This message appears if YESUM file was not initialized. It is followed by: **ENTER CUSTOMER-ID. (_________,?)** This message is displayed during YESUM file initialization. Enter your 9-character customer ID. The following message then appears: **CUSTOMER-ID IS your-customer-id. CONFIRM (Y,N)** This message echoes the customer ID you entered and requests your confirmation. If ? is entered as a response to the customer ID request, the following message appears: **DEPRESS EOT IF ID. NOT AVAILABLE.**
TRACE PERMANENTLY DISABLED. RESTORE? (Y,N) Option TRACE=YES or THRESHOLD was selected, but TRACE has previously been disabled. Respond Y to restore TRACE capabilities. Respond N to ignore this request.
CANNOT HONOR TRACE REQUEST. JOB CANCELLED. Option TRACE=REPEAT was selected, but TRACE has previously been disabled. No TRACE data is available for transmission.
SUMMARY NOT SENT TO TRACE. SEND? (Y,N) Over-threshold has occurred in the last execution of ONUERL, and TRACE data was not transmitted. Transmission must be made before the program can proceed with collection of new data. Respond Y to permit transmission. Respond N to abort program.
INVALID DISKETTE (R,C?) (This message applies only to model 8, 10, and 20 users). The diskette in FDD1 is not a system journal diskette. Respond R to retry after loading the proper diskette. Respond C to cancel the program.

A.6. ONUERL Reports

The reports obtained as a result of running the ONUERL program provide a comprehensive listing of all errors contained in the YELOG file, as defined by the program parameters. The main body of the reports consists of a single line for each error log entry. The entries are sorted chronologically by channel device number. The reports also include a summary of sense bits in error for the device, and a summary of total error count versus total count of valid I/Os by device address.

Note: *ONUERL reports do not include your customer-id unless you add it to the control stream. To do so, use the general editor (EDT). (The ONUERL job control stream resides in the YJCS file located on SYSRES.) When you've read the file into the workspace, look for the // PARAM LOG-ID=*'log id data' *statement, where* 'log id data' *is 29 characters reserved for your installation name followed by your 9-character customer-id. This information must be enclosed in single quotes. When you've saved the file, your customer-id is permanently added to the control stream.*

Figure A-2 shows revised sample ONUERL report formats. Tables A-3 and A-4 define supervisor characteristics and ELOG flags displayed in the system description lines of the ONUERL report. Tables A-5 and A-6 define MSE display and CPU retry stack log-out display of the ONUERL report. Table A-7 defines exigent machine check displays.

The description of each summary report heading follows:

`device name`
Names the device in error.

`CHDV`
The channel, subchannel, and device number representing the did for the device.

`ADDR`
The starting address in main storage for the job running at time of error report.

`SZE`
Main storage assigned to user job.

`CIQ/CAIQ/CEIQ`
Identifies the queue reporting the error: channel interrupt queue (CIQ), channel attention interrupt queue (CAIQ), or channel error interrupt queue (CEIQ).

`SDMA/DMA ICW`
Identifies the error I/O control word (ICW) and the type of channel to which device is connected:

SDMA	(shared direct memory access)
DMA	(direct memory access)

I/O
The number of valid I/Os prior to the error.

RTRY
The number of times the I/O was attempted before reporting the error.

STST
The device and subchannel status of the device in error.

DSX
The device status extension byte.

DSSC
The device state sequence code.

DSW
The directive status word.

SS0-SS11
The sense bytes transmitted by the device to the channel at the time of the error. Up to 12 sense bytes are printed.

OPR
Operator response to error message displayed on console workstation/screen, if any. Possible response is U (unrecovered), I (ignore), C (cancel), or R (retry).

REC
Indicates whether the I/O recovered (Y) or not (N).

date and time
The date and time the error was written to the error log.

J-NME
The name of the job running at time of error, if present in the error log.

VSN
The volume serial number of the media in use at the time of error, if present in the error log.

SENSE BYTE X'80'-X'01'
Summary counts of sense bits set for the device. These counts are taken from the opened error log file and will vary depending upon the OPEN option that was selected.

DEVICE #
The device address (channel, subchannel, and device number) of the device in error.

`DEVICE (did) number of ERROR/S OCCURRED FOR number of VALID I/O'S.`

By device address, a summary of total errors versus total I/Os. These counts are taken from the opened error log file and will vary depending upon the OPEN program option that was selected.

`VSN/ERROR TABLE FOR DVC (did)`

By device address and VSN, a summary of all errors taken from the opened error log file. The count will vary depending upon the OPEN program option that was selected.

Note: *The total of all VSN errors and the device error count will typically be equal only if all device errors are associated with a VSN.*

```
ONUERL-ERROR LOG EDIT. VERSION xx/xx,     site id____________________________          SERIAL# ____ yy/mm/dd  hh.mm  PAGE XX

*****device name*********************    *****device name*********************    *****device name*********************
*    DVC TYPE CODE:                  *    *    DVC TYPE CODE:                  *    *    DVC TYPE CODE:                  *
*    CHAN./DVC ADDR.: 00 00          *    *    CHAN./DVC ADDR.: 00 00          *    *    CHAN./DVC ADDR.: 00 00          *
*    FEATURES:                       *    *    FEATURES:                       *    *    FEATURES:                       *
*    24x80 SCREEN                    *    *    24x80 SCREEN                    *    *    24x80 SCREEN                    *
*    LGERMSK: xxxx, LGEMSK: xxxx     *    *    LGERMSK: xxxx, LGEMSK: xxxx     *    *    LGERMSK: xxxx, LGEMSK: xxxx     *
**************************************    **************************************    **************************************

ONUERL-ERROR LOG EDIT. VERSION 48/26      CUST. ID.: XXXXXXXXX                           yy/mm/dd   16.28  PAGE XX
CUSTOMER ID.: C42722  . OS/3 VER.   13, REV. 00, DATE  yy/mm/dd, TIME  07.27.30, FLAGS 0080,
                                                              CHARACTERSISTICS 33BE0080
SYSRES: 401, PRINTER: 330, READER: 332, $Y$ELOG: 401 MAIN STORAGE SIZE: 00100000, USER MEMORY SIZE:000C2400
 BIT MAP ADDRESS: 01F478, BIT MAP WORD COUNT:   1. BITMAP DATA: (IN BITS. EACH BIT=64K. O=ON LINE, 1=OFF LINE)
00000000 00000000 11111111 11111111

ONUERL-ERROR LOG EDIT. VERSION 48/26      CUST. ID.: XXXXXXXXX                   yy/mm/dd  hh.mm  PAGE    XX

* MSE *  yy/mm/dd, hh.mm.ss, K2: xxxx, CR13:xxxxxxxx, M-SIZE:xxxxxxxx, EL SZE:xxxxxxxx, LPSD TME:    h.mm.ss.mms
   EL.: x, ADR.:xxxxxxxx xxxxxxxx  CNTRL:xxxxxxxx(BIN), SYND.:xxxxxxxx(BIN), COUNT:      x(DEC), BYTE:   x, BIT: x
        x,      xxxxxxxx xxxxxxxx       xxxxxxxx           xxxxxxxx                x              x        x

*** CPU RETRY STACK LOG-OUT ***  yy/mm/dd  hh.mm.ss   K1 FACTOR:xxxx , CR12:xxxxxxxx , CR13:xxxxxxxx , ELAPSED TM    h.mm.ss.mms
   CLASS: xx, RETRY COUNT:   x, HPVSR STATE: xx, CUBA09: xx, CUBA0A: xx, CUBA0B: xx, CUBA0C: xx, WDx, x-x: xx
   CLASS: xx, RETRY COUNT:   x, HPVSR STATE: xx, CUBA09: xx, CUBA0A: xx, CUBA0B: xx, CUBA0C: xx, WDx, x-x: xx
   CLASS: xx, RETRY COUNT:   x, HPVSR STATE: xx, CUBA09: xx, CUBA0A: xx, CUBA0B: xx, CUBA0C: xx, WDx, x-x: xx
   CLASS: xx, RETRY COUNT:   x, HPVSR STATE: xx, CUBA09: xx, CUBA0A: xx, CUBA0B: xx, CUBA0C: xx, WDx, x-x: xx
```

NOTE:

The following EMCIC sample printout is for models 8, 10, 15, and 20 only.

