

Program Logic

IBM System/360 Operating System Control Program With MFT Program Logic Manual

Program Number 360S-CI-505

This publication describes the internal logic of the IBM System/360 Operating System Control Program with MFT. The publication provides an introduction to control program logic and describes the components of the program. It also describes the initialization of the operating system, the functions of the supervisor that differ from those of the PCP and MVT supervisors and the functions of job management that differ from those of PCP and MVT job management.

The appendix contains a description of all routines, major tables, and work areas used by MFT, and flowcharts of the routines of MFT that differ from those of either of the other control programs.

Program Logic Manuals are intended for use by IBM customer engineers involved in program maintenance, and by system programmers involved in altering the program design. Program logic information is not necessary for program operation and use; therefore, distribution of this manual is limited to persons with program maintenance or modification responsibilities.

Restricted Distribution

RESTRICTED DISTRIBUTION: This publication is intended primarily for use by IBM personnel involved in program design and maintenance. It may not be made available to others without the approval of local IBM management.

Fifth Edition (June, 1969)

This is a major revision of, and obsoletes, Y27-7128-3 and Technical Newsletters Y28-2349 and Y28-2376. The text and illustrations have been changed to reflect the following:

- Multiple console support.
- The damage assessment routines of ABEND/ABTERM support.
- The resident reenterable routines facility. The facility allows the user to include both access method routine and other reenterable routines in the resident access method (RAM) area.
- The checkpoint/restart facility.
- The DISPLAY DSNAME and the MODE commands.

In addition the following are included; a revised description of the DEFINE processing routines; updated description of tables and work areas in Appendix A; revised module descriptions in Appendix E; and revised flowcharts in Appendix C.

The section of manual formerly titled Nucleus Initialization Program has been changed to delete the material now covered in the IPL and NIP Program Logic Manual, Form Y28-6661. The section is now titled Initialization of the Operating System and describes only the operation of the master scheduler after completion of the nucleus initialization.

Other changes to the text, and small changes to illustrations, are indicated by a vertical line to the left of the change; changed or added illustrations are denoted by the symbol • to the left of the caption.

This edition applies to release 18 of the IBM System/360 Operating System, and to all subsequent releases until otherwise indicated in new editions or Technical Newsletters. Changes are continually made to the specifications herein; before using this publication in connection with the operation of IBM systems, consult the latest IBM System/360 SRL Newsletter, Form N20-0360, for the editions that are applicable and current.

Requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Programming Systems Publications, Department D58, PO Box 390, Poughkeepsie, N. Y. 12602

Preface

This publication describes the differences in internal logic of the control program that result from the inclusion of multiprogramming with a fixed number of tasks (MFT). It is assumed that the reader of this publication is thoroughly familiar with the basic operation of the control program. Only areas of difference are discussed in this publication.

The manual is divided into four major sections. The Introduction describes control program functions, control program and main storage organization, and control program processing flow. The Initialization of the Operating System section describes differences introduced by MFT into system initialization. The Supervisor section describes supervisor functions including an explanation of task dispatching in MFT.

The Job Management section contains the changes to the job management components made by MFT. Job management is divided into three major components: reader/interpreter, initiator/terminator, and output writer. Also described are the Queue Manager which is used by all three major job management components, the Communications Task which handles operator-system communication, and the Master Scheduler Task which processes operator commands.

Appendix A contains descriptions of major tables and work areas used by MFT. Appendix B contains descriptions of modules used by MFT. Appendix C contains MFT flowcharts.

PREREQUISITE PUBLICATIONS

Knowledge of the information in the following publications is required for a full understanding of this manual.

IBM System/360 Operating System:

Principles of Operation, Form A22-6821

Introduction to Control Program Logic, Program Logic Manual, Form Y28-6605

Fixed-Task Supervisor, Program Logic Manual, Form Y28-6612

MVT Job Management, Program Logic Manual, Form Y28-6660

Initial Program Loader and Nucleus Initialization Program, Form Y28-6661

Planning for Multiprogramming With a Fixed Number of Tasks (MFT), Form C27-6939

The following publications may be useful for reference although they are not prerequisites for this publication.

IBM System/360 Operating System:

Concepts and Facilities, Form C28-6535

Linkage Editor, Form C28-6538

System Programmer's Guide, Form C28-6550

System Generation, Form C28-6554

MVT Control Program Logic Summary, Form C28-6658

Input/Output Supervisor, Program Logic Manual, Form Y28-6616

Contents

INTRODUCTION	9	WTO/WTOR Macro Instruction Processing	42
Functions of the Control Program With		External Interruption Processing . . .	42
MFT	10	Communications Task Modules	43
Job Management	10	Console Attention Interruption	
Task Management	10	Routine (IEECVCRA)	44
Data Management	10	Communications Task Wait Routine	
Control Program Organization	11	(IEECVCTW)	44
Resident Portion of the Control		Communications Task Router (IEECVCTR)	44
Program	11	Console Device Processor Routines	
Nonresident Portion of the Control		(IEECVPMX, IEECVPMC, IEECVPMP)	44
Program	11	Write-to-Operator Routines (IEECVWTO	
Main Storage Organization	11	and IEEVWTOR)	45
Fixed Area	12	External Interruption Routine	
Dynamic Area	12	(IEECVCRX)	45
Theory of Operation	14	Communications Task With Multiple	
INITIALIZATION OF THE OPERATING SYSTEM	18	Console Support	46
Main Storage Preparation	18	Master Scheduler Task	46
SUPERVISOR	21	Multiple Console Support Requirements	46
Interruption Supervision	21	SVC 34 Functions	47
The Dispatcher (Macro IEAAPS)	21	System Initialization	48
STAE Service Routine	28	Partition Definition by the Master	
ABEND and Damage Assessment Service		Scheduler	50
Routine	29	Job Processing	53
Damage Assessment Routines	30	Queue Manager	54
Task Supervision	31	Work Queues	54
The Attach Routine (Macro IEAAAT)	31	Queue Management	54
The Wait Routine (Macro IEAAWT)	32	Job Queue Initialization	54
The Post Routine (Macro IEAAPT)	32	Queue Manager Modules	55
The ENQ/DEQ Routine (IEAGENQ1)	32	Reader/Interpreter	60
Contents Supervision	33	Resident Readers	61
LINK Service Routine (Macro IEAATC)	33	Transient Readers	61
ATTACH Service Routine (Macro IEAAAT)	33	Reader Control Flow	61
LOAD Service Routine (Macro IEAATC)	33	Initiator/Terminator (Scheduler)	62
XCTL Service Routine (Macro IEAATC)	34	Job Selection (IEFSD510)	63
IDENTIFY Service Routine (IEAAID00)	34	Small Partition Scheduling	64
DELETE Service Routine (IEAADL00,		Initiator/Terminator Control Flow	67
IEABDL00)	34	System Output Writers	72
SYNCH Service Routine (IEAASY00)	34	Resident Writers	72
Main Storage Supervision	35	Non-Resident Writers	72
Timer Supervision	35	System Output Writer Modules	72
Timing Procedure	35	System Task Control	73
Timer Pseudo Clock Routine (IEATPC)	36	Initiating System Tasks	74
Comparison of PCP, MFT, and MVT Timer		System Restart	75
Supervision	36	APPENDIX A: TABLES AND WORK AREAS	76
Overlay Supervision	36	Command Scheduling Control Block	
MFT Recording/Recovery Routines	36	(CSCB)	76
Machine-Check Routines	36	Data Set Enqueue Table (DSENQ)	79
Channel-Check Routine	37	Interpreter Work Area (IWA)	79
Systems Without Recording/Recovery		Job Control Table (JCT)	84
Routines	37	Job File Control Block (JFCB) and	
Entry to Recording/Recovery Routines	37	Extension (JFCBX)	87
Checkpoint/Restart Routines	37	Life-of-Task (LOT) Block	87
JOB MANAGEMENT	39	Linkage Control Table (LCT)	87
Job Scheduler Functions	39	Master Scheduler Resident Data Area	87
Communications Task Functions	39	Partition Information Block	92
Master Scheduler Task Functions	39	Small Partition Information List	
Job Management Control Flow	40	(SPIL)	94
Command Processing	42	Step Control Table (SCT)	94
Communications Task	42	Step Input/Output Table (SIOT)	95
		Task Input/Output Table (TIOT)	97

APPENDIX B: MFT MODULES100
Unique MFT Modules100
Major Component Modules101

Module Descriptions106
APPENDIX C: FLOWCHARTS156
INDEX189

Illustrations

Figures

Figure 1. Main Storage Organization in MFT	9	Figure 23. Table Breakup Parameter List	60
Figure 2. Division of Main Storage	13	Figure 24. Scheduling a Problem Program in a Large Partition	64
Figure 3. MFT Theory of Operation (Part 1 of 4)	14	Figure 25. Scheduling a Problem Program in a Small Partition	65
Figure 4. Main Storage During Execution of NIP	19	Figure 26. Scheduling a Writer in a Small Partition	66
Figure 5. Main Storage at Termination of Master Scheduler Initialization	20	Figure 27. Allocate/Terminate Parameter List	69
Figure 6. MFT Supervisor	22	Figure 28. User's Parameter List	70
Figure 7. TCB Queue	23	Figure 29. Scheduling a Writer in a Large Partition	73
Figure 8. Dispatching Communications and Master Scheduler Tasks	25	Figure 30. START Descriptor Table (SDT)	74
Figure 9. Task Switching	26	Figure 31. Command Scheduling Control Block (CSCB) (Part 1 of 2)	77
Figure 10. System Control Block Relationship	32	Figure 32. Data Set Enqueue Table (DSEQ)	79
Figure 11. Recording/Recovery Routines	38	Figure 33. Interpreter Work Area (IWA) (Part 1 of 3)	81
Figure 12. Job Management Data Flow	41	Figure 34. Job Control Table (JCT)	85
Figure 13. Command Processing Flow	42	Figure 35. Job File Control Block (JFCB) and Extension (JFCBX)	86
Figure 14. WTO/WTOR Macro Instruction Processing Flow	43	Figure 36. Life-of-Task (LCT) Block	88
Figure 15. External Interruption Processing Flow	43	Figure 37. Linkage Control Table (LCT)	89
Figure 16. START Command Processing Flow	48	Figure 38. Master Scheduler Resident Data Area (Part 1 of 2)	91
Figure 17. DEFINE Command Processing Flow	51	Figure 39. Partition Information Block (PIB)	93
Figure 18. Master Queue Control Record (Master QCR) Format	54	Figure 40. Small Partition Information List (SPIL)	94
Figure 19. Job Queue Control Record (QCR)	56	Figure 41. Step Control Table (SCT)	96
Figure 20. Logical Track Header (LTH) Record Format	56	Figure 42. Step Input/Cutput Table (SIOT)	98
Figure 21. Sample Job Queue (SYS1.SYSJCBQE) Format After Initialization	57	Figure 43. Task Input/Output Table (TIOT)	99
Figure 22. Input and Output Queue Entries	58		

Tables

Table 1. Responders to Commands After Initial Processing	40	Table 7. Interpreter Modules	103
Table 2. MFT Modules	100	Table 8. Master Scheduler Modules	103
Table 3. ABEND Modules	102	Table 9. Queue Management Modules	103
Table 4. Communication Task Modules	102	Table 10. SVC 34 Modules	104
Table 5. Initiator Modules	102	Table 11. System Output Writer Modules	104
Table 6. I/O Device Allocation Modules	102	Table 12. System Restart Modules	104
		Table 13. System Task Control Modules	104
		Table 14. Termination Modules	105

Charts

Chart 01. Task Dispatcher (Without Time Slicing)156	Chart 16. Communications Task171
Chart 02. Task Dispatcher (With Time Slicing)157	Chart 17. IEFSD518 - Partition Recovery Routine172
Chart 03. ABEND and DAR Control Flow (Part 1 of 2)158	Chart 18. Initiator Control Flow173
Chart 04. ABEND and DAR Control Flow (Part 2 of 2)159	Chart 19. Job Selection Routine (Sheet 1 of 5)174
Chart 05. Small Partition Routine (Part 1 of 4)160	Chart 20. Job Selection Routine (Sheet 2 of 5)175
Chart 06. Small Partition Routine (Part 2 of 4)161	Chart 21. Job Selection Routine (Sheet 3 of 5)176
Chart 07. Small Partition Routine (Part 3 of 4)162	Chart 22. Job Selection Routine (Sheet 4 of 5)177
Chart 08. Small Partition Routine (Part 4 of 4)163	Chart 23. Job Selection Routine (Sheet 5 of 5)178
Chart 09. Master Scheduler Task164	Chart 24. Reader/Interpreter (Sheet 1 of 3)179
Chart 10. Queue Search165	Chart 25. Reader/Interpreter (Sheet 2 of 3)180
Chart 11. Queue Manager Table Breakup Routine166	Chart 26. Reader Interpreter (Sheet 3 of 3)181
Chart 12. Master Scheduler Resident Command Processor167	Chart 27. JCL Statement Processor182
Chart 13. SVC 34 Command Processing (Part 1 of 3)168	Chart 28. Job and Step Enqueue Routine	183
Chart 14. SVC 34 Command Processing (Part 2 of 3)169	Chart 29. Transient Reader Suspend Routine184
Chart 15. SVC 34 Command Processing (Part 3 of 3)170	Chart 30. Transient Reader Restore Routine185
		Chart 31. System Output Writer Control Flow186
		Chart 32. System Output Writer187
		Chart 33. System Task Control188

Introduction

In a single task environment, main storage is divided into two areas: the fixed area, and the dynamic area. In multiprogramming with a fixed number of tasks (MFT), the dynamic area is divided further into as many as fifty-two discrete areas called partitions. Figure 1 shows the division of main storage.

The fixed area, located in the lower portion of main storage, contains the resident portion of the control program, and control blocks and tables used by the system. The size of the fixed area depends on the number of partitions established by the user, and the control program options selected at system generation.

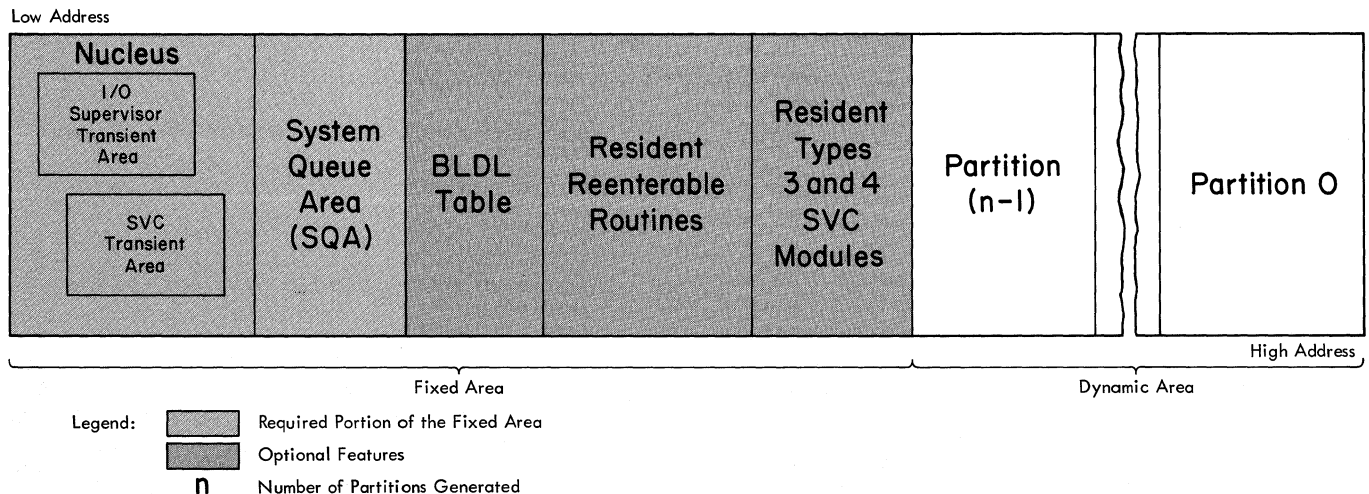
Partitions are defined within the dynamic area, located in the upper portion of main storage, at system generation. The number of partitions may be varied within the number specified at system generation, and the sizes and job classes of partitions may be redefined at system initialization or during operation. (See IEM System/360 Operating System: Planning for Multiprogramming with a Fixed Number of Tasks (MFT), Form C27-6939.) Each partition may be occupied by a processing program, or by control program routines that prepare job steps for execution (job management routines), or handle data for a processing program (access method routines).

Provided the total number of partitions does not exceed 52 and enough computing

system resources are available, MFT provides for the concurrent execution of as many as 15 problem programs, 3 input readers, and 36 output writers, each in its own fixed partition of main storage. The MFT system provides for task switching among the tasks operating in the partitions, and between those tasks and the communications task and master scheduler task in the system area.

Task dispatching in MFT differs from the primary control program (PCP) primarily in that task switching is required, and that certain system functions such as abnormal termination must be carried out so that other, unrelated, tasks are not affected. The dispatching priority of a task is determined by the relative position of the partition used to process the task. The highest-priority partition (P0) is at the highest address in storage. Successively lower partitions (P1 - P51) have correspondingly lower priorities. Control of the CPU is given to the program in the highest-priority partition that is ready.

The integrity of programs operating under MFT is preserved if the storage protection feature is included. MFT uses the 16 protection keys to prevent a user job from modifying the control program or another job; it uses the two operating states of the CPU to restrict the use of control and I/O instructions.



•Figure 1. Main Storage Organization in MFT

Because many components of MFT are similar to those of PCP and multiprogramming with a variable number of tasks (MVT), many of the modules for a given MFT component are the same for the comparable component in either PCP or MVT. Therefore, this publication describes differences between MFT and the other configurations. The corresponding PCP and MVT routines are described in the following IBM System/360 Operating System program logic manuals and are referenced where applicable:

Fixed Task Supervisor, Form Y28-6612

MVT Supervisor, Form Y28-6659

MVT Job Management, Form Y28-6660

Information on modified or new routines for MFT is contained in the three sections that follow this introduction.

The Initialization of the Operating System section describes how the dynamic area of main storage is prepared by the master scheduler task after completion of the Nucleus Initialization Program.

The Supervisor section describes the task management modifications made to the supervisor for MFT. The major area of change has been in the initialization of main storage.

The Job Management section describes modifications and additions to the routines for processing communications with the programmer and the operator. The major changes are in the master scheduler task, and the MFT initiator. Other modifications have been made to the queue manager, the reader/interpreter, system output writer, and system task control routines.

Functions of the Control Program with MFT

As in PCP and MVT, the control program routines of MFT have three major functions: job management, task management, and data management.

JOB MANAGEMENT

Job management is the processing of communications from the programmer and operator to the control program. There are two types of communications: operator commands, which start, stop, and modify the processing of jobs in the system, and job control statements, which define work being entered into the system. Processing of these commands and statements is referred to as command processing and job processing, respectively.

TASK MANAGEMENT

Task management routines monitor and control the entire operating system, and are used throughout the operation of both the control and processing programs. Task management has six major functions:

- Interruption supervision.
- Task supervision.
- Main Storage supervision.
- Contents supervision.
- Overlay supervision.
- Timer supervision.

The task management routines are collectively referred to as the "supervisor."

DATA MANAGEMENT

Data management routines control all operations associated with input/output devices: allocating space on volumes, channel scheduling, storing, naming, and cataloging data sets, moving data between main and auxiliary storage, and handling errors that occur during input/output operations. Data management routines are used by processing programs and control program routines that require data movement. Processing programs use data management routines primarily to read and write required data, and also to locate input data sets and to reserve auxiliary storage space for output data sets of the processing program.

Data management routines are of five categories:

- Input/Output (I/O) supervisor, which supervises input/output requests and interruptions.
- Access methods, which communicate with the I/O supervisor.
- Catalog management, which maintains the catalog and locates data sets on auxiliary storage.
- Direct-access device space management (DADSM), which allocates auxiliary storage space.
- Open/Close/End-of-Volume, which performs required initialization for I/O operations and handles end-of-volume conditions.

The operation of these routines is identical with MVT and is described in the following IBM System/360 Operating System program logic manuals:

Input/Output Supervisor, Form Y28-6616

Sequential Access Methods, Form Y28-6604

Indexed Sequential Access Methods, Form Y28-6618

Basic Direct Access Method, Form Y28-6617

Graphics Access Method, Form Y27-7113

Catalog Management, Form Y28-6606

Direct Access Device Space Management, Form Y28-6607

Input/Output Support (OPEN/CLOSE/EOV), Form Y28-6609

These routines are described in the Supervisor section of this publication, and in the program logic manual IBM System/360 Operating System: Fixed Task Supervisor, Form Y28-6612.

The resident job management routines are those routines of the communications task that receive commands from the operator. The MFT communications task is described in this publication.

The resident data management routines are the input/output supervisor and, optionally, the BLDL routines of the partitioned access method. These routines are described in the following IBM System/360 Operating System program logic manuals:

Input/Output Supervisor, Form Y28-6616

Sequential Access Method, Form Y28-6604

The user may also select resident reenterable routines, which are access method routines from SYS1.SVCLIB, and other reenterable routines from SYS1.LINKLIB. At system generation, the user specifies that he wants such routines resident in main storage. At IPL, he identifies the specific routines desired in the RAM=entry. The selected routines are loaded during system initialization and reside adjacent to the higher end of the system queue area unless the BLDL table is also resident (see Figure 1).

Normally-transient SVC routines (i.e., types 3 and 4 SVC routines) can be made resident through the RSVC option, specified by the user. NIP loads these routines adjacent to the higher end of the resident reenterable routines. If there is no resident BLDL table or resident reenterable routines, the routines are loaded adjacent to the higher end of the system queue area. (See Figure 1.)

Control Program Organization

The control program is on auxiliary storage in three partitioned data sets created when the system is generated. These data sets are:

- The NUCLEUS partitioned data set (SYS1.NUCLEUS), which contains the Nucleus Initialization Program (NIP) and the resident portion of the control program.
- The SVCLIB partitioned data set (SYS1.SVCLIB), which contains nonresident SVC routines, nonresident error-handling routines, and the access methods routines.
- The LINKLIB partitioned data set (SYS1.LINKLIB), which contains other nonresident control program routines and IBM-supplied processing programs.

RESIDENT PORTION OF THE CONTROL PROGRAM

The resident portion (nucleus) of the control program is in SYS1.NUCLEUS. It is made up of those routines, control blocks, and tables that are brought into main storage at initial program loading (IPL) and are never overlaid by another part of the operating system. The nucleus is loaded into the fixed area of main storage.

The resident task management routines include all of the routines that perform:

- Interruption supervision.
- Main storage supervision.
- Timer supervision.

They also include portions of the routines that perform:

- Task supervision.
- Contents supervision.
- Overlay supervision.

NONRESIDENT PORTION OF THE CONTROL PROGRAM

The nonresident portion of the control program comprises routines that are loaded into main storage as they are needed, and which can be overlaid after their completion. The nonresident routines operate from the partitions and from two sections of the nucleus called transient areas (described below).

Main Storage Organization

Main storage in MFT is organized similarly to main storage in MVT, except that the optional resident areas are adjacent to the nucleus.

Main storage may be expanded by including IBM 2361 Core Storage (core storage) units in the system. Main Storage Hierarchy Support for IBM 2361 Models 1 and 2 permits access to either processor storage (hierarchy 0) or core storage (hierarchy 1). Each partition established during system generation is described by a boundary box. The first half of the boundary box describes the processor storage partition segment and the second half describes the core storage partition segment. Any partition segment not assigned main storage in the system has the applicable boundary box pointers set to zero. If a partition is established entirely within hierarchy 1, the processor storage pointers in the first half of the partition's boundary box are set to zero. If a partition segment is not generated in core storage, the core storage pointers in the second half of the partition's boundary box are set to zero. If core storage has been included in the system, but is offline, the second half of the boundary box will contain zeros. If core storage is excluded from the system, the second half of the boundary box is not generated.

FIXED AREA

In MFT (as in PCP and MVT) the fixed area is that part of main storage into which the nucleus is loaded at IPL. The storage protection key of the fixed area is zero so that its contents can be modified by the control program only. The fixed area also contains two transient areas into which certain nonresident routines are loaded when needed: the SVC transient area (1024 bytes) and the I/O supervisor transient area (1024 bytes). These areas are used by nonresident SVC routines and nonresident I/O error-handling routines, respectively, which are read from SYS1.SVCLIB.

Each transient area contains only one routine at a time. When a nonresident SVC or error-handling routine is required, it is read into the appropriate transient area. The transient area routines operate with a protection key of zero, as do other routines in the fixed area.

System Queue Area

The system queue area (SQA) is established by NIP adjacent to the fixed area and provides the main storage space required for tables and queues built by the control program. The SQA must be at least 1600 bytes for a minimum two-partition system. Its storage protection key is zero so that it can be modified by control program routines only. Data in the system queue area indicates the status of all tasks.

DYNAMIC AREA

Figure 2 shows how the contents of each partition in the dynamic area are organized and how they are related to the rest of main storage. Routines are brought into the high or low portion of an MFT partition similarly to the way routines are brought into the entire dynamic area of PCP. Job management routines, processing programs, and routines brought into storage via a LINK, ATTACH, or XCTL macro instruction, are loaded at the lowest available address. The highest portion of the partition is occupied by the user parameter area and user save area. The next portion of the partition is occupied by the task input/output table (TIOT) which is built by a job management routine (I/O Device Allocation routine). This table is used by data management routines and contains information about DD statements.

Each partition may be used for a problem program as well as for system tasks (readers, initiators, and writers). When the control program requires main storage to build control blocks or work areas, it obtains this space from the partition of the processing program that requested the space. Access method routines and routines brought into storage via a LCAD macro instruction are placed in the highest available locations below the task input/output table.

Working storage and data areas are assigned from the highest available storage in a partition.

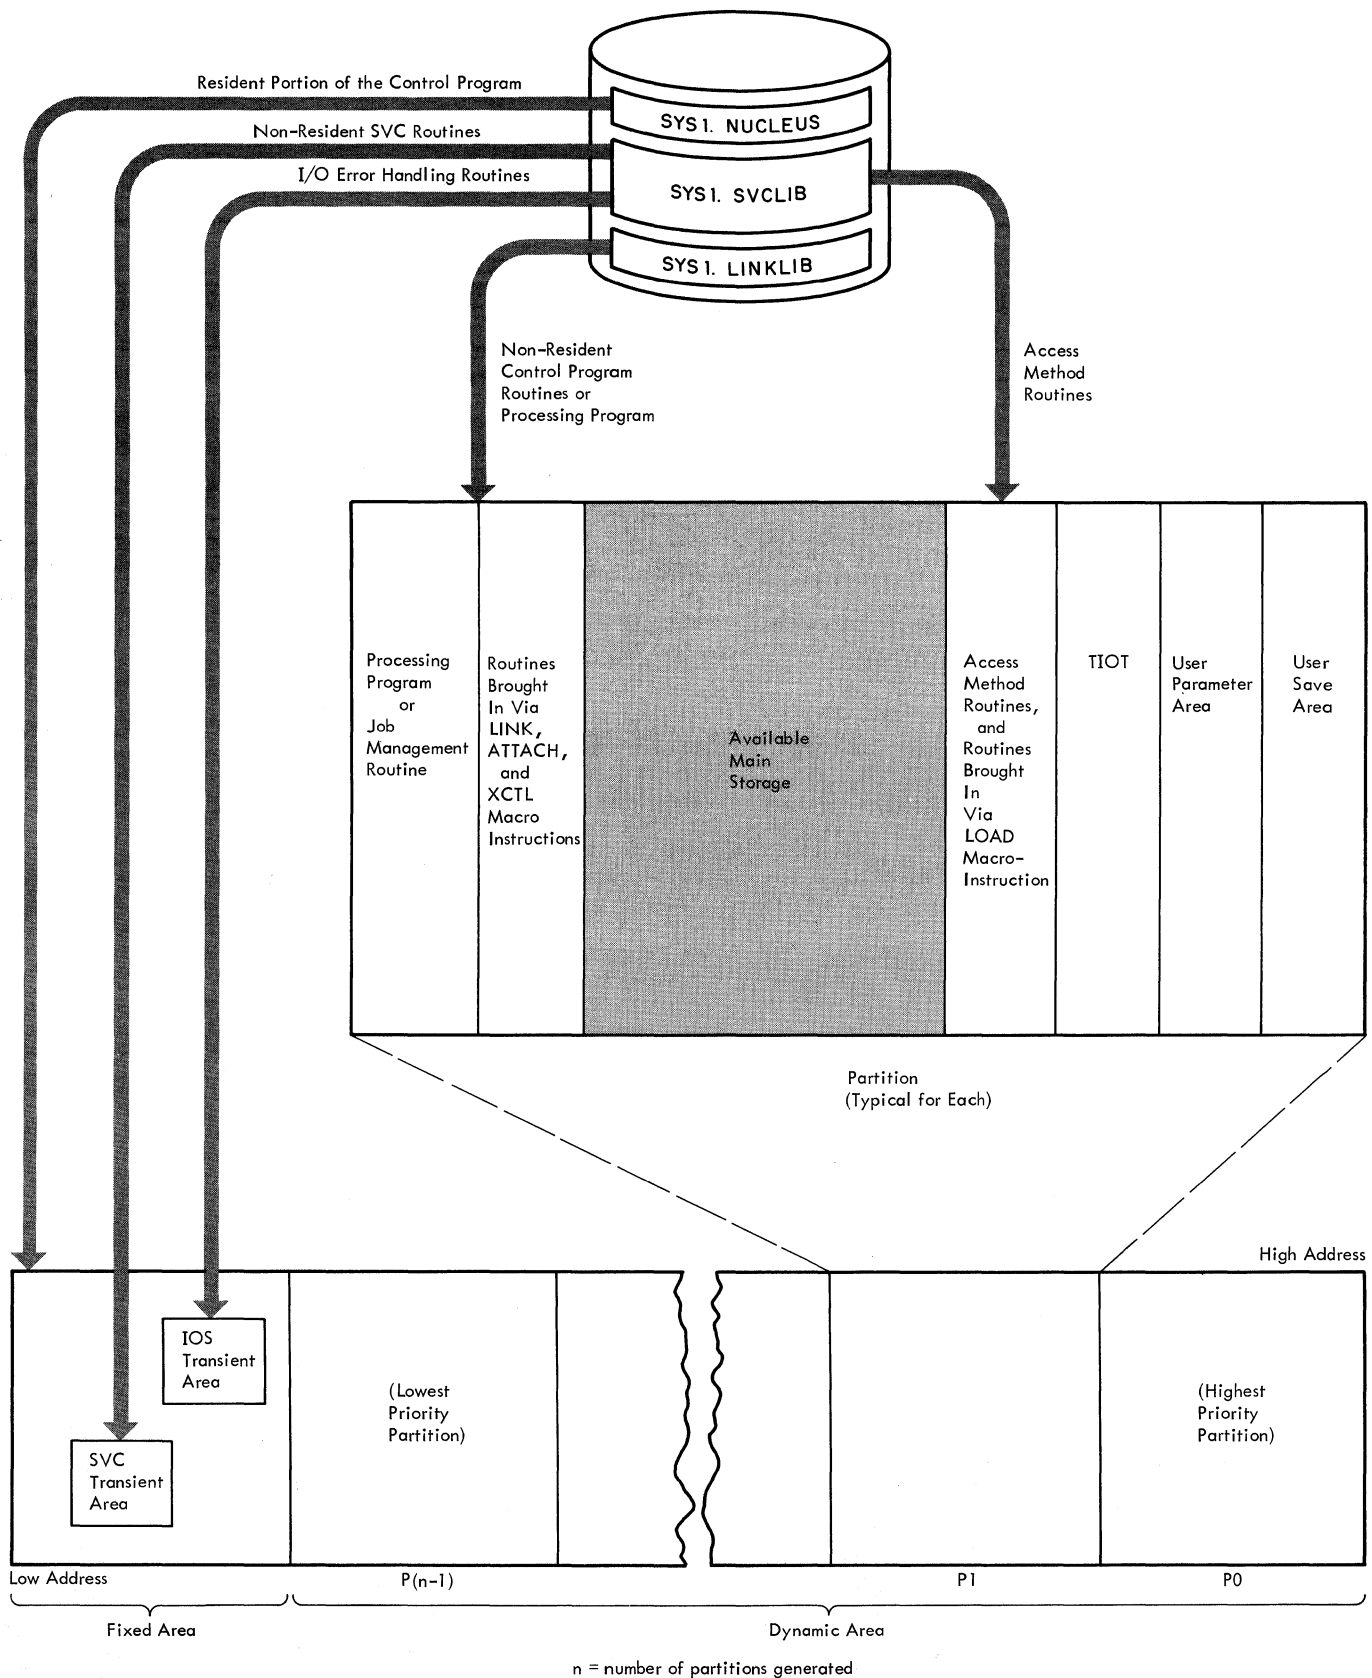
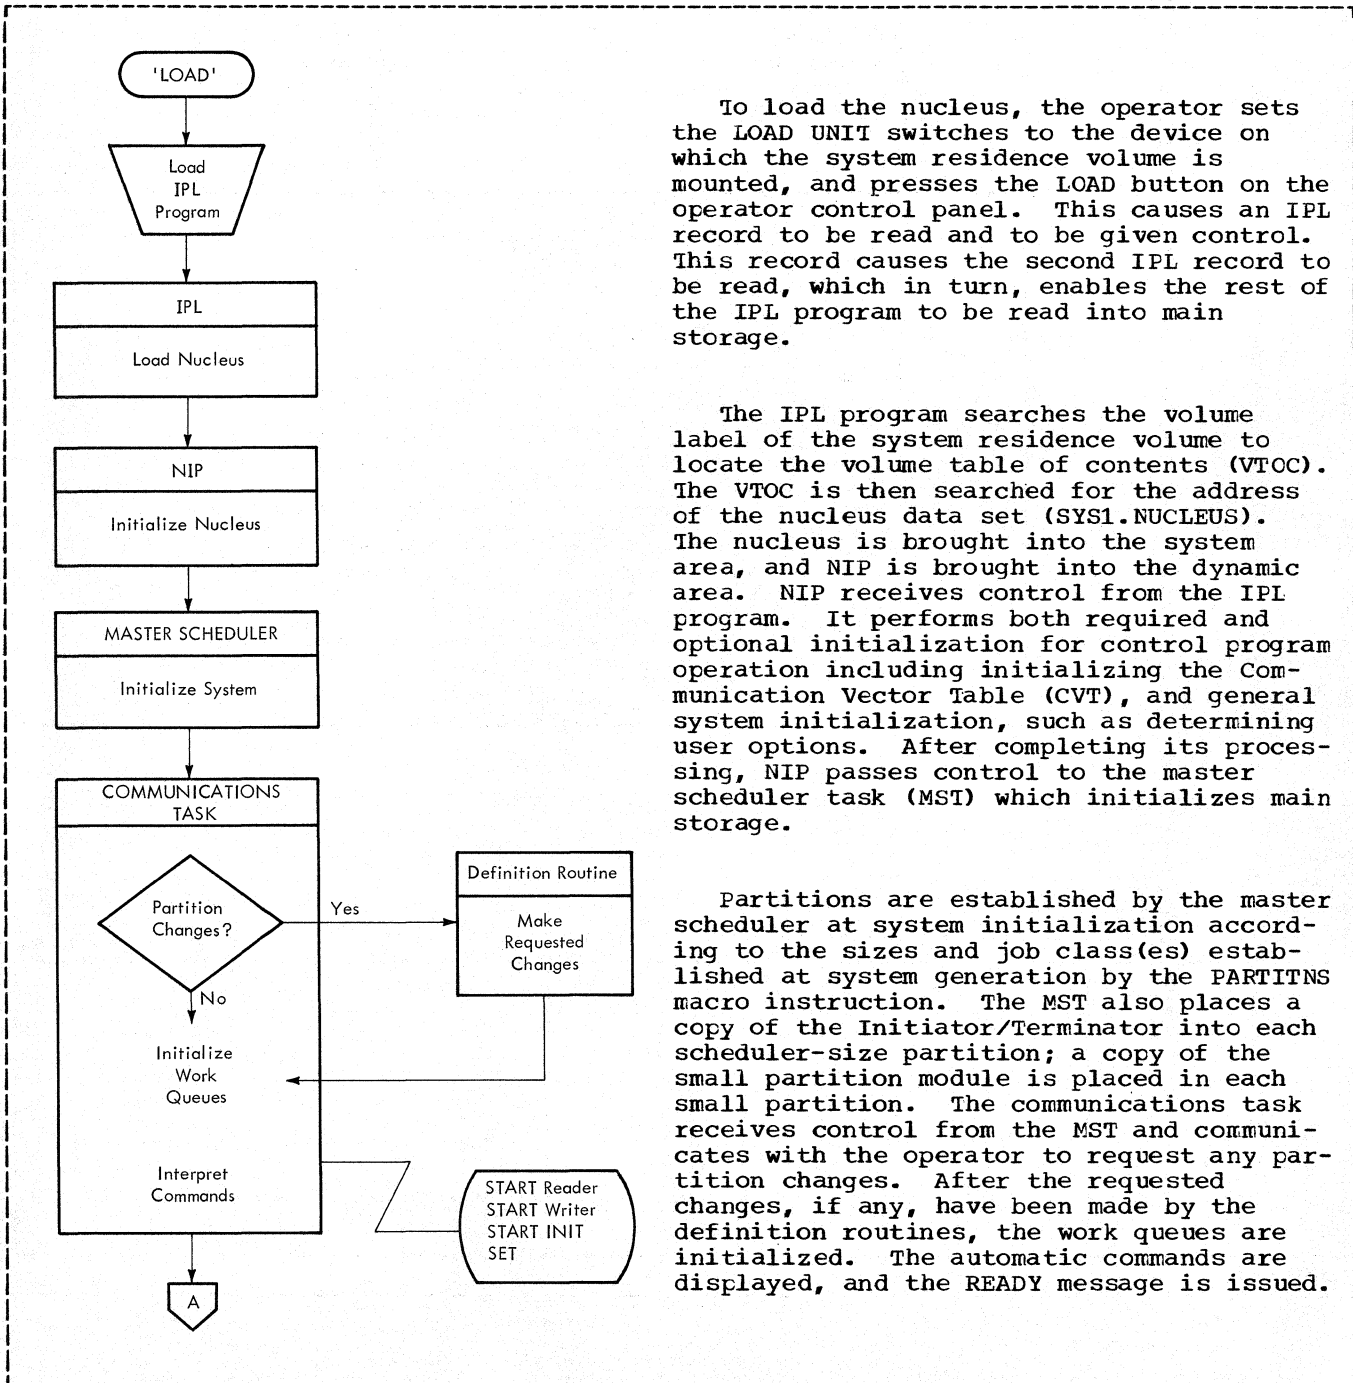


Figure 2. Division of Main Storage

Theory of Operation

Figure 3 describes the overall processing flow through each job cycle. These paragraphs describe the processing performed by various components of the control program as it loads the nucleus, reads control statements, initiates the job step, causes processing to begin or end in other partitions, and terminates the job step.