Figure A-2. Sample Error Log Report Format (Part 1 of 11)

```
ONUERL-ERROR LOG EDIT. VERSION 48/26        CUST. ID.: XXXXXXXXX                              yy/mm/dd  15.28  PAGE  XX

* EMCIC *  yy/mm/dd  20.40.28  JOB:          .  ADDR: 000000.  SIZE: 000000.
MCIC: 0800005D 20000000.   PSW: E0010000 00000000.   TOD: 9786DAC6 33CE8000.
SPV RGS:00: 00000000, 01: 0000D7C0, 02: 000261B8, 03: 600261D2, 04: 000030D2, 05: 00008860, 06:046FAD42, 07: 0000E240,
        08: 00001720, 09: 0000E240, 10: 400034FA, 11: 00000400, 12: 0000DE94, 13: 40026IBA, 14: 000034F4, 15: 0000320A,
MMU LOG #0   0: 00000000 00000000,  1: 00000000 00000000,  2: 00000000 00000000,
             3: 00000000 00000000,  4: 00000000 00000000,  5: 00000000 00000000,
ED MCIC: 0800005D 20000000.  ED ER0: 00008001.  ED ER1: 8040A004.
MMU LOG #1   0: 00000000 00000000,  1: 00000000 00000000,  2: 00000000 00000000,
             3: 00000000 00000000,  4: 00000000 00000000,  5: 00000000 00000000,
BPU LOGOUT: 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
            00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
            00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
            00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 9786DAC6 33CE8000
LOW MEMORY (X'310'-'313'): 00000000
MACHINE CHECK INTERRUPT CODES:
SUBCLASS: EXTERNAL DAMAGE
AREA CODE: I/O EXTERNAL DAMAGE: CHC AND IOP ABNORMALITY
IOP ADDRESS: 2

ONUERL-ERROR  LOG  EDIT.  VERSION  48/26     CUST. ID.: XXXXXXXXX                           yy/mm/dd  16.28  PAGE  XX

device name. CHDV: xxxx, ADDR: xxxxxx. SZE: xxxxxx
     ICW/CCW                                      I/O   RTRY  STST  DSX  DSSC DSW SS0-1  2-3 4-5  6-7  8-9  10-11 OPR  REC
yy/mm/dd hh.mm.ss J.NME:________  VSN:______

xxxxxxxxxxxxxxxxxxxxxxxx     xxxxxxxxxxxxxxxxxxxxxxx  xxxxx  xx   xxxxx xxxx  xx   xx xxxxx xxxx xxxx xxxx xxxx xxxx   x    x

ONUERL-ERROR LOG EDIT. VERSION 48/26       CUST. ID.: XXXXXXXXX                           yy/mm/dd  15.28  PAGE  XX

U-14 TAPE DR., CHDV: 180
            CCW/TCW                             I/O    RTRY STST DSX DSSC DSW SS0-1 2-3 4-5 6-7 8-9 10-11 OPR REC
   yy/mm/dd 18.57.53 J-NME: DUMPTAPE, VSN: TAPE1
C308A80C60002900 0108A80C20022900                2  2 0E00 8000 35  00  0840 0084 0000 0000 0000 0000      Y
CMD: WRITE           XPCT BLK:  0001: RCVD BLK:  0001                          VRC SPEED CHK
```

Figure A-2. Sample Error Log Report Format (Part 2 of 11)

```
SENSE BYTE   X'8Ø'   X'4Ø'   X'2Ø'   X'1Ø'   X'Ø8'   X'Ø4'   X'Ø2'   X'Ø1'   DEVICE # xxx
     Ø        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     1        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     2        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     3        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     4        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     5        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     6        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     7        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     8        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     9        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     A        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     B        xxx     xxx     xxx     xxx     xxx     xxx     xxx     xxx
     x ERROR/S OCCURRED FOR    xxxxxxx VALID I/O's

VSN/ERROR TABLE FOR DVC 1B1.
VSN:TAPE2  :      24 ERRORS FOR    7475 I/O'S
VSN:SCRT   :      33 ERRORS FOR   12612 I/O'S
VSN:TESTØ1 :       6 ERRORS FOR    6283 I/O'S
VSN:OLMTAP :       1 ERRORS FOR       5 I/O'S
VSN:R512Q1 :      22 ERRORS FOR      25 I/O'S

ONUERL-ERROR LOG EDIT. VERSION 48/26     CUST. ID.: XXXXXXXXX                yy/mm/dd  15.28  PAGE  XX

847Ø DISC DR., CHDV: 29Ø, ADDR: ØØØØØØ, SZE: ØØØØØØ
                CCW                  I/O    RTRY STST DSSC   Ø-3    4-7    8-11    12-15    16-19    2Ø-23   OPR REC
   yy/mm/dd 14.Ø5.ØØ J-NME: ________, VSN: RELO8S
ØØØØØØØØØØØØØØØØ 13ØØ516860000008       1   1   Ø2ØØ   35 ØØØØØØØØ ØØØØØØ38 ØØØØØØØØ ØØØØØØØØ ØØØØØØØØ ØØØØØØØØ      Y
CMD: RECALIB.      CLY: 768,    HD:N/A .
```

Figure A-2. Sample Error Log Report Format (Part 3 of 11)

```
SENSE BYTE    X'80'    X'40'    X'20'    X'10'    X'08'    X'04'    X'02'    X'01'    DEVICE # xxx
     Ø         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     1         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     2         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     3         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     4         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     5         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     6         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     7         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     8         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     9         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     A         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     B         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     x ERROR/S OCCURRED FOR    xxxxxxx VALID I/O's

VSN/ERROR TABLE FOR DVC xxx.
VSN:REL08S :        x ERRORS FOR    xxxx I/O'S  CYL     HEAD    ERR
                                                 8Ø      N/A      Ø4
                                                768      N/A      Ø1
                                                 2Ø      N/A      Ø3

ONUERL-ERROR LOG EDIT. VERSION 48/26      CUST. ID.: XXXXXXXXX                            yy/mm/dd  15.28  PAGE  XX

8418 DISC DR., CHDV: 2A1
               CCW                                        I/O    RTRY STST DSX DSSC DSW SSØ-1 2-3 4-5 6-7 8-9 1Ø-11 OPR REC
CMD: SEEK-READ      CYL: 711,   HD:N/A .                                                  DTA FLD ECC
   yy/mm/dd 13.34.32 J-NME: PREP18  , VSN: SCRT
ØØØØØØØØØØØØØØØØ 17ØØ5BØØ6ØØØØØØ8                                1   1  ØEØØ ØØØØ 35  ØØ  Ø811 Ø22Ø ØCØ6 8ØØØ ØØØØ ØØØØ    Y
CMD: SEEK-READ      CYL: 711,    HD:N/A .                                                 DTA FLD ECC
```

Figure A-2. Sample Error Log Report Format (Part 4 of 11)

```
SENSE BYTE   X'80'    X'40'    X'2Ø'    X'1Ø'    X'Ø8'    X'Ø4'    X'Ø2'    X'Ø1'    DEVICE # xxx
    Ø          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    1          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    2          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    3          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    4          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    5          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    6          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    7          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    8          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    9          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    A          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    B          xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
    x ERROR/S OCCURRED FOR    xxxxxxx VALID I/O's

VSN/ERROR TABLE FOR DVC xxx.
VSN:RELØ8S :        x ERRORS FOR    xxxx I/O'S  CYL   HEAD   ERR
                                                 8Ø    N/A    Ø4
                                                768    N/A    Ø1
                                                 2Ø    N/A    Ø3

ONUERL-ERROR LOG EDIT. VERSION 48/26       CUST. ID.: XXXXXXXXX                          yy/mm/dd  15.28  PAGE  XX

8420/22  DSKTE, CHDV: C2Ø, ADDR: ØØØØØØ, SZE: ØØØØØØ
               ICW                                       I/O   RTRY  STST DSX DSSC DSW SSØ-1 2-3 4-5 6-7 8-9 1Ø-11 OPR REC
   yy/mm/dd 2Ø.57.17 J-NME: ________, VSN:________
56ØØ886A2ØØØØØ8ØØØØØØ7Ø ØØØØØØØØØØØØØØØØØØØØØØØØ       1     1   ØØØØ ØØØØ 35   Ø3  ØØØØ ØØØØ ØØØØ ØØØØ ØØØØ ØØØØ      N
CMD: READ VOL. ID
                                                                  I/O TERMINATION UNKNOWN
```

Figure A-2. Sample Error Log Report Format (Part 5 of 11)

```
SENSE BYTE    X'8Ø'    X'4Ø'    X'2Ø'    X'1Ø'    X'Ø8'    X'Ø4'    X'Ø2'    X'Ø1'    DEVICE # xxx
     Ø         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     1         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     2         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     3         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     4         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     5         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     6         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     7         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     8         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     9         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     A         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     B         xxx      xxx      xxx      xxx      xxx      xxx      xxx      xxx
     x ERROR/S OCCURRED FOR    xxxxxxx VALID I/O's

VSN/ERROR TABLE FOR DVC xxx.
VSN:ID16JØ :        x ERRORS FOR        x I/O'S  TRK     SCTR     ERR
                                                  Ø1      Ø1        Ø5
VSN:JHPDD1 :        x ERRORS FOR     xxxx I/O'S  TRK     SCTR     ERR
                                                  36      Ø6        Ø5
                                                  36      Ø7        Ø5
VSN:JHPDD2 :        x ERRORS FOR      xxx I/O'S  TRK     SCTR      ERR
                                                  36      Ø8        Ø5

DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
DEVICE xxx:  xxxx ERRORS OCCURRED FOR   xxxxxxx VALID I/O'S
```

Figure A-2. Sample Error Log Report Format (Part 6 of 11)

```
ONUERL-ERROR LOG EDIT. VERSION 48/26.     CUST. ID.: XXXXXXXX                           yy/mm/dd   15.28   PAGE  XX

OS/3 VER.  13/00    CUST. ID.: XXXXXXXXX      yy/mm/dd       14.31.15
FLAGS 0080, CHAR. 73B30080, MAIN STOR. 00600000
SYSRES: 2A0, PRNTR: 005, READER: C20, $Y$ELOG: 2A0
BIT MAP DATA (IN BITS. ONE BIT=64K. 0=ON LINE, 1=OFF LINE) :
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000

ONUERL-ERROR LOG EDIT. VERSION 48/26      CUST. ID.: XXXXXXXX                           yy/mm/dd   15.28   PAGE   XX

VER. 48/26   - SENSE SUMMARY DATA AT:   15.28

DVC LIST:C12,C14,C18,. . . . . . . .

SUMMARY SPAN:  yy/mm/dd. TO yy/mm/dd.

U-RCD ADDR          390
DEVICE             0770
THRESH. VALUE      .050
TOT. I/O COUNT       15
RECOV. COUNT          4
UNRECOV. COUNT        5
OPR RTRY              3
BUF LOAD CHECK        9

TAPE ADDR            5D3     480     481
DEVICE            U26/28    BT32    BT32
THRESH. VALUE       .050    .050    .050
TOT I/O COUNT      58582   68231   81323
NOISE                        114      42
SKEW CHCK              6
END CHK                        1
RD/WR VRC CHK.               174     211
MTE/LRC CHK.                   3       6
ENV/EDD CHK.          29     101      46
```

Figure A-2. Sample Error Log Report Format (Part 7 of 11)

```
DISC ADDR.         1A1       1A2       1A3
DEVICE            8494      8494      8494
THRESH. VALUE     .050      .050      .050
TOTAL I/O COUNT   6267      8746     12089

DISC ADDR.         1B6
DEVICE            8419
THRESH. VLAUE     .050
TOT. I/O COUNT     142
RECOV. COUNT        12
UNRECOV. COUNT      68
OPR RTRY             2
DVC CHCK            74
ID FLD DATA CHK      5
STOP STATE           2

DSKT ADDR          C21
DEVICE            8420/22
THRESH. VALUE     .050
TOT. I/O COUNT      19
RECOV. COUNT         7
UNRECOV. COUNT       4
DVC. NOT RDY         5
NO DTA SEP            6
```

Figure A-2. Sample Error Log Report Format (Part 8 of 11)

DELETION

Figure A-2. Sample Error Log Report Format
(Part 9 of 11) was deleted.

```
BETWEEN yy/mm/dd AND yy/mm/dd, OVER-THRESHOLD DEVICES:
* EMCIC *
DATE: yy/mm/dd TIME: 00.00.15   MCIC: 2020000D 000BA9D0
    ED MCIC: 00000000 00000000   PSW: E0040000 9006825C
    ED R0: 00000000  ED R1: 00000000
DATE: yy/mm/dd TIME: 00.00.15   MCIC: 2020000D 000BBD60
    ED MCIC: 00000000 00000000   PSW: E0040000 9006825C
    ED R0: 00000000  ED R1: 00000000

ERROR ACTIVITY BY VSN.
 DVC   VSN        TIC       TOC      TIC/TOC

2A0   REL08X        24     67921  0.000353
CYL,ER#: 155,   2; 536,   1; 466,   4;  28,  2;
 162,  10; 132,   2; 140,   1; 465,   1; 776,  1;

      REL082        22      68105  0.000323
CYL,ER#: 462,   3; 466,    4; 464,  2; 163,  1;
 162,   1; 543,   2; 542,    1; 140,  2; 781,  1;
  28,   2; 155,   3;

2A1   DATAPK        89      39856  0.002233
CYL,ER#:  54,   2;  69,   12;  81,   7; 145, 12;

ONUERL-ERROR LOG EDIT. VERSION 48/26      CUST. ID.: XXXXXXXXX                     yy/mm/dd   15.28  PAGE  XX

  70,  18;  71,  11;  00,   8;  74,   9;  55,    9;
   53,  1;
      SCRT          49    51091  0.000959
CYL,ER#:  11,   4; 424,   8; 453,  11; 711,  12;
 764,  14;

290   REL08S         8      3095  0.002585
CYL,ER#:  80,   4; 768,   1;  20,   3;
C20   ID16J0         5         6  0.833333
TK,SCTR,ER#:  01, 01,   5;
      JHPDD2         5       674  0.007418
```

Figure A-2. Sample Error Log Report Format (Part 10 of 11)

```
TK,SCTR,ER#:  36, 08,     5;
       JHPDD1       10      7277  0.001374
TK,SCTR,ER#:  36, 06,     5; 36, 07,    5;

C21   0S3PCB         2         10  0.2
TK,SCTR,ER#:  16, 01,    2;

- END OF DATA
NOTES: + SIGNS IN TOTAL I/O COUNT INDICATE THAT THE I/O COUNTER HAS WRAPPED.
       NUMBER SHOWN IS MINIMUM AND MAY BE GREATER BY 32768 BYTE INCREMENTS.

       ** IN OTHER COUNTS INDICATE THAT ERROR THRESHOLD HAS BEEN EXCEEDED.

       ONLY COUNTERS WITH NON-ZERO VALUES ARE DISPLAYED.
        I/O COUNT INCLUDES THE ONE ON WHICH ERROR HAS INCURRED.

ONUERL-ERROR LOG EDIT. VERSION 48/26      CUST. ID.: XXXXXXXXX                       yy/mm/dd   15.28  PAGE XX

SYSTEM DEFINITION FILE

DEVICE ADDRESS          C10
TYPE/FEATURE            WSC
MICROCODE NAME          WSC2-400

DEVICE ADDRESS          C13
TYPE/FEATURE            3560
MICROCODE NAME          WS1DJ000
```

NOTE:

Each device in the device table is printed here in this report.

```
BASE OS/3 RELEASE   -- 13.0.0
```

Figure A-2. Sample Error Log Report Format (Part 11 of 11)