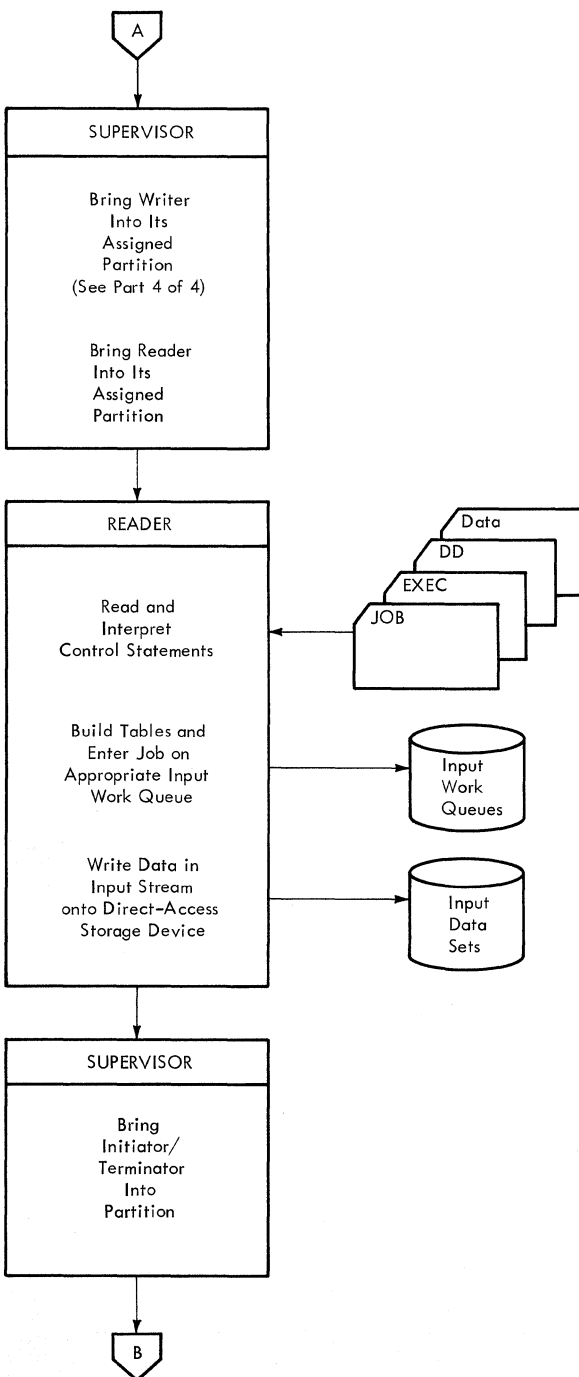


To load the nucleus, the operator sets the LOAD UNIT switches to the device on which the system residence volume is mounted, and presses the LOAD button on the operator control panel. This causes an IPL record to be read and to be given control. This record causes the second IPL record to be read, which in turn, enables the rest of the IPL program to be read into main storage.

The IPL program searches the volume label of the system residence volume to locate the volume table of contents (VTOC). The VTOC is then searched for the address of the nucleus data set (SYS1.NUCLEUS). The nucleus is brought into the system area, and NIP is brought into the dynamic area. NIP receives control from the IPL program. It performs both required and optional initialization for control program operation including initializing the Communication Vector Table (CVT), and general system initialization, such as determining user options. After completing its processing, NIP passes control to the master scheduler task (MST) which initializes main storage.

Partitions are established by the master scheduler at system initialization according to the sizes and job class(es) established at system generation by the PARTITNS macro instruction. The MST also places a copy of the Initiator/Terminator into each scheduler-size partition; a copy of the small partition module is placed in each small partition. The communications task receives control from the MST and communicates with the operator to request any partition changes. After the requested changes, if any, have been made by the definition routines, the work queues are initialized. The automatic commands are displayed, and the READY message is issued.

Figure 3. MFT Theory of Operation (Part 1 of 4)



When the required SET command is entered, the communications task calls the master scheduler command scheduling routine to have the command executed. An automatic START reader command or a subsequent operator entered START reader command causes a copy of the Reader/Interpreter (reader) to be brought into its appropriate partition. If a START writer command is entered, a copy of a writer is also brought into the specified partition(s).

When the reader gets control, it reads control statements and data from the input job stream. Information from the JOB, EXEC, and DD statements controls the execution of each job step. This information is placed in the following tables:

- Job control table (JCT) for the job being read.
- Step control table (SCT) for the step being read.
- Data set enqueue table (DSEQ) for the job being read.
- Job file control block (JFCB) and step input/output table (SIOT) for each data set being used or created by the job step.
- Volume table (VOLT) containing each volume serial number to be used by the job.

Information from these tables and control blocks is updated with information in the data control block (DCB) and data set control block (DSCB) or volume label when a data set is opened during step execution.

The reader then places these updated control blocks into the input work queue corresponding to the CLASS parameter on the JOB statement. Data sets in the input stream are written onto a direct-access storage device for later use by the problem program.

After the reader has completed processing all input for a job and has entered the job on an input work queue, all initiators that are waiting for that job class are posted. If the job is for a small partition, the small partition module is also posted.

Figure 3. MFT Theory of Operation (Part 2 of 4)

After receiving control, the initiator/terminator prepares to initiate the highest priority job in its primary input work queue. Using information which the reader extracted from the DD statement, the initiator/terminator processes the user accounting routine, in addition to the following:

Locates Input Data Sets: The Allocation routine, running as a subroutine of the initiator/terminator, determines the volume containing a given input data set by examining the JFCB, or by searching the catalog. This search is performed by a catalog management routine entered from allocation. (A description of the routines that maintain and search the catalog is given in IBM System/360 Operating System: Catalog Management, Program Logic Manual, Form Y28-6606.)

Allocates I/O Devices: A job step cannot be initiated unless there are enough I/O devices to fill its needs. Allocation determines whether the required devices are available, and makes specific assignments. If necessary, messages are issued to the operator to request the mounting of volumes.

Allocates Auxiliary Storage Space: Direct access volume space required for output data sets of a job step is acquired by the allocation routine, which uses the Direct Access Space Management (DADSM) routines. (A description of the operation of the DADSM routines is given in the publication IBM System/360 Operating System: Direct Access Device Space Management, Program Logic Manual, Form Y28-6607.)

The JFCB, which contains information concerning the data sets to be used during step execution, is written on auxiliary storage. This information is used when a data step is opened, and when it is closed, the job step is terminated.

The initiator causes itself to be replaced by the problem program it is initiating (if for a large partition), or initiates the job in a small partition.

The problem program can be an IBM-supplied processor (e.g., COBOL, linkage editor), or a user-written program. The problem program uses control program services for operations such as loading other programs and performing I/O operations.

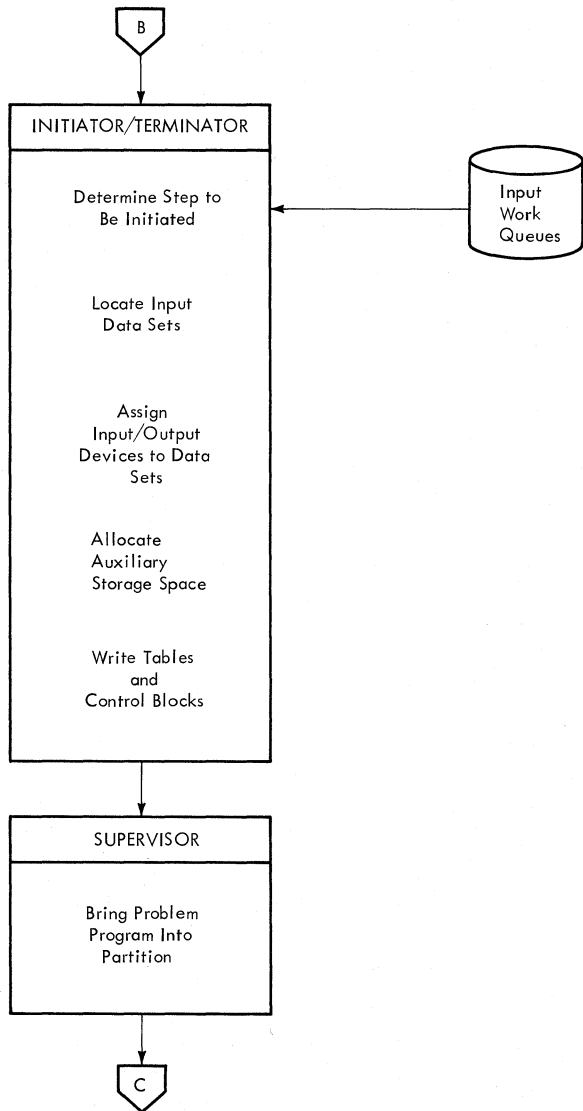
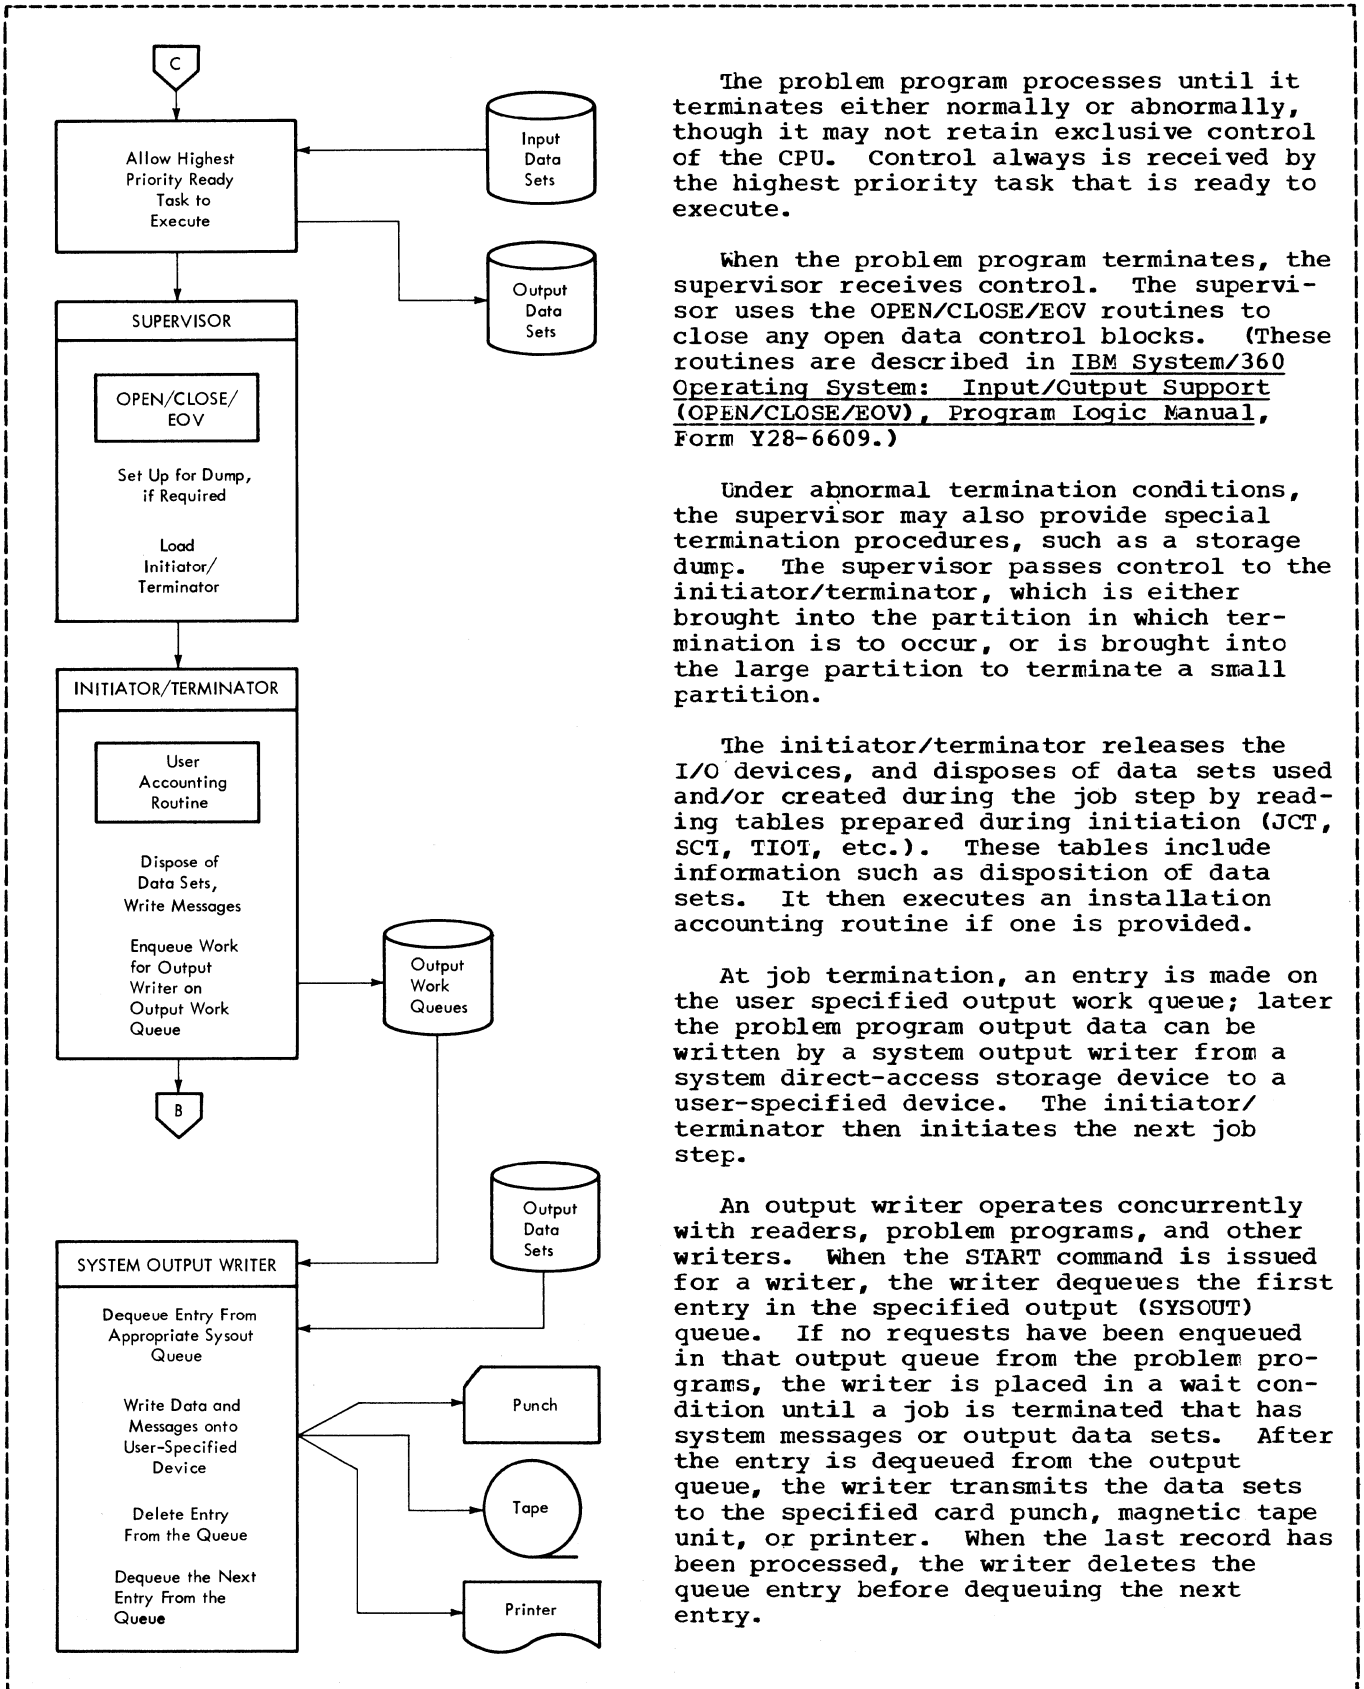


Figure 3. MFT Theory of Operation (Part 3 of 4)



The problem program processes until it terminates either normally or abnormally, though it may not retain exclusive control of the CPU. Control always is received by the highest priority task that is ready to execute.

When the problem program terminates, the supervisor receives control. The supervisor uses the OPEN/CLOSE/EOV routines to close any open data control blocks. (These routines are described in IBM System/360 Operating System: Input/Output Support (OPEN/CLOSE/EOV), Program Logic Manual, Form Y28-6609.)

Under abnormal termination conditions, the supervisor may also provide special termination procedures, such as a storage dump. The supervisor passes control to the initiator/terminator, which is either brought into the partition in which termination is to occur, or is brought into the large partition to terminate a small partition.

The initiator/terminator releases the I/O devices, and disposes of data sets used and/or created during the job step by reading tables prepared during initiation (JCT, SCT, TIOT, etc.). These tables include information such as disposition of data sets. It then executes an installation accounting routine if one is provided.

At job termination, an entry is made on the user specified output work queue; later the problem program output data can be written by a system output writer from a system direct-access storage device to a user-specified device. The initiator/terminator then initiates the next job step.

An output writer operates concurrently with readers, problem programs, and other writers. When the START command is issued for a writer, the writer dequeues the first entry in the specified output (SYSOUT) queue. If no requests have been enqueued in that output queue from the problem programs, the writer is placed in a wait condition until a job is terminated that has system messages or output data sets. After the entry is dequeued from the output queue, the writer transmits the data sets to the specified card punch, magnetic tape unit, or printer. When the last record has been processed, the writer deletes the queue entry before dequeuing the next entry.

Figure 3. MFT Theory of Operation (Part 4 of 4)

Initialization of the Operating System

When the system is loaded, routines perform required and optional initialization of functions needed for control program operation. (These routines are described in IBM System/360 Operating System: Initial Program Loader and Nucleus Initialization Program, Program Logic Manual, Form Y28-6661.)

When the Nucleus Initialization Program (NIP) has defined the fixed area, it then assigns the rest of main storage to the master scheduler task to be prepared as the dynamic area for control program operation.

Main Storage Preparation

When NIP completes its functions it constructs a request block (RB) and an XCTL macro instruction (specifying master scheduler initialization routine IEFSD569) at the low address of the temporary master scheduler area. NIP places the address of this RB in master scheduler task TCB field TCBRBP. (The original contents of TCBRBP are saved and passed to IEFSD569 in a parameter list along with the original master scheduler task boundary box contents.) NIP sets master scheduler task TCB field TCBFLGS to make the master scheduler task dispatchable, and then branches to the dispatcher.

The dispatcher gives control to the master scheduler task causing execution of the XCTL instruction which NIP placed in the temporary master scheduler area. The master scheduler initialization routine is brought into the temporary master scheduler area and begins executing. Figure 4, excluding the medium shaded area, illustrates main storage at completion of NIP before branching to the dispatcher. Figure 4, excluding the light shaded area, illustrates main storage when the master scheduler initialization routine receives control from the dispatcher.

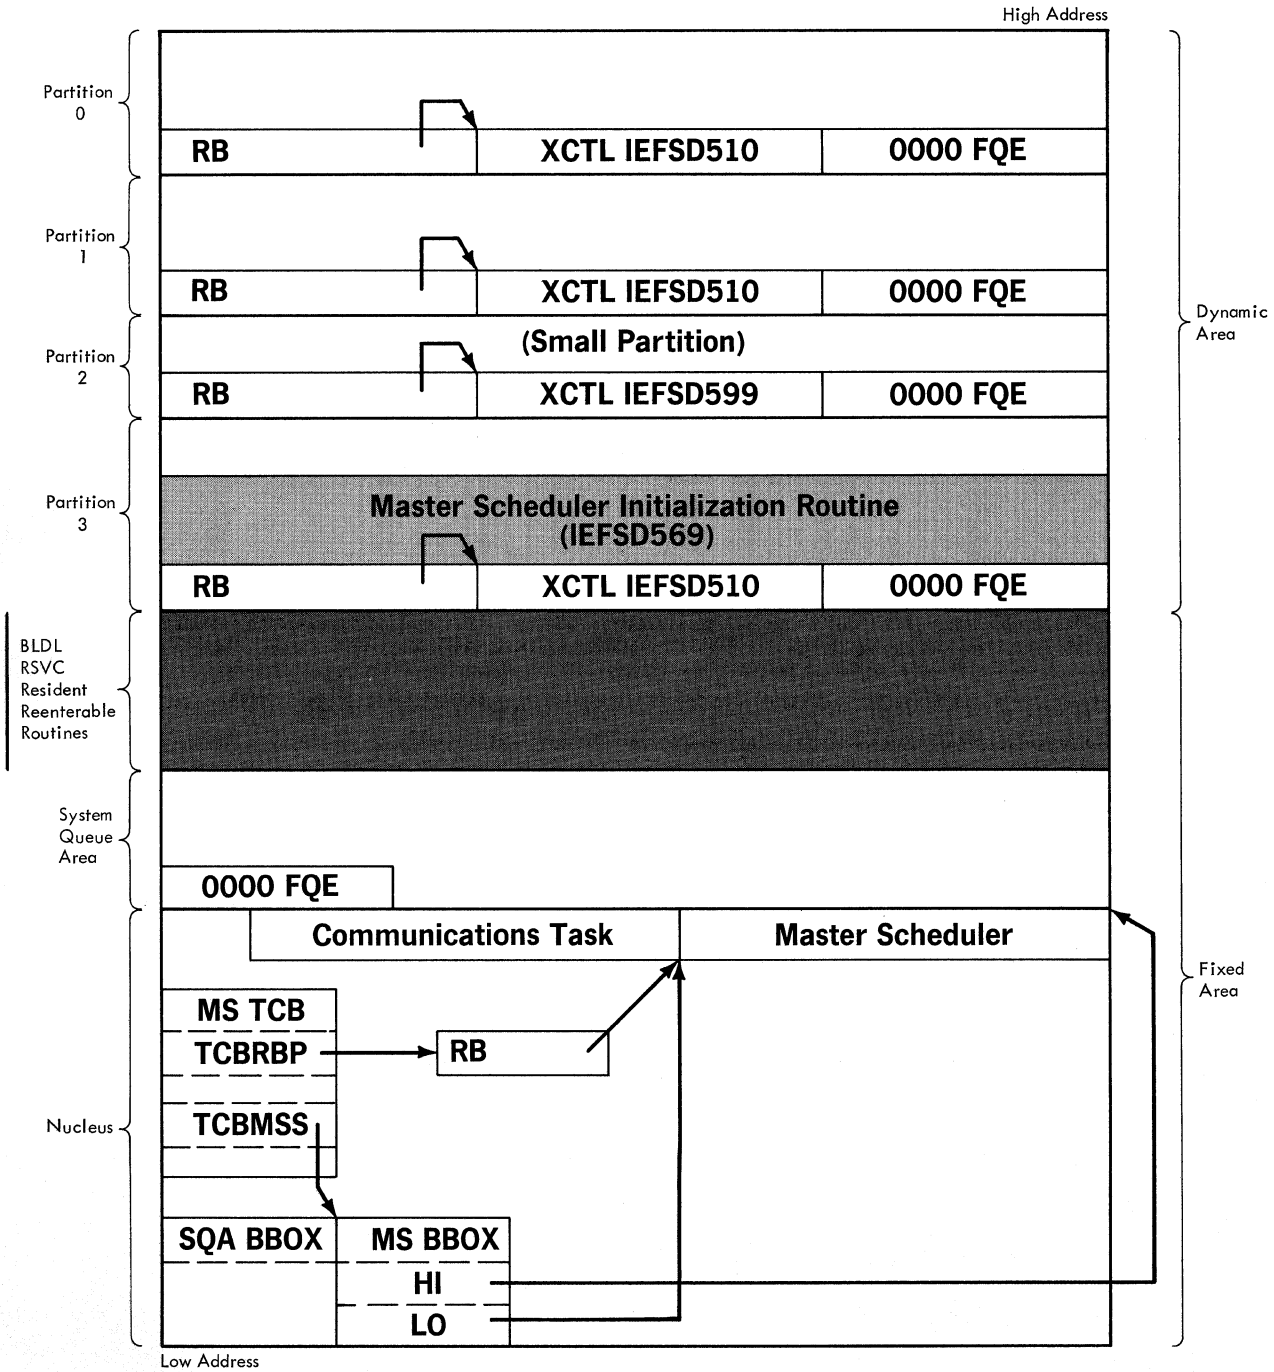
For a description of the master scheduler initialization routine see "Master

Scheduler Task" in the Job Management section. Figure 5 illustrates main storage (four partition example) at completion of master scheduler initialization. When the initialization routine completes processing, it branches to the dispatcher.

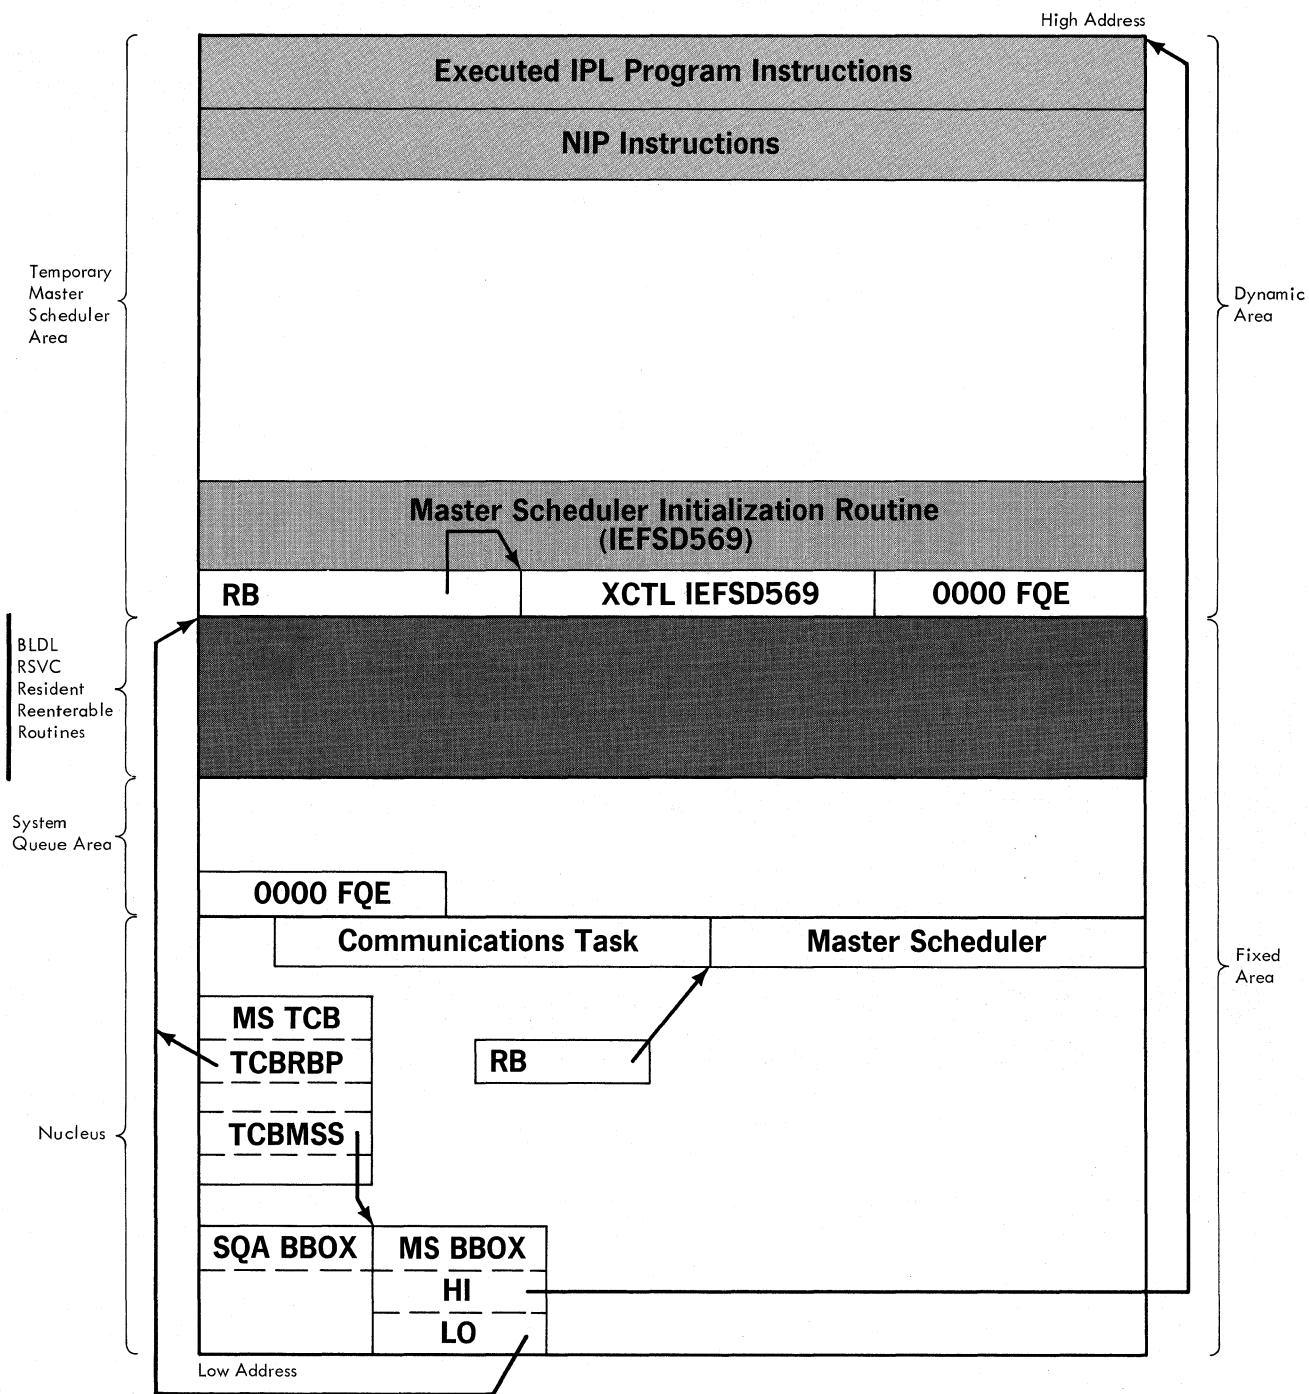
Initializing the Partitions

During master scheduler initialization the operator must accept automatic START commands or enter START commands manually. When a START command is processed, the partition number specified in the command is determined, and a CSCB is built. The CSCB (see Appendix A) is used for communication between the command scheduling routines (SVC 34) and the command execution routines. The address of the CSCB is placed in the partition information block (PIB) of the specified partition, and the partition is posted. The PIB for each partition contains information used by command processing and scheduler routines. (See Appendix A for a description of the PIB, and "Initiator/Terminator" in Job Management for a discussion of its use.)

After the initialization routine completes processing, the dispatcher gives control to the master scheduler router routine. When this routine completes processing, it returns to the dispatcher which begins searching the TCB queue. The highest priority task posted through START command processing receives control. The XCTL macro instruction addressed by the partition's RB is executed and the Job Select module (IEFSD510) or Small Partition module (IEFSD599) is brought into the partition. When an interruption occurs and the partition can no longer retain control, the dispatcher gives control to the next posted partition. This process continues, enabling all posted partitions to receive control and to execute the XCTL instruction placed in them by the initialization routine.



•Figure 4. Main Storage During Execution of NIP



- Legend:
- Contents of the Dynamic Area During IPL and NIP.
 - Contents of the Dynamic Area After The Master Scheduler Task Receives Control on Completion of NIP.
 - Optional Features

•Figure 5. Main Storage at Termination of Master Scheduler Initialization

Supervisor

The MFT Supervisor manages the operation of the control program and processing programs. Job management selects jobs for execution, allocates devices and storage to the step to be executed, and gives control to the program that represents the step. After receiving control, a program is known as a task and becomes the responsibility of the Supervisor. As many as 15 job-step tasks may operate in the system concurrently with system tasks. Each task must be isolated so it does not interfere with any other task. To do this, each job-step task operates in its own partition in main storage. If the system has the optional storage protection feature, each partition is assigned a unique protection key (1-15). The resident portion of the control program, including some supervisor routines, occupies a fixed area of main storage and operates under a protection key of zero.

To maintain control of the computing system, the supervisor must perform many services. Routines within the supervisor are grouped into general categories depending upon the services which they perform. These categories are:

Interruption Supervision: All supervisor activity begins with an interruption. The five types of interruptions are: supervisor call, timer/external, input/output, program, and machine. When an interruption occurs, the interruption handling routine for the type of interruption that occurred gains control. The interruption handling routine then passes control to those parts of the control program that perform the services required as a result of the interruption. Many of the services which must be performed are included in other general categories of the supervisor.

Task Supervision: The supervisor maintains control information including the current status of program and interruption request blocks, task control blocks, and event control blocks.

Contents Supervision: The supervisor keeps records of the status and characteristics of all programs in each partition of main storage, initiates program fetch for the dynamic loading of programs, and maintains the active request block queue.

Main Storage Supervision: Within each partition, the supervisor allocates and releases main storage space for a task on request, and maintains a record of all free storage space within each partition.

Timer Supervision: The supervisor sets and maintains a clock, and honors requests for time intervals and exact time.

Overlay Supervision: The supervisor monitors the flow of control between segments of a program operating in an overlay structure established by the user through the linkage editor.

Interruption Supervision

With the exception of the dispatcher and the ABEND routines which are described below, the interruption supervisor of MFT functions as described in IBM System/360 Operating System: Fixed-Task Supervisor, Program Logic Manual, Form Y28-6612.

When an interruption occurs and is serviced, the task which had been executing may relinquish control of the CPU. Control must always be given to the highest priority ready task. The transfer of control from one task to another is called task switching and is accomplished by the task dispatcher. When an interruption handling routine completes processing an interruption, it branches to the task dispatcher rather than returning control to the interrupted program. Type 1 EXIT is the only interruption handling routine which may return control directly to the interrupted program. Figure 6 illustrates how the task dispatcher receives control after an interruption has been serviced.

THE DISPATCHER (MACRO IEAAPS)

The dispatcher gives control to the highest priority task ready to execute. It uses information located by communication vector table (CVT) fields CVTHEAD and CVTTCBP, and if the time-slicing feature is in the system, field CVTTSCE.

Field CVTHEAD addresses a queue of task control blocks (TCBs). This TCB queue is arranged in dispatching priority order beginning with the highest priority task. The highest priority TCB is the optional log task TCB, followed by the communication task TCB, the master scheduler task TCB, and one TCB for each of the partitions generated in the system (in ascending order by partition number). Figure 7 illustrates the TCB queue.

INTERRUPTIONS

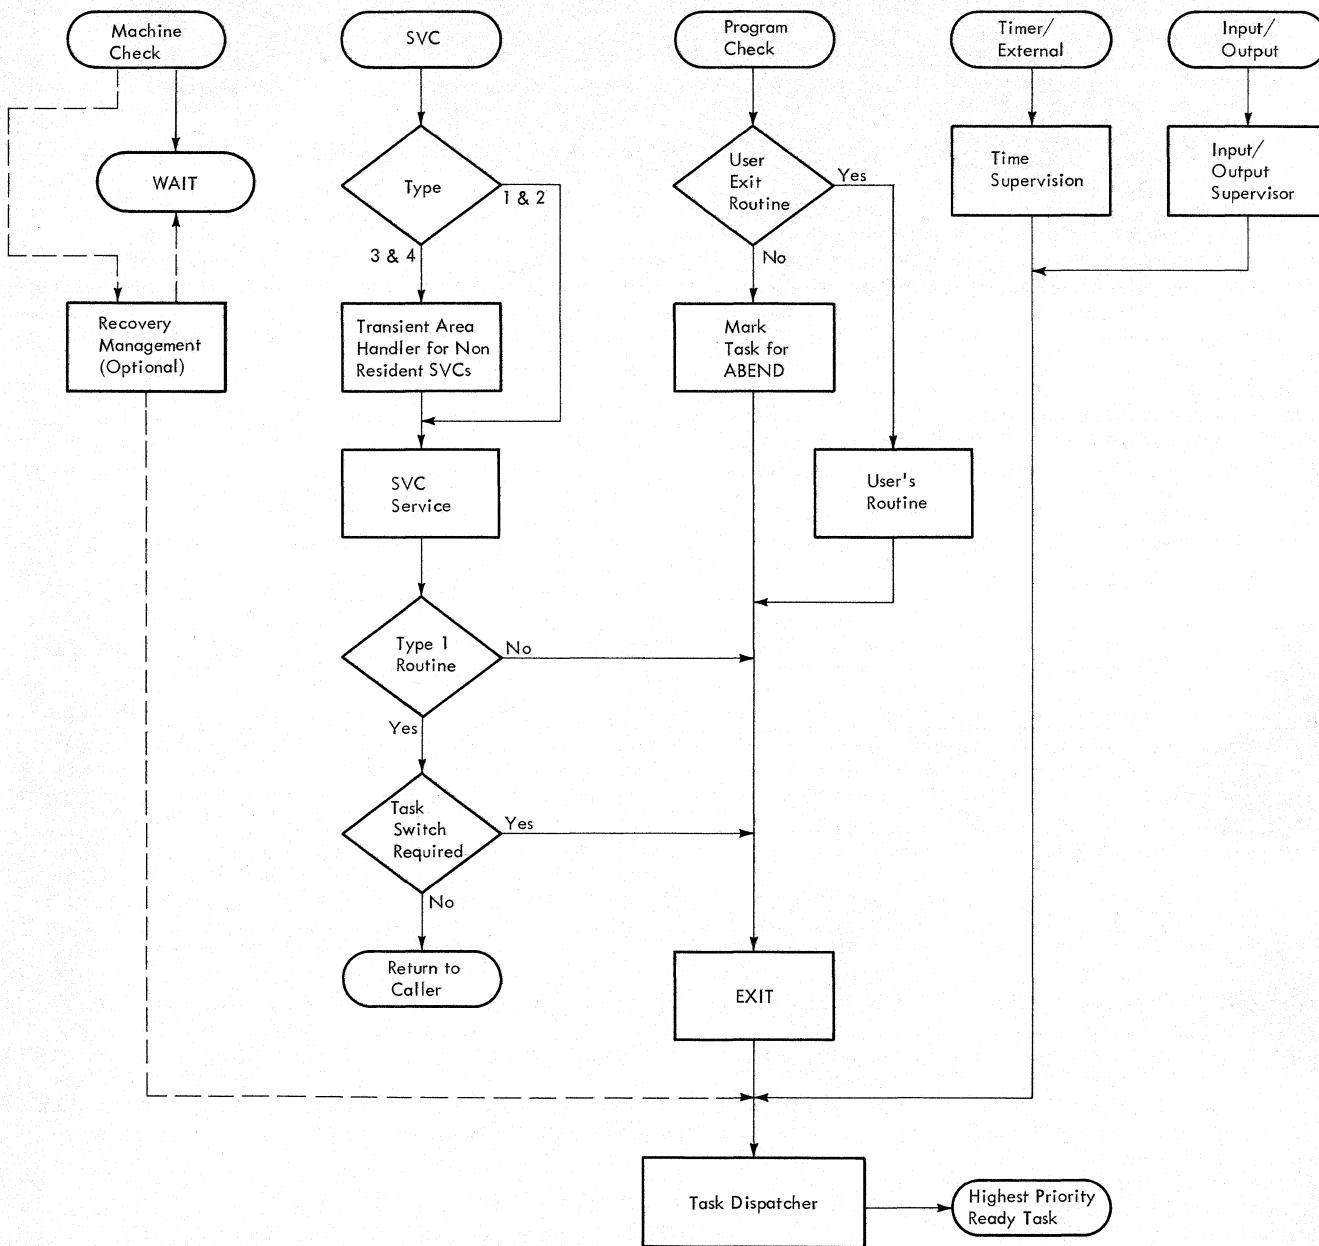


Figure 6. MFT Supervisor

Any number of partitions (up to 52) may be specified during system generation. Partitions must be numbered consecutively beginning with zero. Note that in Figure 7 there is a TCB for partition 1, but partition 1 is assigned no storage space. This illustrates a partition which was specified at system generation but which has been made inactive. If a partition is not specified during system generation, no TCB is constructed. If, for example, only 3 partitions (0 through 2) are specified at system generation, then only three TCBs are constructed and partitions 3 through 51 do not exist.

All of the TCBs in the system are chained together through TCB field TCBCB. In each TCB, this field contains the address of the next TCB on the queue. The TCBCB field of the last TCB on the queue contains zero.

CVT field CVTTCBP addresses two full words called NEW and OLD. The first word (NEW) contains either zero or the TCB address for the task to be given control. The second word (OLD) contains the TCB address for the task currently in control. NEW can be set by any of the supervisory

routines associated with task switching (WAIT, POST, ENQ/DEQ, Manual Purge). When a supervisory routine determines that the task currently in control can no longer retain control, it sets NEW to zero. When a supervisory routine determines the new task to be given control, it inserts the TCB address for that task in NEW.

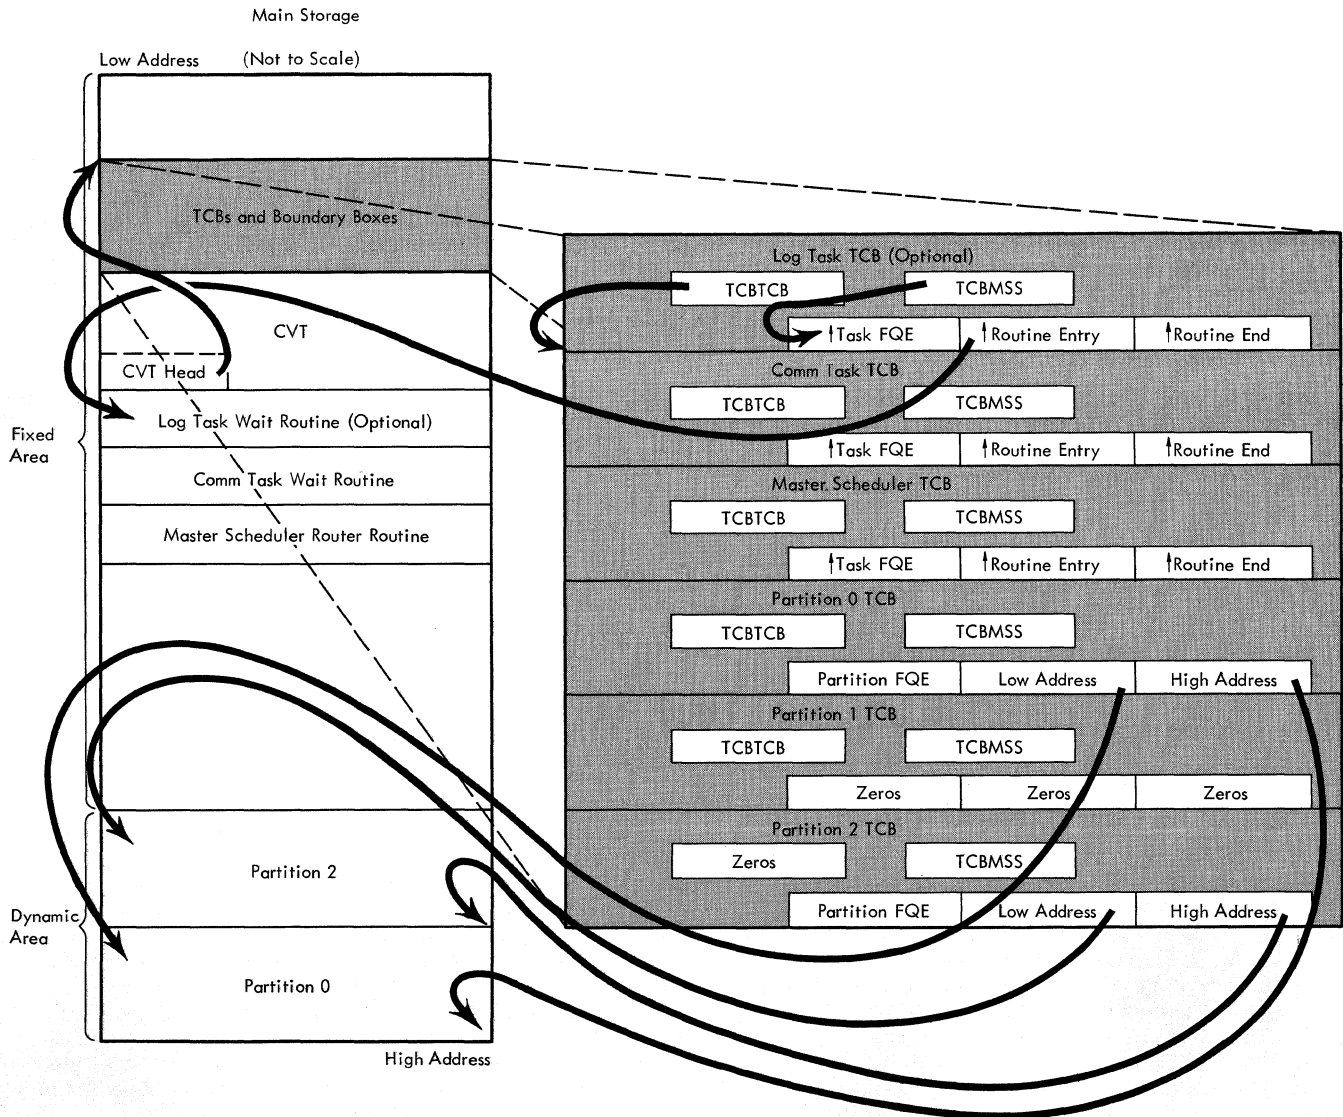
CVT field CVTTSCE contains the address of the time-slice control element (TSCE). This field is used by the dispatcher in determining the next time-slice task to receive control, providing time-slicing was specified as a system generation option. The format of a TSCE is explained later in this section.

When the interval timer is in use and a user accounting routine is supplied, the

dispatcher accumulates the total amount of time used to execute a job step. Each time a new job step is dispatched, the dispatcher stores the time from the hardware timer in the PTIMER field of IEATPC (pseudo clock area). When control is returned to the dispatcher, it calculates the elapsed time by subtracting the stored value from the current value of the hardware timer. The dispatcher adds the result to the TCBTCT field in the task's TCB. Time is not calculated for the job step if it is dispatched in a wait state.

Dispatching a Task

When the dispatcher receives control, it first schedules any requests for system asynchronous exit routines. Then it determines if NEW equals OLD (see Chart 01). If



• Figure 7. TCB Queue

so, no task switch is indicated. If necessary, the dispatcher enqueues timer elements for the task. It then returns to the task currently in control.

If NEW does not equal OLD, a task switch is indicated. If job/step CPU timing is included in the system, the dispatcher calculates the job step time for OLD, and increments the job time accumulator in the TCB. If necessary, the dispatcher dequeues timer elements associated with the task currently in control. Then it determines if NEW equals zero.

If NEW does not equal zero, it contains the TCB address for the task to be given control. The dispatcher sets OLD equal to NEW, and enqueues timer elements if necessary. Additionally, if job/step CPU timing is included in the system, the dispatcher stores the interval timer value in the pseudo timer field of IEATPC. Control then passes to the new task.

If NEW equals zero, the dispatcher must examine the TCB queue to determine which task should be given control. This examination begins with the TCB addressed by OLD. (For a task of higher priority than OLD to receive control, the address of its TCB must be inserted in NEW by a supervisory routine.)

When examining a TCB to determine if its associated task should be given control, the dispatcher first determines if the request block (RB) of the program executing under the TCB is waiting. This is done by examining field XREWT in the RB addressed by TCB field TCBRBP. If the RB is not waiting, the dispatcher examines TCB field TCBFLGS to determine if the task is dispatchable. If so, the dispatcher sets NEW and OLD to the address of the TCB and enqueues timer elements (if necessary). Additionally, if job/step CPU timing is included in the system, the dispatcher stores the interval timer value in the pseudo timer field of IEATPC. Control then passes to the new task.

In one case, the dispatcher does not pass control directly to the new task. If TCB field TCBRBP for the task to be given control addresses an SVRB for a transient SVC routine, a check is made to determine the contents of the double word XCNTCC (in IEAATA00) which contains the name of the routine presently in the SVC transient area. If the routine names in XCNTCC and the SVRB are identical, the dispatcher passes control to the new task. If they are not identical, the Transient SVC Refresh routine (IEAARF00) brings the required routine into the SVC transient

area and then returns to the dispatcher. Since NEW and OLD have already been set equal, the dispatcher need only enqueue timer elements if necessary and pass control to the new task.

If the RB for a task is waiting or the task is nondispatchable, the task is not ready to receive control. The dispatcher examines TCB field TCBTCB to obtain the address of the next TCB on the queue. The dispatcher then examines this TCB to identify whether it is ready to receive control. This process continues until a ready task is found or until the end of the queue is reached (indicated by a zero in TCBTCB).

If no task is able to receive control, the dispatcher sets the resume PSW wait bit of the TCB addressed by OLD. This PSW is then loaded, placing the CPU in a wait condition. The resume PSW is located in field XRBPSW of the RB addressed by TCB field TCBRBP.

Figures 8 and 9 illustrate how control is switched assuming a three partition system in which P1 is inactive (see Figure 7). All tasks are dispatchable except task P1. Initially, only the communications task and master scheduler task are waiting. Because task P0 is the highest priority task which is dispatchable and not waiting, it is given control. Task P0 has already enqueued and received exclusive control of a resource which task P2 will later enqueue (see Figure 9).

Dispatching the Communications Task and Master Scheduler Task

Figure 8 illustrates how control passes to the communications task and master scheduler task through the dispatcher. In the example illustrated, the communications task receives control in order to read a DEFINE command from the operator console.

Initially, the task in P0 has received control from the dispatcher and is executing. The operator presses the REQUEST key to indicate that he wishes to enter a command from the console. An I/O interruption is generated and control passes to the I/O supervisor which identifies the interruption as an attention signal. The I/O supervisor then passes control to the console interruption routine which issues a POST macro instruction. The PCST routine posts the attention ECB and sets the communications task RB to a non-wait condition. Because the communications task is of higher priority than the task in partition 0, the POST routine places the address of the communications task TCB in location NEW. Control then passes to the dispatcher.

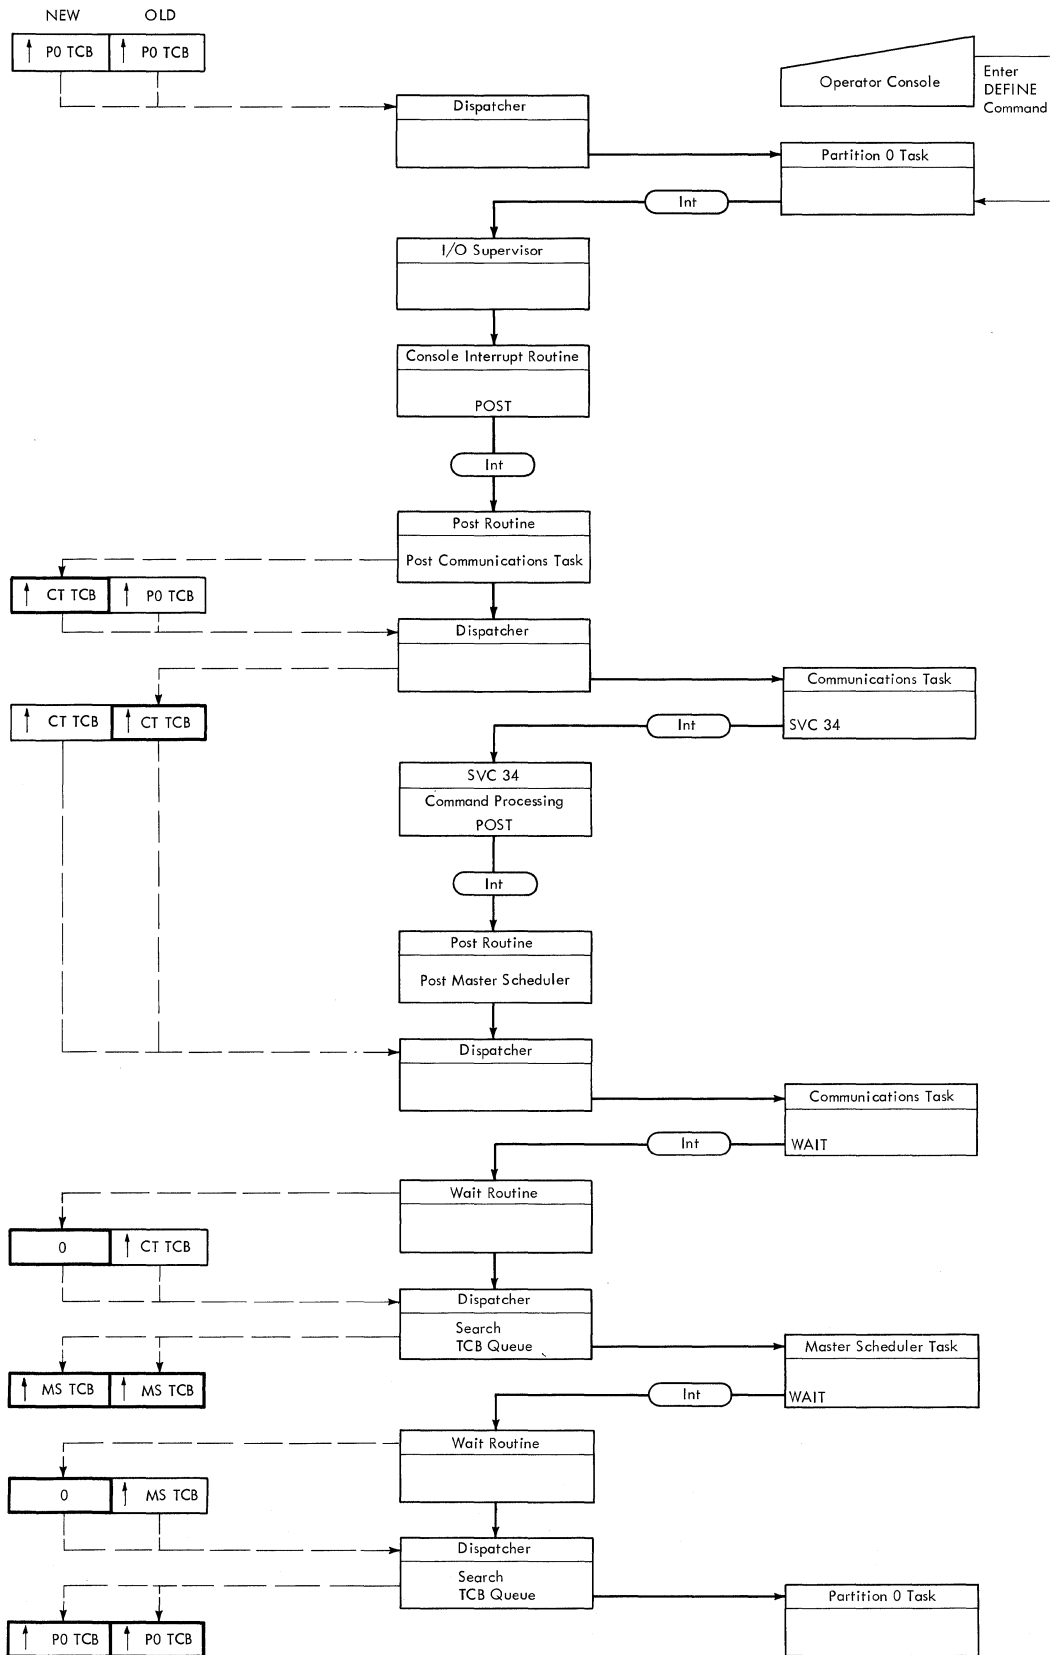


Figure 8. Dispatching Communications and Master Scheduler Tasks

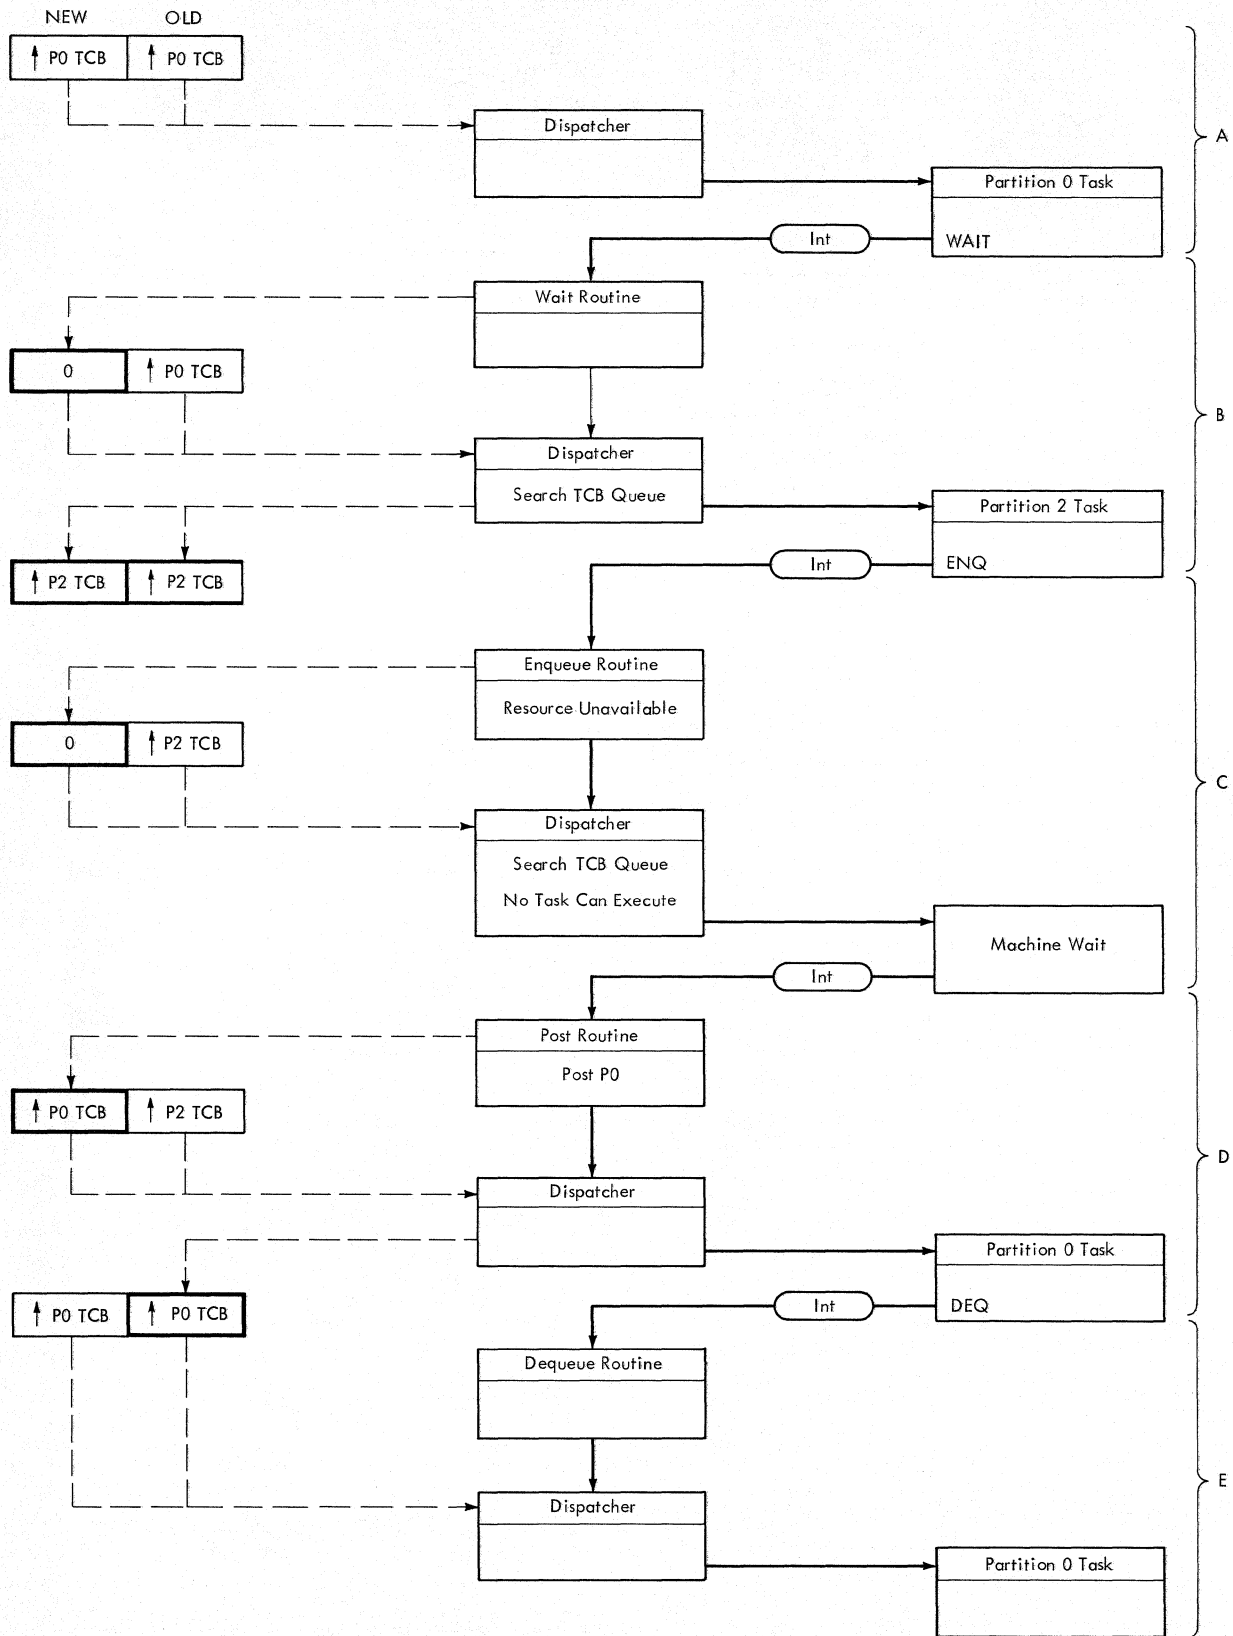


Figure 9. Task Switching

The dispatcher gives control to the communications task which passes control to resident device-support routines or issues SVC 72 for transient device-support routines. The device-support routines read the console's command and then issue SVC 34 to process the command. SVC 34 processes some commands completely but must pass control to the master scheduler resident command processor routine to complete processing the DEFINE command. (See "Command Processing" in the Job Management section for a complete description of SVC 34 and the master scheduler task.) SVC 34 issues a POST macro instruction to post the master scheduler task. The POST routine sets the master scheduler RB to a non-wait condition and gives control to the dispatcher. Because the master scheduler task is of lower priority than the communications task, locations NEW and OLD remain unchanged and the dispatcher returns control to the communications task.

The communications task issues a WAIT macro instruction and waits on an ECB. The WAIT routine sets the communications task RB in a wait state and sets location NEW to zero. The dispatcher then receives control and searches the TCB queue. Since the master scheduler task is the next ready task on the TCB queue, the address of the master scheduler TCB is placed in locations NEW and OLD, and the dispatcher passes control to the master scheduler.

The master scheduler completes processing the DEFINE command and then issues WAIT. The WAIT routine sets location NEW to zero and passes control to the dispatcher which searches the TCB queue until it finds a task ready to receive control. In Figure 8, control returns to the task which was executing before the operator entered the DEFINE command.

Dispatching Tasks by Partition Priority

Figure 9 illustrates task switching among tasks executing in partitions.

- A. The task in partition P0 (task P0) is the highest-priority ready task and is given control by the dispatcher. When task P0 issues a WAIT on an ECB, an interruption occurs and control passes to the WAIT routine.
- B. The WAIT routine places the RB for partition 0 in a wait condition and sets location NEW to zero. It then passes control to the dispatcher which searches the TCB queue beginning with the TCB for partition 0. Since task P0 is waiting and task P1 is non-dispatchable, the dispatcher passes control to task P2, the highest priority task ready to execute. When

task P2 attempts to enqueue a resource through use of the ENQ macro instruction, an interruption occurs and control passes to the ENQ routine.

- C. The resource is unavailable because task P0 has already enqueued it. Therefore, task P2 cannot continue executing. The enqueue routine places zero in location NEW and then passes control to the dispatcher which searches the TCB queue. Since task P2 is the last task on the queue, the dispatcher sets the wait bit in the resume PSW of task P2. The dispatcher passes control to task P2, placing the CPU in a machine wait condition.
- D. While the CPU is waiting, an interruption occurs signifying the completion of the event for which task P0 was waiting. The POST routine receives control and posts the ECB for task P0 which is now able to resume control. The POST routine places the TCB address for task P0 in location NEW and gives control to the dispatcher. The dispatcher sets OLD equal to NEW and gives control to task P0. Task P0 executes and when finished using the resource it has enqueued, it issues a DEQ macro instruction.
- E. An interruption occurs and the DEQ routine receives control. The queue element for task P0 is removed from the resource queue. The next element on the resource queue is for task P2. The resource is assigned to task P2 and its RB is placed in a non-wait condition. The DEQ routine then compares the priority of the task which has been in control with the priority of the task which is now ready. Because task P0 has a higher priority than task P2, location NEW remains unchanged. The DEQ routine passes control to the dispatcher which returns control to task P0.

Dispatching a Task (with Time Slicing)

If time slicing was selected as a system generation option, the user can select a number of contiguous partitions to be a time-slice group. The tasks executing in time-sliced partitions have equal priority. Each ready task in the time-slice group executes for a selected amount of time, the time-slice length, and then loses control to the next ready task in the time-slice group. The time-slice group is supervised through use of a time-slice control element (TSCE) shown following.

FIRST - Address of the first time-slice TCB on the TCB queue	4
LAST - Address of the last time-slice TCB on the TCB queue	4
NEXT - Address of the next time-slice TCB to be dispatched	4
LENGTH - Time-slice length (in milliseconds)	4

When time-slicing is selected, the dispatcher performs functions in addition to those explained in the preceding paragraphs. The following text describes the additional dispatcher functions, and parallels the flow of data shown in Chart 02.

NEW EQUALS OLD: The dispatcher first determines if NEW equals OLD. If it does, the dispatcher further determines if the task represented by OLD is a time-slice task.

OLD a Time-Slice Task: If OLD is a time-slice task, the dispatcher determines if the time-slice interval has expired; i.e., if the time-slice queue element (TQE) has been removed from the timer queue.

If the interval has expired, the next ready time-slice task must be dispatched. The dispatcher searches the time-slice group beginning with the TCB addressed by TSCE NEXT (see preceding explanation of TSCE fields). When the TCB addressed by TSCE LAST is reached, the dispatcher checks the TCB addressed by TSCE FIRST, until a ready task is found or until all time-slice TCBs have been checked.

When a ready task is found, TSCE NEXT is updated, the time-slice TQE is enqueued, and the ready task is dispatched. If no time-slice tasks are ready, the dispatcher searches the TCB queue for the highest-priority ready task.

If the interval has not expired, i.e., the time-slice TQE has not been dequeued, control is returned to the interrupted task.

OLD Not a Time-Slice Task: If OLD is not a time-slice task, control is returned to the interrupted task.

NEW NOT EQUAL TO OLD: If NEW does not equal OLD, the dispatcher determines if OLD is a time-slice task.

OLD Time-Slice Task -- NEW Equal Zero: If OLD is a time-slice task and NEW equals zero, the time-slice TQE is dequeued for the current task. The dispatcher then searches (using the TSCE) for the next ready TCB in the time-slice group. If no time-slice TCBs are ready, the dispatcher searches the TCB queue for the highest-priority ready task.

OLD Time-Slice Task -- NEW Not Equal to Zero: If OLD is a time-slice task and NEW does not equal zero, the dispatcher determines if NEW is a time-slice task.

If NEW is a time-slice task, the task represented by OLD, if ready, is redispatched. (The time-slice TQE remains on the queue.) If the task represented by OLD is not ready, the time-slice TQE is dequeued, and the dispatcher searches (using the TSCE) for the next ready time-slice task. If no time-slice tasks are ready, the dispatcher searches the TCB queue for the highest-priority ready task.

If NEW is not a time-slice task, the time-slice TQE is dequeued and the NEW task is dispatched.

OLD Not a Time-Slice Task: If OLD is not a time-slice task, the dispatcher finds the next highest-priority ready task. It does this by either obtaining the TCB address from NEW or, if NEW is zero, by scanning the TCB queue. If the highest-priority ready task is not a time-slice task, it is dispatched. If the highest-priority ready task is a time-slice task, the dispatcher finds (using the TSCE) the next ready task in the time-slice group. The time-slice TQE is enqueued, and the task is dispatched.

STAE SERVICE ROUTINE

The STAE service routine is a type 3 SVC routine which prepares the task to intercept scheduled abnormal termination (ABEND) processing. When the STAE macro instruction (resulting in an SVC 60) is issued, the STAE service routine is invoked. The STAE service routine creates a 16-byte STAE control block (SCB), which contains the addresses of a user-written STAE exit routine and parameter list. When the task becomes scheduled for abnormal termination, the ABEND/STAE interface routine (ASIR) is given control by the ABEND routine. ASIR returns control to the user at the STAE exit routine address. After the STAE exit routine has been executed, control is returned to ASIR. ABEND processing continues for the task as previously scheduled unless the STAE exit routine has requested that a STAE retry routine be scheduled. If a STAE retry routine is provided by the

user, ASIR reestablishes the task scheduled for ABEND processing and exits, giving control to the dispatcher so that the STAE retry routine is executed next. See IBM System/360 Operating System: System Programmer's Guide, Form C28-6550, for further explanation of the STAE macro instruction.

The five modules which perform the functions of the STAE macro instruction are the STAE service routine (IGC00060) and the four ABEND/STAE interface modules (IGCOB01C, IGCOC01C, IGCOD01C, and IGC0E01C). These modules perform the same functions as in MVT, (see IBM System/360 Operating System: MVT Supervisor, Form Y28-6659) with the exception both IGCOC01C and IGC0E01C pass control via the XCTL macro instruction to the ABEND module IEAGTMOA to purge the WTOR queue before giving control to the next ASIR module (IGCOD01C).

ABEND AND DAMAGE ASSESSMENT SERVICE ROUTINE

ABEND is a type 4 SVC routine that is used for both normal and abnormal task termination. ABEND terminates the task under which it is running, resets the partition, and passes control to the job management routines for continued processing.

ABEND can be entered directly from the problem program or system task via an ABEND macro instruction, or indirectly through the ABTERM service routine. (ABTERM schedules the execution of ABEND for system routines that detect an error but cannot issue an ABEND macro instruction.) The SVC SLIH (second level interruption handler) fetches the first load module of ABEND and passes control to it. Control is passed from one ABEND load module to the next via an XCTL instruction (SVC 7). The flow of control between modules for normal and abnormal termination is shown in Charts 03 and 04.

The Damage Assessment Routines (DAR) process, and attempt to recover from the following failures:

- System tasks (log, communication, or master scheduler).
- Tasks in "must complete" status.
- Tasks experiencing invalid ABEND recursion.

A record of the failures is provided in a core image dump. A primary DAR recursion results when a failure occurs while writing the main storage image dump. A secondary DAR recursion results when a failure occurs during partition recovery.

The ABEND and DAR functions provided for MFT are similar to those provided for PCP. ABEND modules IEAGTMOA, IEAGTM00, IEAGTM05,

IEAGTM06, and DAR modules IEAGTM08 and IEAGTM09 exist only in MFT and are described below. The remaining modules are also used in PCP and are explained in IBM System/360 Operating System: Fixed Task Supervisor, Form Y28-6612. (For a brief description of all ABEND modules used in MFT, see "MODULE DESCRIPTIONS," IN Appendix B of this publication.)

ABEND STAE Test Routine (IEAGTMOA)

IEAGTMOA first sets a bit to prohibit asynchronous exits for this task and tests the TCB for evidence of invalid recursion. If a primary DAR recursion or an invalid ABEND recursion is found, control is passed to DAR module IEAGTM08. If a secondary DAR recursion is found, control is passed to DAR module IEAGTM09.

The routine then tests for normal end of task. If this is a normal end, the normal completion code is stored in the TCB, and control passes to ABEND module IEAGTM00. If the task is abnormally terminating, IEAGTMOA determines if STAE (specify task asynchronous exit) processing is indicated for this task (via a user-issued STAE macro instruction).

If a STAE was issued, i.e., TCB field TCBNSTAE does not equal zero, IEAGTMOA checks for a valid STAE and performs as follows:

- If the ABEND was issued by the Purge routine during STAE processing, i.e., the purge bit in TCB field TCBNSTAE equals one, the resume PSW of the Purge RB routine is set to the address of an EXIT instruction (SVC 3). IEAGTMOA then issues an EXIT instruction.

Note: If ABEND was not entered from the Purge routine, i.e., the purge bit in TCB field TCBNSTAE equals zero, IEAGTMOA stores the abnormal completion code in the TCB.

- If the Task is being abnormally terminated because of a timer expiration or an operator cancel, STAE processing is bypassed and IEAGTMOA exits to ABEND module IEAGTM00. (Since task timing and cancelations from the console are directly controlled by the user, the ABEND was intentional and should not be handled by the STAE routines.)
- If STAE processing is already in progress, regular ABEND processing continues, since the STAE routine can process only once per STAE issued. IEAGTMOA thus exits to ABEND module IEAGTM00.

- If this is a valid request for STAE processing, i.e., none of the above conditions are true, IEAGTMOA tests TCB field TCBPIE for zero. If it is not zero, IEAGTMOA frees the PIE and zeros TCBPIE. Thus, subsequent program checks will not be handled by a user routine which may not be designed for such a program check. IEAGTMOA then exits to STAE module IEAATMOB.

If a STAE has not been issued, or if all STAEs have been processed, i.e., TCB field TCBNSTAE equals zero, IEAGTMOA determines if this is a graphics or an ABEND recursion and if this is a graphics job with a Graphics Abend Exit routine.

- If this is a recursion, or if the task abnormally terminating is not a graphics job, IEAGTMOA exits to ABEND module IEAGTM00.
- If this is a graphics job, IEAGTMOA passes control to the Graphics Abend routine. At the completion of this routine, control is returned either to the caller via an SVC 3 (if the task is resumable, i.e., the ABEND was issued by a user program or caused by a program check in a user routine, and if he wishes to resume processing) or to IEAGTMOA, which exits to ABEND module IEAGTM00.

ABEND Initialization Routine (IEAGTM00)

This routine prepares for ABEND processing of a task by canceling the task timer element and by dequeuing all interruption queue elements (IQEs) belonging to the task. If the routine was entered from STAE, it also purges the WTOR queue elements for the task and passes control to IEAATMOD. If the routine was entered as the result of a normal end of task, it passes control to IEAGTM05. If it was entered as the result of an abnormal end of task, it passes control to IEAGTM06.

ABEND Input/Output Purge Routine (IEAGTM06)

IEAGTM06 purges I/O requests and I/O operations via a macro instruction version of the SVC Purge Routine, which is assembled within this module. (See "SVC PURGE ROUTINE" in Input/Output Supervisor, Form Y28-6616.) This prevents errors that can cause recursion to the ABEND routine. (Since ABEND frees main storage, an I/O operation that is not halted can cause information to be read into, or an ECB to be posted in main storage that may have been relocated, thus destroying data or programs.) RQEs (request queue elements) removed from the request queue are returned to a list of available RQEs for reuse by the I/O supervisor.

IEAGTM06 also dequeues, from the SIRB (system interruption request block), IQEs (interruption queue elements) representing requests for the use of I/O error handling routines. The routine passes control to DAR module IEAGTM08 when Recovery Management Support is not the caller, and the failing task is a task in "Must Complete" status or if the failing task is a system task. Otherwise IEAGTM06 passes control to ABEND module IEAATM01.

ABEND Termination Routine (IEAGTM05)

ABEND termination routine IEAGTM05 is the final ABEND module for both normal and abnormal termination. For normal termination it is entered from IEAGTM06. On abnormal termination it may have been entered from any of the previous modules of ABEND except initialization routine IEAGTM00.

If a dump message is required, i.e., if ABDUMP has been initiated but has failed to complete, IEAGTM05 causes message IEA002I, "ABEND/ABDUMP ERROR" to be printed. IEAGTM05 issues a CLOSE macro instruction for any open data sets. The timer queue and the main storage supervisor queue are purged, and fields in the TCB are reset so that a new task may be initiated. If an indicative dump is provided by IEAATM03, it is moved to the upper boundary of the partition.

IEAGTM05 does not directly transfer control to a job management routine. In the first 72 bytes of the problem program partition, IEAGTM05 establishes a dummy PRB, an XCTL parameter list, and a set of instructions including an XCTL to a step deletion routine. The XRBLNK field of the dummy PRB contains a pointer to the TCB. The dummy PRB therefore becomes the only RB queued for this task.

The dummy PRB is then placed at the beginning of the RB queue. This ensures that the XCTL instruction will be the next operation executed for this task after IEAGTM05 has completed. For scheduler-size partitions, step deletion routine IEFSD515 gains control, at entry point GO. For small partitions, control is passed to entry point SMALLGO in small partition routine IEFSD599.