Table A-3. CHARACTERISTICS Field Displayed in ONUERL Report

Byte No.	Bit No.	Definition
0	0	Minimum supervisor
	1	Software performance monitor
	2	Online diagnostics
	3	Timer
	4	Daily build supervisor
	5	Do not terminate job for EC39
	6	Suppress R06 message display
	7	Memory consolidation
1	0	Error logging
	1	User files locked
	2	Job accounting
	3	Spooling
	4	Roll out/roll in
	5	Input spooling
	6	Shared code
	7	Tape block numbers
2	0	Console log
	1	Retain console log
	2	File lock = share
	3	File lock = granule
	4	Interactive OS/3
	5	Maximum CPU time configuration
	6	Maximum wall time configuration
	7	Workstation support
3	0	RC symbiont available
	1	CDI supported
	2	DTFs supported
	3	Ignore DM SFTs
	4	Dynamic buffer management
	5	Resident loader
	6	Setime support
	7	N/A

Table A-4. FLAGS Field Displayed in ONUERL Report

Byte No.	Bit No.	Definition
0	N/A	N/A
1	0	Error log effective
	1	Error log active
	2	Logging machine check HPR
	3	ELOG not recovered at IPL
	4	Exigent machine check to be logged
	5	Retry stack machine check to be logged
	6	Main storage machine check to be logged
	7	Resident ELOG in control

Table A-5. MSE Display of ONUERL Report

Item	Definition
ADR	Address range in hexadecimal
CNTRL bits	0: ENC (ENable Count) When 1, permits the incrementation of the near-valid count field in this MSE log entry. 1: ENI (ENable Incrementation) When 1, permits the incrementation of the total near-valid count field in control register 13. 2: SR (Source Report) When 1, indicates the near-valid count field in this MSE log entry has been incremented. 3: MS (Multiple Syndrome) When 1, indicates that one or more syndrome changes have been detected while logging into this MSE log entry. 4–6: Reserved 7: CO (Carry Out) When 1, indicates that near-valid count in this MSE log entry has wrapped around ($FFFF_{16}$ to 0000_{16}). Multiple occurrences of wraparound will be merged.
SYND. (SYNDrome)	Contains the most recent syndrome recorded, if any, in this MSE log entry. This identifies the byte and bit position (X-coordinate) at which the single-bit failure (corrected) occurred.
COUNT	Contains a count of the occurrences of logging into the particular MSE log entry.

Table A-6. CPU Retry Stack Log-Out Display of ONUERL Report

Item	Definition
error CLASS	00: Not used 01: Internal CPU 02: COS 03: Main storage 04: M-bus 05: D-bus 06: MSP detected unconnectable ECC error 07: MSP detected error (non-ECC related) 08: I/O structure fault 09: Channel check 0A: Machine check hardware mask (MC bit) set 0B–FF : Reserved
RETRY COUNT	Indicates the number of retries that occurred in the process of recovering from the failure described. Exceeding 255 retries results in an exigent machine check.
HyPerViSoR State indicator (currently active slot)	01: DBUS extension 02: DMA 03: MLCM 04: SDMA 05: QUER 06: Maintenance panel 07: S-2
WD2 and WD3	The format of words 2 and 3 depends upon the error class: Error Class 01, 02 Word 3, bits 16–31 indicate the failing microinstruction or COS data address. Error Class 03, 06, 07 Word 3, bits 0–7 indicate the absolute main storage address. Error Class 04 Words 2 and 3 contain the error report from failing unit.

Table A-7. Exigent Machine Check Displays

Item	Definition
MCIC	Machine check interrupt code
PSW	Program status word
TOD	Time of day
SPV RGS	Supervisor registers
ED MCIC	External damage MCIC
ED ERO, ER1	External damage registers 0,1

A.7. ONELAN Program

The ONELAN program provides error log analysis and functions as part of ONUERL. You use ONELAN to display accumulated error data in the form of graphs and tables. Analyzing errors in this way provides you with:

- Equipment deviation
- Equipment trends and performance
- Significant error counts
- Serious equipment error problems

You can obtain this information for individual devices (device error count) or for all devices of a particular type (subsystem error count). The error counts are based on monthly totals. Each graph displays up to five months of accumulated errors.

Before running ONELAN for the first time, ONUERL version 48/14 must be run to properly initiate the YESUM file.

To execute ONELAN, enter the following command.

```
RV△ONELAN
```

ONELAN then displays the following message:

```
OS/3 ERROR LOG ANALYSIS. VER. 48/1
DO YOU NEED HELP RUNNING THIS PROGRAM? (Y,N)
```

Your answer to this message determines which mode you will use: standard or tutorial. Respond with N for standard mode, Y for tutorial.

- Standard Mode

  If you chose standard mode, the following message is displayed:

  ```
  ENTER REQUEST (LST,DOC,ALL,DVC,SUB,END,?)
  ```

  Your answers may be free form but must be separated by commas. If you enter '?', an explanation of the requests is displayed and the original request is repeated:

  ```
  AVAILABLE REQUESTS:
  LST:  DISPLAY LIST OF AVAILABLE GRAPHS.
  DOC:  PRINT OPERATING INSTRUCTIONS.
  ALL:  PRINT ALL GRAPHS.
  END:  TERMINATE PROGRAM.
  DVC:  PRINT DEVICE GRAPH.
  SUB:  PRINT SUB-SYSTEM GRAPH.
  ENTER REQUESTS (LST,DOC,ALL,DVC,SUB,END,?)
  ```

If more than one request is entered on a line (for example: LST,DOC,DVC), the requests will be honored one at a time. The END request, if entered, is always the last request honored. If no END request has been entered, new requests will be solicited by the original request message.

When you request a device graph (DVC), the following message appears:

```
ENTER DEVICE ADDRESS (CSU(,CSU)...,?)
```

You must now enter the channel, subchannel, and device address. Multiple addresses must be separated by commas. To display a sample address, enter a question mark.

When you request a subsystem graph, the following message appears:

```
ENTER SUB-SYSTEM ADDRESS (CS(,CS)...,?)
```

You must now enter the channel/subchannel address. Multiple addresses must be separated by commas. To display a sample address, enter a question mark.

After the device or subsystem has been validated, the requested graphs are directed to the line printer either directly or via spooling.

To terminate ONELAN, enter the END request when the original request message (ENTER REQUEST) appears.

A sample ONELAN dialog in standard mode is shown in Figure A-3.

```
11  ONELAN - OS/3 ERROR LOG ANALYSIS. VER. 48/01
12 ?ONELAN - DO YOU NEED HELP RUNNING THIS PROGRAM? (Y,N)
12   N
13 ?ONELAN - ENTER REQUEST (LST,DOC,ALL,DVC,SUB,END,?)
13   LST,ALL,SUB
14  ONELAN - GRAPHS AVAILABLE : 311,312,313,314,330,332,370,100,101,103
15  04,105,320,321
16 ?ONELAN - ENTER SUB-SYSTEM ADDRESS (CS(,CS)...,?)
16   31,32,33,10,37
17 ?ONELAN - ENTER;REQUEST;(DOC,ALL,DVC,SUB,END,?)
17   DVC,END
18 ?ONELAN;-;ENTER;DEVICE;ADDRESS;(CSU(,CSU)...,?)
18   103,104,321
```

Figure A-3. Sample ONELAN Dialog, Standard Mode

- Tutorial Mode

  If you chose tutorial mode, ONELAN prompts you for the following requests:

  ```
  DO YOU NEED OPERATING INSTRUCTIONS? (Y OR N)
  DO YOU WANT A LIST OF AVAILABLE GRAPHS? (Y OR N)
  DO YOU WANT ALL DEVICES GRAPHED? (Y OR N)
  DO YOU WANT SUB-SYSTEM GRAPH(S)? (Y OR N)
  ```

  If you answer Y to the last request, the following request appears:

  ```
  ENTER SUB-SYSTEM ADDRESS (CS(,CS)...,?)
  ```

  Enter the channel and subchannel addresses of the devices you want graphed, or enter a question mark (?) to display an example:

  ```
  EXAMPLE: 10,37
  ```

  The subsystem address request is then repeated. After you respond with the addresses, ONELAN asks you if you want device graphs:

  ```
  DO YOU WANT DEVICE(S) GRAPH(S)? (Y OR N)
  ```

  A response of Y is followed by the next request for the channel address (C), the subchannel address (S), and the device address (U):

  ```
  ENTER DEVICE ADDRESS (CSU(,CSU)...,?)
  ```

  Enter ? to display an example:

  ```
  EXAMPLE: 100,101,370
  ```

  The device address request is repeated after the example. After you enter the addresses (or if you answered N to the device graph request), the following request is displayed:

  ```
  TERMINATE PROGRAM? (Y OR N)
  ```

  Respond with Y to terminate ONELAN. If N is entered, the program will again make requests, starting with the subsystem request.