DAMAGE ASSESSMENT ROUTINES

The damage assessment routines are used to eliminate wait states due to system failures. When a system failure occurs, the damage assessment routines provide an image of main storage at the time of failure and

reinitialize the failing task. The routines also advise the operator of the failure and subsequent reinstatement of the task.

There are two routines that are unique to MFT: Damage assessment core image dump routine IEAGTM08, and damage assessment task reinstatement routine IEAGTM09.

DAR Core Image Dump Routine (IEAGTM08)

The DAR core image dump routine IEAGTM08 writes on the SYS1.DUMP data set an image of main storage at the time of failure. When entered, the routine sets all tasks except the failing task and the communications task nondispatchable. The routine then writes the image of main storage and passes control to the DAR task reinstatement routine IEAGTM09.

If the SYS1.DUMP data set has not been allocated, the routine informs the operator via a WTO. If the routine is entered as a result of a primary DAR recursion, which is caused by a failure to write the image of main storage, the routine does not try to rewrite but informs the operator of the failure via a WTO. In both cases the routine passes control to IEAGTM09.

If the communications task is the failing task, messages are queued pending reinstatement of the communications task by IEAGTM09.

DAR Task Reinstatement Routine (IEAGTM09)

If the DAR task reinstatement routine IEAGTM09 is entered as a result of a failing system task, the routine attempts to reinstate the task. It points the resume PSWs of all but the highest level RB of the task's TCB to an SVC 3 instruction in the CVT. It points the highest level RB to entry point IEECIR50 for the Master Scheduler task, entry point IEECIR45 for the Communications task, or entry point IEEVLIN for the Log task. The routine then passes control to the dispatcher via a branch instruction.

If the routine is entered as a result of a secondary DAR recursion, which is caused by a failure to reinstate the failing task, the routine informs the operator via a WTO, sets all tasks dispatchable except the failing task, and passes control to the dispatcher via a branch instruction.

If the failing task is in "Must Complete" status, the task reinstatement routine issues a message to the operator listing the major and minor names of the enqueued resources that have caused the "Must Complete" condition and asking the operator to reply whether the resources are

critical. If the reply indicates the resources are critical, processing is identical to the processing of a secondary DAR recursion described above. If the reply indicates the resources are not critical, the "Must Complete" status is removed, and the resources are designated as shareable. The task is processed as a failing non-system task as described below.

If the routine is entered as a result of a failing non-system task, it sets indicators showing that a dump has been taken by DAR and issues a message to the operator indicating that the system has been reinstated. The routine then sets all tasks dispatchable and passes control to IEAGTM05.

Task Supervision

The task supervisor maintains the status of tasks within the system. Task supervision service routines:

- Maintain task control blocks.
- Enter tasks into the wait state.
- Post completed events in the event control block (ECB).
- Maintain control levels indicated by request blocks.

The routines which accomplish these functions are WAIT, POST, ENQ, and DEQ.

Each task within the operating system has an associated task control block (TCB). The TCB contains task-related information and pointers to additional control blocks containing task-related information. The control blocks used by MFT are the same as those used by PCP except for the addition of the partition information block (PIB) which is described in Appendix A. The last three bytes of the word at displacement 124 (decimal) of each partition TCB contain the address of the associated PIB. Figure 10 shows the major control blocks maintained by the supervisor and their relationship to the TCB.

Task supervision is described in IBM System/360 Operating System: Fixed-Task Supervisor, Program Logic Manual, Form Y28-6612. Additional information applicable to MFT is presented in the following paragraphs.

THE ATTACH ROUTINE (MACRO IEAAAT)

In MFT, the ATTACH and LINK macro instructions are handled identically. An RB is created for the requested program, the program is brought into the requesting task's partition, and its RB is chained to the RB

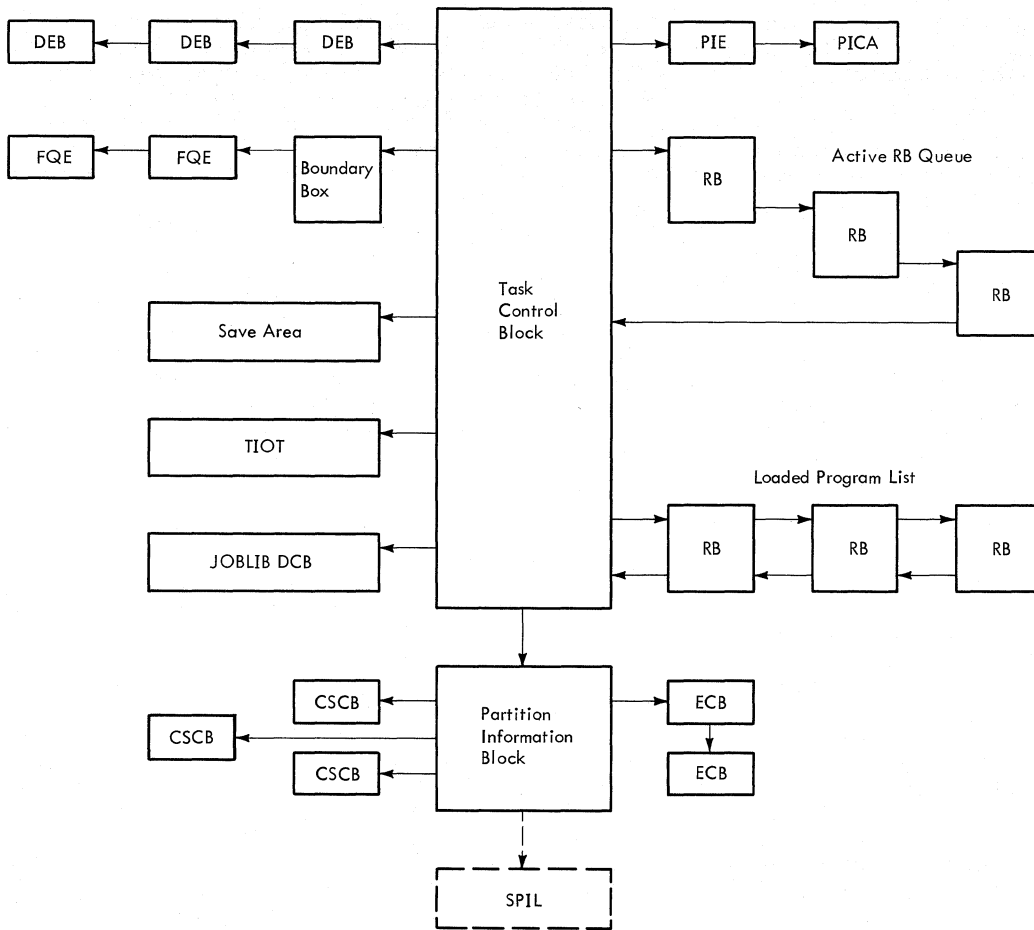


Figure 10. System Control Block Relationship

queue for that partition. See IBM System/360 Operating System: Supervisor and Data Management Services, Form C28-6646 for further explanation of the ATTACH macro instruction with MFT.

THE WAIT ROUTINE (MACRO IEAAWT)

The WAIT routine is not changed from that described in Fixed-Task Supervisor, Program Logic Manual. However, the user should remember the effect the optional validity checking feature has on WAIT. If validity check is included in the system and the program issuing the WAIT macro instruction is not in supervisor mode, the WAIT routine checks that:

1. The boundary alignment of the ECBs is correct.
2. The storage protection key of the ECBs is that of the issuing program.
3. The addresses specified do not exceed main storage boundaries of the machine.

Because of point 2, it is not possible for one partition to WAIT on an ECB within another partition.

THE POST ROUTINE (MACRO IEAAPT)

The POST routine, like the WAIT routine, is unchanged from that described in Fixed-Task Supervisor, Program Logic Manual. Validity checking applies to PCST in the same way it applies to WAIT.

THE ENQ/DEQ ROUTINE (IEAGENQ1)

The ENQ/DEQ routine provides a means of controlling serially reusable resources. This is done by assigning unique names consisting of a Qname and an Rname to each serially reusable resource. The ENQ/DEQ routine controls access to resources by building resource queues consisting of a queue control block (QCB) for each Qname and Rname specified in an ENQ macro instruction and a queue element (QEL) to represent each actual request. ENQ/DEQ is

fully described in IBM System/360 Operating System: MVT Supervisor, Program Logic Manual, Form Y28-6659. ENQ/DEQ for MFT is identical to MVT except as described below.

In MFT, resource queues are located in the system queue area (subpool 255). Location IEACQCB0 in the ENQ/DEQ routine contains the address of the first queue control block in the queue. There is only one TCB for each job step in MFT. Therefore, the "must complete" function of ENQ/DEQ applies only to the system, not to job steps. If "system must complete" is specified by a task, all other tasks in the system are set non-dispatchable until the task which specified "system must complete" completes its processing.

Contents Supervision

Contents supervision routines determine the location of requested programs and fetch them into main storage if necessary. They also maintain records of all programs in main storage. Programs requested via LINK or XCTL macro instructions are scheduled for use by placing a request block (RB) for each program on the requesting task's active request block queue.

Programs requested via LOAD macro instructions are represented by RBs on the loaded program list.

There are six types of request blocks in MFT:

- Program Request Block (PRB) -- represents a nonsupervisory routine that must be executed in the performance of a task. PRBs are created by the contents supervision routines that perform the LINK or XCTL functions.
- Supervisor Request Block (SVRB) -- represents a supervisory routine. SVRBs are created by the SVC interruption handling routines.
- Interruption Request Block (IRB) -- controls a routine that must be executed in the event of an asynchronous interruption. IRBs are created in advance of an interruption by the CIRB routine at the user's request, but not placed on an RB queue until an interruption actually occurs.
- System Interruption Request Block (SIRB) -- used only for the system I/O error task. There is only one SIRB in the system.

- Loaded Program Request Block (LPRB) -- controls programs brought in by a LOAD macro instruction. LPRBs also control sections of programs that are specified by the IDENTIFY macro instruction. LPRBs are created by the contents supervision routines that perform the LOAD function.
- Loaded Request Block (LRB) -- a shortened form of LPRB and controls load modules that have the "load only" attribute. It is invalid to issue ATTACH, LINK, or XCTL macro instructions to these load modules. LRBs are created by the routines that perform the LOAD function.

Contents supervision routines alter the active RB queue and the loaded program list, and bring nonresident programs into the problem program partitions in response to LINK, ATTACH, LOAD, and XCTL macro instructions. Additional contents supervision services are provided by the use of IDENTIFY, DELETE, and SYNCH macro instructions. IDENTIFY and DELETE alter the loaded program list. SYNCH alters the active request block queue. The routines that service these macro instructions are described below.

LINK SERVICE ROUTINE (MACRO IEAATC)

The LINK service routine determines if the RB of the requested routine is on the loaded program list. If it is and is inactive, LINK places the RB on the active RB queue. If the requested RB is not on the loaded program list (or if it is on the list, but is active), and the resident re-entrant routine option was selected at system generation, the routine searches the resident area. If the module is found in the area, a load list element for the module is placed on the loaded program list, and processing continues as if the module were originally found on the load list. If the module is not found, the LINK routine constructs an RB for the requested routine, places the RB on the active RB queue, and fetches the requested routine into main storage.

ATTACH SERVICE ROUTINE (MACRO IEAAAT)

The ATTACH macro instruction is handled as a LINK macro instruction. For a complete explanation, see "The ATTACH Macro Instruction" under the topic Task Supervision.

LOAD SERVICE ROUTINE (MACRO IEAATC)

The LOAD service routine first determines if the requested routine is in the resident

reenterable routine area (if the resident routine option was specified at system generation). If so, the entry point of the routine is passed to the requesting routine in register zero. If the routine is not resident, LOAD searches the loaded program list for the RB of the requested routine. If it is found, the LOAD routine increments the RB use count by one and returns the entry point of the requested routine in register zero.

If the requested routine is not found on the loaded program list, the LOAD routine branches to the FINCH routine to load the requested routine into storage. On return from the FINCH routine, the LOAD routine initializes the requested routine's RB and places it on the loaded program list, sets the RBs use count to one and branches to the LINK routine to issue the SVC EXIT instruction.

XCTL SERVICE ROUTINE (MACRO IEAATC)

The XCTL service routine first determines if XCTL was issued by a transient SVC routine. It then determines if the resident SVC (RSVC) option was chosen at system generation and determines if the requested SVC routine is an RSVC routine. If it is, the routine need not be brought into main storage. If the requested routine is not an RSVC, the XCTL routine branches to the FINCH routine to locate the routine on the SVC library and to bring it into the SVC transient area. The XCTL routine initializes the routine's RB and executes an SVC EXIT instruction.

If the XCTL macro instruction was not issued by a transient SVC routine, the XCTL routine dequeues the primary RB and each minor RB of the issuer from the active RB queue. The routine which issued the XCTL macro instruction and its RB are removed from storage unless the routine was brought in via a LCAD macro instruction. If the requested routine is on the loaded program list and is inactive, the XCTL routine branches to the LINK routine to place the RB on the active queue and to issue an SVC EXIT instruction.

If the XCTL routine determines that the scheduler has issued an XCTL macro instruction to branch to the problem program, the XCTL routine zeros out the TCBTCT field of the TCB so that the optional Job/Step CPU Timing entry can be made.

If the RB of the requested routine was not found inactive on the loaded program list, and the resident reenterable module option was selected at system generation, the routine searches the resident area. If the module is found in the area, a load

list element for the module is placed on the loaded program list, and processing continues as if the module were originally found on the load list. If the module is not found, the XCTL routine branches to the FINCH routine to bring in the routine. On return from the FINCH routine, the XCTL routine branches to the LINK routine to place the RB on the active queue and issue an SVC EXIT instruction.

IDENTIFY SERVICE ROUTINE (IEAAID00)

The IDENTIFY service routine builds and initializes a minor request block to describe a routine specified in the parameters of the IDENTIFY macro instruction. The IDENTIFY routine chains this minor RB to the loaded program list and to the RB of the routine which contains the identified routine. The IDENTIFY routine returns to the issuer by issuing an SVC EXIT instruction.

DELETE SERVICE ROUTINE (IEAADL00, IEABDL00)

The DELETE service routine determines if the routine specified in the DELETE macro instruction is resident. If it is, the DELETE routine exits immediately. If the routine is not resident, the DELETE routine finds the routine's RB on the loaded program list and decrements the use count in the RB by one. If the use count reaches zero, the DELETE routine dequeues the routine from the loaded program list and issues a FREEMAIN macro instruction to release the storage occupied by the specified routine and its RB. On return from the FREEMAIN routine, the DELETE routine repeats the deleting process for each minor RB belonging to the specified routine. The DELETE routine returns by branching to the type 1 SVC exit.

SYNCH SERVICE ROUTINE (IEAASY00)

The SYNCH service routine uses GETMAIN to obtain 32 bytes of main storage from the lower end of the partition for the creation of a program request block (PRB). The PSW in the PRB is initialized by the SYNCH routine to address the location specified in register 15 by the issuer of the macro instruction. The SYNCH routine sets the PSW completely enabled in problem program mode, with the protection key recorded in the task control block. After the PRB is created and initialized, the SYNCH routine queues it on the active request block queue below the SVRB for SYNCH, and returns by issuing an SVC EXIT instruction.

Additional information describing PCP and MFT Contents Supervision can be found in Fixed-Task Supervisor, Program Logic Manual.

Main Storage Supervision

In MFT, the main storage supervisor:

1. Allocates space via the GETMAIN SVC.
2. Deallocates space via the FREEMAIN SVC.
3. Allocates space in the system queue area.
4. Checks validity of requests that are to be serviced.
5. Maintains the pointers and control blocks necessary to supervise main storage.

Each job is assigned to a partition in which it must operate. Each partition has an associated TCB which contains a pointer (TCBMSS field) to the main storage boundary box for that partition. The main storage supervisor, in response to GETMAIN macro instructions, obtains storage from either the problem program partition or the system queue area. Obtaining storage space from the system queue area is the basic difference in main storage supervision between MFT and PCP. In MFT, a system task can issue a GETMAIN macro instruction specifying subpool 255 and the required storage will be allocated from the system queue area. The system queue area is used to obtain space for system control blocks which might be destroyed by problem programs if they were placed in problem program partitions. The system tasks which request storage space from subpool 255 are:

- The CSCB creation module of SVC 34, for CSCBs.
- The ENQ/DEQ processing routines of task supervision, for all control blocks associated with ENQ/DEQ.
- The communications task, for write-to-operator (WTO) buffers if all WTO buffer storage space specified during system generation is unavailable.

Note: Although subpools are not created in MFT (as in PCP and MVT), problem programs and system tasks may specify subpools in the GETMAIN macro instruction. However, all main storage requests from problem programs are allocated from the highest available main storage in the partition which issued the GETMAIN.

The boundary box for the system queue area is located in master scheduler resident data area IEESD568 (see Appendix A). The master scheduler resident data area is addressed by the CVTMSER field in the Communications Vector Table.

When problem programs issue GETMAIN macro instructions specifying a subpool from 0 through 127, storage is allocated from the high-address portion of the partition in which the GETMAIN macro instruction was issued. When problem programs attempt to issue a GETMAIN macro instruction specifying a subpool from 128 through 255, the program is abnormally terminated.

When system tasks issue a GETMAIN macro instruction specifying a subpool from 0 through 127, storage is allocated from the low-address portion of the partition; when specifying a subpool from 128 through 254, storage is allocated from the high-address portion of the partition. Subpool 255 is handled as a special case as described in preceding paragraphs.

For a complete description of main storage supervisor functions, see Fixed-Task Supervisor, Program Logic Manual.

TIMER SUPERVISION

Timer supervision routines are an optional feature of MFT. If selected, the user may request timer services through the TIME, STIMER, and TTIMER macro instructions. The TIME service routine (IEA0RT00) determines the date and time of day. The STIMER service routine (IEA0ST00) sets a user specified interval, and the TTIMER service routine (IEA0TT00) determines the amount of time remaining in a previously specified interval. Whenever a timer interval is requested in an STIMER macro instruction, a timer queue element (TQE) is constructed. These elements are chained together in a timer queue. The queue is ordered so that the TQE representing the next interval to expire is always at the top of the queue. When a requested interval expires, a timer interruption occurs and the Supervisor Timer Interruption Handling Routine (IEA0TI00) takes appropriate action, depending on the type of interval which has expired. If job/step CPU timing is included in the system, IEA0TI00 adjusts the pseudo timer field in IEATPC in the same manner it adjusts the hardware timer.

Timer Supervision

The System/360 interval timer is a 32 bit word in lower main storage which continually decrements as long as the system is running and the interval timer switch is on. The timer supervision routines use this hardware timer to accomplish their functions. The timer supervision routines can set the hardware timer to any interval between zero and six hours. An interruption occurs when the hardware timer decre-

ments to zero. Since the hardware timer never exceeds six hours, four values are needed to maintain elapsed time for a full day. These values are:

- Hardware timer.
- Six Hour Pseudo Clock (SHPC).
- Twenty-four Hour Pseudo Clock (T4PC).
- Local Time Pseudo Clock (LTPC).

The SHPC is used to time intervals up to six hours; the T4PC is used to time intervals up to twenty-four hours. The LTPC contains the local time of day entered by the operator during system initialization.

When an STIMER macro instruction is issued, the STIMER supervisory routine adjusts the time interval requested relative to the intervals in the hardware timers and pseudo clocks. This enables the supervisory routines to place the newly requested timer element in the correct place on the timer queue.

TIMER PSEUDO CLOCK ROUTINE (IEATPC)

The timer pseudo clock routine (IEATPC) contains all variable information that would normally be included in the resident timer routines. This information includes:

- Pseudo clocks.
- Work space used for incrementing CVT date.
- Accumulator for the job/step CPU timing feature.

COMPARISON OF PCP, MFT, AND MVT TIMER SUPERVISION

Requests for timer services in PCP, MFT, and MVT are made using the same macro instructions. Timer requests are enqueued on the timer queue in the same way in all three systems. There is one difference between PCP and MFT timer supervision. Because there is only one partition in PCP, the timer completion exit routine receives control as soon as a requested task time interval expires. When a timer interval expires in MFT, the timer completion exit routine does not receive control until the task which requested the timer interval is the highest priority ready task in the system. In MVT, the maximum amount of time permitted to complete a job step or cataloged procedure may be specified on the EXEC card. This facility is not provided in MFT.

For a complete description of timer supervisor, see Fixed-Task Supervisor, Program Logic Manual, and MVT Supervisor Program Logic Manual.

Overlay Supervision

The routines which supervise loading of overlay program segments and assist flow of control between segments of the overlay program are identical in operation for PCP and MFT. A complete description of PCP and MFT overlay supervision can be found in Fixed-Task Supervisor, Program Logic Manual.

MFT Recording/Recovery Routines

Operating System Recording/Recovery routines are optional control program routines which may be selected during system generation. They handle two types of equipment malfunctions:

- Malfunctions of the central processing unit (CPU), which cause machine-check interruptions.
- Malfunctions in a channel, which cause input/output interruptions.

Operating System Recording/Recovery routines are divided into two groups: System Environment Recording and Recovery Management.

System Environment Recording includes:

- System Environment Recording 0 (SER0), described in IBM System/360 Operating System: Fixed-Task Supervisor, Program Logic Manual, Form Y28-6612.
- System Environment Recording 1 (SER1), also described in the Fixed-Task Supervisor PLM.

Recovery Management includes:

- Machine-Check Handler (MCH), described in IBM System/360 Operating System: Machine-Check Handler for IBM System/360 Model 65, Program Logic Manual, Form Y27-7155.
- Channel-Check Handler (CCH), described in IBM System/360 Operating System: Input/output Supervisor, Program Logic Manual, Form Y28-6616.

MACHINE-CHECK ROUTINES

There are three machine-check routines.

The recording routines:

SER0, which records information about the error and then places the system in a wait state.

SER1, which records information about the error and attempts to associate the error with a task. If it can do this, it abnormally terminates the task and allows the system to continue operation.

The recovery routine:

MCH, which records information about the error and attempts complete recovery from it, including retry of the instruction that caused the error.

For the Model 65, any one of these three routines may be selected during system generation. For the Model 40, 50, 75, and 91, either SER0 or SER1 may be selected. If no routine is selected, either SER0 or SER1 is used by default. The version used by default depends on the model (or models) specified, and on the size of the system (see IBM System/360 Operating System: System Generation, Form C28-6554).

CHANNEL-CHECK ROUTINE

There is only one channel-check routine:

CCH, which aids recovery from channel errors by:

- Providing channel error information to IBM-supplied device dependent error recovery procedures (ERP).
- Building a record entry which is later written on SYS1.LOGREC by the statistical data recorder of the I/O supervisor.

This routine may be selected during system generation for the Model 65, 75, and 91 only.

SYSTEMS WITHOUT RECORDING/RECOVERY ROUTINES

A machine check or I/O interruption caused by an equipment malfunction places in a wait state those IBM System/360 models that do not have Recording/Recovery routines (See Figure 11). A message is issued on the console telling the operator to load the System Environment Recording, Editing, and Printing (SEREP) program. SEREP is a model-dependent, stand-alone diagnostic program. Its use is described in IBM

System/360 Operating System: Operator's Guide, Form C28-6540.

ENTRY TO RECORDING/RECOVERY ROUTINES

When a machine-check interruption occurs, the machine-check new PSW is loaded. This causes control to pass directly to the Recording/Recovery routine which was selected during system generation (see Figure 11).

When an I/O interruption occurs because of a channel error, the I/O new PSW is loaded. This causes control to pass to the I/O FLIH and then to the I/O Supervisor.

If the Channel-Check Handler option was not selected during system generation, the I/O Supervisor enters the SER Interface subroutine (SERR04) within the I/O Supervisor. This routine loads the machine-check new PSW (See Figure 11).

If the Channel-Check Handler was selected during system generation, the I/O Supervisor enters the Channel-Check Handler Interface (SERR04) within the I/O Supervisor (see Figure 11).

Checkpoint/Restart Routines

The checkpoint/restart routines used by MFT allow a job to restart after an abnormal termination. The checkpoint routine (SVC 63) is used by the programmer to create a record of the job's main storage region at selected points during the execution of a job step. The routine is identical with the PCP checkpoint routine described in IBM System/360 Operating System: Fixed-Task Supervisor, Program Logic Manual, Form Y28-6612.

The restart routine (SVC 52) allows jobs to restart at a checkpoint. If the restart is automatic, it will occur at the last valid checkpoint taken by the job before it abnormally terminated. If the restart is deferred, it will occur at the checkpoint specified by the job statement. Processing of the restarting job is discussed in the Job Processing section of this manual. The restart routine is described in IBM System/360 Operating System: Fixed Task Supervisor, Program Logic Manual, Form Y28-6612.

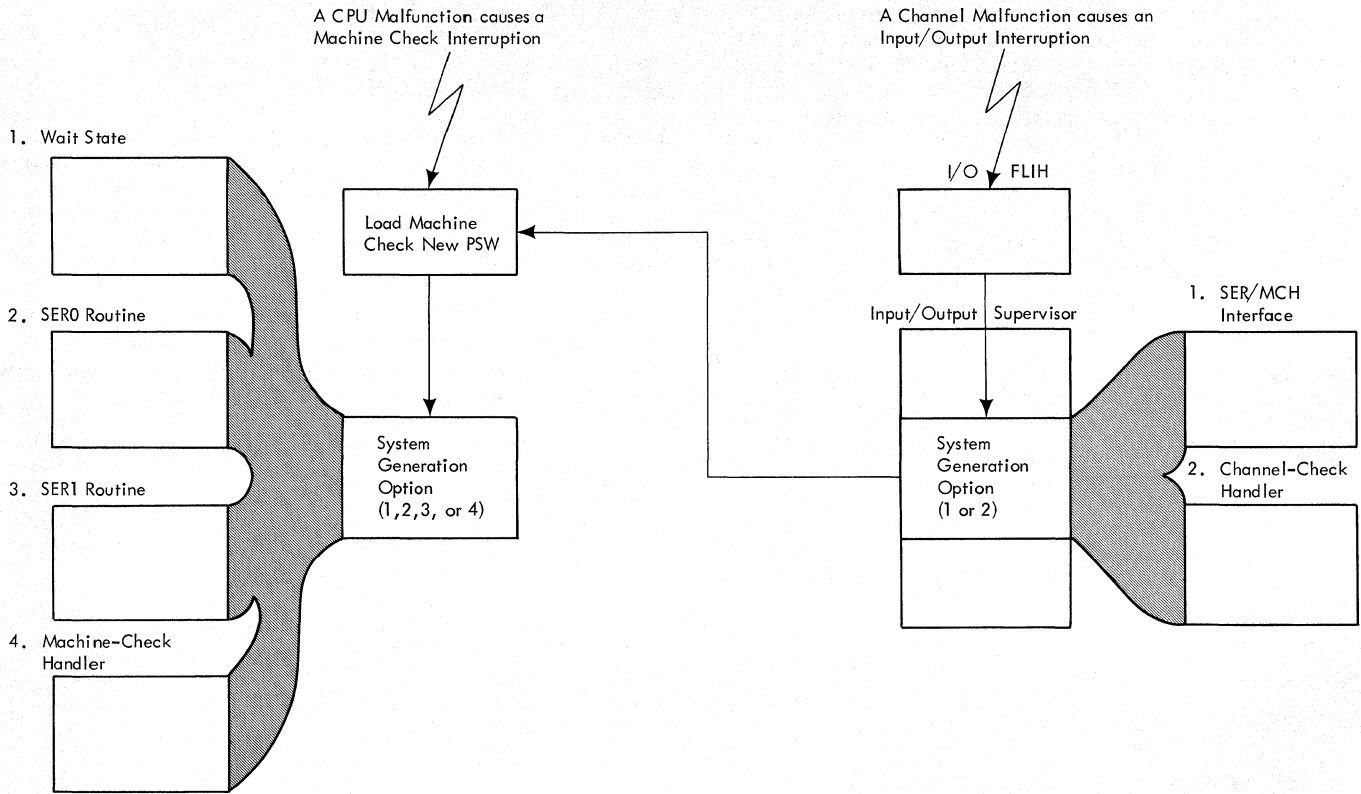


Figure 11. Recording/Recovery Routines

Job Management

The primary job management function is to prepare job steps for execution and, when they have been executed, to direct the disposition of data sets created during execution. Prior to step execution, job management:

- Reads control statements from the input job stream
- Places information contained in the statements into a series of tables.
- Analyzes input/output requirements.
- Assigns input/output devices.
- Passes control to the job step.

Following step execution, job management:

- Releases main storage space occupied by the tables.
- Frees input/output devices assigned to the step.
- Disposes of data sets referred to or created during execution.

Job management also performs all processing required for communication between an operator and the control program. Major components of job management are the job scheduler, which introduces each job step to the system (job processing), and the communications and master scheduler tasks, which handle all operator-system communication (command processing).

JOB SCHEDULER FUNCTIONS

The job scheduler includes three programs: the reader/interpreter, the initiator/terminator, and the system output writer. The functions of the reader/interpreter are similar to the MVT reader; additional information can be found in IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.

After all control statements for a job have been processed, all initiators that are waiting for that job class are posted and the initiator residing in the highest priority partition is given control. The MFT initiator is described in the Job Management section of this publication; for information on allocation and termination, refer to IBM System/360 Operating System:

MVT Job Management, Program Logic Manual, Form Y28-6660.

When the job step has been executed, control is returned to the initiator/terminator which performs data set dispositions and releases input/output (I/O) resources. If the entire job is to be terminated, the terminator enqueues all data sets on the appropriate system output (SYS-OUT) queues.

When the system output writer receives control, it dequeues a job from an output queue, and transcribes the data sets to the user-specified output device. (See IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660, for further information on the system output writer.)

COMMUNICATIONS TASK FUNCTIONS

The routines of the communications task process the following types of communication between the operator and the system:

- Operator commands, entered through a console.
- Write-to-operator (WTO) and write-to-operator with reply (WTOR) macro instructions.
- Interruptions caused when the INTERRUPT key is pressed, to switch functions from the primary console/master console to its alternate console.
- If the system has Multiple Console Support, the communications task processes the delete operator message (DCM) macro instruction and provides buffer management for all console devices.

MASTER SCHEDULER TASK FUNCTIONS

The master scheduler task consists of SVC 34 and the master scheduler resident command processor routines. The SVC 34 command scheduler routines process all commands initially. The job queue manipulation and partition definitions, which are not fully processed by SVC 34, are passed to the master scheduler resident command processor. Table 1 lists the commands used in MFT and indicates the routine which responds to the commands after initial processing.

Table 1. Responders to Commands After Initial Processing

Command	Responder
CANCEL (active jobs)	Initiator
CANCEL (job in queue)	Master Scheduler
DEFINE	Master Scheduler
DISPLAY STATUS, JOBNAMES, DSNAME	Initiator
DISPLAY A,Q,N,jobname, CONSOLES	Master Scheduler
DISPLAY R	Master Scheduler
DISPLAY SPACE	I/O Device Allocation
DISPLAY T	Timer Maintenance Routine *
HALT	Statistics Update Routine *
HOLD	Master Scheduler
LOG	System Log
MODE	Master Scheduler
MODIFY	Writer
MOUNT	Master Scheduler
RELEASE	Master Scheduler
REPLY	Master Scheduler
RESET	Master Scheduler
SET CLOCK, DATE	Timer Maintenance Routine *
SET PRCC, Q, AUTO	Master Scheduler
START/STOP Reader	Reader/Interpreter
START/STOP Writer	Writer
UNLOAD	Initiator
VARY	Initiator
WRITELOG	System Log*

* See the publication IBM System/360 Operating System: MVT Supervisor, Program Logic Manual, Form Y28-6659.

JOB MANAGEMENT CONTROL FLOW

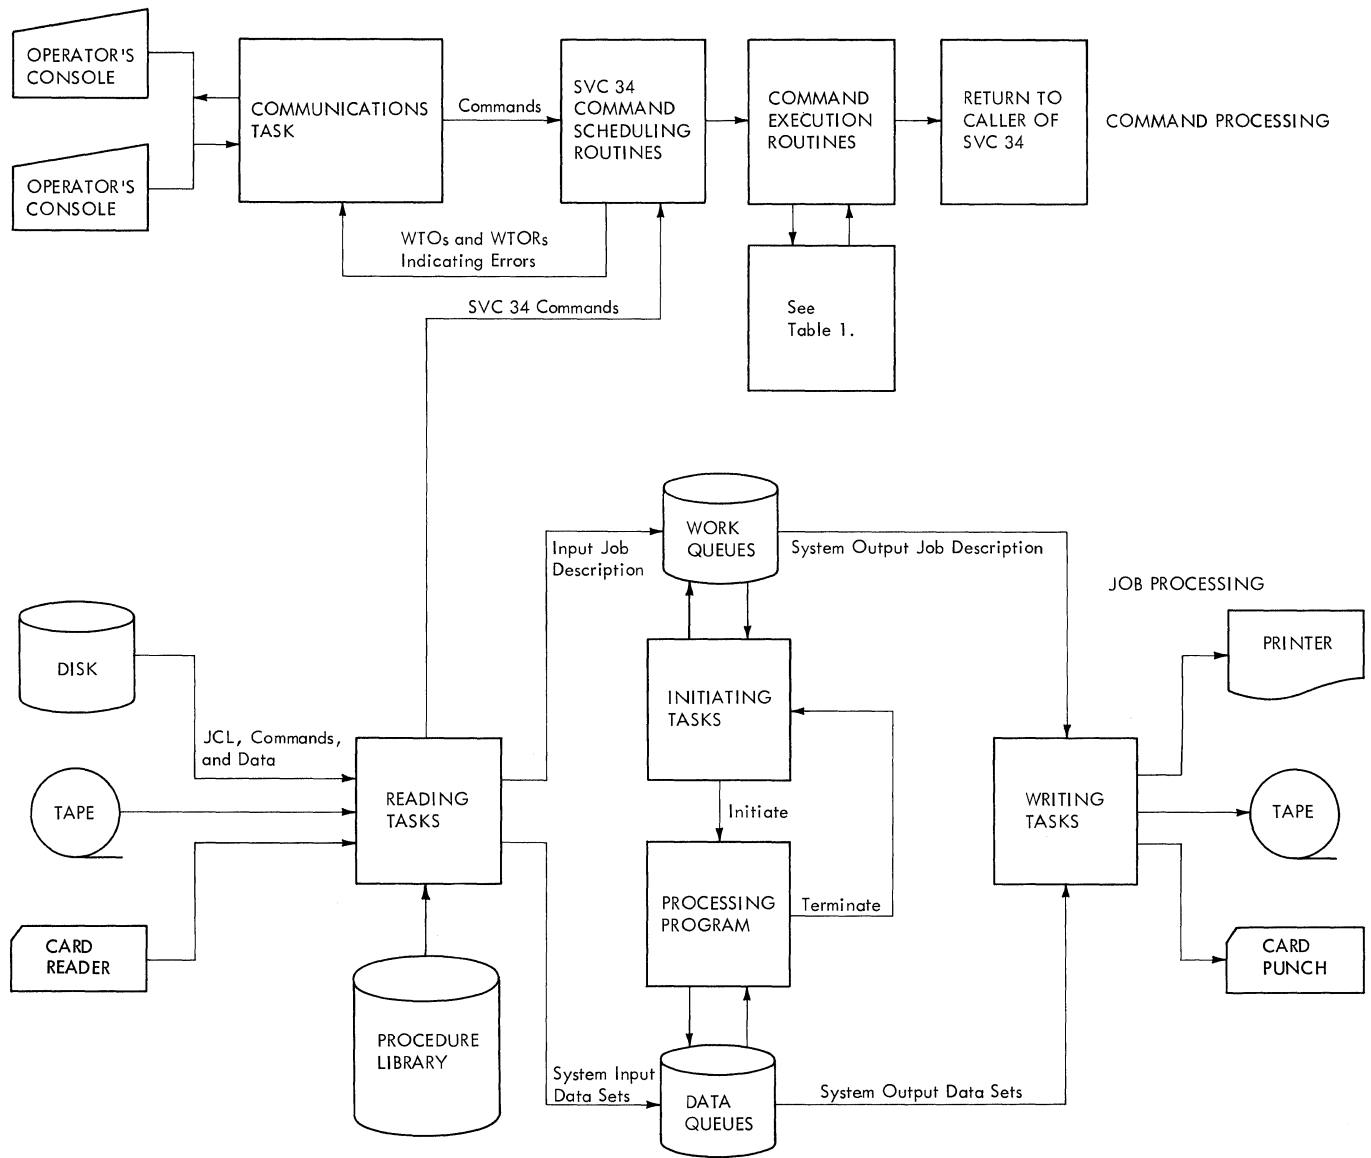
Figure 12 shows the major components of job management and the general flow of control.

Control is passed to job management whenever the supervisor finds that there are no program request blocks in the request block queue. This can occur for two reasons: either the initial program loading (IPL) procedure has just been com-

pleted, or a job step has just been executed.

Entry to Job Management Following Initial Program Loading

Following IPL, certain actions must be taken by the operator before job processing can begin. Therefore, control passes to the communications task which issues a message to the operator instructing him to



•Figure 12. Job Management Data Flow

enter commands, or to redefine the system. If he chooses to redefine the system, control passes to the master scheduler task to handle the redefinitions. If the system is not to be redefined, the initialization commands (a SET command, a START reader command, a START writer command, and a START INIT command) are issued (either automatically by the master scheduler task or by the operator performing the IPL), and job processing begins.

Entry to Job Management Following Step Execution

Following step execution, control is passed to the step termination routine of the initiator/terminator. If no further job steps are to be processed, control is also

passed to the job termination routine of the initiator/terminator. Both routines are described in the topic "Initiator/Terminator."

MFT job management is similar in many respects to MVT job management. However, certain major differences in logic exist. These differences are described in two major topics. "Command Processing" includes the communications task and master scheduler task. "Job Processing" includes:

- Queue Management.
- Reader/Interpreter.
- Initiator/Terminator.
- System output writer.
- System task control.
- System restart.

References to IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660 are made in the topics where the logic is the same as in MVT.

Tables and work areas used by MFT, MFT module descriptions, and MFT flowcharts are included in the appendixes.

Command Processing

Operator commands control system operation and modify system tasks. Command processing in MFT is handled by the communications task and the master scheduler task. With the exception of DEFINE, and HALT, commands can be entered into the system through the console or the input job stream. The DEFINE and HALT commands can be entered only through the console. Commands entered through the console are read by the communications task and routed to the master scheduler (see Figure 13). The communications task also communicates between the system and the operator; it handles WTO/WTOR macro instructions, assigns message identifiers (including partition numbers), and maintains reply queue elements. It also deletes messages from the CRT display of the Model 85 operator console via the DOM macro instruction.