Notes:

1. *A response of Y to any request causes the appropriate graph / display to occur before the next request is displayed.*

2. *Before running ONELAN for the first time, ONUERL version 48 / 14 must be run. This is required to properly initiate YESUM, which is used by ONELAN.*

3. *It is recommended that ONUERL be run at least twice between executions of ONELAN. This results in more meaningful ONELAN output.*

A sample ONELAN dialog in tutorial mode is shown in Figure A-4.

```
11  ONELAN - OS/3 ERROR LOG ANALYSIS. VER. Ø1/Ø1.
12 ?ONELAN - DO YOU NEED HELP RUNNING THIS PROGRAM ?(Y,N)
12   Y
13 ?ONELAN - DO YOU NEED OPERATING INSTRUCTIONS ?(Y OR N)
13   N
14 ?ONELAN - DO YOU WANT A LIST OF AVAILABLE GRAPHS ?(Y OR N)
14   Y
15  ONELAN GRAPHS AVAILABLE : 33Ø,37Ø,1ØØ
16 ?ONELAN - DO YOU WANT ALL DEVICES GRAPHED ?(Y OR N)
16   Y
17 ?ONELAN - DO YOU WANT SUB-SYSTEM GRAPH(S) ?(Y OR N)
17   Y
18 ?ONELAN - ENTER SUB-SYSTEM ADDRESS (CS(,CS)...,?)
18   ?
19  EXAMPLE: 1Ø,37
2Ø ?ONELAN - ENTER SUB-SYSTEM ADDRESS (CS(,CS)...,?)
2Ø   37,33,1Ø
21 ?ONELAN - DO YOU WANT DEVICE(S) GRAPH(S) ?(Y OR N)
21   Y
24 ?ONELAN - ENTER DEVICE ADDRESS (CSU(,CSU)...,?)
24   ?
25  ONELANN - EXAMPLE: 1ØØ,1Ø1,37Ø
26 ?ONELAN - ENTER DEVICE ADDRESS (CSU(,CSU)...,?)
26   1ØØ
27 ?ONELAN - TERMINATE PROGRAM ?(Y OR N)
27   N
28 ?ONELAN - DO YOU WANT SUB-SYSTEM GRAPH(S  ?(Y OR N)
28   N
29 ?ONELAN - DO YOU WANT DEVICE(S) GRAPH(S) ?(Y OR N)
29   Y
3Ø ?ONELAN - ENTER DEVICE ADDRESS (CSU(,CSU)...,?)2
3Ø   37Ø
31 ?ONELAN - TERMINATE PROGRAM ?(Y OR N)
31   Y
32   1K JCØ2  JOB JPELAN    TERMINATED NORMALLY
   BR CN
```

Figure A-4. Sample ONELAN Dialog, Tutorial Mode

The following console workstation or console messages may be displayed during execution of the ONELAN program:

`INVALID ENTRY`

The request entered is invalid, probably due to incorrect spelling. The 'ENTER REQUEST' message will be redisplayed. Reenter the request.

```
TOO MANY ENTRIES. REQUEST TRUNCATED.
```

The request buffer cannot accommodate the number of requests entered. The leftover requests must be reentered when the accepted requests have been honored.

```
NO RECORD AVAILABLE FOR DVC xxx.
```

A request was made for a graph for the specified device, but there is no summary data for that device. The request is ignored.

```
$Y$ESUM nnnn error. CD$FNMC: sense data
```

An error was detected during accessing of YESUM, where nnnn is OPEN, CLOSE, INPUT, or OUTPUT. Sense data is returned by the system. ONELAN terminates.

Figures A-6 and A-7 show sample device error log analysis graphs for a 3561 workstation (address 312) and an 8420 diskette (address 321), respectively. The graph on the left is the device error count for the months of November and December. The graph on the right shows the ratio of the total incident count (device error count) to the total output count (TIC/TOC ratio). The exact error counts and ratios are given in the table at the bottom of each display.

Note: *As a new month's counts become available, the graphs will expand to the right until five months are shown. Thereafter, the leftmost month is truncated automatically.*

The following dialog was used to obtain the displays shown in Figures A-6 and A-7:

```
ENTER REQUEST (LST,DOC,ALL,DVC,SUB,END,?)
DVC,END
ONELAN - ENTER DEVICE ADDRESS (CSU(,CSU)...,?)
312,321
```

A sample subsystem error log analysis for diskette is provided in Figure A-7. The device error count and TIC/TOC ratio graphs show the device addresses for all diskettes present on the system, and chart their errors for the months of November and December, with each month appearing on a separate page. The exact totals and TIC/TOC ratios are given in the table at the bottom of each display. The third part of the subsystem error log analysis charts the total subsystem errors by month. The following dialog was used to obtain the subsystem error log analysis graphs for disk:

```
ONELAN - ENTER REQUEST (LST,DOC,ALL,DVC,SUB,END,?)
SUB
ONELAN - ENTER SUB-SYSTEM ADDRESS (CS,(,CS)...,?)
10
```

When the ENTER REQUEST message reappears, you may request additional displays or END.

```
ONELAN-ERROR LOG ANALYSIS. VER. 48/1.                        88/12/Ø7  12.Ø2  PAGE 1
                  SYSTEM-8Ø - 3561 WK-STA    ADDRESS: Ø312
     DEVICE ERROR COUNT                           DEVICE TIC/TOC RATIO
     25Ø -                                             1.Ø -
         I                                                 I
         I                                                 I
         I                                                 I
         I                                                 I X
     2ØØ -                                              .1 - X
         I                                                 I X
         I                                                 I X
         I                                                 I X
         I                                                 I X
     15Ø -                                             .Ø1 - X
ERROR    I X                             TIC/TOC           I X
         I X                                               I X
COUNT    I X                              RATIO            I X
         I X                                               I X
     1ØØ - X                                          .ØØ1 - X
         I X                                               I X
         I X                                               I X
         I X                                               I X
         I X                                               I X
      5Ø - X                                         .ØØØ1 - X
         I X                                               I X
         I X                                               I X
         I X                                               I X
         I X                                               I X
       Ø - X    X                                   .ØØØØ1 - X    X
        --I----I--                                         --I----I--
          NOV  DEC                                           NOV  DEC
                  MONTH         NOV        DEC
                  TIC            14Ø         Ø
                  TOC            5Ø6         3
                  TIC/TOC:   Ø.27668    Ø.
```

Figure A-5. Device Error Log Analysis for 3561 Workstation

```
ONELAN-ERROR LOG ANALYSIS. VER. 48/1.                  88/12/Ø7  12.Ø2  PAGE 2
                    SYSTEM-8Ø - 842Ø    DISKETTE ADDRESS: Ø312
     DEVICE ERROR COUNT                          DEVICE TIC/TOC RATIO
       25Ø -                                            1.Ø -
           I                                                I
           I                                                I
           I                                                I
           I                                                I
       2ØØ -                                             .1 -
           I                                                I
           I                                                I
           I                                                I
           I                                                I
       15Ø - X                                          .Ø1 - X    X
ERROR      I X                                  TIC/TOC     I X    X
           I X                                              I X    X
COUNT      I X                                  RATIO       I X    X
           I X                                              I X    X
       1ØØ - X                                         .ØØ1 - X    X
           I X                                              I X    X
           I X                                              I X    X
           I X                                              I X    X
           I X                                              I X    X
        5Ø - X                                        .ØØØ1 - X    X
           I X    X                                         I X    X
           I X    X                                         I X    X
           I X    X                                         I X    X
           I X    X                                         I X    X
         Ø - X    X                                  .ØØØØ1 - X    X
           --I----I--                                       --I----I--
            NOV  DEC                                         NOV  DEC
                      MONTH        NOV        DEC
                      TIC            15Ø          43
                      TOC          15Ø49        4357
                      TIC/TOC    Ø.ØØ9967    Ø.ØØ9869
```