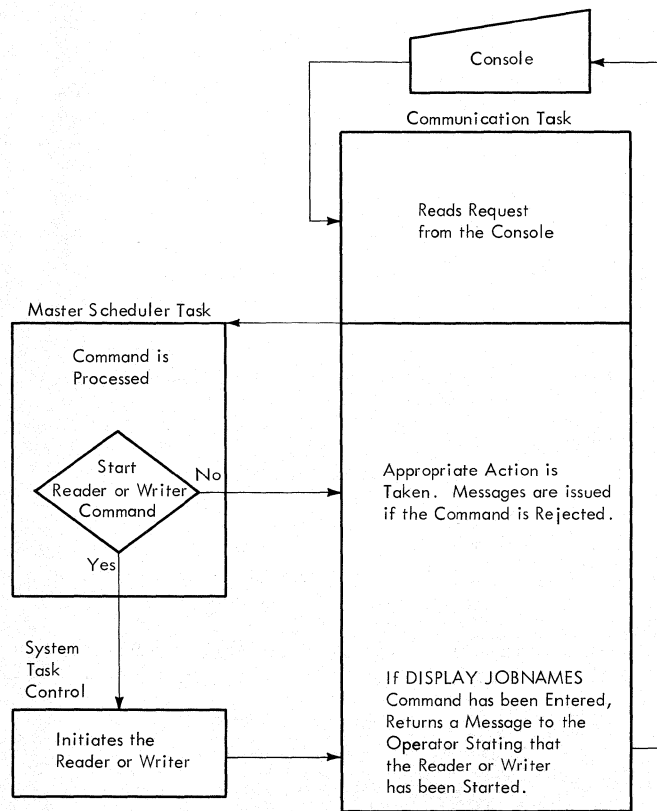


Figure 13. Command Processing Flow

When a command is encountered in the input stream, the reader/interpreter passes control to SVC 34 to process the command. SVC 34 processes most commands completely and returns control to the interrupted routine.

The commands accepted and processed by MFT are the following:

CANCEL
 DEFINE
 DISPLAY
 HALT
 HOLD
 LOG
 MODE
 MODIFY
 MOUNT
 RELEASE
 REPLY
 RESET
 SET
 START
 STOP
 UNLOAD
 VARY
 WRITELOG

The format and syntax of these commands can be found in IBM System/360 Operating System: Operator's Guide, Form C28-6540.

Communications Task

The routines that handle operator-system communication are contained in the communications task. Communication may take either of two forms: commands, which allow the operator to change the status of the system or of a job or job step, and WTO or WTOR macro instructions, which allow problem programs or system components to issue messages to the operator. The communications task routines also switch functions from the primary console device to an alternate console device when the INTERRUPT key is pressed.

WTO/WTOR MACRO INSTRUCTION PROCESSING

Whenever a WTO or WTOR macro instruction is issued, a supervisor call (SVC) interruption occurs. The supervisor identifies the type of interruption and passes control to the communications task to issue messages and/or to read replies. (See Figure 14.)

EXTERNAL INTERRUPTION PROCESSING

When the operator presses the INTERRUPT key, an external interruption occurs. The communications task then switches from the primary console/master console to its alternate device. (See Figure 15.)

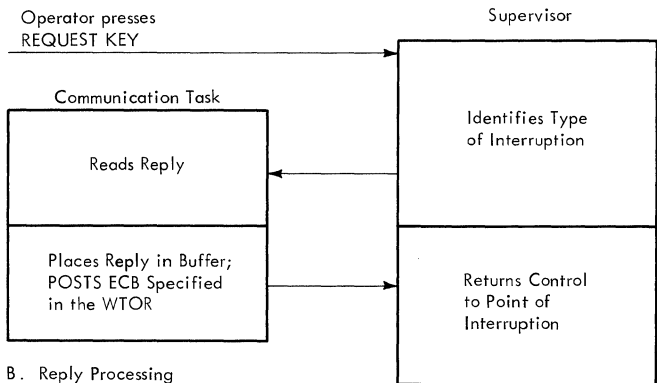
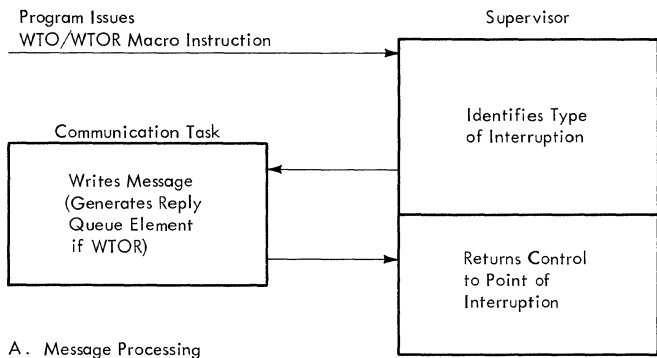


Figure 14. WTO/WTOR Macro Instruction Processing Flow

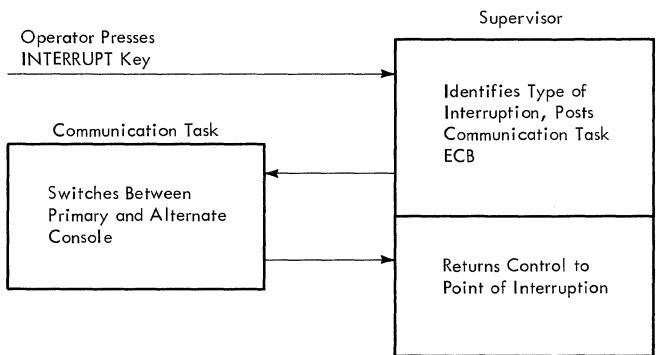


Figure 15. External Interruption Processing Flow

Communications Task Modules

The communications task (Chart 16) receives control through interruptions which occur when commands are entered or messages are written. The following paragraphs describe the seven major routines of the communications task.

Console interruption routine (IEECVCRA): notifies the communications task wait routine that a console read has been requested.

Communications task wait routine (IEE-CVCTW): waits for all WTC/WTCR requests and console interrupts and calls the communications task router routine.

Communications task router routine (IEE-CVCTR): determines the type of request or interruption that occurred and passes control to the appropriate processing routine.

Console device processor routines (IEE-CVPM): performs console read and write operations and error checking.

Write-to-operator routine (IEECVWTO): manages WTO buffers.

Write-to-operator with reply routine (IEE-VWTOR): manages WTOR buffers.

External interruption routine (IEECVCRX): switches to the alternate console device when an external interruption occurs.

Commands are issued through the console device or the input reader. Before entering commands through the console device, the operator must cause an I/O interruption by pressing the REQUEST key. When he does, control is given to the supervisor, which recognizes the interruption and passes control to the I/O supervisor. The I/O supervisor determines that the interruption is an attention signal and passes control to the communications task console interruption routine in the nucleus. The console interruption routine posts the attention event control block (ECB) in the unit control module (UCM) and sets the attention flag in the UCM list entry corresponding to the device from which the interruption came. Posting of the attention ECB causes the communications task wait routine to be dispatched.

The communications task wait routine waits on all communication ECBs associated with WTO/WTOR. The wait routine issues a multiple WAIT macro instruction on a list of ECBs contained in the UCM. When one of the ECBs is posted, as by attention or external interruptions, the wait is satisfied and the communications task thus becomes ready. When it becomes the active task, it issues SVC 72. This SVC includes the console communication service routines and the router.

The communications task serves a number of purposes. The first segment of SVC 72, called the router, distinguishes among these purposes and establishes the order of response. When a posted ECB is found by the router, the router passes control to the specified processor routine via an XCTL macro instruction.

The console-device processor routines read and write using the EXCP macro instruction. The processor routines consist of a routine to service external interruptions and three device-oriented routines: 1052 Printer-Keyboard routine, card reader routine, and printer routine. Each of the three console input/output processor routines is associated with an OPEN/CLOSE support routine, which provides data management and input/output supervisor control blocks. The specified processor routine reads the input message into a buffer area and calls the master scheduler task via an SVC 34.

The write-to-operator routine moves the text from the requesting program's area to a buffer area within the nucleus and posts the communication ECB for write-to-operator.

The write-to-operator with reply routine generates a message ID, including a partition identifier, and creates a reply queue element (RPQE) to handle the operator's reply.

The external interruption routine, residing in the nucleus, switches to an alternate console device when the operator presses the INTERRUPT key on the console.

CONSOLE ATTENTION INTERRUPTION ROUTINE (IEECVCRA)

The console attention interruption routine (IEECVCRA), operating in privileged mode, posts the communications task attention ECB to request reading of the console. Input/output interruptions are disabled without destroying register contents, and without macro access to supervisor services. Using the address of the UCB (found in register 7), the UCB address is matched to a UCM entry. The attention flag for the entry is turned on. Control then passes to the POST routine, indicating the attention ECB in the UCM. The address in register 14 is used for return to the input/output supervisor (IOS).

COMMUNICATIONS TASK WAIT ROUTINE (IEECVCTW)

Upon entry from the dispatcher, the communications task wait routine (IEECVCTW) issues a WAIT (with a count of one) specifying the list of ECBs whose address is contained in the Event Indication List (EIL). Thus the communications task can respond to a variety of events since the posting of any one ECB satisfies the wait. The POST macro instruction issued in the console attention interruption routine satisfies the wait, causing the TCB to be placed on the ready queue. When next dis-

patched, the wait routine issues an SVC 72 which results in creation of a supervisor request block (SVRB), and fetching of the first segment of the console processor routines into the system transient area.

COMMUNICATIONS TASK ROUTER (IEECVCTR)

The communications task router (IEECVCTR) is the first segment of SVC 72 brought into the transient area. Because the communications task serves a number of purposes, and many service requests may be pending, the router establishes the order of response. The order is: external interruption, input/output list completion, attention (console interruption), and WTC/WTCR. Multiple attentions are treated in order of appearance in the UCM. Multiple input/output completions are treated in order of first use of the device. The router responds to an attention by building a parameter list in the SVRB extended save area. The parameter list consists of a remote XCTL parameter list, the address of the appropriate UCM entry, and the address of (contents of CVTCUCB) the UCM. The router then passes control to a processor routine by issuing an XCTL macro instruction to the remote parameter list, using the name obtained from the unit control block (UCB) entry. The flag signifying the request to be serviced by the processor routine is turned off by the routine. Consequently, processor routines return control to the router by issuing an XCTL macro instruction to allow the router to schedule service for other requests. If no requests are pending, the router exits to the wait routine using the address in register 14.

In addition to distinguishing the output request from other requests, the router selects the device to which the message is to be sent. The router establishes the output device by checking UCB entry attribute indicators. The appropriate entry is the first active UCB entry that supports WIO. As before, the router builds a remote interface for, and passes control to, a processor routine via an XCTL macro instruction.

CONSOLE DEVICE PROCESSOR ROUTINES (IEECVPMX, IEECVPMC, IEECVPMP)

Control flow in a processor routine is determined by the setting of flags in the router-selected UCM entry. The close flag is tested first. If this flag is on, any pending input/output activity is suspended by issuing a WAIT macro instruction. Control is then passed to an associated OPEN/CLOSE support routine via an XCTL macro instruction for release of various control blocks. If the close flag is off, the busy

flag is tested to determine input/output status. If there is outstanding input/output activity, error checking and buffer disposition occur if the activity has been posted complete. Otherwise, any attention request is temporarily abandoned (as are output requests), and control returns to the router via an XCTL macro instruction. If the busy flag is off, the attention flag is tested; if it is on, the status of the device is examined. If the device has not been opened, control passes to an associated OPEN/CLOSE support routine via an XCTL macro instruction to obtain storage for a DCB and access-method dependent control blocks, and for execution of the OPEN macro instruction.

When return is made from the OPEN/CLOSE support routine, a response to the attention flag is prepared. A fixed buffer in the UCB is reserved and an access-method dependent interface is constructed. Input/output activity is initiated by issuing an EXCP macro instruction for a 1052, and by issuing a READ macro instruction for a unit record device. In no case does the processor routine await completion of this activity. Control immediately returns to the router via an XCTL macro instruction.

Control flow within the processor routine is as described previously up to the point at which the output request flag is tested. If the flag is on, the processor routine obtains the address of an output buffer from the UCM. The element is not removed from the queue at this time; this occurs only on successful completion of input/output activity. This preserves a means of retrying the message if an external interruption intervenes before the message is successfully presented to the current device. Since output buffers are always selected from the top of the queue, the initiation of output to an alternate device is unaffected by previous attempts to present the message to the primary device.

Having selected a buffer, the processor routine establishes data management and input/output supervisor (IOS) control block linkages. The routine then issues an EXCP macro instruction for a 1052, or a WRITE macro instruction for a printer. Without awaiting completion of the input/output, the processor routine returns to the router via an XCTL macro instruction.

WRITE-TO-OPERATOR ROUTINES (IEECVWTO AND IEEVWTOR)

The write-to-operator routine (SVC 35) writes operator messages on the console when a WTO or WTOR macro instruction is

issued by system component programs or problem programs. Messages and replies are buffered; the period of time between issuing the message and receiving the reply is available for processing. Issuance of either macro instruction causes an SVC interruption. When the SVC interruption is handled, the supervisor causes the write-to-operator routine to be loaded into the transient area of the nucleus and passes control to it.

There are two console queues: the buffer queue and the reply queue. The extent of both queues is defined by specifying the number of buffers at system generation. An attempt to exceed this value results in the requesting task being placed on a queue to wait for service; i.e., the task is placed in a wait condition. Each WTO and WTOR macro instruction results in the addition of a WTO Queue Element (WQE) to the buffer queue; each WTOR results in the addition of a Reply Queue Element (RPQE) to the reply queue. SVC 35 (IEECVWTO) sets up the problem program message. If it is a WTOR, the write-to-operator-with-reply routine (IEEVWTOR) inserts the message identification (ID) in addition to a partition identifier. The same message ID (which the operator must use for his reply) is placed in the RPQE with other information to insure passing the reply, when received, to the proper area. WTO messages are always written; a WTOR message may be purged (removed from the queue) if the issuing task terminates while the message is on the buffer queue. Therefore, an RPQE differs from a WQE in that it contains the address of the issuing task's TCB. The buffer queue is accessed through the entry UCMWTOQ in the UCM.

The reply queue contains RPQEs for operator replies to WTOR. Like WTOR elements in the buffer queue, RPQEs contain a TCB address to permit their being purged from the queue if the issuing task is abnormally terminated.

For a REPLY (to WTOR), the processor issues SVC 34 (see "Master Scheduler Task"). The SVC routine determines that the incoming command is a REPLY, processes the reply, posts the user's ECB and branches back to the processor.

EXTERNAL INTERRUPTION ROUTINE (IEECVCRX)

The external interruption routine assigns functions performed by the primary console device to an alternate console device. When the operator presses the INTERRUPT key on the console, an external interruption occurs and control passes to the supervisor. The supervisor identifies the interruption and passes control to the external

interruption routine which switches consoles and returns control to the supervisor. Console functions may later be reassigned to the primary console device, if the operator causes another external interruption.

Communications Task with Multiple Console Support

The MFT communications task with Multiple Console Support (MCS) is similar to the MVT communications task except that MFT does not obtain buffers dynamically. The MCS communications task receives control as a result of an external interruption, an operator console attention, an I/O interruption for a console, or a WTO (R) or DOM macro instruction. The following paragraphs describe the communications task routines with MCS (for a detailed description of these modules see IBM System/360 Operating System: MVT Supervisor, Form Y28-6659):

Communications Task Router Routine (IEECMAWR): waits for the posting of an external, attention, I/O, WTO(R), or DOM ECB. Control is passed to the appropriate routine to handle the posted ECB, to provide console switching, or to provide buffer management.

Communications Task Device Interface Routine (IEECMDSV): passes control to the device support routine for the device on which I/O is to be performed, or consolidates system and console output queues.

Communications Task Console Switch Routine (IEECMCSW): performs console switching as a result of an external interruption, an unrecoverable I/O error, or a VARY command. It also switches the hard copy log to the master console when both log data sets are full.

Communications Task WTO(R) Routine (IEECMWSV): marks WTO queue elements to appropriate console output queues.

Communications Task DOM Routine (IEECMDOM): marks WTC queue elements on the system output queue to be purged.

Console Device Support Routines: provide read and write functions for the associated console devices.

The following modules remain unchanged with MCS:

Write-to-operator (IEECVWTO)
Write-to-operator with reply (IEEVWTOR)
External Interrupt (IEECVCRX)
Console Interrupt (IEECVCRA)

Note: The routines that support the Model 85 integrated operator's console with CRT display are identical with those used with MVT. For a complete description of these routines, see IBM System/360 Operating System: MVT Supervisor, Program Logic Manual, Form Y28-6659.

Master Scheduler Task

The MFT master scheduler task (MST) processes all commands, and initializes main storage at system initialization. It is composed of the SVC 34 routines and the master scheduler resident command processor routines. SVC 34 processes all commands directly except HOLD, RELEASE, RESET, CANCEL (inactive jobs), DISPLAY (A,Q,N, jobname), WRITELOG, and DEFINE. SVC 34 calls the resident command processor to complete the processing of all but the WRITELOG command. When a WRITELOG command is found, SVC 34 stores it and posts the System Log task ECB.

The master scheduler resides in the nucleus and operates under control of its own TCB. The master scheduler TCB is always dispatchable and is of higher priority on the TCB queue than the TCBs for the partitioned area (the problem program area) of storage. Therefore, when a command is issued, the master scheduler always gains control of the CPU after the communications task for processing the command.

When processing commands, interruptions are disabled so that command processing may be completed before any other interruptions are serviced. Although commands are processed when issued, the command may not take effect immediately. An example of this is the STOP writer command. The master scheduler marks a command scheduling control block (CSCB) which is checked by the writer between jobs. The command does not take effect until the writer completes the job it was processing when the command was issued.

MULTIPLE CONSOLE SUPPORT REQUIREMENTS

In systems that include Multiple Console Support (MCS), a hard copy of all operator and system messages is required when there is an active graphic console or more than one active non-graphic console. Because of this requirement, a system log function is provided which may be specified as the hard copy log. In MFT, the System Log operates under its own TCB created at system generation. The System Log task is the highest priority task in the operating system. The master scheduler routine IEFSD569 calls the log initialization routine IEEVLIN which initializes control blocks and obtains

storage for the Log Control Area and the log buffer. The Log Support routines in an MFT environment function similarly to those in an MVT environment. For a further description of the system log and the Log Support routines with MCS, see IBM System/360 Operating System: MVT Supervisor, Form Y28-6659.

SVC 34 FUNCTIONS

SVC 34 (Charts 13, 14, and 15) is called to process all commands. As previously noted, it processes some of these commands completely and calls the resident command processor to process the remaining commands. The commands processed completely by SVC 34 are:

- START
- STOP
- MODE
- MODIFY
- CANCEL (active jobs only)
- HALT
- MOUNT
- VARY
- UNLOAD
- REPLY
- DISPLAY (JOBNAMES, R, SPACE, DSNAME, T, or STATUS)

For CANCEL (inactive jobs), HOLD, RELEASE, RESET, DISPLAY (A, Q, N, jobname), and DEFINE commands, SVC 34 does preliminary processing before passing control to the resident command processor. If the resident command processor is processing a DEFINE command, SVC 34 will queue all commands until the DEFINE command has been completely processed.

For the log command, SVC 34 issues a WTL (SVC 36) to have the LOG command processed in manner similar to a Write-to-log macro instruction issued from a problem program.

The same routines are used in the MFT Command Processor as are used in the MVT Command Processor with two exceptions. DEFINE, MOUNT, and CANCEL command processing is performed in module IEESD571, and STOP INIT and START command processing is performed in module IEESD561. In addition, the Validity Check Command routine (IEE0403D) passes control to the MFT routines rather than their MVT counterparts when operating in an MFT system. The following paragraphs describe the two MFT routines within SVC 34.

DEFINE, MOUNT, and CANCEL Routine (IEESD571)

This routine processes the DEFINE command by setting the necessary indicators in the

master scheduler resident data area. It then posts the ECB for the resident command processor IEECIR50. This routine also processes the CANCEL command (for active jobs), and the MOUNT command.

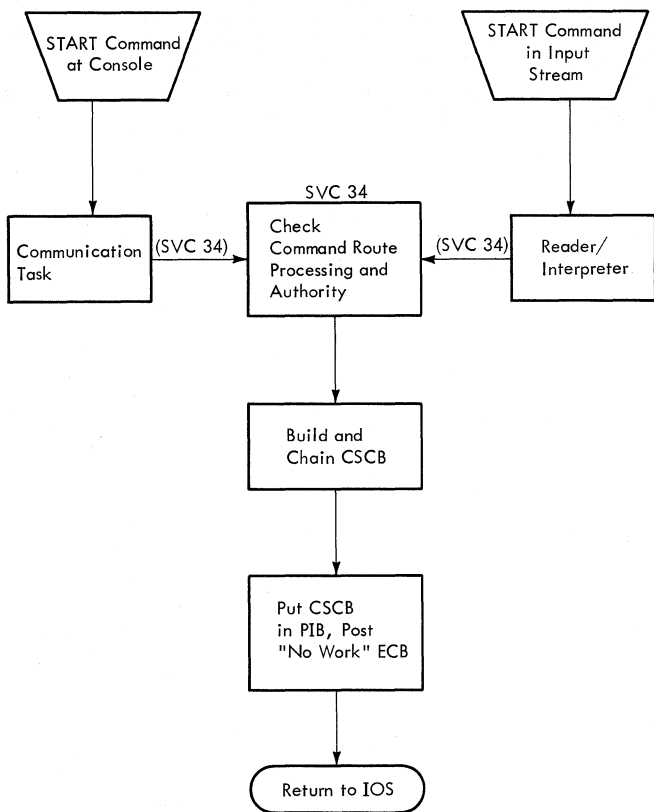
MOUNT processing parallels that of PCP by building a parameter list for, and issuing an XCTL macro instruction to the PCP master command EXCP routine (IGC0103D).

Canceling of an active job is handled by scanning the CSCBs for a jobname compare. If the compare is equal and the CSCB is marked cancelable, IEESD571 issues a BALR to ABTERM with the job's TCB address and proper completion code dump indication. If the CSCB is not marked cancelable, the CSCB is marked canceled and is posted. If the job is not found, IEESD571 passes control to the CSCB creation routine (IEE0803D) via an XCTL macro instruction, to CANCEL the jobname on the job queue. (See IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660, for a description of IEE0803D.)

STOP INIT and START Commands Routine (IEESD561)

This routine processes all the START commands and the STOP INIT command. (The STOP commands that deal with console displays of job names, data set names, and space available are processed by IGC0703D.) When processing a START command, the routine first examines the command parameters. If anything other than a system reader or writer is to be started, the routine determines the number and status of the partition. The routine then builds and chains a CSCB, passes the address of the CSCB to the partition's PIB, and posts the partition. If a system reader is to be started, the routine searches for a scheduler-size problem partition which is inactive; if a system writer is to be started, the routine searches for any inactive problem partition. If a partition is located, the routine builds and processes a CSCB as stated above. If a partition cannot be found, the routine issues a message to the operator stating that the command has failed.

To process a STOP INIT command the routine determines which partition contains the initiator to be stopped, verifies that the partition contains an initiator, and sets the STOP INIT indicator in the partition's PIB. If the routine cannot process the command, it will issue a message to the operator stating that the command has failed.



•Figure 16. START Command Processing Flow

Machine Status Control Routines (IGF2603D and IGF2703D)

These routines process the MODE command which is valid only for the Model 85. The operator obtains a display of machine status by issuing the MODE command with the parameter STATUS. This parameter is processed by machine status control routine IGF2603D. If the parameter is other than STATUS, machine status control routine IGF2703D processes the parameter to allow the operator to control the mode of recording soft machine checks, to reactivate previously deleted buffer sectors or the high speed multiply circuitry, and to restore machine status to the system reset condition. Exit from either routine is via an SVC 3.

The Machine Status Control modules are part of the Master Scheduler but are stored with modules relating to RMS/85 and given the RMS/85 identification of IGF rather than the Master Scheduler identification of IEE.

For a more detailed description of the MODE command, see IBM System/360 Operator's Guide, Form C28-6540. For a more detailed description of the Machine Status Control modules, see IBM System/360 Machine Check Handler for Model 85, Form Y27-7128.

SYSTEM INITIALIZATION

The master scheduler task (Chart 09) performs the function of initializing main storage. In MVT this is done by NIP. In MFT it is done by the master scheduler to facilitate redefinition of main storage. The following paragraphs describe the action of the master scheduler in defining main storage at system initialization.

The master scheduler task is loaded with the nucleus. Its task control block (TCB) points to the master scheduler request block (RB) in the nucleus. NIP saves the RB address and the contents of the boundary box describing the normal master scheduler task partition, for later use by the master scheduler initialization routine IEFSD569. (Note: IEFSD569 is brought into main storage by the macro instruction SGIEEOVV generated during system generation.)

The boundary box (BBX) is then changed by NIP to describe a partition including all of storage except the nucleus. The address of an RB at the low address of this partition is placed in the master scheduler TCB. NIP then creates the RB. The RB points to an XCTL to IEFSD569. NIP then sets the master scheduler task dispatchable and branches to the dispatcher.

The master scheduler initialization routine is given control to perform scheduler initialization. First it passes control to the communications task initialization routine (IEECVCTI) via a LINK macro instruction. After the communications task is initialized, the master scheduler initialization routine passes control to the definition routine, IEEDFIN1, via a LINK macro instruction. IEEDFIN1 communicates with the operator, or prepares the partition as it was described at system generation. IEFSD569 then issues the READY message, and if the system log was requested, passes control to IEEVLIN to initialize the system log. It then types the automatic commands, and issues a WAIT macro instruction.

When the operator presses the REQUEST key, control is given to the supervisor which recognizes the interruption and passes control to the input/output supervisor. The input/output supervisor determines that the interruption is an attention signal and passes control to communications task console attention interrupt routine (described above). The interrupt routine posts the communications task attention ECB to request reading of the console. The operator enters a SET command. SVC 34

posts the WAIT and places the parameters of the SET command in the master scheduler resident data area. The master scheduler initialization routine then regains control to continue processing. Control blocks for the job queue and procedure library are created. To format the job queue, the routine passes control to queue initialization routine IEFSD055 via a LINK macro instruction which places a queue control record (QCR) in the nucleus after the DCB and DEB. Control then passes to queue manager formatting routine IEFORMAT which formats the job queue and returns control to the queue initialization routine. (For a discussion of these two modules, see the topic "Queue Manager".) After return from the queue manager initialization routine, the master scheduler initialization module displays and processes any automatic commands.

The master scheduler initialization routine then establishes partitions based on information in the TCBS. It constructs an RB in each partition, with an XCTL macro instruction addressing job selection module IEFSD510 (for large partitions), or small partition module IEFSD599 (for small partitions). The master scheduler initialization routine then readjusts the pointers to the master scheduler area, and returns to the dispatcher. The dispatcher returns control to the master scheduler task, but the TCB now points to master scheduler router routine IEECIR50, in the nucleus.

Master Scheduler Router Routine (IEECIR50)

Resident master scheduler router routine IEECIR50 waits on an ECB which is posted by SVC 34 when a command has been scheduled for processing. This router (Chart 12) scans the CSCB chain for any outstanding commands to be processed. If a command is found, the CSCB is removed from the chain. The router routine then passes control to syntax check routine IEESD562 via a LINK macro instruction, passing the address of the CSCB.

After all commands are processed, or if none are found, the router routine determines if a DEFINE command has been entered. If so, the router routine passes control to IEEDFIN1, the first module of the definition routines, via a LINK macro instruction. If no DEFINE command has been issued, the router routine returns to wait on its ECB. No test is made for DEFINE command scheduling until all other commands have been processed.

Syntax Check Routine (IEESD562)

Syntax check routine IEESD562 checks the syntax of the command parameter in the CSCB (Chart 10). If a search of the input work

queues (SYS1.SYSJOBQE) is required for processing the command, the syntax check routine sets internal codes for the queue search, issues a GETMAIN to obtain storage, and constructs an event control block (ECB) and an input/output block (IOB). Control is then passed to queue search setup routine IEESD563. If the command was a DISPLAY A command, control is passed to DISPLAY A routine IEESD566. If it was a DISPLAY CONSOLES command, control is passed to DISPLAY CONSOLES routine IEEXEDNA.

Queue Search Setup Routine (IEESD563)

Queue search setup routine IEESD563 determines which of the queues is to be searched and reads the queue control record (QCR) for that queue. If the queue must be searched, the queue search setup routine establishes parameters for the search. The queue search setup routine then passes control to queue search routine IEESD564 via an XCTL macro instruction. When the queue search setup routine regains control, the QCR is scanned and if any information in the record has been changed, the updated information is rewritten on SYS1.SYSJOBQE. The queue search setup routine then establishes a parameter list and passes control to service routine IEFSD565 via an XCTL macro instruction.

Queue Search Routine (IEESD564)

Queue search routine IEESD564 reads the entries of a queue based on the parameter information passed by setup routine IEESD563. If the command processing requires changes in the chaining information in a queue entry or control record, the updated information is written on the queue. Action indicators are passed as parameters when control returns to setup routine IEESD563.

Service Routine (IEESD565)

Based on the information passed by the calling routine, service routine IEESD565 performs the following:

1. Passes control to queue manager enqueue routine IEFQMNQC via a LINK macro instruction to enqueue an entry or QCR.
2. Issues a FREEMAIN macro instruction to free the ECB/IOB which was used to read SYS1.SYSJOBQE.
3. Passes control to the master scheduler message module (IEE0503D) via a LINK macro instruction to write a message.

4. If another queue needs to be searched, it passes control to queue search set up routine IEESD563 via an XCTL macro instruction.

After the requested processing has been performed, the service routine transfers control to router routine IEECIR50.

DISPLAY A Routine (IEESD566)

DISPLAY A routine IEESD566 receives control from syntax check routine IEESD562 when the DISPLAY A (active) command is entered. This routine constructs WTO messages containing the active job and stepnames. The DISPLAY A routine returns control to the router routine.

DISPLAY CONSOLES Routine (IEEXEDNA)

DISPLAY CONSOLES routine IEEXEDNA receives control from the Syntax Check routine IEESD562 when the DISPLAY CONSOLES command is entered. This routine issues a header message that describes the status message. It then constructs and issues a message describing the status of the hard copy log (if one exists) and each console in the system, both active and inactive. When the message is issued, it returns to the Master Scheduler Router routine IEECIR50.

PARTITION DEFINITION BY THE MASTER SCHEDULER

The master scheduler uses the DEFINE command processing routines (shown in Figure 17) to initialize or change partition definitions in MFT. These routines handle:

- Commands from the operator via a console, issued after nucleus initialization, to change the size and description of any partition while processing continues in unaffected partitions.
- Commands from the system at IPL time to prepare the partition as it was described at system generation.

All transfers of control among the processing routines are accomplished via an XCTL macro instruction.

DEFINE Command Initialization Routine (IEEDFIN1)

The master scheduler passes control to DEFINE command initialization routine IEEDFIN1 whenever a DEFINE command is entered

by the operator. The routine also receives control from the master scheduler during system initialization, after the nucleus initialization program (NIP) completes its preparation of the system. In either case the routine builds the DEFINE data area containing the size and description (job classes A-0, or R or W) of each partition. If Main Storage Hierarchy Support is included in the system, the data area contains the size of the partitions in terms of hierarchies. Hierarchy 0 represents processor storage and hierarchy 1 represents 2361 Core Storage.

If the time-slicing feature is included in the system, the data area also contains a doubleword of time-slicing information, including the first and last partition numbers in the time-slicing group and the time interval (in milliseconds) assigned to the group of partitions. This data is used at completion of DEFINE processing to define the partitioning of main storage.

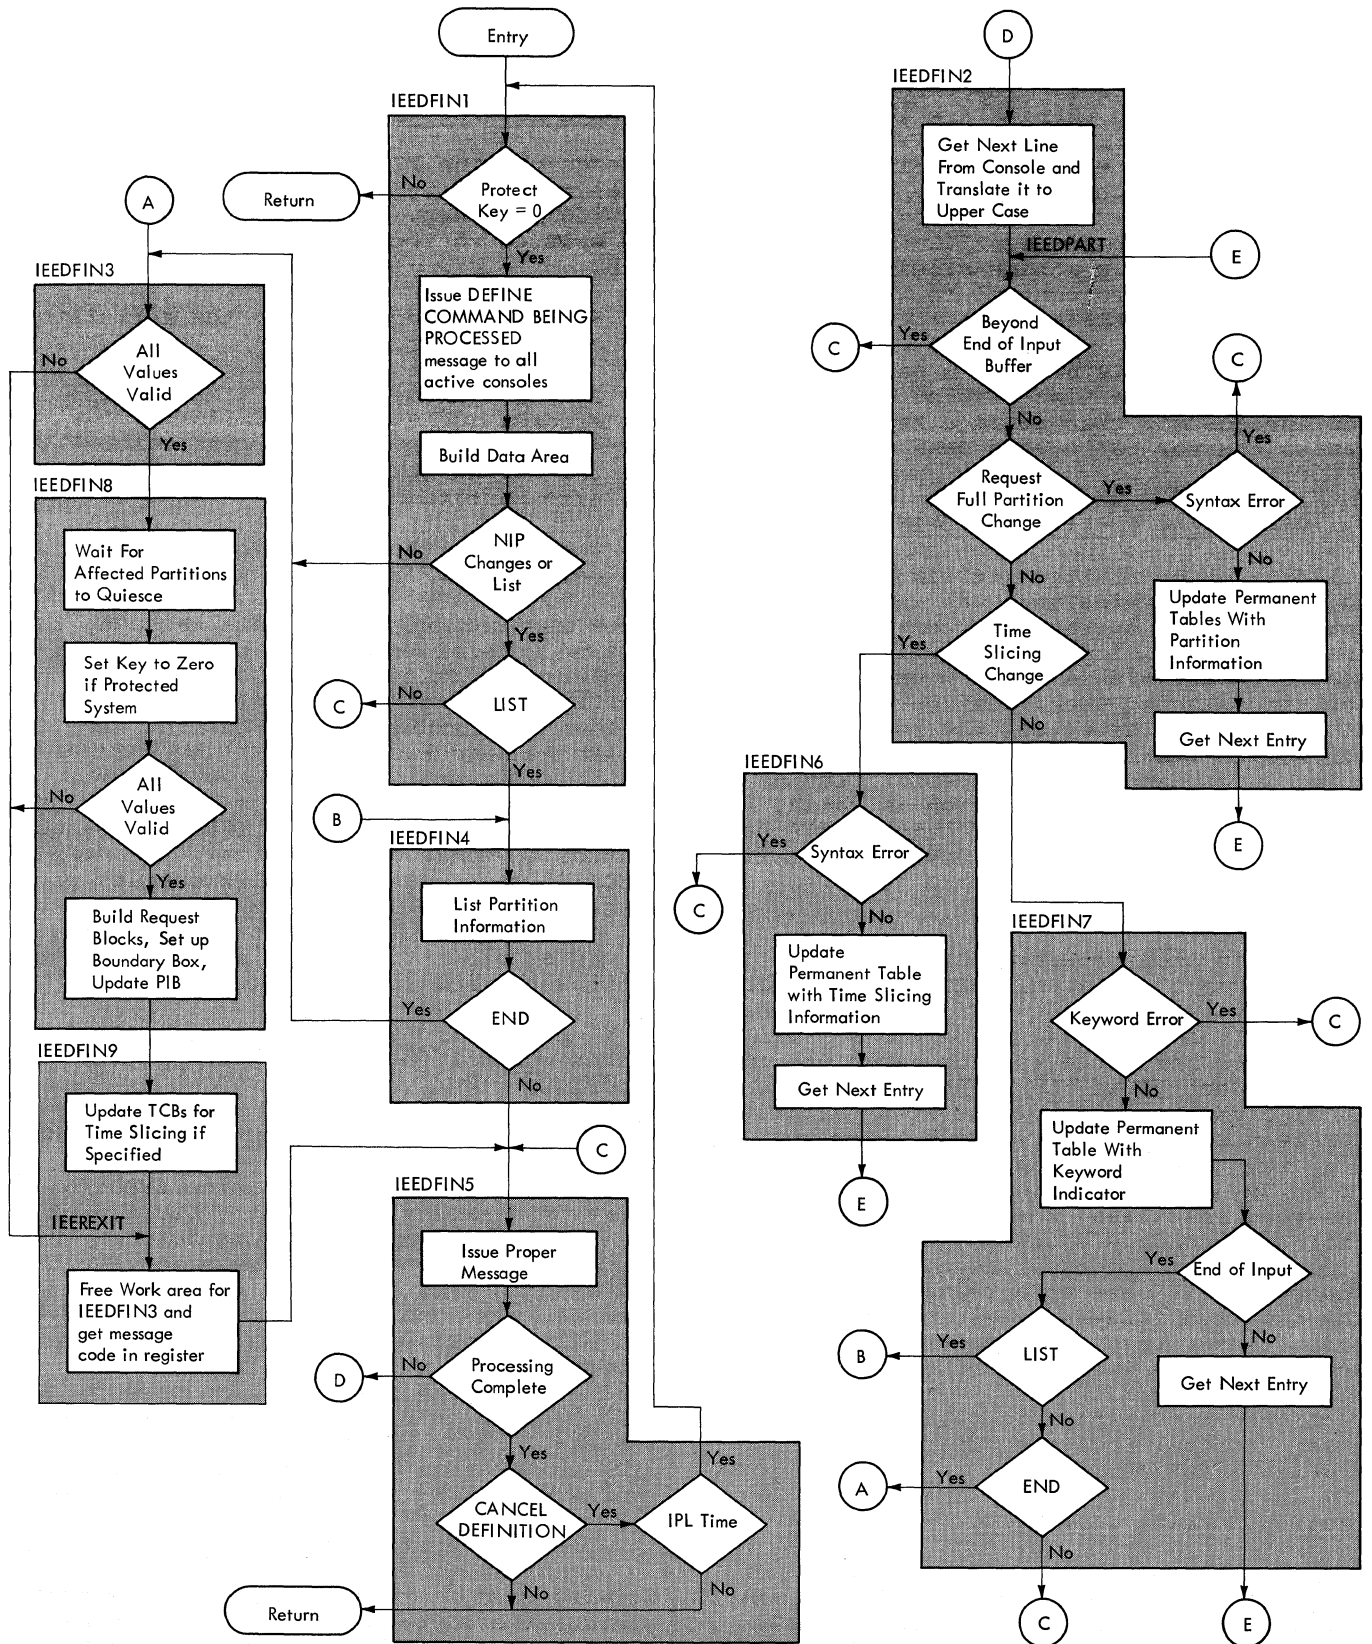
If the DEFINE command initialization routine was entered as the result of a DEFINE command, the routine issues a DEFINE COMMAND BEING PROCESSED message to all active consoles. It then determines whether LIST was specified and if so, passes control to listing routine IEEDFIN4. If not, the routine passes control to message routine IEEDFIN5 for issuance of an ENTER DEFINITION message.

If the DEFINE command initialization routine was entered during the system initialization, the routine also issues a DEFINE COMMAND BEING PROCESSED message to all active consoles. It then determines whether partition redefinition or LIST was specified by the operator, and if not, passes control to validity check routine IEEDFIN3. If either LIST or partition redefinition was specified, the routine continues processing as if a DEFINE command had been entered by the operator.

Syntax Check Routine (IEEDFIN2)

When syntax check routine IEEDFIN2 receives control at primary entry point IEEDFIN2, it translates the statements entered by the operator to upper case. When the routine receives control at secondary entry point IEEDPART, this operation is bypassed.

The statement is scanned and each entry in the statement -- a partition definition, a time-slicing change, or a keyword -- is processed separately.



•Figure 17. DEFINE Command Processing Flow

If the entry is a partition definition, the routine checks the entry for syntax errors. If a syntax error is found, the routine passes control to message routine IEEDFIN5 for issuance of the appropriate syntax error message. The erroneous entry and all following entries are ignored. If the syntax is correct, IEEDFIN2 updates the DEFINE data area with the partition information and gets the next entry for processing.

If the entry is a time-slicing change, the routine passes control to time-slice check routine IEEDFIN6.

If the entry is neither a partition definition, nor a time-slicing change, the routine assumes that it is a keyword and passes control to keyword scan routine IEEDFIN7.

Validity Check Routine (IEEDFIN3)

Validity check routine IEEDFIN3 makes final checks to determine that the information entered by the operator is correct (e.g., that the definition changes which have been requested are within legal bounds or that the time-slicing specification is valid). If an error is detected, the routine passes control to IEEREXIT, a secondary entry point in command final processor routine IEEDFIN9. If the information is valid, the routine determines the partitions affected by the DEFINE command and constructs a list of PIB pointers (one for each affected active partition).

If Main Storage Hierarchy Support is included in the system, IEEDFIN3 determines if a partition has been defined in two segments. If both H0 and H1 size have been reduced to zero, the routine marks the partition inactive in the DEFINE data area. It also checks to determine if a partition has been specified for excess bytes resulting from a redefinition in either H0 or H1 of an adjacent partition. If no partition has been specified, the routine passes control to secondary entry point IEEREXIT in command final processor routine IEEDFIN9. Otherwise, it sets up a message indicating the number of excess bytes, the partition, and the hierarchy to which they have been added. It then passes control to IEEREXIT.

If the information is valid, IEEDFIN3 passes control to system reinitialization routine IEEDFIN8.

Listing Routine (IEEDFIN4)

Listing routine IEEDFIN4 lists partition definitions and job classes. If the time-slicing feature is in the system, it also lists the time-slicing attributes. After performing the listing function, the rou-

tine determines whether an END keyword has been read from the console, and if so, passes control to validity check routine IEEDFIN3. If not, it passes control to message routine IEEDFIN5.

Message Routine (IEEDFIN5)

Message routine IEEDFIN5 handles the messages required by the DEFINE command processing routines. These messages, which are written to the operator, are concerned with:

- Entering and continuing the definition of partitions.
- Syntax, parameter, and time-slicing errors.
- Illegal number of partitions or over-size partitions.
- Completing the definition of partitions.

After issuing the appropriate message, the routine determines whether processing is complete and if so, issues a DEFINITION COMPLETED message to all active consoles. It then determines if a DEFINITION CANCELLED message has previously been issued and if so, tests to see if the system is being initialized. If the message has been issued and it is IPL time, IEEDFIN5 passes control to command initialization routine IEEDFIN1 to repeat the DEFINE command processing. If the DEFINITION CANCELLED message has not been issued, or if it has been issued at other than IPI time, the routine returns control to the caller.

If processing is not complete, IEEDFIN5 passes control to syntax check routine IEEDFIN2.

Time-Slice Syntax Check Routine (IEEDFIN6)

Time-slice syntax check routine IEEDFIN6 checks the time-slicing entry for syntax errors. If a syntax error is found, the routine passes control to message routine IEEDFIN5 for issuance of a PARAMETER ERROR message. It ignores the erroneous entry and all following entries. If there are no syntax errors, the routine updates the DEFINE data area with the time-slicing information, gets the next entry in the statement being processed, and passes control to secondary entry point IEEDPART in syntax check routine IEEDFIN2.

Keyword Scan Routine (IEEDFIN7)

Keyword scan routine IEEDFIN7 determines whether the entry being processed is a valid keyword. If it is not a valid keyword, the routine passes control to message

routine IEEDFIN5 for issuance of a PARAMETER ERROR message. It ignores the erroneous entry and all following entries. If a valid keyword is found, the routine sets the appropriate keyword indicator in the DEFINE data area.

If there are more entries to be processed, the routine gets the next entry and passes control to secondary entry point IEEDPART in syntax check routine IEEDFIN2.

If there are no more entries to be processed (end of input), the routine determines whether a LIST keyword has been entered and if so, passes control to listing routine IEEDFIN4. If LIST was not specified, a check for the END keyword is made. If an END entry is found, the routine passes control to validity check routine IEEDFIN3. If an END entry is not found, the routine passes control to message routine IEEDFIN5 for issuance of a CONTINUE DEFINITION message.

System Reinitialization Routine (IEEDFIN8)

System reinitialization routine IEEDFIN8 places the ECB that must be posted by the affected partition in the PIB of the partition. If a partition has been marked inactive (i.e., no H0 or H1 size is contained in the DEFINE data area), IEEDFIN8 sets the partition's TCB nondispatchable. If any partition being redefined contains a system writer, the routine posts the STOP ECB in the Start Parameter List to stop the writer as if a "Stop Writer" command had been issued from the console. Therefore the operator must issue a "Start Writer" command for any writer partition involved in the redefinition.

The routine then issues the WAIT macro instruction for the posting of the ECB list. When the ECB list is posted, IEEDFIN8 sets the protection key to zero if the system is protected. It makes one final check to determine that no more than 15 problem program partitions have been defined. If an error is found, the routine passes control to secondary entry point IEEREXIT in command final processor routine IEEDFIN9.

If no error is found, IEEDFIN8 uses the information in the DEFINE data area to build request blocks and boundary boxes for the defined partition. The routine then passes control to IEEDFIN9 at its primary entry point, IEEDFIN9.

Command Final Processor Routine (IEEDFIN9)

Command final processor routine IEEDFIN9 updates the time-slice control element and the task control blocks affected by time-slicing if this feature is specified. The

routine then passes control to message routine IEEDFIN5 for issuance of the DEFINITION COMPLETED message.

If the routine receives control from validity check routine IEEDFIN3, it frees the work area obtained by its caller. It passes control to IEEDFIN5 for issuance of the appropriate error message specified by its caller (IEEDFIN3 or IEEDFIN8).

Job Processing

Job processing is accomplished by three types of tasks:

- Reading tasks, which control the reading of input job streams and the interpreting of control statements in these input streams.
- Initiating tasks, which control the initiating of job steps whose control statements have been read and interpreted. (Terminating procedures are also part of initiating tasks.)
- Writing tasks, which control the transferring of system messages and user data sets from direct-access volumes on which they were written initially to some other external storage medium.

These tasks are created in response to START commands entered for readers, initiators, and writers. Whenever a START reader or writer command is entered, the resulting command processing brings a reader or writer into the associated partition. Initiators are brought into all scheduler-size partitions at system initialization, and after a START INIT command has been issued following partition redefinition. An initiator is also brought into a partition that is specified in a STOP INIT command to terminate the initiator.

There may be more than one of each of the job processing tasks so long as the total does not exceed 52. Input job streams may be read simultaneously from three input devices by issuing a START reader command for each input stream. System messages or data sets may be written to as many as 36 output devices by issuing a START command for each device. Up to 15 initiating tasks can exist concurrently. Each initiating task is created in response to a START INIT command issued for a specific partition, or a START INIT.ALL command. (See IBM System/360 Operating System: Operator's Guide, Form C28-6540.)

This section is divided into six topics, including the three major tasks discussed above, and three other areas associated with the major tasks: Queue Manager, System Task Control, and System Restart.

QUEUE MANAGER

MFT uses the MVT Queue Manager. However, to reduce possible interlocks due to unavailability of requested tracks, the assign routine (IEFQASGQ) has been modified, and a new module (IEFSD572) has been added. A table breakup routine (IEFSD514) has also been added to subdivide variable size tables located in main storage into 176-byte data records on disk. Descriptions of some MVT modules have also been included to provide a more complete explanation of the relationship of these modules to the entire system.

WORK QUEUES

An MFT system contains 54 work queues which form the job queue data set (SYS1.SYSJOBQE). These 54 work queues are:

- Automatic SYSIN blocking queue.
- HOLD queue.
- Remote job entry (RJE) queue.
- 36 output class queues.
- 15 input job class queues.

The job entries are enqueued in priority order within each job class on the appropriate job class queue. Jobs are selected for processing according to the job class designation of the partition requesting work.

Queue Manager

Queue Manager is a general term describing a group of routines used by various system components, such as the reader/interpreter, initiator/terminator, and output writer. The queue manager performs some common functions for all system components. It performs all input/output for accessing the job queue data set and keeps track of all space on this queue. The queue manager assigns space on the job queue in logical track increments for control blocks, tables, and system messages built by the scheduler. When the control blocks and tables have been created, the reader/interpreter enqueues (ENQs) the job using the queue manager. After the job is enqueued, the initiator dequeues (DEQs) the job for execution when a partition that is assigned to service that job class becomes

available for work. The terminator places control information needed by the system output writer on the job queue. At job termination, the terminator enqueues the output work description. The writer then dequeues the output work according to output class and priority within the class, and transcribes it to the appropriate device, specified by the user.

At system generation, the space for the job queue data set is allocated. The device upon which the job queue resides is considered a non-demountable system residence volume.

JOB QUEUE INITIALIZATION

At system initialization, queue initialization routine IEFSD055 receives control from the SET command processor to construct a data control block (DCB) in the nucleus, and to issue an OPEN macro instruction which causes a data event block (DEB) to be built for accessing SYS1.SYSJOBQE. It also places a queue manager master queue control record (master QCR) in the nucleus after the DCB and DEB. (See Figure 18 for the format of the master QCR.) Control then passes to queue formatting routine IEFORMAT.

0 (0)	8 byte disk address of the Master QCR MBBCCHHR		8
8 (8)	1 Reserved	2 Displacement of first track of the free queue	1
12 (C)	2 Number of logical tracks in the job queue data set	2 Number of logical tracks in the free-track queue	2
16 (10)	2 Number of tracks reserved for cancelling of job steps when queue full	2 Number of tracks reserved for any initiator	2
20 (14)	2 Displacement of last available logical track	2 Displacement of first track containing only job queue records	2
24 (18)	2 Number of QCRs per physical track	2 Number of job queue records per physical track	2
28 (1C)	2 Number of records per logical track	2 Number of logical tracks for each Problem Program partition	2
32 (20)	2 Number of QCRs on the mixed track	2 Address of first record on first track containing only job queue records	2
36 (24)			

•Figure 18. Master Queue Control Record (Master QCR) Format

The queue formatting routine divides the job queue data set into a control record area and a logical track area. The control record area contains a copy of the master QCR, a control record for the automatic SYSIN batching (ASB) queue, a control record for the HCLD queue, a control record for the Remote Job Entry (RJE) queue, a control record for each of the 36 SYSOUT writer classes, and a control record for each of the 15 input work queues. (See Figure 19 for the format of an input queue control record.)

Note: The first position of the job queue control record (job QCR) contains zeros if no work exists. The job QCR contains a minimum of two entries if work exists for at least one priority.

The job class specified by the user (on the JOB statement or in a START command) is converted by the system to match the system-assigned job class identifiers. The user-assigned job class and corresponding system job class identifiers are:

User-assigned job class	System-assigned identifier (hexadecimal)
A	28
B	29
C	2A
D	2B
E	2C
F	2D
G	2E
H	2F
I	30
J	31
K	32
L	33
M	34
N	35
C	36

The logical track area length is variable. Logical tracks are used instead of physical tracks so that the job queue can reside on different device types. Each logical track contains a 20-byte header record (LTH) (as shown in Figure 20) which includes a pointer to the next track. The header record is used to chain all tracks of a job together. When the job is enqueued, the header record is used to chain jobs first-in/first-out (FIFO) according to priority. All jobs of the same job class are chained together. Following the header record are a variable number of 176-byte data records. The number of records per logical track is determined at system generation and may range

from 10 to 255 records. The number may be modified within this range at IPL. All tables, control blocks, and system messages are in 176-byte increments.

At system initialization, all tracks are members of the free track queue. The free track queue is a list of logical tracks available for assignment to work queues. As tracks are needed, they are taken from the free track queue. When the system is finished with tracks, they are returned to the free track queue. After system initialization, SYS1.SYSJOBQE appears as shown in Figure 21. Figure 22 illustrates typical input and output work queues. Each input and output QCR contains the address of the last entry in each priority queue.

QUEUE MANAGER MODULES

As jobs are read into the system, they are placed into each job class queue according to priority (established by the PRTY parameter on the JOB statement). When the reader/interpreter reads a job or establishes a new queue for an output class, it establishes a queue entry. This is done by Assign/Start Routine IEFQASGT.

Assign/Start Routine (IEFQASGT)

The Assign/Start routine takes the first track from the available track pool and establishes it as the first track for a job. The queue manager parameter area (QMPA) is updated accordingly. (See IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660, for a description of QMPA.) An IOB and an ECB are created for subsequent input/output operations. The actual reserving of tracks is done by the assign routine, IEFQASGQ.

Note: MFT does not support the track-stacking facility of MVT.

Assign Routine (IEFQASGQ)

The assign routine assigns record space on the job queue, and determines whether the requested blocks can be assigned to the current track. If so, the record addresses are placed in the external parameter list of the QMPA, and the records-available field of the QMPA is decremented to reflect this assignment. If additional logical tracks must be assigned, this routine issues an ENQ macro instruction on the master QCR to prevent concurrent access by other tasks. The master QCR is read into main storage.

0 (0)	Address of last LTH of highest priority entry on queue.		2	14	2
4 (4)	13		2	12	2
8 (8)	11		2	10	2
12 (C)	9		2	8	2
16 (10)	7		2	6	2
20 (14)	5		2	4	2
24 (18)	3		2	2	2
28 (1C)	1		2	0	2
32 (20)	Hold Queue	Highest Priority	1	Address of ECB for first task requesting work	

Addresses of last LTH of latest entry having indicated priority.

Figure 19. Job Queue Control Record (QCR)

0 (0)	Reserved				4
4 (4)	Reserved				4
8 (8)	Reserved	1	First logical track of the job	2	Reserved
12 (C)	Next logical track of the job		2	Number of tracks assigned	1
16 (10)	Reserved	1	Jobclass of the job	1	Type*
20 (14)	Last logical track of the next job				2

interpreter taking all the space and making it impossible for jobs to be initiated or terminated, two limit values have been added: the number of tracks reserved for initiating a job, and the number of tracks reserved for terminating a job.

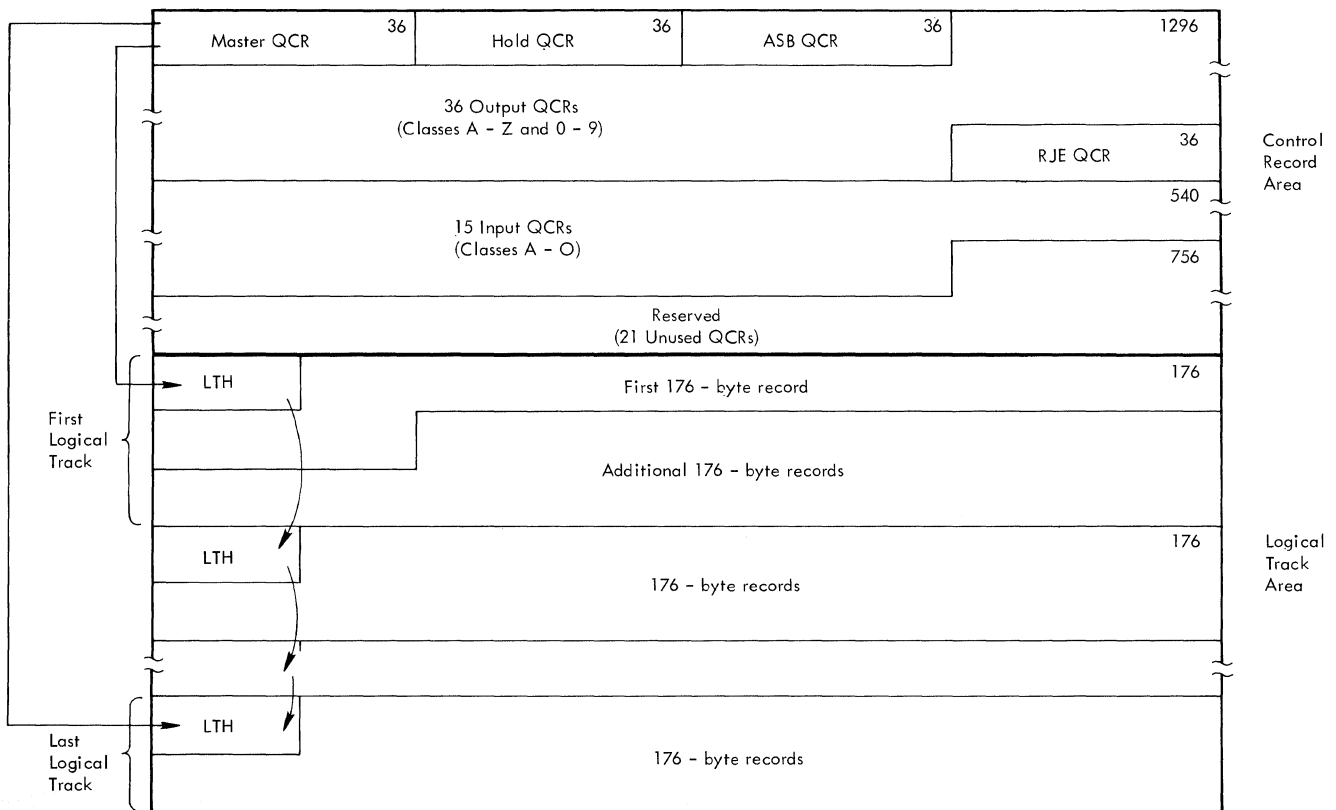
If logical tracks are available, the requested tracks are acquired. The address of the first available logical track is updated and the newly assigned tracks are chained to the tracks assigned to the job. The master QCR is written to the control record area of the job queue data set. A DEQ macro instruction is issued to make the master QCR available to the next user.

If there are no available logical tracks, and the requesting routine is a reader/interpreter, the assign routine passes control to queue manager/interpreter interlock routine IEFSD572. If the reader/interpreter is resident, control returns to the assign routine to wait for tracks to become available. If the reader/interpreter is transient, IEFSD572 issues a message to the operator requesting him to reply "WAIT" or "CANCEL". If the reply is WAIT, control returns to the assign routine, otherwise control is passed to the ABEND routines to cancel the reader/interpreter.

- Type :
- 1 = HOLD queue
 - 2 = ASB queue
 - 3-38 = Output class queues
 - 39 = RJE queue
 - 40-54 = Input work queues

•Figure 20. Logical Track Header (LTH) Record Format

The primary user of this assign routine is the reader/interpreter, although the initiator/terminator also uses it. To prevent the possibility of the reader/



• Figure 21. Sample Job Queue (SYS1.SYSJOBQE) Format After Initialization

If there are no available logical tracks and the requesting routine is an initiator/terminator, the assign routine issues a message to the operator stating that queue space has been exceeded and passes control back to the initiator/terminator to cancel the job.

When the requesting routine is assigned the record TTRs, it can read and write records on the job queue. The master QCR is written, and a DEQ macro instruction is issued to make the master QCR available to the next user. The record addresses in storage and TTR pointers are contained in the external parameter list of the QMPA. When available space on the job queue becomes critical, a warning is sent to the requesting task. Logical tracks are removed from the pool of available tracks and assigned to the job.

If the reply is CANCEL, the interlock routine deletes all queue space assigned to the job, cancels the job, and returns control to the assign routine. Normal initiator operation recovers the partition for further use.

Interpreter/Queue Manager Interlock Routine (IEFSD572)

When the reader/interpreter requests tracks for the job it is processing, and no space is available, IEFQASGQ passes control to interlock routine IEFSD572 to identify whether an interlock can occur. If the reader is transient, the possibility exists that space needed by the reader/interpreter can be provided only by the termination routines, which must operate in the partition that the reader occupies. Because the requested space is not available, the routine issues a message to the operator requesting a reply of 'WAIT' or 'CANCEL'. If the reply is WAIT, this routine returns to the assign routine to wait for available space. (If the reader requesting space is a resident reader, no message is issued, and a reply of WAIT is assumed.)

If the reply is CANCEL, control passes to delete routine IEFQDELQ to delete all queue space assigned to the job being processed (if any space had already been assigned). When control returns, the interlock routine abnormally terminates the job with a job-canceled code of 222. Normal initiator operation recovers the partition for further use.

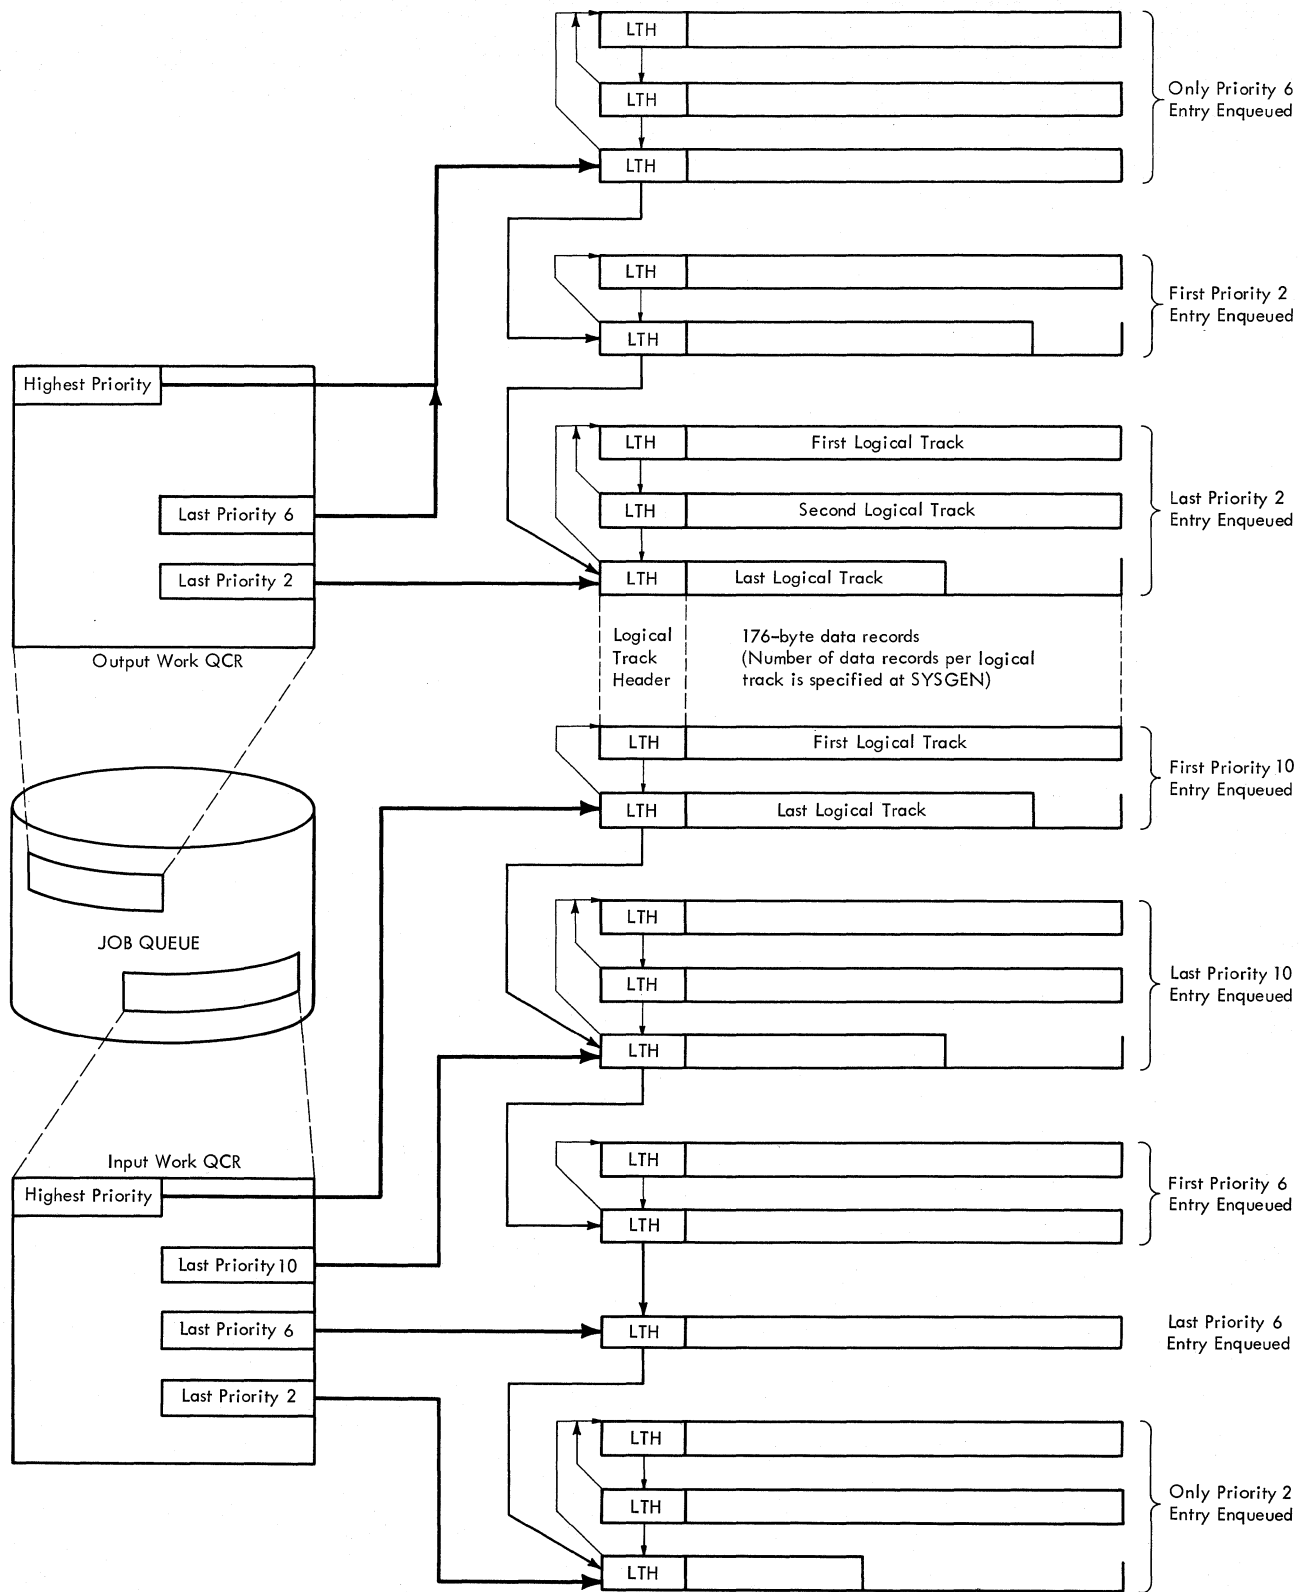


Figure 22. Input and Output Queue Entries

Queue Manager Enqueue Routine (IEFQMNOQ)

After all control blocks for a job have been written, the job is eligible for selection by an Initiator. Declaring a job ready for selection (enqueueing) is done by Queue Manager Enqueue routine IEFQMNOQ.

When an interpreter has completed the processing of a job, (all records generated by the interpreter have been written on the queue), it uses this routine to enqueue the job, in priority order, on the appropriate job class input work queue. When a job completes processing, the terminator uses this routine to enqueue output data sets, in priority order, on the appropriate output work queues.

To prevent concurrent updates, this routine issues an ENQ macro instruction for the queue control record (QCR) of the proper queue. When the QCR becomes available, it is read into main storage. The enqueue routine then places the new queue entry after the last entry with the same priority as shown in Figure 22. The address of the new entry is then placed in the track header of the prior entry (maintaining a chain), and in the QCR position for that priority. The job control table (JCT) is written. The updated QCR is written on the job queue. A DEQ macro instruction is issued making the QCR available. Control is then returned to the calling routine.

Dequeue Routine (IEFQMDQQ)

In addition to dequeuing a job from the input queue for an initiator, the dequeue routine (IEFQMDQQ) removes the output data from an output queue for processing by a system output writer.

The routine issues an ENQ macro instruction on the QCR of the selected queue. When the QCR becomes available, the dequeue routine reads it into main storage. The QCR is examined for a job belonging to the same job class as the partition. Upon finding a job, this routine adjusts the chain. If none is found, the requesting task tries the next job class. If no work is found on any of the selected queues (up to three), the requester places itself in a wait state. In the case of an output writer, a pointer to the "no work" ECB is placed in the QCR. If a pointer already exists, the ECB is chained to the last ECB waiting for that output class. Then the updated QCR is written and a DEQ macro instruction is issued making the QCR available.

Once a job has completed processing, or the output writer has written all records for a job, the tracks are returned to the system. This is known as deleting a job

and is handled by the queue manager delete routine IEFQDELQ.

Delete Routine (IEFQDELQ)

The Delete routine first issues an ENQ macro instruction on the master QCR of the free chain of tracks. After control is returned, the record is updated to reflect the new available tracks. The prior last track of free storage is updated to point to the new set of free tracks. After the master QCR is updated, it is written and a DEQ macro instruction is issued against it. The ECB indicating wait-for-space is posted.

Table Breakup Routine (IEFSD514)

When a reader must be suspended, the job scheduler must prevent the destruction of variable size tables in main storage. To do this, it calls the queue manager table breakup routine, IEFSD514, (Chart 10) which subdivides tables in main storage and writes them on disk as 176-byte data records. The data records are written in a queue entry related to the caller. The job scheduler calls IEFSD514 to retrieve the 176-byte data records and to reconstruct the tables in main storage. Whether reading or writing tables, the caller must build a parameter list (see Figure 23) and place the address of the list in general register 1 before calling the TBR.

When the tables are written initially, the TBR parameter list must contain the address of a QMPA specifying the queue entry into which the tables are to be written. The function code field (QMPCP) of QMPA must specify a write operation. The TBR parameter list must also contain the address, subpool, and size of each table to be written. The last word of the TBR parameter list must be zero. The TBR returns a Head TTR address which locates the beginning of the tables on disk. This TTR must be saved for subsequent retrieval of the tables.

The initial write establishes disk data records for the tables for the duration of the associated queue entry (i.e., until the entry is deleted). Therefore, further write requests must specify the Head TTR in the TBR parameter list. Before issuing a write request, the caller must retrieve any previously written tables to prevent their being overlaid by the new write request.

If the request is for output of tables, (transferring from main storage to direct-access device), the Head TTR (passed in the parameter list) is used to read the first table queue control record (TQCR). If the Head TTR is zero, the assign routine, IEFQASGQ, is called to assign space for a new

TQCR. The TQCR is a 176-byte record containing a 4-byte forward-chain pointer and space for 43 TTRs. These spaces are filled in as the tables are written, using the assign routine to assign the TTRs, and the Read/Write routine, IEFQMRW, to write the tables in 176-byte segments. If more than 43 records are required to hold the tables, a new TQCR is chained to the first, and processing continues. The low-order byte of the last TTR used in writing the tables is set to 'FF' (hexadecimal) to indicate end-of-tables. After these TTRs are assigned, they are used each time the table breakup routine is called to write tables, as long as the Head TTR is preserved by the caller.

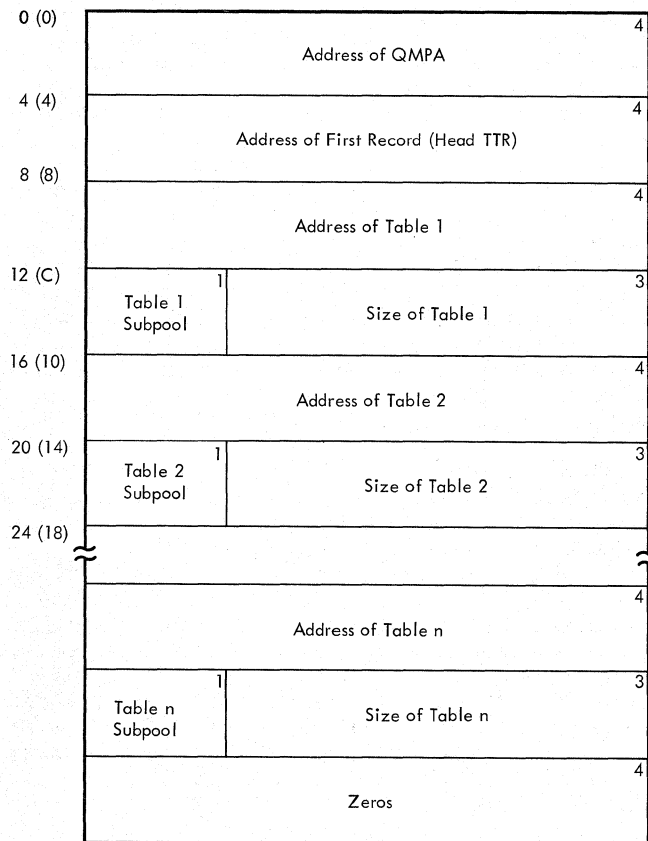


Figure 23. Table Breakup Parameter List

Once a queue entry has been deleted, a caller must issue another initial write request (Head TTR is zero in the table breakup routine parameter list) to establish a new string of table data records. IEFSD514 does not free table storage areas.

In retrieving tables, the TBR parameter list must contain the address of an associated QMPA. The function code (QMPOP) field must specify a read operation. The TBR parameter list must also contain the Head TTR address. Sufficient space must be allowed for the TBR to return the new main

storage address of each table, and the subpool and size of each table as specified when they were written by the TBR.

If the request is for input (reading into storage) of tables, the first TQCR is read into storage using the Head TTR passed in the parameter list. The first record of the first table is read, using the first record in the TQCR. This record contains the size of the table and the number of the desired subpool. IEFSD514 issues a GETMAIN specifying the subpool and the amount of storage required for the table. The remainder of the table is then read into the storage obtained, using read/write routine IEFQMRW. Each table specified in the parameter list is processed in this manner until 'FF' (hexadecimal), indicating end-of-tables, is found. As each table is read into main storage, the parameter list is updated with the main storage address of that table. When all tables have been read, control is returned to the caller. The address of the updated parameter list is returned in register 1. Tables are always written in the same sequence that they appear in the TBR parameter list, beginning with the Head TTR. They are retrieved, totally, in the same sequence; they cannot be read selectively.

Reader/Interpreter

MFT uses the MVT reader/interpreter (reader). However, because of job class, possible MFT interlocks, and the capability of using transient readers, some modifications have been made to the MVT modules, and six new modules have been added. These modifications and additions are described below.

MFT allows as many as three input readers to execute concurrently with problem programs and writers. Resident readers operate in previously defined reader partitions, and transient readers operate in problem program partitions large enough to accommodate them. Input stream data for the step being read is transcribed onto direct-access storage where it is held until execution of the associated job begins. Problem programs retrieve this data directly from the storage device.

In MFT there are three types of system input readers:

- Resident reader.
- User-assigned transient reader.
- System-assigned transient reader.

Resident and transient readers may operate in the same system, provided no more than one system-assigned reader is specified, and the total number of readers does not

exceed three. The primary difference between the user-assigned and system-assigned transient readers is the manner in which the transient reader resumes operation after it is suspended.

RESIDENT READERS

A resident reader operates in a partition designated as such at system generation (by replacing the job class identifier with R), or during system initialization or partition definition (by specifying RDR for the job class identifier). A resident reader reads its input stream, enqueueing jobs until the input stream reaches end-of-file or until it is terminated by a STOP command entered for that partition.

Note: The STOP command does not take effect until the current job is completely read.

TRANSIENT READERS

A transient reader operates in a problem program partition large enough to accommodate it. A transient reader can be terminated by issuing a STOP command or by reaching end-of-file, as can the resident reader. In addition, a transient reader is suspended when a job is enqueued either for the partition occupied by the reader, or for a small partition. (Note that this is possible only when a reader completes reading an entire job.)

If a transient reader is started in a specific partition by including the partition assignment in the START command, it always resumes operation in that same partition, and only when that partition becomes free. This type of transient reader is referred to as user-assigned. If 'S' is substituted for the partition number in the START command, the system assigns the reader to any available large problem program partition. This type of transient reader is called system-assigned.

READER CONTROL FLOW

After a START command is entered to activate a reader, master scheduler routine IEE-CIR50 determines if the size of the requested partition is large enough, and posts the partition. Job selection routine IEFSD510 determines that a START command has been entered, and passes control to system task control (STC) syntax check routine IEEVSTAR. The syntax check routine validates the syntax of the START command, builds job control language tables, and retrieves the reader cataloged procedure

specified in the START command. Each reader is assigned to an input device specified in the START command. Control is then passed to interface routine IEFSD533 which sets up an interpreter entrance list (NEL) for a reader. It also allocates job queue space for a transient reader by issuing a dummy WRITE macro instruction. Control is then passed to linkage routine IEFSD537 which issues a LINK macro instruction to reader initialization routine IEFVH1 to begin reading the input job stream (Chart 24-26).

When reader initialization routine IEFVH1 receives control, it reads its input stream using QSAM, and translates job processing information into convenient form for subsequent processing by an initiator and system output writer. Each job read in by the readers is converted into tables that are placed in the appropriate job class input work queue specified by the CLASS parameter on the JCB statement. One input work queue exists for each of the fifteen problem program job classes (A through O).

For systems that include Multiple Console Support (MCS), the PARM field on an EXEC statement includes a command authority code. This code is included in the option list created by interface routine IEFSD533, and placed in the interpreter work area (IWA) by reader initialization routine IEFVH1. This code is passed by the reader when it issues an SVC 34 due to a command read in the input stream.

After the reader has completed reading a job, control passes to queue manager enqueue routine IEFQMNQO which enqueues the job on the appropriate input work queue according to the PRTY parameter on the JOB statement (see "Queue Management" in this section).

Note: If the reader is being used as a subroutine by a problem program, it does not enqueue the job on the input work queue, but returns control to the problem program passing the addresses of the JCT constructed for that job, and the QMPA associated with that input queue entry.