Figure A-6. Device Error Log Analysis for 8420 Diskette

```
ONELAN-ERROR LOG ANALYSIS. VER. 48/1.                  88/12/10  18.04  PAGE 1
                     SYSTEM-80 -  DISK SUB-SYSTEM.
    DEVICE ERROR COUNT            MONTH:  DEC     DEVICE TIC/TOC RATIO
       250 -                                    1.0-
           I                                       I
           I                                       I
           I                                       I
           I                                       I
       200 -                                    .1 -
           I                                       I
           I                                       I
           I                                       I
           I                                       I
       150 -                                   .01 -
ERROR      I                           TIC/TOC     I
           I                                       I
COUNT      I                            RATIO      I                  X
           I                                       I                  X
       100 -                                   .001-                  X
           I                                       I                  X
           I                                       I                  X
           I                                       I                  X
           I                                       I              X   X
        50 -                                  .0001-              X   X
           I                                       I              X   X
           I                  X                    I              X   X
           I                  X                    I              X   X
           I              X   X                    I              X   X
         0 -              X   X       X      .00001-              X   X       X
           --I---I---I---I---I---I---I---I-        --I---I---I---I---I---I---I---I-
            105 106 107 100 101 102 103 104         105 106 107 100 101 102 103 104

DVC ID.    105        106        107        100        101        102        103
TIC            0          0          0         10         26          0          0
TOC            0          0          0      86236       5339          0          0
TIC/TOC:   0.         0.         0.      0.000116   0.00487     0.         0.         0.
```

Figure A-7. Disk Subsystem Error Log Analysis (Part 1 of 3)

```
ONELAN-ERROR LOG ANALYSIS. VER. 48/1.                  88/12/10  18.04  PAGE 2
                            SYSTEM-80 - DISK SUB-SYSTEM.
    DEVICE ERROR COUNT              MONTH:  NOV              DEVICE TIC/TOC RATIO
       250 -                         X               1.0-                 X
           I                         X                  I                 X
           I                         X                  I                 X
           I                         X                  I                 X
           I                         X                  I                 X
       200 -                         X               .1 -                 X           X
           I                         X                  I                 X           X
           I                         X                  I                 X           X
           I                         X                  I                 X           X
           I                         X                  I                 X           X
       150 -                         X              .01 -                 X           X
ERROR      I                         X      TIC/TOC     I                 X           X
           I                         X                  I                 X           X
COUNT      I                         X       RATIO      I                 X       X   X
           I             X           X                  I                 X       X   X
       100 -             X           X              .001-                 X       X   X
           I             X           X                  I                 X       X   X
           I             X           X                  I                 X       X   X
           I             X           X                  I                 X       X   X
           I             X           X                  I             X   X       X   X
        50 -             X           X             .0001-             X   X       X   X
           I             X           X   X              I             X   X       X   X
           I             X           X   X              I             X   X       X   X
           I             X           X   X              I             X   X       X   X
           I             X           X   X              I             X   X       X   X
         0 -             X   X       X   X        .00001-             X   X       X   X
           --I---I---I---I---I---I---I---I-            --I---I---I---I---I---I---I---I-
            105 106 107 100 101 102 103 104              105 106 107 100 101 102 103 104

DVC ID.    105        106      107        100       101       102      103
TIC              0         0         0     105         3       0     366
TOC              0         0         0  413005         2        0  103354
TIC/TOC   0.         0.         0.       0.000254   0.        0.       0.003541
```

Figure A-7. Disk Subsystem Error Log Analysis (Part 2 of 3)

```
ONELAN-OS/3     ERROR ANALYSIS PROGRAM.       MM/DD/YY    HH.MM       PAGE 3
                        SYSTEM-80      DISK SUB-SYSTEM
                          *-----------------------*
                          **  SUB-SYSTEM TOTALS  **
                          *-----------------------*
             SUB-SYSTEM ERROR COUNT                SUB-SYSTEM TIC/TOC RATIO
      1250 -                                 1.0 -
           I                                     I
           I                                     I
           I                                     I
           I                                     I
      1000 -                                  .1 -
ERROR      I                                     I
           I                         TIC/TOC     I
           I                          RATIO      I
           I                                     I
       750 -                                 .01 -
COUNT      I                                     I
           I                                     I
           I                                     I
           I                                     I
       500 -                                .001 -
           I                                     I
           I                                     I X
           I                                     I X    X
           I           X                         I X    X
       250 - X         X                   .0001 - X    X              X
           I X         X    X                    I X    X              X
           I X    X    X    X                    I X    X    X         X
           I X    X    X    X    X               I X    X    X    X    X
           I X    X    X    X    X               I X    X    X    X    X
         0 - X    X    X    X    X        .00001 - X    X    X    X    X
           --I----I----I----I----I               --I----I----I----I----I
            NOV  DEC  JAN  FEB  MAR               NOV  DEC  JAN  FEB  MAR

MONTH        NOV      DEC      JAN     FEB       MAR
TIC           XX       XX       XX      XX        XX
TOC          XXX      XXX      XXX     XXX       XXX
TIC/TOC:   .XXXX   .XXXXX   .XXXXX   .XXXX    .XXXXX
```

Figure A-7. Disk Subsystem Error Log Analysis (Part 3 of 3)

Appendix B
Supervisor Modification Procedures

For special processing requirements, your system administrator may tell you to modify the selections made for the supervisor during SYSGEN. Tables B-1 and B-2 list the procedures you use to modify the supervisor. Respond with Y to the MODIFY SUPERVISOR= IPL statement, then proceed with the appropriate operator action described in Table B-1.

Note: *For a description of system output messages and how to respond to them (via solicited input messages), see Section 8.*

Table B-1. How to Modify List of Resident Shared Code Modules

IPL Output Message	Operator Response
IS THE SHARED CODE DIRECTORY INDEX TO BE BUILT? (Y,N)	Y Specifies an index is to be built for the shared code library. The system will do a binary search on the index rather than a serial directory search. This allows faster loading of shared code modules into main storage and faster processing between two or more shared code modules. The directory itself uses approximately 4000 bytes of main storage. N No shared code directory index is to be built.
ANY RESIDENT SHARED MODULES TO ADD OR DELETE? (Y,N)	N Terminates output messages for modifying shared code module list. Specifies no modification is required to SYSGEN list of shared code modules to be made resident. Proceed with operator action described in Table B–2. Y Specifies SYSGEN list of shared code modules requires modification. Output message requesting a function is displayed.
FUNCTION? (ADD,DEL,LIST,END, HELP)	END Terminates output messages for modifying shared code module list. Specifies no further supervisor modification is required. The SYSGEN list of shared code modules is updated by the changes previously specified with the ADD and DEL functions. Modules in updated list are loaded and made resident until next IPL. Proceed with operator action described in Table B–2. HELP Displays brief description of the ADD, DEL, LIST, and END functions, followed by an output message requesting another function. LIST Displays a list of the individual shared code modules currently specified to be made resident, followed by an output message requesting another function.

continued

Table B-1. How to Modify List of Resident Shared Code Modules (cont.)

IPL Output Message	Operator Response
	$\left\{\begin{matrix}\texttt{ADD}\\ \texttt{DEL}\end{matrix}\right\}$ Specifies an individual module or group of modules is to be added to or deleted from the current list of shared code modules. Displays the following list of shared code module groups, followed by an output message requesting a group or module name. POSSIBLE SYSTEM GROUPS ARE: SF - SCREEN FORMAT SERVICES ISB - BASIC INTERACTIVE SERVICES ISF - FULL INTERACTIVE SERVICES EDT - EDITOR DP - DIALOGUE PROCESSOR RPG - RPG EDITOR DDPR - REMOTE DISTRIBUTED DATA PROCESSING DDPL - LOCAL DISTRIBUTED DATA PROCESSING ESC - ESCORT ESCF - ESCORT FULL BAS - BASIC
WHICH GROUP TO $\left\{\begin{matrix}\texttt{ADD?}\\ \texttt{DELETE?}\end{matrix}\right\}$	$\left\{\begin{matrix}\texttt{group-name}\\ \texttt{module-name}\end{matrix}\right\}$[,L] Specifies the group of shared code modules or the individual shared code module to be added or deleted. Provides option to display (L) individual module names as they are added or deleted from the current list. After processing the addition or deletion, an output message requests another function. NOTE: Run the job stream SCLIST to get a list of all system shared code modules filed in YCLOD.