If data is encountered in the input stream, control is passed to interpreter CPO routine IEFVHG to transcribe the data onto direct-access storage for later retrieval by the problem program. If there is no space for the data, control passes to interpreter operator message routine IEFSD536 to issue a DISPLAY active command and a WTOR message. The operator replies with either 'WAIT' or 'CANCEL'. If 'WAIT' is specified, the reader waits for space to become available. If 'CANCEL' is specified, the reader is canceled and a READER

CLOSED message is issued. IEFSD536 then sets indicators which cause cleanup of the current job, and control to be passed to interpreter termination routine IEFVHN to terminate the reader.

After a reader enqueues each job, control passes to transient-reader suspend tests routine IEFSD532. This routine decides whether to 1) terminate the reader, 2) suspend the reader, or 3) have the reader continue reading the job stream. (The decision to suspend the reader would never be made if the reader is resident.) If the reader is to be terminated, control passes to termination routine IEFVHN. If the reader is to be suspended, control passes to transient reader suspend routine IEFSD530. Otherwise, control returns to job and step enqueue routine IEFVHH to continue reading the job stream.

Transient Reader Suspend Routine (IEFSD530)

When a transient reader is suspended, transient reader suspend routine IEFSD530 (Chart 29) writes the tables and work areas used by the reader onto the work queue data set (SYS1.SYSJOBQE).

The routine closes the reader and procedure library. Data needed to restore the reader is temporarily saved in the interpreter work area (IWA). The IWA is then written to the work queue data set. When a user-assigned transient reader is suspended, the address of the reader space on the work queue is placed in the partition information block (PIB). When a system-assigned transient reader is suspended, the address of the IWA is placed in the master scheduler resident data area (IEFSD568). (See Appendix A for the format of IEFSD568.) The work queue data set is later used by transient reader restore routine IEFSD531 to restore the reader when the assigned partition becomes available after job termination. "No work" ECBs for problem program partitions are posted (see "Job Selection"), and suspend routine IEFSD530 returns control to system task control.

Transient Reader Restore Routine (IEFSD531)

Once a partition is again free for the reader, transient reader restore routine IEFSD531 (Chart 30) receives control and issues a GETMAIN for the IWA, Local Work Area (LWA), reader DCB, and procedure library DCB. The direct-access device address of the IWA is retrieved from the PIB if a user-assigned reader is to be restored, or from the master scheduler resident data area, if a system-assigned reader is to be restored. The IWA is then read in from the job queue. The TIOT is read into storage and the TCB pointer is

updated; other tables and work areas necessary to restore the reader are reset from the information saved in the IWA. The reader and procedure library DCBs are opened and the reader resumes operation to start reading at the point in the job stream where it was suspended. Control is then passed to interpreter routine IEFVHCB to continue reading the job stream.

Initiator/Terminator (Scheduler)

To provide independent scheduling, schedulers operate in any problem program partition of sufficient size. A partition large enough to accommodate the scheduler is referred to as a "large partition." A partition not large enough to accommodate the scheduler is referred to as a "small partition". Within a given large partition, a scheduler operates independently of schedulers in other large partitions. Because small partitions cannot accommodate the scheduler, they rely on large partitions to perform their initiation, allocation, and termination operations. Scheduling for small partitions is described in "Small Partition Scheduling" in this section.

An MFT initiator (Chart 18) dequeues a job (entry) for its partition based on a job class designated for the partition. Once dequeued, the job is scheduled according to the information contained in the entry.

During allocation and termination of each job step, the allocation and termination routines place messages and output data set pointer blocks in a specified output queue. The queue entry is created by the reader/interpreter. (The output queue entry becomes input to an output writer when the job is completed.)

An initiator functions as a control program for the scheduling process, using the allocation and termination functions as closed subroutines. The MFT initiator is composed of the following routines:

- Job Selection
- Small Partition
- Job Initiation
- Data Set Integrity
- Step Initiation
- Problem Program Interface
- Step Deletion
- ENQ/DEQ Purge Routine
- Alternate Step Deletion
- Job Deletion

JOB SELECTION (IEFSD510)

The job selection routine (Charts 19-23) acts as the control routine for the MFT initiator. The routine is brought into all large problem program partitions by the master scheduler at system initialization, by the job deletion routine when a job has terminated, or by system task control when a writer has been scheduled for a small partition or a reader has been suspended.

Job selection first waits on a "no work" ECB in the PIB. This ECB is posted complete by the command processing routines, the job deletion routine, system task control, or the small partition module when a small partition needs scheduler services.

When the "no work" ECB has been posted complete, the job selection routine checks the PIB to determine if a life-of-task (LOT) block exists (see Appendix A for a description of the LOT block). If not, it creates one for the task.

Job selection then checks the PIB for a small partition information list (SPIL) pointer (see Appendix A for a description of SPIL). If one exists, scheduling is performed for the small partition by passing control to IEFSD599. If no SPIL pointer exists, the PIB is checked to determine if the partition is involved in partition redefinition; if the partition is to be changed, the PIB is checked further. If a job is queued on the checkpoint/restart internal queue it is processed; if a restart reader is pending, it is started. If neither exists, no further scheduling is allowed in the partition and the partition can be redefined. (See "Master Scheduler Task".)

If the partition in which the initiator is operating is not part of a partition redefinition, a test is made for a pending Restart Reader command. If no command is pending, a test is made to determine if a system-task reader or writer is to be started. If a restart reader or a system-task reader or writer is to be started, control passes to system task control which initiates readers and writers. If a restart reader is being started, and a user-assigned reader had been rolled out of the partition, the PIB is marked accordingly.

If no small partition is requesting service, no reader or writer is to be started, and the partition is not part of a redefinition operation, a final check is made to determine if a START INIT command has

been issued; if so, job selection attempts to dequeue work from the input work queue (see Figure 24). If a STOP INIT command has been issued, the attempt to dequeue a job is bypassed.

A threshold check is then made to determine if enough logical tracks are available on SYS1.SYSJOBQE to start the initiator. If not, message IEF427I COMD REJECTED FOR INITIATOR 'ident' - INSUFFICIENT QUEUE SPACE is sent to the operator and job selection again waits on the "no work" ECB.

The job selection routine obtains storage for the job control table (JCT) and checks to determine if a job is queued on the checkpoint/restart internal queue. If a job exists, dequeue by jobname routine (IEFLOCDQ) is used to remove it from the hold queue for processing. If no job is on the internal queue, the routine then uses the queue manager dequeue routine (IEFQMDQQ) to obtain work from one of the input job queues according to the job class assignment of the partition. If work is found, IEFQMDQQ constructs a CSCB for the job and an IOB to be used when reading or writing the input queue. The CSCB is constructed in the system queue area and the address of the CSCB is placed in the ICT. The address of the IOB is placed in QMGRI. When a user accounting routine is supplied, the job selection routine sets the ICT fields ICTMWRK and ICTMWRK+4 to zero. These fields are used in calculating the execution time of a job step. Job selection then branches to job initiation routine IEFSD511.

If the search for work for the partition is unsuccessful (i.e., no work has been enqueued for any of the job classes assigned to the partition) tests are made to determine if a transient reader is to be restored in the partition or if a START command has been entered for a system-assigned transient reader. If so, system task control is called. If a reader is to be restored in the partition, job selection passes control to special entry point IEE534SD in system task control.

Command Processing Services

In response to system commands entered in the input stream or from a console, the command processing routines request a service by storing information in the PIB of the affected partition or in the master scheduler resident data area for START and STOP commands issued for system-assigned transient readers and writers. The job selection routine recognizes these requests and takes one of the following actions:

- Inhibits further job scheduling for the partition in preparation for the processing of a DEFINE command. (The DEFINE command can be entered only from a console.)
- Prevents execution of problem programs in large partitions in response to a STOP INIT command.
- Passes control to system task control in response to a START reader or START writer command.
- Schedules problem program execution in response to a START INIT command.

dependent on a large partition while scheduler services are being performed. Scheduling for a small partition is independent of scheduling for other small partitions in the system.

The small partition module interfaces with job selection module IEFSD510 to schedule a problem program, or with system task control to schedule a writer in a small partition. Communication between the small partition module and job selection or system task control is maintained through a small partition information list (SPIL). (The format of a SPIL is shown in Appendix A.)

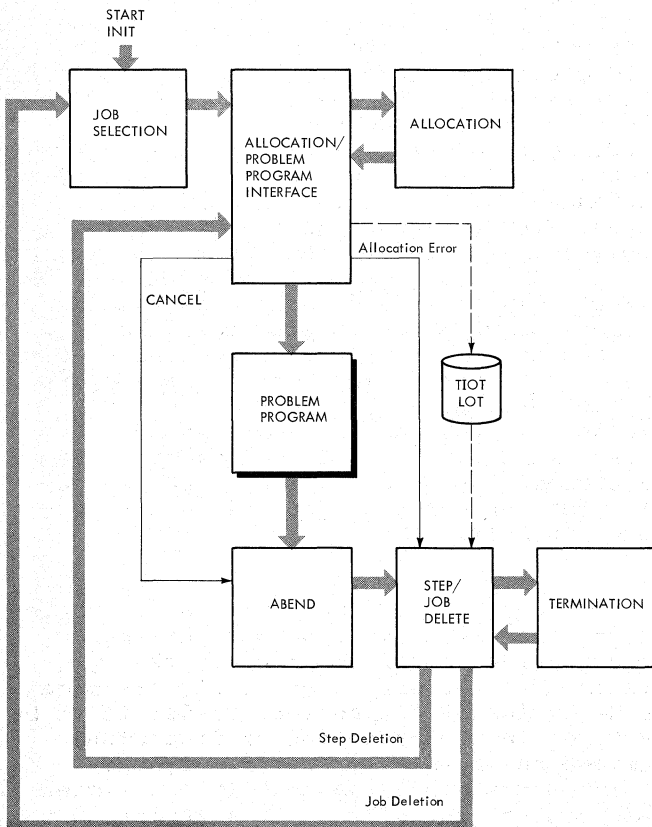


Figure 24. Scheduling a Problem Program in a Large Partition

SMALL PARTITION SCHEDULING

A partition is defined as "small" when its size is at least 8K bytes but less than the job scheduler generated for the system. Small partition scheduling is performed by an initiator in a scheduler-size partition at the request of small partition module IEFSD599 (IEFSD599 is described later in the topic "Small Partition Module"). The small partition is therefore temporarily

Small partition module IEFSD599 requests the scheduling function by placing the address of a SPIL in the partition information block (PIB) of each scheduler-size partition in the system. Each time job selection is entered between jobs, the PIB is checked for a non-zero SPIL address. If the PIB contains a valid address, the SPIL is analyzed, the job class queues for small partitions are searched for work, and control is passed to one of the following:

- Job Initiation (IEFSD511), if work has been found for a small partition.
- Step Deletion (IEFSD515), if a small partition is waiting for termination.
- System Task Control (IEEVSTAR), if a writer is to be started in the small partition.

These routines perform the requested service in the large partition and use the SPIL to indicate their action to IEFSD599. When the requested service has been performed, these routines return to IEFSD510.

Initiating a Problem Program

As shown in Figure 25, initiation of a problem program in a small partition is performed by a large partition. If a small partition is waiting for work, job selection module IEFSD510 dequeues a job from an input work queue that the small partition is assigned to service. The large partition posts a completion code in field ECBA of the SPIL when initiation services have been performed.

A completion code of one indicates that no work was found for the small partition. The small partition then waits on the ECB list in the SPIL. The posting of any of the listed ECBA causes the small partition to request initiation services.

A completion code of zero indicates that initiation services have been performed and the problem program job step is ready to be executed. The small partition, using the allocate parameter list (APL), moves the

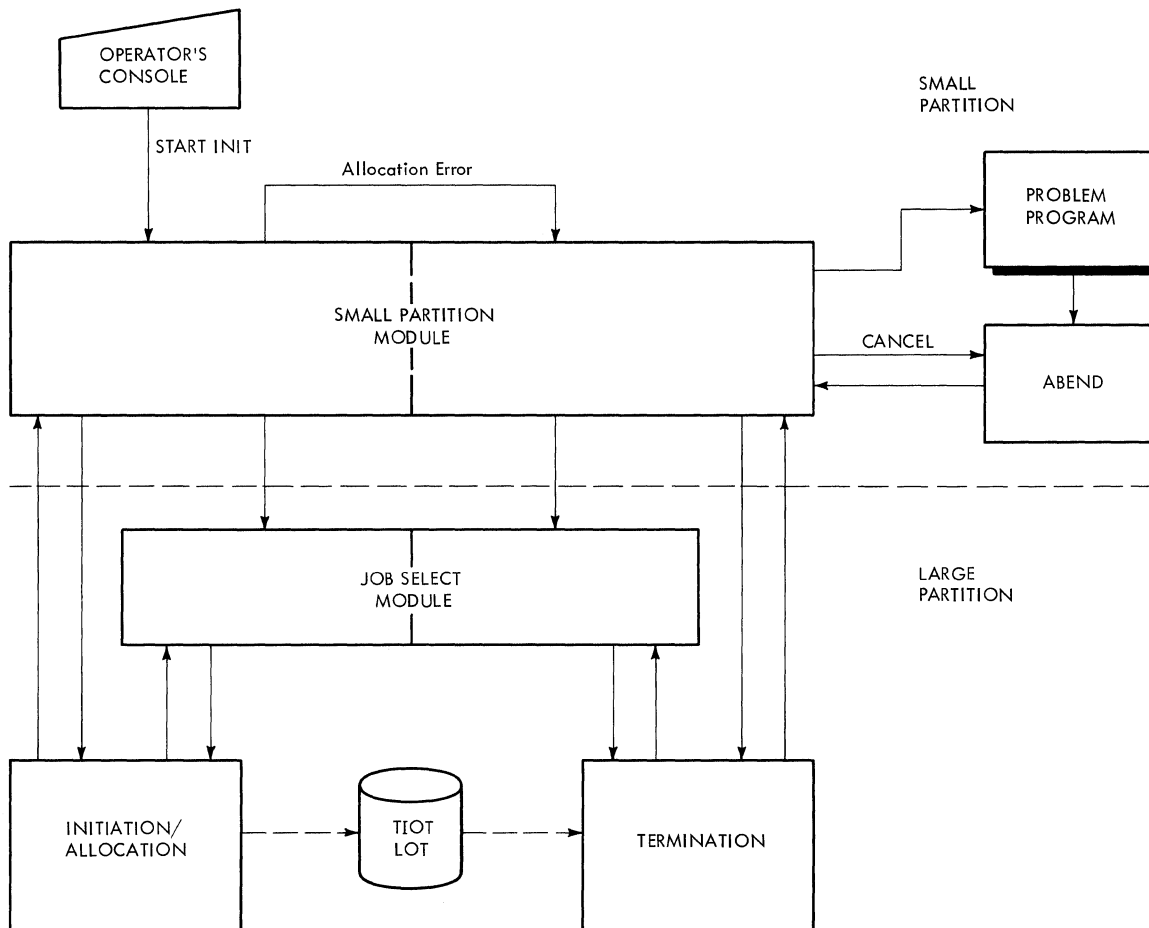


Figure 25. Scheduling a Problem Program in a Small Partition

task input/output table (TIOT) and life-of-task (LOT) block from the large partition, opens required DCBs, and establishes problem program mode. (If the system has the storage protection feature, the protection key is set.) If the job has not been canceled, control passes to the problem program, thus freeing the large partition to continue processing.

Initiating a Writer

As shown in Figure 26, if a writer is to be started in the small partition, small partition module IEFSD599 requests initiation of the writer by system task control. A large partition responds to the request by bringing system task control routine IEEVSTAR into the large partition. IEEVSTAR initiates the small partition to the point of calling in the writer. IEEVSTAR then posts ECBA in the SPIL with a completion code of zero to indicate to IEFSD599 that initiation services have been performed, and the writer is ready to be executed. Small partition module IEFSD599, using the link parameter list (LPL), moves the TIOT

from the large partition to the small partition. ECBC in the SPIL is posted, thus freeing the large partition to continue normal processing. Problem program mode is established, the SPIL is freed, and control passes to the writer via an XCTL macro instruction.

Terminating the Small Partition

When the job step is completed, or a writer is stopped, small partition module IEFSD599 is brought back into the partition and entered at special entry point SMALLGC. A check is made to determine whether a scheduler ABEND occurred. If it did, a message is issued to the operator with a completion code, and all CSCBs associated with that job are removed from the CSCB chain. Control then passes to the normal entry point of IEFSD599. If no scheduler ABEND occurred, the SPIL is created, and a status bit is set indicating that termination services are requested. The small partition module then begins a search for a large partition to perform the job termination services or writer end-of-job processing.

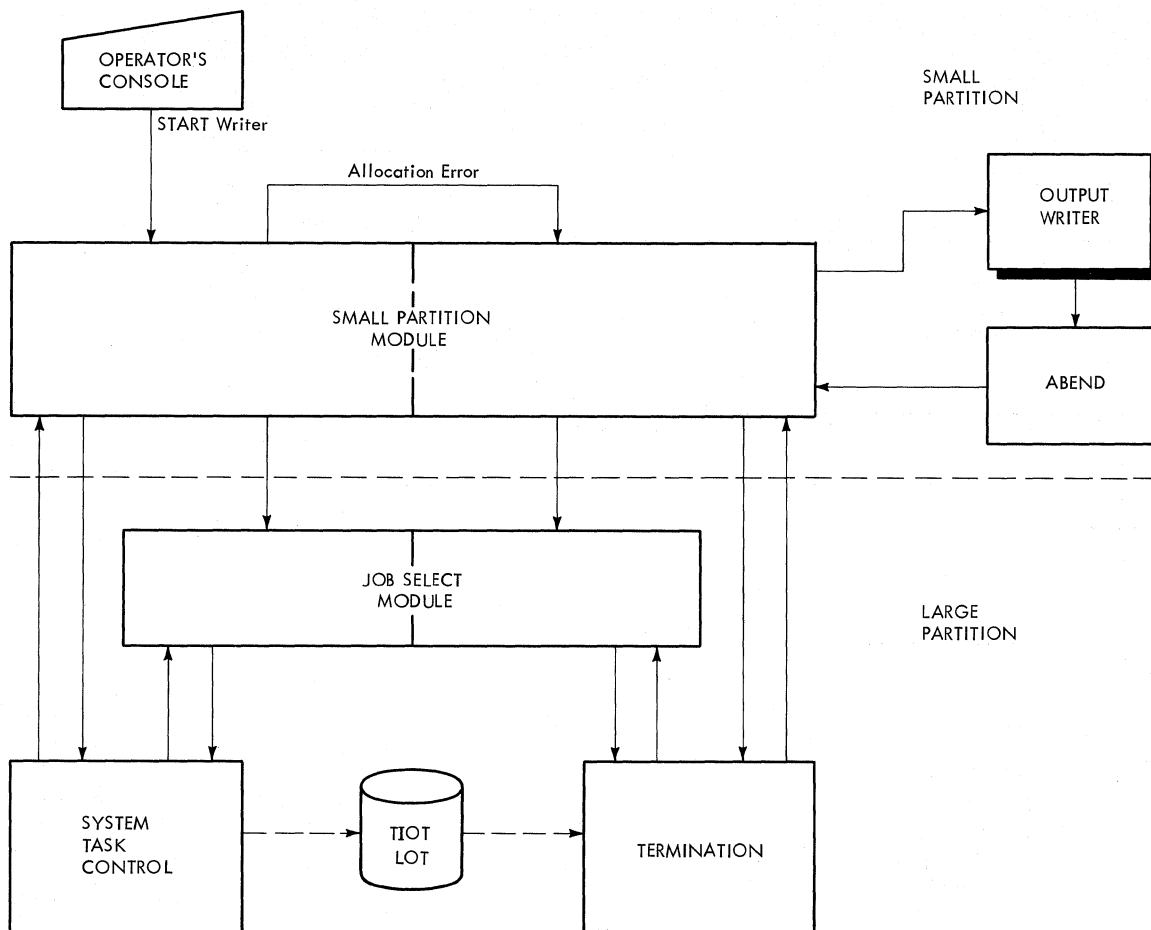


Figure 26. Scheduling a Writer in a Small Partition

After an initiator in a large partition has performed the termination services, ECBA in the SPIL is posted with a completion code of two to indicate that job termination has taken place. A check is made to determine if the small partition is involved in a redefinition operation. If it is, the small partition is made quiescent. If the small partition is not associated with a redefinition operation, it requests additional services from an initiator in a large partition.

Note: If the initiator in a large partition performs step termination instead of job termination, the next step of the job in the small partition is scheduled before the initiator schedules a job into its partition, or before it performs scheduling services for another small partition.

Small Partition Module (IEFSD599)

Small partition module IEFSD599 (Charts 05-08) is entered from the redefinition routines at system initialization or when a DEFINE command is issued or from the master

scheduler. The module is entered at special entry point SMALLGC from the ABEND routines when a step has completed execution. IEFSD599 first waits on a "no work" ECB located in the partition's PIB. When this ECB is posted complete, the PIB is checked to determine if a SPIL has been created. If not, one is created and an indicator is set in the PIB. The PIB is checked to determine if the partition is involved in a redefinition operation. If a redefinition is pending, the internal job queue of checkpoint/restart jobs is checked and any jobs on the queue are processed before the partition redefinition. If there is nothing on the internal job queue and redefinition is pending, assigned tracks are deleted, the SPIL is freed, and pending CSCBs are freed. The 'DEFINE' ECB in the PIB is posted to indicate that the partition has been made quiescent, and a return is made to wait on the "no work" ECB.

If no redefinition operation is pending, the PIB is checked to determine if a writer is to be started in the partition. If so, an indicator is set in the SPIL, assigned

tracks are deleted, and a request for scheduling is made to a large partition (described below). If a writer is not to be started, the STOP INIT bit in the PIB is checked. If this bit is on, assigned tracks are deleted, the SPIL is freed, and a return is made to wait on the 'no work' ECB. If the STOP INIT bit is not on, the PIB is checked for track assignment. If needed, tracks are assigned and indicated in the PIB. The SPIL is updated to indicate a request for initiation of a problem program.

A request is made for a large partition to service the small partition based on the contents of the SPIL. First, an exclusive ENQ macro instruction is issued to prevent concurrent service requests by small partitions. Interruptions are disabled to prevent interference with the address of the SPIL in the large partition's PIB. IEFSD599 then searches for a scheduler-size partition. The TCBS are tested for problem program status; when a scheduler-size partition is found, a determination is made of whether the small partition is involved in a DEFINE operation.

If the small partition is involved in a DEFINE operation, the test for the large partition involved in a DEFINE operation is bypassed. If the small partition is not involved in a DEFINE operation, the large partition is tested to determine if it is involved in a DEFINE operation. If so, the large partition is bypassed and the TCB search is continued.

The address of the SPIL is stored in the PIB of the large partition, thus constituting a request. An indication is made when storing occurs. If a large partition is waiting on its 'no work' ECB (in its PIB), the large partition is posted and the large partition routine clears the SPIL addresses in the other large partition PIBs. When a large partition is posted, or all applicable TCBS are checked, interruptions are enabled.

If no SPIL pointers were stored during the search, a DEQ macro instruction is issued (to allow other small partitions to make requests), and a WAIT macro instruction is issued on a 'dormant' ECB in the small partition's PIB. (When later posted by the command processing routines, the small partition module will repeat its search). If at least one SPIL pointer was stored, a WAIT macro instruction is issued on ECBB in the SPIL. This allows a large partition, immediately upon recognition of the request, to post the ECB complete. The small partition module may then issue a DEQ macro instruction to release the SPIL pointer field so other small partitions may make requests.

Next, a WAIT macro instruction is issued on ECBA (in the SPIL) to delay the small partition until the requested service has been performed. When ECBA is posted complete by the large partition, the completion code is tested to determine the action which occurred. If the completion code is two, job termination occurred and return is made to the point of determining the DEFINE status of the small partition. If the completion code is one, 'no work' was found for the small partition and a return is made to WAIT on the ECB list in the SPIL. If the completion code is zero, the large partition is at the point of calling either the problem program or a writer. The large partition is waiting on ECBC (in the SPIL) to allow transfer of information into the small partition by the small partition module.

If a problem program is to be initiated, IEFSD599 uses the allocate parameter list (APL) to move the TIOT and user parameter area into the small partition. It then posts ECBC (freeing the large partition), and opens Fetch and/or JCBLIB DCBs if required. The partition is established in problem program protection mode. The SPIL is freed. If the program to be initiated is the DSDR processing step of a checkpoint restart, IEFSD599 uses the APL to move the TIOT and user parameter area into the small partition, and posts ECBC. The routine moves the job QMPA and the SYSCUT QMPA from the LOT to the CSCB, and bypasses opening the JOBLIB and FETCH DCBs. The routine also bypasses setting the storage protection key but frees the SPIL.

A check is made to determine if the job has been canceled. If so, an ABEND macro instruction is issued. If the job has not been canceled, an XCTL macro instruction is issued to call the problem program into the small partition (the problem program passes control to ABEND at completion of its execution).

ABEND recalls small partition module and enters at special entry point SMALLGC. The small partition protection key is changed to zero and a SPIL is created. A termination request is indicated in the SPIL, and IEFSD599 begins the search for a large partition to service the request.

If a writer is to be initiated, the control flow is the same as described above in "Initiating a Writer".

INITIATOR/TERMINATOR CONTROL FLOW

There are no terminator routines that are unique to MFT; the modules used in MFT task termination are described in IBM System/360

Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.

In addition to IEFSD510 and IEFSD599, several other initiator routines are unique to MFT. These are described in the following paragraphs. Descriptions of the MVT allocation and step initiation routines that have not been modified by MFT can be found in IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.

Job Initiation Routine (IEFSD511)

Job initiation routine IEFSD511 issues a GETMAIN specifying subpool 0 to obtain space for the system output class directory (SCD). The SCD is then read into the area and the contents of the SCD are used to initialize QMGR2 in the LOT block. (QMGR2 is the queue manager parameter area which is used for referencing the output data set.) After QMGR2 has been initialized, the storage obtained for the SCD is freed. A GETMAIN is then issued to obtain storage for IOB2, the IOB used in conjunction with QMGR2. A GETMAIN is issued (specifying subpool 253) to obtain space for the step control table (SCT). The SCT is read into the area thus obtained. Job initiation then branches to data set integrity routine IEFSD541.

Data Set Integrity Routine (IEFSD541)

The data set integrity routine is entered only once per job, from job initiation routine IEFSD511. It first determines whether data set integrity processing is required.

If the JCT indicates a 'failed' job or if there are no explicit data sets (DSNAME parameter in a DD statement) for the job, processing is bypassed and exit is made to step initiation routine IEFSD512. If data set integrity processing is required, the DSENG table records are read from the job's entry in the input job queue (SYS1.SYSJOBQE). Duplicate DSNAMEs are eliminated from the table and each unique DSNAME is placed in a minor name list. The most restrictive attribute (exclusive or share) is chosen for each DSNAME placed in the minor name list. After this processing is complete, an ENQ supervisor list is constructed which contains an entry for each DSNAME in the minor name list. Each entry is initialized with the following:

- RET=TEST option of ENQ.
- SYSTEM option of ENQ.
- Attribute (E/S) of the corresponding DSNAME.

- Address of the common major name 'SYSDSN'.
- Address of the corresponding DSNAME (considered the minor name) in the minor name list.

The DSNAME (minor name) length is contained in the first byte of each DSNAME field in the minor name list.

When the ENQ supervisor list is constructed, the system is disabled and an ENQ supervisor call is issued against the list to test the availability of the DSNAMEs. If the DSNAMEs are available, the ENQ supervisor list is updated so that each entry reflects the RET=NONE option of ENQ. A second ENQ supervisor call is issued against the list to reserve DSNAMEs for the job. The system is enabled and exit is made to step initiation routine IEFSD512.

If the DSNAMEs are unavailable for the job (already reserved with conflicting attributes by other task(s) in the system), the operator is notified of the condition. In notifying the operator, the return code field of each entry in the ENQ supervisor list is tested for a non-zero setting. If the setting is non-zero, the associated DSNAME (minor name) is identified to the operator as unavailable. The operator is given the following reply options:

- RETRY, in case the resources have been freed by the other task(s) (processing is delayed until the operator replies).
- CANCEL the job.

If RETRY is entered by the operator, processing continues at the initial ENQ supervisor call to again test the availability of the DSNAMEs. The operator is again notified, and he can reply either RETRY or CANCEL. If the job is canceled by the operator, the 'job fail' bit in the JCT is set and exit is made to step initiation routine IEFSD512.

Step Initiation Routine (IEFSD512)

Step initiation routine IEFSD512 first issues a GETMAIN specifying subpool 253 to obtain storage for an allocate register save area (ARSA) and an allocate parameter list (APL). The APL (Figure 27) is initialized containing addresses of the LOT, JCT, and SCT, and two words of zeros.

The step initiation routine checks the current step to determine if it is either the checkpoint/restart data set descriptor record (DSDR) processing step or the restart step. If the step is a DSDR processing step being scheduled for a small

partition containing less than 12K bytes, the PIB of the partition containing the step initiation routine will be tagged to indicate that the DSDR step is to execute in that partition. The step initiation routine will place the address of its TCB and PIB in the LOT and pass control to allocation via a LINK macro instruction. If the DSDR step is to be processed in a large partition, normal processing is continued.

0 (0)	Address of the LCT	4
4 (4)	Address of the JCT	4
8 (8)	Address of the SCT	4
12 (C)	Address of the TIOT List	4
16 (10)	Zeros	4
20 (14)		

Figure 27. Allocate/Terminate Parameter List

If the step is the restart step, the step initiation routine will pass control to partition recovery routine IEFSD518 via a LINK macro instruction. If the return code from IEFSD518 is a zero, normal processing is continued; if the return code from IEFSD518 is a four, the address of the LOT is placed in register 1 and control is passed to job selection IEFSD510 via an XCTL macro instruction.

The step initiation routine then passes control to allocation via a LINK macro instruction. Allocation returns the addresses of a task input/output table (TIOT) list (which points to the TIOT) in the first word of zeros in the APL. On return from allocation, the return code is tested to determine if allocation was successful. If not, step initiation branches to alternate step deletion routine IEFSD516 via an XCTL macro instruction. If allocation was successful, the ARSA is freed, and the "step started" bit in the SCT is turned on. The address of the job's CSCB is stored in the APL (in the last word of the list).

Step initiation then uses queue manager read/write routine IEFQMRW to write the JCT and SCT back on the input queue. The disk addresses of the JCT and SCT are saved in the LCT. A GETMAIN specifying subpool 253 is issued for the table breakup routine

(TBR) parameter list and register save area. The TBR parameter list is initialized with the address, size, and subpool specifications for the TIOT and LOT block. The TIOT and LOT are then written into the job's entry in the job queue, and the Head TTR is saved in the JCT. The storage obtained for the TBR parameter list and register save area, IOB1, and IOB2 is freed. The JCT is then written out. Step initiation then passes control to problem program interface routine IEFSD513 via an XCTL macro instruction.

Problem Program Interface Routine (IEFSD513)

The problem program interface routine prepares the partition for execution of the job step. A test is made to determine if scheduling was performed for a small partition. If so, this routine tests its partition's PIB to determine whether a checkpoint/restart data set descriptor record (DSDR) is to be processed. If the DSDR step is to be processed, the SPIL pointer in the LOT is ignored; otherwise the address of the APL is placed in the SPIL, ECBA in the SPIL is posted to indicate that scheduling is complete, and a WAIT is issued on ECBC. This WAIT allows the small partition module to copy tables and work areas into the small partition. When the tables have been copied, ECBC is posted complete, and the interface routine frees all storage obtained for tables and work areas except for the LOT block, which is retained. The address of the LCT block is placed in register 1 and this routine passes control to job selection, IEFSD510, via an XCTL macro instruction.

If scheduling was not performed for a small partition, a test is made to determine if the job has been canceled. If so, exit is made by issuing an ABEND macro instruction.

If the job has not been canceled, the LOT block is freed, the TIOT is moved to the lowest possible location (subpool 0) in the partition, and a GETMAIN macro instruction specifying subpool 253 is issued for the user's parameter list (UPL). The UPL (Figure 28) is initialized from the SCT. Another GETMAIN macro instruction (subpool 253) is issued to create a register save area for the user's problem program. If STEPLIB, JOBLIB, and/or FETCH have been specified, their DCBs are created (but not opened) in subpool 253. The JCT, SCT, and APL are now freed, the STEPLIB or JOBLIB and FETCH DCBs are opened, and the TIOT is then moved to subpool 253. A single DCB is used for STEPLIB or JOBLIB, with STEPLIB overriding JOBLIB if both are present.

If the job being started in the partition is a checkpoint/restart data set descriptor record (DSDR) processing job, the routine moves the DSDR step QMPAs to the CSCB. The routine bypasses opening the STEPLIB, JOBLIB, and FETCH DCBs and also bypasses setting the storage protection key.




Figure 28. User's Parameter List

Note: The use of subpools, and the order in which control blocks and tables are created, moved, or deleted, follows a particular sequence even though this handling occurs within different modules. This is done to prevent fragmenting main storage within the partition.

After the TIOT has been moved to the highest available position within the partition, the task control block (TCB) is updated and the problem program's protection key is set (if the system has storage protection). The problem program interface routine then passes control to the problem program via an XCTL macro instruction.

Step Deletion Routine (IEFSD515)

Step deletion routine IEFSD515 is entered at the end of step execution to prepare the partition for continued execution of the job, to interface with the termination subroutine, to prepare for the initiation of the next step, or to branch to job deletion if there are no more steps in the current job. When step deletion is entered, a check is made to determine whether the routine was entered due to an ABEND with the scheduler in control. If so, a message is issued to the operator and all CSCBs are removed from the CSCB chain. Control passes to IEFSD510.

If an ABEND did not occur, the step deletion routine branches to ENQ/DEQ purge routine IEFSD598 via a BALR instruction to remove any control blocks which were

enqueued, but not dequeued, by the problem program step.

Step deletion then issues a series of GETMAIN requests to obtain storage for queue manager IOBs (IOB1 and ICB2), a temporary QMPA, and a register save area and parameter list for the table breakup routine. These blocks and tables are initialized and step deletion branches to queue manager table breakup routine IEFSD514, to read in the TIOT and LCT blocks for the job step. The addresses in these blocks are restored and the storage obtained for the temporary work areas is freed.

A GETMAIN (subpool 253) is issued to obtain storage for the SCT and JCT. The SCT is read into storage from the job queue, the JCT from its temporary area. The JCT is updated with the address of the next SCT and written back on the job queue. Storage is obtained for a terminate register save area and a terminate parameter list. The terminate parameter list is initialized with addresses of control blocks (LOT, JCT, SCT, and TICT list) and the step deletion routine branches to the termination subroutine via a BALR instruction. When termination returns control, step deletion frees the terminate register save area and terminate parameter list and then checks the return code. If the partition was executing the DSDR step for a small partition, step deletion places the addresses of the small partition's TCB and PIB in the LOT.

If the return code indicates that the job is to be suspended, step deletion passes the address of the LCT block in register one to job suspension module IEFSD168 via a BALR instruction. If the return code indicates that job termination was entered, step deletion branches to job deletion routine IEFSD517. If job termination was not entered, the SCT for the next step of the job is read from the job queue, and step deletion passes control to IEFSD512 via an XCTL macro instruction.

Note: If a small partition is requesting termination, entry to the step deletion routine is made at special entry point SMA-LTERM. Entry at this point causes pointers to the SPIL and the small partition's TCB to be established before the step deletion routine invokes ENQ/DEQ Purge routine IEFSD598.

ENQ/DEQ Purge Routine (IEFSD598)

At job termination, this routine purges all ENQ/DEQ control blocks associated with the TCB address passed in Register 4 by the caller. If step termination was completed instead, this routine purges all ENQ/DEQ

control blocks except the data set integrity blocks associated with the major name SYSDSN.

When a given resource is dequeued for the subject TCB, a task switch may occur for a higher priority requestor whose wait count becomes zero, due to availability of the resource. (This purge routine operates in a disabled state to prevent concurrent updating of the ENQ/DEQ control blocks.)

Alternate Step Deletion Routine (IEFSD516)

Alternate step deletion routine IEFSD516 is entered from step initiation routine IEFSD512 when allocation for a step has not been successful. Using the APL and ARSA (created by the step initiation routine) as the terminate parameter list and terminate register save area, this routine branches to termination subroutine IEFSD22Q via a BALR macro instruction. When control is returned from termination, the storage used for the parameter list and register save area is freed and a test is made to determine if job termination was entered. If so, this routine branches to job deletion routine IEFSD517. If job termination was not entered, the SCT for the next job step is read from the job queue and this routine branches to step initiation routine IEFSD512.

Job Deletion Routine (IEFSD517)

The job deletion routine is called at job termination to delete the job from the input queue and to prepare the partition for initiation of the next job. The routine sets the high-order byte of the LCTTCBAD field of the LCT to '80' (hexadecimal) to indicate to the ENQ/DEQ purge routine that it is job termination instead of step termination. The routine then branches to ENQ/DEQ purge routine IEFSD598 to purge the control blocks. On return from the purge routine, the high-order byte is reset to '00'.

The job deletion routine then deletes the job from the input queue, using queue manager delete routine IEFQDELQ. All areas of storage in the partition which were used for the job (except the LOT block) are freed, and the job's CSCB is freed by issuing an SVC 34. The PIB fields used for the disk address of the TIOT and the LOT block are set to zero. If termination was for a small partition, ECBA in the SPIL is posted with a code of two (indicating job termination for the small partition). If termination was for a large partition (or after ECBA has been posted) the "no work" ECB in the PIB is posted and the job deletion routine branches to job selection routine IEFSD510.

Partition Recovery Routine (IEFSD518)

Partition recovery routine IEFSD518 determines the location of main storage required for a checkpoint restart. If the partition being scheduled for the job to be restarted contains the required main storage, the routine returns to the step initiation routine for normal processing. If the nucleus has expanded past the lower boundary of the partition containing the required main storage, the routine sets the job fail bit in the JCT, issues a message stating that main storage is not available for the job, and returns to the step initiation routine IEFSD512 with a return code of zero.

If the partition being scheduled does not contain the required main storage, the routine places the job on the hold queue, updates the SCD and places the SCD back on the job queue. The job's CSCB is unchained and the space containing the CSCB and the ECB/IOBs is freed. The routine then branches to ENQ/DEQ purge routine IEFSD598.

Upon return from ENQ/DEQ purge routine, if a problem program partition exists that contains the required main storage, this routine will create an internal queue element and chain it to the partition's PIB. The partition's "no work" ECB will be posted and a message will be issued stating that the job will start in the partition. If an existing partition contains the required main storage and is defined as a reader or writer partition, this routine issues a message indicating that the partition must be redefined to accept the desired jobclass. If no partition contains the required main storage or the partition that contains the required main storage is about to be redefined, this routine issues a message stating the length and displacement of the required main storage. If the partition being scheduled was a large partition its no-work ECB is posted; if it was a small partition, the SPIL is posted indicating job termination. The partition recovery routine frees the JCT and SCT areas of the partition and returns control to step initiation routine IEFSD512 with a return code of four.