Table B-2. How to Modify Buffer Sizes Used by Dynamic Buffer Management

IPL Output Message	Operator Response
ANY CHANGES TO DYNAMIC BUFFER MANAGEMENT PARAMETERS? (Y,N)	N Terminates output messages for modifying buffer sizes. Specifies no changes are required to the expansion region and resident buffer sizes specified at SYSGEN. Proceed with step 5 of the interactive IPL procedure (3.3.2 for models 4 and 6 or 6.3.2 for models 8 thru 20). Y Specifies SYSGEN-specified expansion region and resident buffer sizes require modification. Output message requesting new expansion region size is displayed.
ENTER NEW EXPANSION REGION SIZE? (CURRENTLY decimal-byte-size)	new-size Specifies new size*, in decimal bytes, of expansion region and terminates output messages for modifying buffer sizes. Expansion region size is changed until next IPL. Proceed with step 5 of the interactive IPL procedure (3.3.2 for models 4 and 6 or 6.3.2 for models 8 thru 20). Ø Specifies no expansion regions are to be allocated until the next IPL. Dynamic buffer management must use resident buffer pool. Output message requesting new resident buffer size is displayed.
ENTER NEW RESIDENT BUFFER SIZE? (CURRENTLY decimal-byte-size)	new-size Specifies new size*, in decimal bytes, of resident buffer pool and terminates output messages for modifying buffer sizes. Resident buffer size is changed until next IPL. Proceed with step 5 of the interactive IPL procedure (3.3.2 for models 4 and 6 or 6.3.2 for models 8 thru 20).

*The new decimal size can be specified with or without a comma, or as a multiple of K. For example:

200000	is 200,000 decimal bytes
200,000	is 200,000 decimal bytes
200K	is 200 x 1024 = 204,800 decimal bytes

Appendix C
Disk Cache Facility Supportive Information

This appendix provides you with background and illustrative information to help you use the disk cache facility optimally. Also refer to the DCF section (Section 12).

C.1. Cache Modules

There are two modules that support disk cache, the cache symbiont (SL@@CM) and the cache manager (CAxx0000). The cache symbiont is responsible for loading the cache manager, initializing the cache buffer, and processing unsolicited cache commands. Once the manager is loaded and initialized, the symbiont becomes inactive until an unsolicited cache command is entered. The manager is responsible for managing the cache buffer and handling I/O requests to the disk.

C.2. Main Storage Layout with Disk Cache

When disk cache is initialized at IPL, the cache buffer is established at the high end of main storage. The cache buffer contains the cache manager, a search/read buffer, and cache segments. The search/read buffer is used by search/read high or equal commands, which will be explained later.

Note: *8470, 8494, and M9720 disks have equivalent logical track sizes.*

The size of the search/read buffer is that of an 8470 or 8417 disk track depending on whether or not an 8470 (or equivalent) disk is configured. The manager is located in the beginning of the cache buffer and is followed by the search/read buffer. If the cache buffer size is greater than 511K, DCF automatically allocates two search/read buffers on models 8-20. Also, if the segment size configured is greater than or equal to the largest track size for a disk on the system, DCF does not build separate search/read buffers (models 8-20 only). DCF uses any available cache segments for search/read operations.

The rest of the cache buffer is divided into fixed size segments. These segments are used by disk cache to read blocks of data. Each segment is associated with a disk track and device address. It may take several segments to contain a track of data, depending on the segment size. If the segment size is larger than a track, then part of the segment will be unused.

Segments are not reserved for specific disks on the system, but are available for use by any disk, based on a least recently used (LRU) algorithm. Section C.6 explains how the segment size is determined.

Use the following formula to calculate the number of segments in the cache buffer for models 8-20.

$$\#Segments = \frac{(cache\ buffer\ size * 1024) - (cache\ manager + search/read\ buffer(s))}{(segment\ size * 256)}$$

The following example applies this formula. The cache manager is approximately 8K.

If:

cache buffer size	=	700K
segment size	=	48 sectors
search buffer	=	24576 bytes (8470)

Then:

#Segments = ((700 * 1024) - (8000 + (2 * 24576))) / (48 * 256)
= (716800 - 32576) / 12288
= 684224 / 12288
= 55

In this example there are 55 segments containing 48 sectors of disk data in the cache buffer.

When disk cache is initialized manually (via the CM command) the cache buffer cannot be placed in high main storage because that area is reserved for the job region. The cache buffer is obtained from an available buffer in middle main storage.

C.3. Cached I/O Commands

The cached I/O commands are read, write, and search/read high/equal commands. Search/read high/equal commands are primarily issued by data management, while searching through the index of a MIRAM file, as a fast method to locate the index record(s) that point(s) to the user data record. The search is based on a key argument that is supplied by the initiator of the search.

Search/read equal commands are not cached due to the infrequency of use. If data on the disk is accessed infrequently and is read into the cache, it displaces other data and increases disk cache overhead. Search/read equal commands are used mainly in VTOC searches by disk space management and in YSHR searches by data management during open and close processing.

C.4. I/Os Not Cached

Disk cache does not cache all read and search/read high/equal commands. It is important to know which I/Os are not cached because they may cause the cache statistics to reflect lower hit rates.

The following I/Os are not cached:

- Search/read or read multitrack
- Read with no data transfer
- Read larger than half of the segment size

A MIRAM file with IRAM characteristics or an IRAM file has an index that is more than one track. Data management issues multitrack searches and reads to this index.

A read with no data transfer is commonly used by data management when issuing a write command with a read verification option.

A read larger than half of the segment size is not cached, so that a program issuing large reads bypasses the caching of data, reduces disk cache overhead, and avoids the displacement of other disk data. If the disk data is accessed infrequently, then a large I/O buffer or the use of the selective caching options (see 12.3) is recommended. If the data is accessed frequently, it is recommended that the I/O buffer be less than or equal to the cache segment size, so that the data is cached.

C.5. Monitoring DCF Effectiveness

Generally, the larger the cache buffer, the greater the performance. However, there is an exception. When the production mix is random, there is a significantly large amount of cache overhead with a very low hit rate. In this situation, it is better to run without DCF.

The most effective way to determine the optimal cache buffer and segment sizes is to use the cache statistics to monitor DCF operations, while experimenting with different cache and segment sizes. See Section 12 for information on DCF statistics commands.

C.6. Fine-Tuning DCF

C.6.1. Specifying Segment Size

Disk cache segment size is specified by the following methods:

- System generation (SYSGEN)
- Unsolicited cache command (CM RESEG)
- Manual cache loading (CM SIZ and CM SEG)
- Default value

For information on specifying segment size at system generation, see the *Installation Guide* (UP-8839). See Section 12 for information on the CM#RESEG, SIZ, and SEG commands and default values.

The default segment sizes may not be optimal for every user. You may not have enough main storage to run with a large cache buffer to get large segments. Specifying segment size allows you to run with large segments even with a small cache buffer and, conversely, small segments with a large cache buffer.

If a segment size is not specified, the default values specified in the following table are used:

	Number of 256-Byte Sectors	
Cache Buffer Size	**Models 8-20 Segment**	**Models 4/6 Segment**
64/160 - 255K	12	10
256 - 511K	24	30
512 - 1,024K	48	60
1,025 - 8,192K	48	-

Note that 8417 and 8419 disks have 60 and 50 sectors per track respectively. The maximum default segment size is 48 sectors on models 8-20. Thus, with a 48-sector segment, an 8419 track would use one segment of 48 sectors and a second segment with 2 sectors (46 sectors of the second segment are unused). An 8417 track would use two segments of 48 and 12 sectors (36 sectors of the second segment are unused).

It is recommended that models 8-20 users with predominantly 8419 and/or 8417 disks run with a segment size that minimizes unused segment space (i.e., full track segment size).

The following lists the System 80 disks and the corresponding number of segments per track:

8430/8433	33 sectors (256-byte sectors)
8416/8418	40 sectors
8419	50 sectors
8417	60 sectors
8470/8494/M9720	96 sectors

Note: *8430/33 disks are variable-sectored disks. When you choose a segment size, consider the fact that smaller sector sizes increase track overhead.*

C.6.2. Optimizing Disk Cache Performance

Generally, sequential processing performs better with a small number of larger segments, and random processing performs better with a larger number of small segments. By experimenting with different sizes, you can find the combination best suited to your production environment. Figure C-1 shows the effect of segment size on run time for models 8-20.

You may need to remove a disk from disk cache for performance reasons. If data accessed infrequently is read into the cache buffer, then other disk data is displaced, a lower hit rate is obtained, and disk cache overhead is increased. For information on the remove and activate commands, see Section 12.

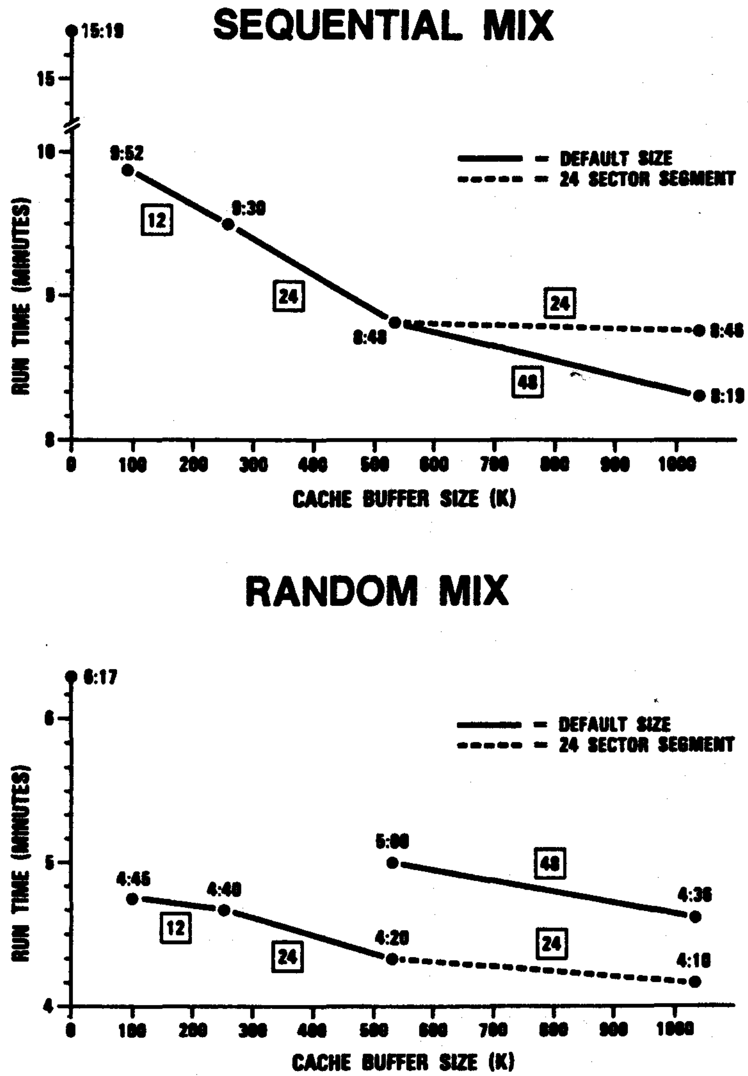

Figure C-1. Effect of Segment Size on Run Time for Models 8-20

Index

D

E

K

L

M

N

O

P

Q

R

S

T

U

V

W

NOTES

NOTES

NOTES

NOTES

NOTES

UNISYS

USER COMMENTS

We will use your comments to improve subsequent editions.

NOTE: Please do not use this form as an order blank.

(Document Title)

(Document No.) *(Revision No.)* *(Update Level)*

Comments:

From:

(Name of User)

(Business Address)

Fold on dotted lines, and mail. (No postage is necessary if mailed in the U.S.A.)
Thank you for your cooperation

CUT

FOLD

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 21 BLUE BELL, PA.

POSTAGE WILL BE PAID BY ADDRESSEE

Unisys Corporation
E/MSG Product Information Development
PO Box 500 — E5-114
Blue Bell, PA 19422-9990

FOLD

UNISYS

USER COMMENTS

We will use your comments to improve subsequent editions.

NOTE: Please do not use this form as an order blank.

(Document Title)

(Document No.) *(Revision No.)* *(Update Level)*

Comments:

From:

(Name of User)

(Business Address)

Fold on dotted lines, and mail. (No postage is necessary if mailed in the U.S.A.)
Thank you for your cooperation

CUT

FOLD

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 21 BLUE BELL, PA.

POSTAGE WILL BE PAID BY ADDRESSEE

Unisys Corporation
E/MSG Product Information Development
PO Box 500 — E5-114
Blue Bell, PA 19422-9990

FOLD

Help Us To Help You

Publication Title

Form Number　　　　Date

Unisys Corporation is interested in your comments and suggestions regarding this manual. We will use them to improve the quality of your Product Information. Please check type of suggestion:

☐ Addition　　☐ Deletion　　☐ Revision　　☐ Error

Comments

Name

Title　　　　Company

Address (Street, City, State, Zip)

Telephone Number

Help Us To Help You

Publication Title

Form Number　　　　Date

Unisys Corporation is interested in your comments and suggestions regarding this manual. We will use them to improve the quality of your Product Information. Please check type of suggestion:

☐ Addition　　☐ Deletion　　☐ Revision　　☐ Error

Comments

Name

Title　　　　Company

Address (Street, City, State, Zip)

Telephone Number

Help Us To Help You

Publication Title

Form Number　　　　Date

Unisys Corporation is interested in your comments and suggestions regarding this manual. We will use them to improve the quality of your Product Information. Please check type of suggestion:

☐ Addition　　☐ Deletion　　☐ Revision　　☐ Error

Comments

Name

Title　　　　Company

Address (Street, City, State, Zip)

Telephone Number

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

First Class Permit No. 21 Blue Bell, PA

Postage Will Be Paid By Addressee

Unisys Corporation
OS/3 Systems Product Information Development
PO Box 500 - E5-114
Blue Bell, PA 19422-9990

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

First Class Permit No. 21 Blue Bell, PA

Postage Will Be Paid By Addressee

Unisys Corporation
OS/3 Systems Product Information Development
PO Box 500 - E5-114
Blue Bell, PA 19422-9990

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

First Class Permit No. 21 Blue Bell, PA

Postage Will Be Paid By Addressee

Unisys Corporation
OS/3 Systems Product Information Development
PO Box 500 - E5-114
Blue Bell, PA 19422-9990

Help Us To Help You

Publication Title

Form Number Date

Unisys Corporation is interested in your comments and suggestions regarding this manual. We will use them to improve the quality of your Product Information. Please check type of suggestion:

☐ Addition ☐ Deletion ☐ Revision ☐ Error

Comments

Name

Title Company

Address (Street, City, State, Zip)

Telephone Number

Help Us To Help You

Publication Title

Form Number Date

Unisys Corporation is interested in your comments and suggestions regarding this manual. We will use them to improve the quality of your Product Information. Please check type of suggestion:

☐ Addition ☐ Deletion ☐ Revision ☐ Error

Comments

Name

Title Company

Address (Street, City, State, Zip)

Telephone Number

Help Us To Help You

Publication Title

Form Number Date

Unisys Corporation is interested in your comments and suggestions regarding this manual. We will use them to improve the quality of your Product Information. Please check type of suggestion:

☐ Addition ☐ Deletion ☐ Revision ☐ Error

Comments

Name

Title Company

Address (Street, City, State, Zip)

Telephone Number

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

First Class Permit No. 21 Blue Bell, PA

Postage Will Be Paid By Addressee

Unisys Corporation
OS/3 Systems Product Information Development
PO Box 500 - E5-114
Blue Bell, PA 19422-9990

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

First Class Permit No. 21 Blue Bell, PA

Postage Will Be Paid By Addressee

Unisys Corporation
OS/3 Systems Product Information Development
PO Box 500 - E5-114
Blue Bell, PA 19422-9990

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

First Class Permit No. 21 Blue Bell, PA

Postage Will Be Paid By Addressee

Unisys Corporation
OS/3 Systems Product Information Development
PO Box 500 - E5-114
Blue Bell, PA 19422-9990