Dequeue by Jobname Interface Routine (IEFSD519)

Dequeue by jobname interface routine (IEFSD519) builds a parameter list used by dequeue by jobname routine IEFLOCDQ to locate a job named on the checkpoint/restart internal job queue. When a checkpoint/restart job is indicated by an entry in the internal job queue pointer in the FIB being processed by job selection routine IEFSD510, job selection branches to IEFSD519 which builds the seven-word parameter list required by IEFLOCDQ. When the

job is dequeued, IEFLOCDQ returns control to IEFSD519.

The interface routine marks the job as ready and returns to job selection with a code of zero in register 15, indicating that the job has been found, and a pointer to the LCT in register 1. If the job is not found by IEFLOCDQ, a return code of four is returned in register 15 to job selection. (A description of IEFLOCDQ is in IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.)

System Output Writers

MFT uses the MVT system output writer (Charts 31-32) with minor changes to five of the modules. As in MVT, the user may have up to 36 system output writers operating concurrently in the system. Each output writer can handle eight output classes; output classes may be shared by writers. However, in MFT, system output writers are classified as either resident or non-resident. A resident writer operates in its own partition. A non-resident writer operates in any problem program partition large enough to accommodate it.

RESIDENT WRITERS

Resident output writer partitions are designated in the TCB by a setting of '10' in the first two bits of the pointer to the partition information block (PIB). This designation is made at system generation by assigning W to the partition in place of the job class or by redefining a partition and assigning WTR to it.

A resident writer is activated by issuing a START command specifying a partition designated previously as a writer partition. A resident writer can be terminated only by issuing a STOP command specifying the device assigned to that writer.

NON-RESIDENT WRITERS

A non-resident system output writer may be started in a problem program partition large enough to hold the writer by issuing a START command specifying either that partition or by replacing the partition number with an 'S' to specify a system-assigned non-resident writer. This causes a "command pending" flag to be set in the partition's PIB.

When the writer has started, it executes in the same way as a resident writer and must be terminated by a STOP command to allow processing of problem programs to be resumed in the partition.

SYSTEM OUTPUT WRITER MODULES

The following five MVT system output writer modules are modified for MFT.

- IEFSD070 - Data Set Writer Linkage Routine.
- IEFSD079 - Linkage to Queue Manager Delete Routine.
- IEFSD084 - Wait Routine.
- IEFSD085 - Data Set Block (DSB) Handler Routine.
- IEFSD087 - Standard Writer Routine.

Descriptions of all other system output writer modules can be found in IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.

Data Set Writer Linkage Routine (IEFSD070)

This routine passes control to the appropriate writer routine via a LINK macro instruction. The normal linkage is to the standard writer, IEFSD087. If a special user-written output writer routine is requested, this routine passes control to that writer. Upon return from either writer, the routine passes control to data set delete routine IEFSD171 via an XCTI macro instruction which deletes the output data sets from the output queue.

Linkage to Queue Manager Delete Routine (IEFSD079)

Upon completion of a job, linkage module IEFSD079 passes control to queue manager delete routine IEFQDELQ via an XCTI macro instruction to delete all control blocks and SMBS associated with the output job from the job queue. Following deletion, the routine then posts all reader ECBs that are waiting for space to indicate that space is now available. (The reader ECB chain address is obtained from the master scheduler resident data area.) When all ECBs have been posted, control is returned to main logic routine IEFSD082.

Wait Routine (IEFSD084)

This routine serves as a multiple WAIT when there is no work in any of the output classes associated with the writer. It issues a WAIT macro instruction on the ECB list created by class name setup routine IEFSD081. When the system output writer enters a wait state, the wait routine issues a message informing the operator that the writer is waiting for work. Any posting (such as a command, or work for the writer) causes control to be given to IEFSD082.

DSB Handler Routine (IEFSD085)

DSB handler routine IEFSD085 is the setup module for printing data sets. It issues a GETMAIN macro instruction for the input DCB if it was not obtained before, and constructs a new TIOT containing an entry for the input data set. It also sets up any user-written output writer program. A check is then made to determine if a pause is required between data sets or only at forms change. If a special form is to be used, the routine writes a message to the operator telling him what form to put in the output device. The form change only occurs if the output device is unit record. This routine then passes control to linkage routine IEFSD070 via an XCTL macro instruction.

Standard Writer Routine (IEFSD087)

This routine first issues an OPEN macro instruction to open the output data set. If the data set was not opened by the problem program, no attempt is made to process the data set. After OPEN, a test is made to check for machine control characters. A switch is set that is interrogated by PUT routine IEFSD089. The writer then passes control to transition routine IEFSD088 which creates header and trailer records. Upon return from IEFSD088, the writer routine checks the CANCEL ECB in the CSCB to determine if a CANCEL command has been issued for this writer. If the CANCEL ECB has been posted complete, control passes to transition routine IEFSD088 to create a trailer record. When control is returned from IEFSD088, the writer is closed. Control is then returned to linkage routine IEFSD078 via a RETURN macro instruction.

If the writer is not to be canceled, the writer routine issues a GET macro instruction to read a record and checks for a control character. If no control character exists, the writer puts one in which causes the printer to skip one line or the punch to feed into the normal pocket. If the printer has overflowed, a skip is made to the next page.

The writer then adjusts the pointer to the record so that it points to the first data character (instead of control character) and passes control to transition routine IEFSD088 for trailer records. It then issues a CLOSE macro instruction to close the input data set, a FREEPOOL macro instruction to free the buffers, and returns control to linkage module IEFSD078 via a RETURN macro instruction.

System Task Control

System task control (STC) (Chart 33) initiates all tasks except the initiator (START INIT). When the master scheduler determines that a START command with an identifier operand has been issued, it checks the validity of the partition specified in the command, builds and chains a CSCB, places a pointer to the CSCB in the partition's PIB, and posts the partition.

Note: If the procedure being started is for a system-assigned reader or writer, the CSCB pointer is placed in the master scheduler resident data area. (See Appendix A for the format of the master scheduler resident data area).

As shown in Figure 29, job selection module IEESD510 responds when the partition is posted, and calls STC when a START command for a reader or writer is recognized. If a reader or system output writer is to be started, STC must process a job description similar to a user's job description.

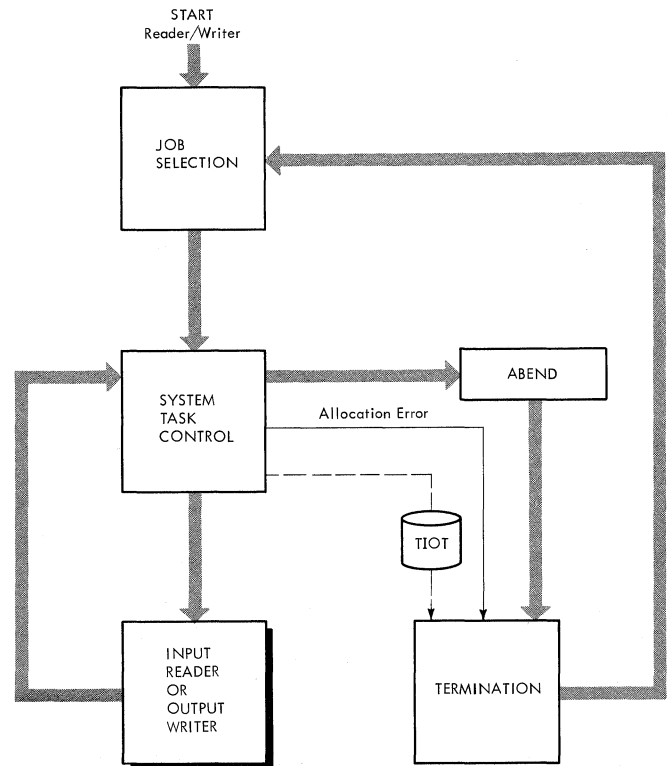


Figure 29. Scheduling a Writer in a Large Partition

The job description information for a reader or writer comes from three sources: the procedure library, Job Control Language (JCL) statements, and the operator. The procedure library contains standard descriptions of a reader and writer. JCL

statements (corresponding to input stream JCL) are stored internally; these statements invoke and modify the reader or writer procedure. The operator furnishes additional information in the operand of the START command; this information is edited into the internally stored JCL statements before they are used to invoke and modify the procedure.

INITIATING SYSTEM TASKS

When initiator job selection routine IEESD510 determines that a START command for a reader or writer has been entered, it passes control to START syntax check routine IEEVSTAR via an XCTL macro instruction.

START Syntax Check Routine (IEEVSTAR)

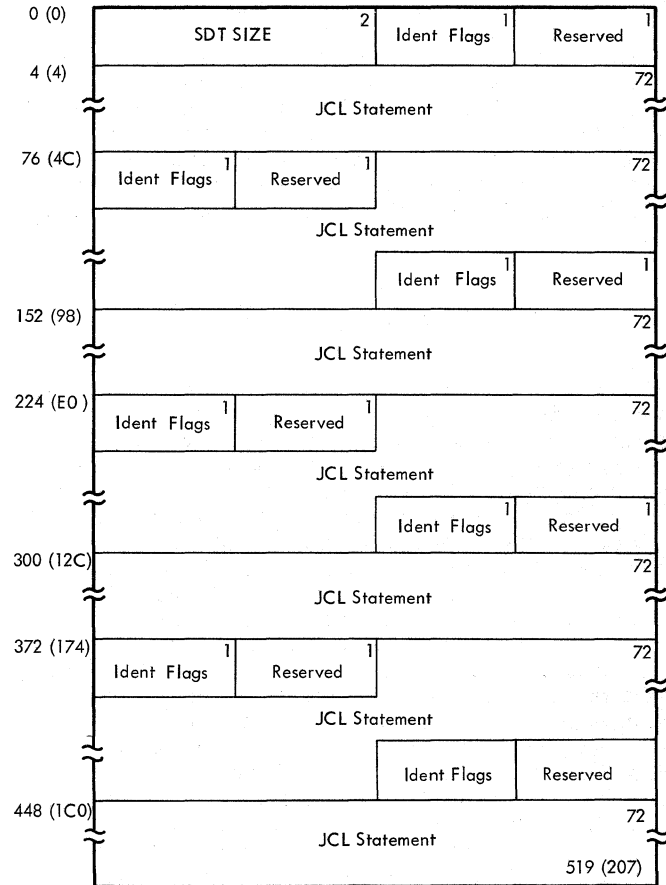
The START syntax check routine gets main storage for, and builds, the start descriptor table (SDT) (see Figure 30). Seven entries are provided in the SDT: the first contains the JOB statement, the second contains the EXEC statement that calls the procedure specified in the START command, the remaining entries are provided for a DD statement and continuations of the EXEC and DD statements. Each entry contains a one-byte identification flags field, whose bits, when set to one, have the following meanings:

- Bit 0 indicates a JOB statement.
- Bit 1 indicates an EXEC statement.
- Bit 2 indicates a DD statement.
- Bit 3 indicates a DD statement continuation.
- Bit 4 indicates an EXEC statement continuation.
- Bits 5 through 7 are reserved.

The routine generates the JOB, EXEC, and DD statements that are placed in the SDT. The keyword parameters in the START command are compared with a list of keyword parameters that are allowable in a DD statement; they are not compared with DD subparameters. If the keyword corresponds to a member of the list, the routine stores it in the DD statement in the SDT. This DD statement overrides the IEFDRD DD statement in the procedure specified in the START command. If the keyword does not correspond to a member of the list, it is assumed to be a symbolic parameter keyword and is placed in the EXEC statement in the SDT.

Finally, the Syntax Check routine passes control to the JCL Edit routine (module IEEVJCL), which builds the job control language set (JCLS). Using the information in the SDT, the JCL Edit routine puts the JCL

in the form appropriate for the interpreter. Each statement is built in an 88-character buffer (obtained with a GETMAIN macro instruction). A pointer to the first buffer is placed in the CSCB associated with the START command. Each buffer contains a pointer to the next buffer, 4 bytes of reserved space, and a "card image" of the statement in the last 80 bytes.



•Figure 30. START Descriptor Table (SDT)

Reader Control Routine (IEEVRCTL)

Reader control routine IEEVRCTL then receives control and builds the interpreter entrance list (NEL), option list, and exit list. The interpreter entrance list contains the address of the JCLS in its third word. The reader control routine passes control to the reader via a LINK macro instruction.

The reader, used as a closed subroutine, is the same routine that performs the reading task. The non-zero value of the third word of the entrance list indicates that the input stream is an internal data set. Since the input stream is internal, the reader issues a pseudo OPEN macro instruction to bring a special access method (a

modified QSAM) into storage and places a pointer to the access method in the input DCB. This special access method reads the JCLS; it is entered from the expansion of the standard GET macro instruction.

The internally-stored job control language statements, and the statements from the procedure library are analyzed and combined. The standard job description tables are built, and an input queue entry is constructed; however, because bit 7 of the option switches field of the option list is off, the entry is not enqueued, and the reader or writer "job" cannot be selected by an initiator. If errors are detected during reader processing, appropriate messages are placed in system message blocks, which are enqueued in the message class queue. When processing is complete, the reader places the main storage address of the job control table (JCT) in the NEL and returns control to the reader control routine with a code that indicates whether processing was successful. The reader control routine then passes control to allocation interface control routine IEEVACTL.

Allocation Interface Control Routine (IEEVACTL)

The reader control routine passes control to allocation interface control routine IEEVACTL, with an indication of whether the reader had encountered errors. The allocation interface control routine uses error message routine IEEVMSG1 to issue the WTO macro instruction to inform the operator of any errors that have been found. The routine then constructs the required allocate parameter list, and passes control to the I/O device allocation routine via a LINK macro instruction.

I/O device allocation routine IEFSD21Q uses the JCT to find the appropriate tables in the input queue, allocates the necessary devices to the reader or writer, and issues any necessary mounting messages. The allocation recovery routines issue WTO macro instructions to inform the operator of any errors found during allocation. When allocation is complete, or if allocation cannot be performed, control is returned to the allocation control interface routine.

Allocation control interface routine IEEVACTL determines if the routine to be given control is an authorized routine and then transfers control to Write TIOT routine IEESD590.

Note: A list of "authorized" routines is contained in a table in link-table routine IEEVLKNT.

Write TIOT on Disk Routine (IEESD590)

Write TIOT on disk routine IEESD590 checks that a reader has not been started in a small partition, writes the TIOT which is used for job selection, and checks for a small partition writer. If a writer is to be started in a small partition, this module issues a POST macro instruction and a WAIT macro instruction for the SPIL and then passes control to job selection routine IEFSD510 via an EXIT macro instruction. If it is not for a small partition writer, control is transferred to linkor routine IEESD591.

Linkor Routine (IEESD591)

The linkor routine passes control to the requested routine via a LINK macro instruction. When the reader or writer stops, it returns control to the linkor routine, which checks for a small partition writer. If a small partition writer returned control to the linkor routine, control then passes to IEFSD510. If a resident reader or large partition writer returned control, termination interface routine IEEVTCTL is given control via an XCTL macro instruction. If a transient reader was suspended, IEFSD591 returns to job selection routine IEFSD510.

POST Routine (IEESD592)

POST routine IEESD592 checks the CSCB to determine if it has been freed; if not, it is freed. It also checks for a small partition. The valid condition is posted in the SPIL or the PIB. The post routine then passes control to IEFSD510 via an EXIT macro instruction.

System Restart

The system restart functions may be requested at any time that a system restart becomes necessary; e.g., end-of-day, end-of-shift, system malfunction, power failure. This feature provides a means whereby a maximum amount of information concerning input work queues, output work queues, and jobs in interpretation, initiation, execution, or termination can be preserved. System restart permits reinitialization, rather than a complete reformatting, of the job queue data set (SYS1.SYSJOBQE).

MFT uses the MVT system restart modules. For a complete description of these modules, and how they function, see IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.

Appendix A: Tables and Work Areas

This appendix contains descriptions and format diagrams of the major tables and work areas that are used by MFT job management. The tables and work areas are in alphabetical order, as shown below:

- Command Scheduling Control Block (CSCB)
- Data Set Enqueue (DSEQ) Table
- Interpreter Work Area (IWA)
- Job Control Table (JCT)
- Job File Control Block (JFCB)
- Job File Control Block Extension (JFCBX)
- Life-of-Task Block (LOT)
- Linkage Control Table (LCT)
- Master Scheduler Resident Data Area
- Partition Information Block (PIB)
- Small Partition Information List (SPIL)
- Step Control Table (SCT)
- Step Input/Output Table (SIOT)
- Task Input/Output Table (TIOT)

Tables and work areas are shown four or eight bytes wide for convenience, but are not necessarily drawn to scale. Tables that are stored in work queue entries are limited, by convention, to a length of 176 bytes.

The names of most fields are sufficient to describe the fields; those that require further explanation are described in the text accompanying the table. Where a macro instruction may be used to include a DSECT of a table in routines using the table, the name of the mapping macro instruction is also given. The displacement of each field is shown to the left of each table; the values in parentheses show the hexadecimal displacement.

COMMAND SCHEDULING CONTROL BLOCK (CSCB)

Description: A command scheduling control block (CSCB) (Figure 31) is an area for communications between the command scheduling routine (SVC 34) and the command execution routines. Input CSCBs are created by several system routines. When an input CSCB is created, it is placed in a chain of CSCBs by the command scheduling routine. It remains in the chain until it is deleted from the chain by the command scheduling routine, which may also free the main storage occupied by the CSCB. An input CSCB is created under the following circumstances:

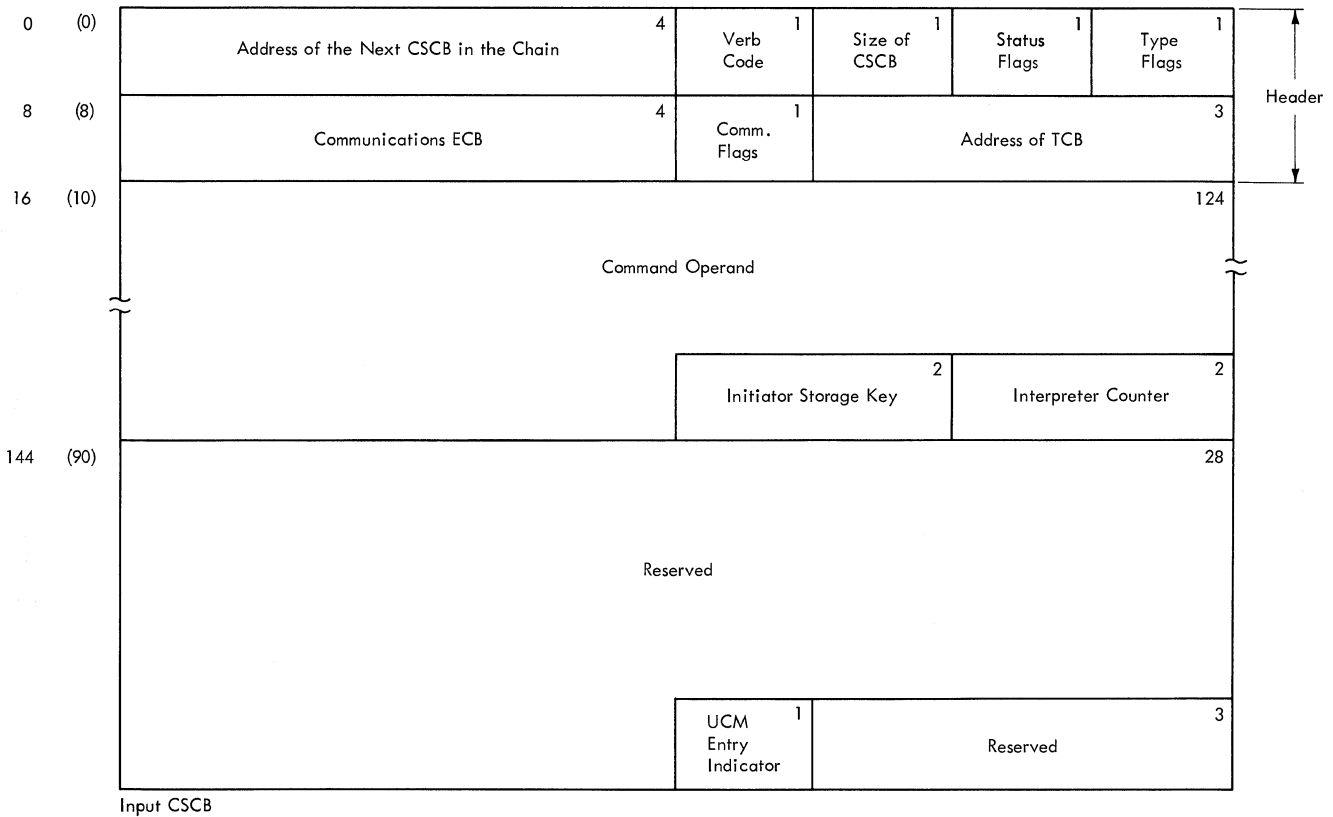
- A CSCB is created by the command scheduling routine each time a task-creating command is encountered. If the task is a reading or writing task, the CSCB is deleted from the chain, and its main storage released, when the task terminates.
- A CSCB is created by the queue management dequeue routine each time the initiator dequeues a job. This CSCB is deleted from the chain, and its main storage released, when the last step of the job has terminated.
- A CSCB is created by a system output writer each time it encounters a DSB that was not preceded by another DSB in the current queue entry. The CSCB serves as a communication area, allowing the cancelation (by operator command) of the subtasks established by the writer. The CSCB is deleted from the chain, and its main storage released, when the writer encounters an SMB (or the last block in the current queue entry).

A control CSCB is updated (and changed to the control format if necessary) by the command scheduling routine when a CANCEL jobname (job selected), CANCEL writer device, MODIFY, or STCP command is encountered.

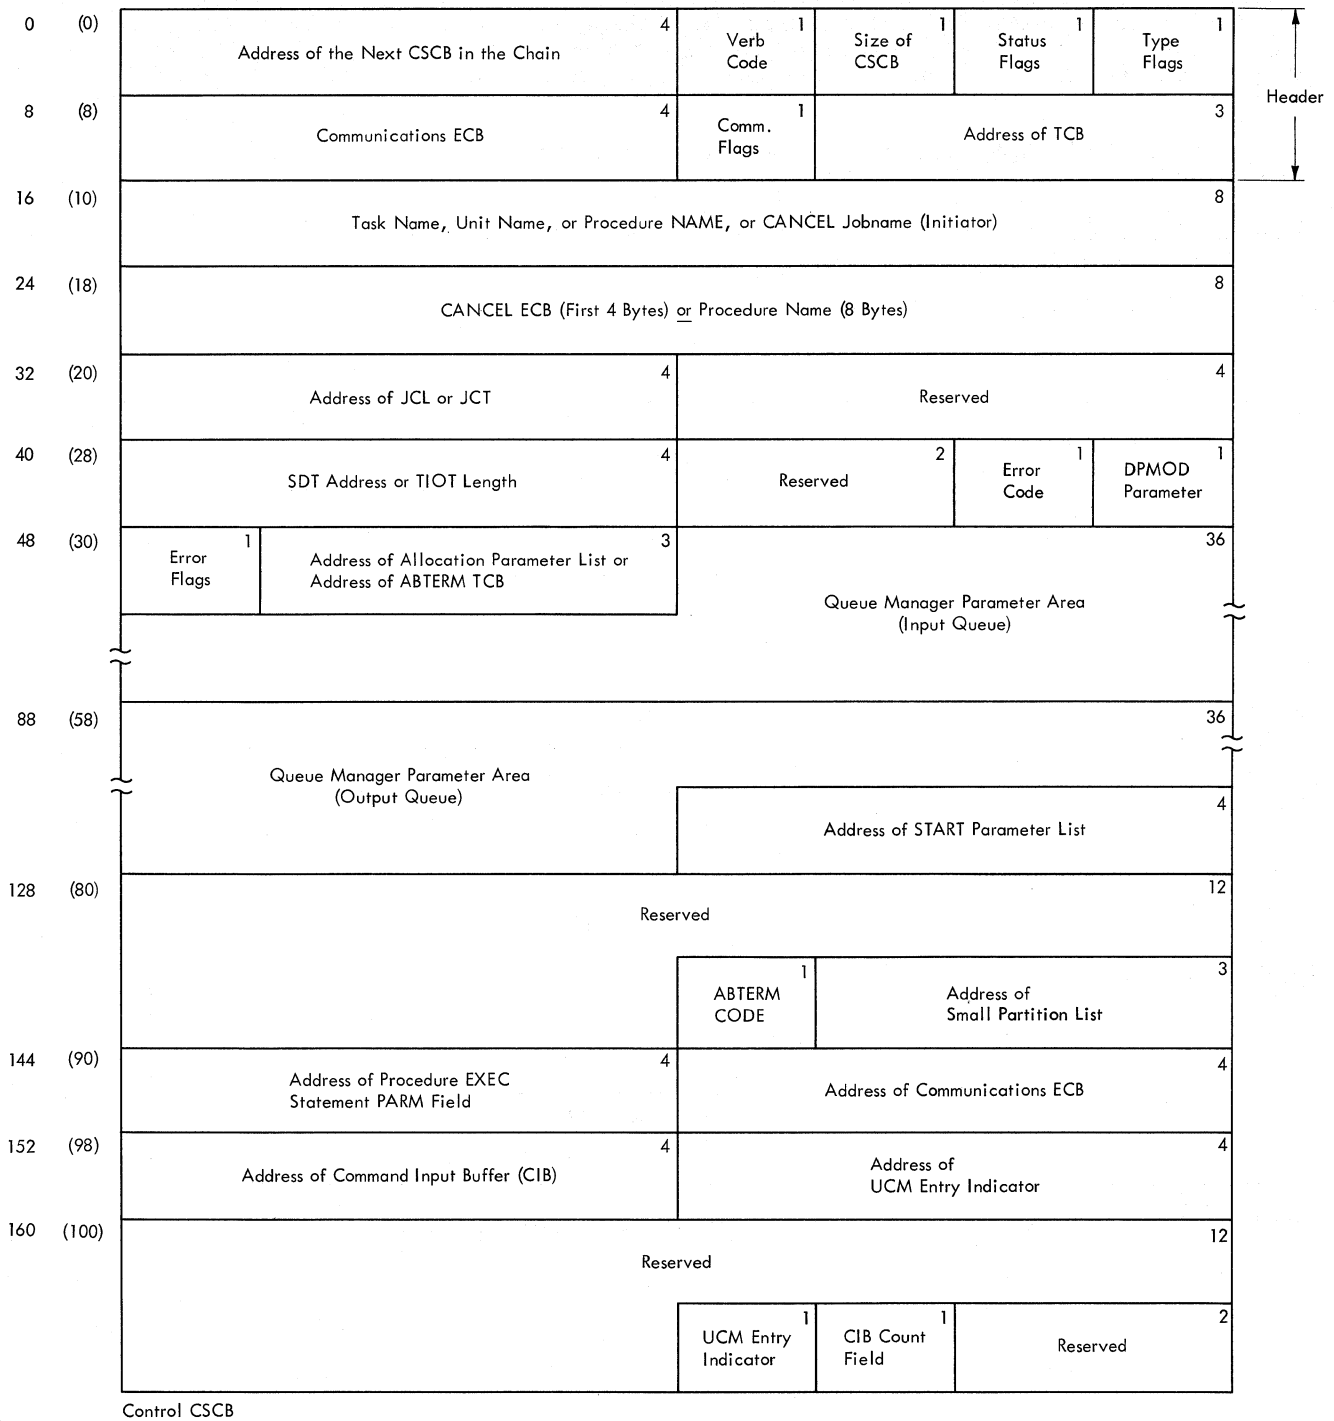
Although most of the fields are self-explanatory, the following require further description:

- **Status Flags:** This byte indicates the status (pending/not pending) of the CSCB, and the action to be taken by the command scheduling routine. In addition to command processing, the command scheduling routine may be entered to add the CSCB to the chain, delete it, free its main storage, or to branch to the abnormal termination routine.
- **Type Flags:** This byte indicates the type of activity with which the CSCB is associated.
- **Communication Flags:** This byte indicates the function to be performed by the command processing routine.

Mapping Macro Instruction: IEECHAIN




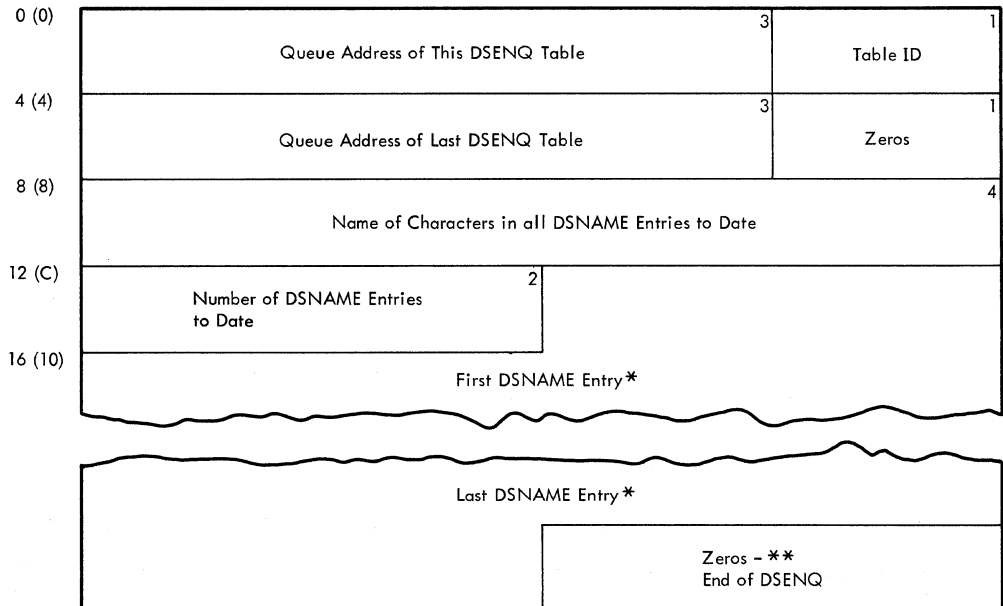
•Figure 31. Command Scheduling Control Block (CSCB) (Part 1 of 2)



•Figure 31. Command Scheduling Control Block (CSCB) (Part 2 of 2)

DATA SET ENQUEUE TABLE (DSENQ)

Description: The data set enqueue table (DSENQ) (Figure 32) is built by the DD statement processor routine of the interpreter, and is used by the initiator to construct an ENQ macro instruction parameter list to prevent routines performing different tasks from using the same exclusive data sets concurrently. The table contains an entry for each data set (except temporary data sets) required for a job.



** If the last entry uses the last available space in the tables but no overflow occurs, the zero bytes are omitted.

Figure 32. Data Set Enqueue Table (DSENQ)

INTERPRETER WORK AREA (IWA)

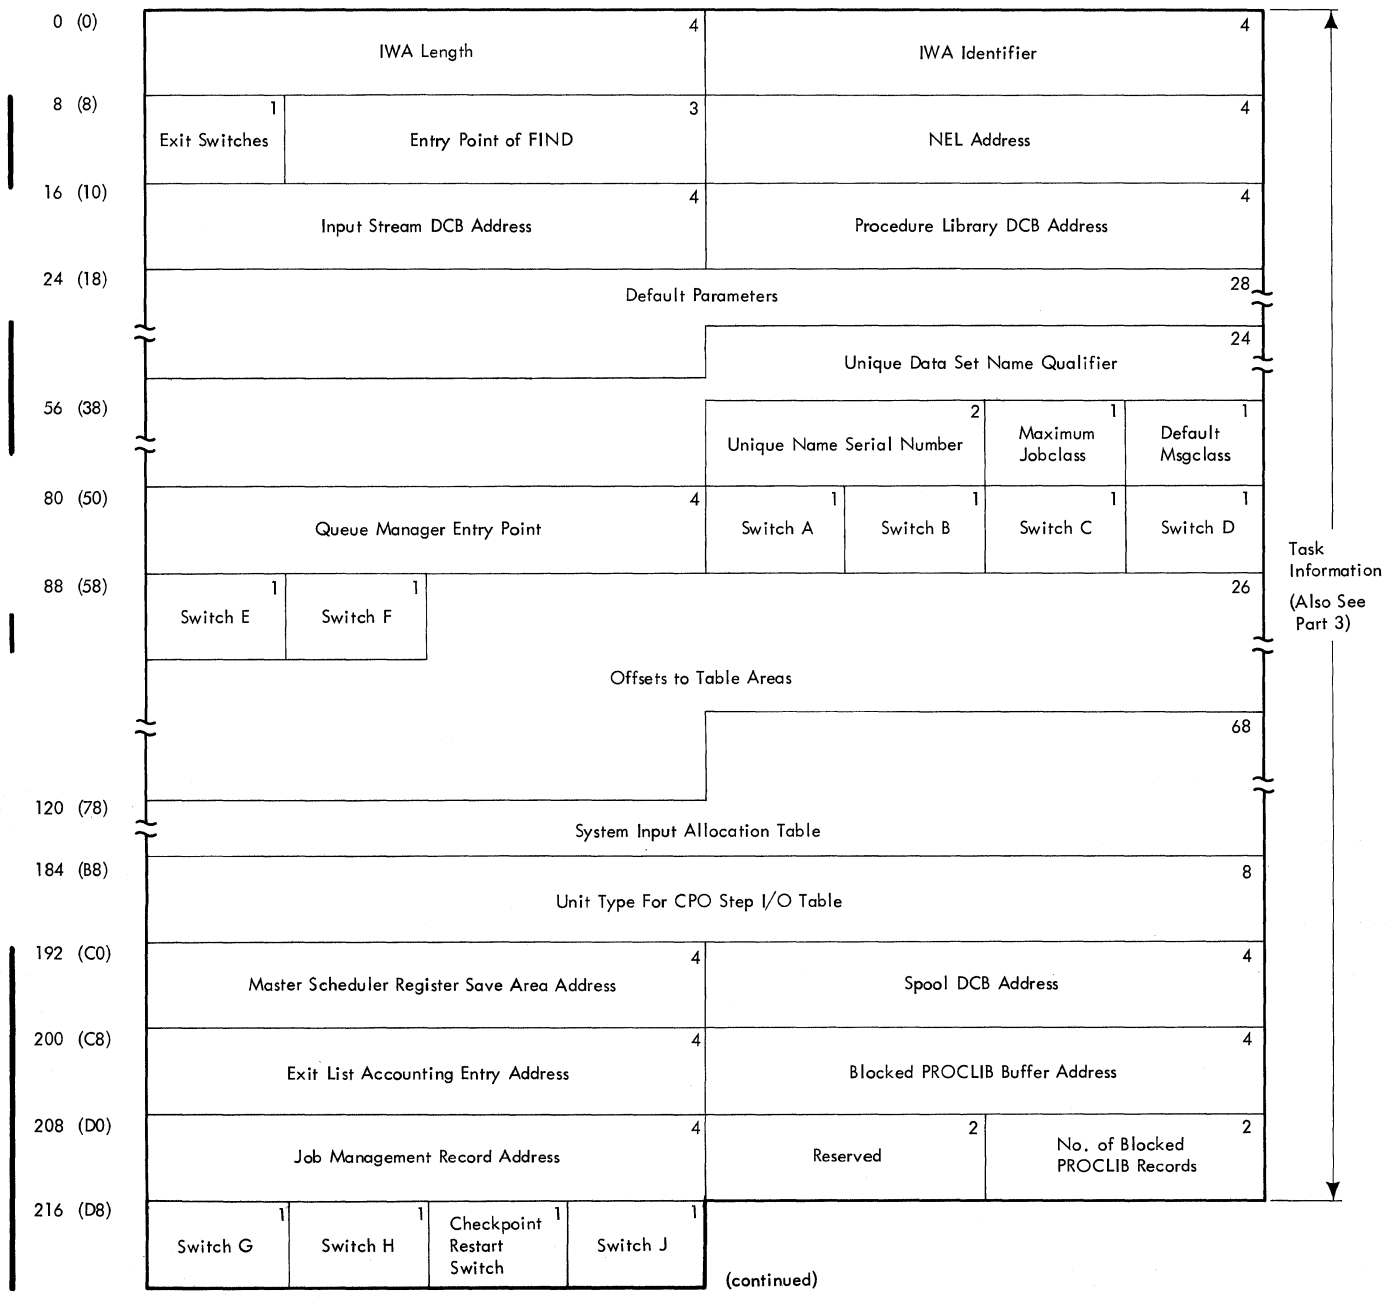
Description: The 2044-byte interpreter work area (IWA) (Figure 33) is obtained from subpool zero by a GETMAIN macro instruction in the interpreter initialization module (IEFVH1). The IWA contains information used by the interpreter routines; it is the area in which job description tables are built before they are placed in the work queues.

Although most of the fields in the interpreter work area are self-explanatory, the following require further description:

- **Default Parameters:** The PARM field of the EXEC statement in the reader procedure contains parameters to be used when no explicit specification is made. These parameters specify whether the installation requires a programmer's name or account number on each JOB statement, the priority to be assigned to a job if no priority has been specified, whether commands in the input stream should be processed (or ignored), and the device, primary quantity, and secondary quantity to be allocated to system output data sets.
- **Switches A-J:** These fields contain internal switches used for communicating status information among the interpreter routines.

- Switch K: This field contains the Priority Change Value for the CHAP macro instruction.
- Switch L: This field contains the Default Allocation level in MSGLEVEL.
- Switch M: This field contains the Default JCL level in MSGLEVEL.
- Switch N: This field contains the length of the fixed part of the message for symbolic parameter substitution.
- Switch X1: This field is set to X'80' for a search of the DDNAME reference table or to X'40' for SYSOUT.
- Checkpoint/Restart Switch: This field contains switches that communicate checkpoint/restart status information to the interpreter routines.
- System Input Allocation Table: This area contains a list of pointers to the UCBS corresponding to units available for allocation to system input data sets.
- Queue Address Table: This area contains the addresses (in TTR form) of the next two records assigned to the job's input queue entry, and the addresses (in TTR form) of the first joblib SIOT, the first scan dictionary record, and the DD override table.
- Input Stream Parameter List: This area describes the statement last encountered in the input stream, and contains a pointer to the field currently being processed.
- Procedure Library Parameter List: This area describes the statement last read from the procedure library, and contains a pointer to the field currently being processed.
- Procedure Library Merge Control Data: This area contains information used in merging statements from the input stream with statements from the procedure library. The information includes the statement names, the step names, and the names of the previous and next procedure steps.

Mapping Macro Instruction: IEFVMIWA



•Figure 33. Interpreter Work Area (IWA) (Part 1 of 3)

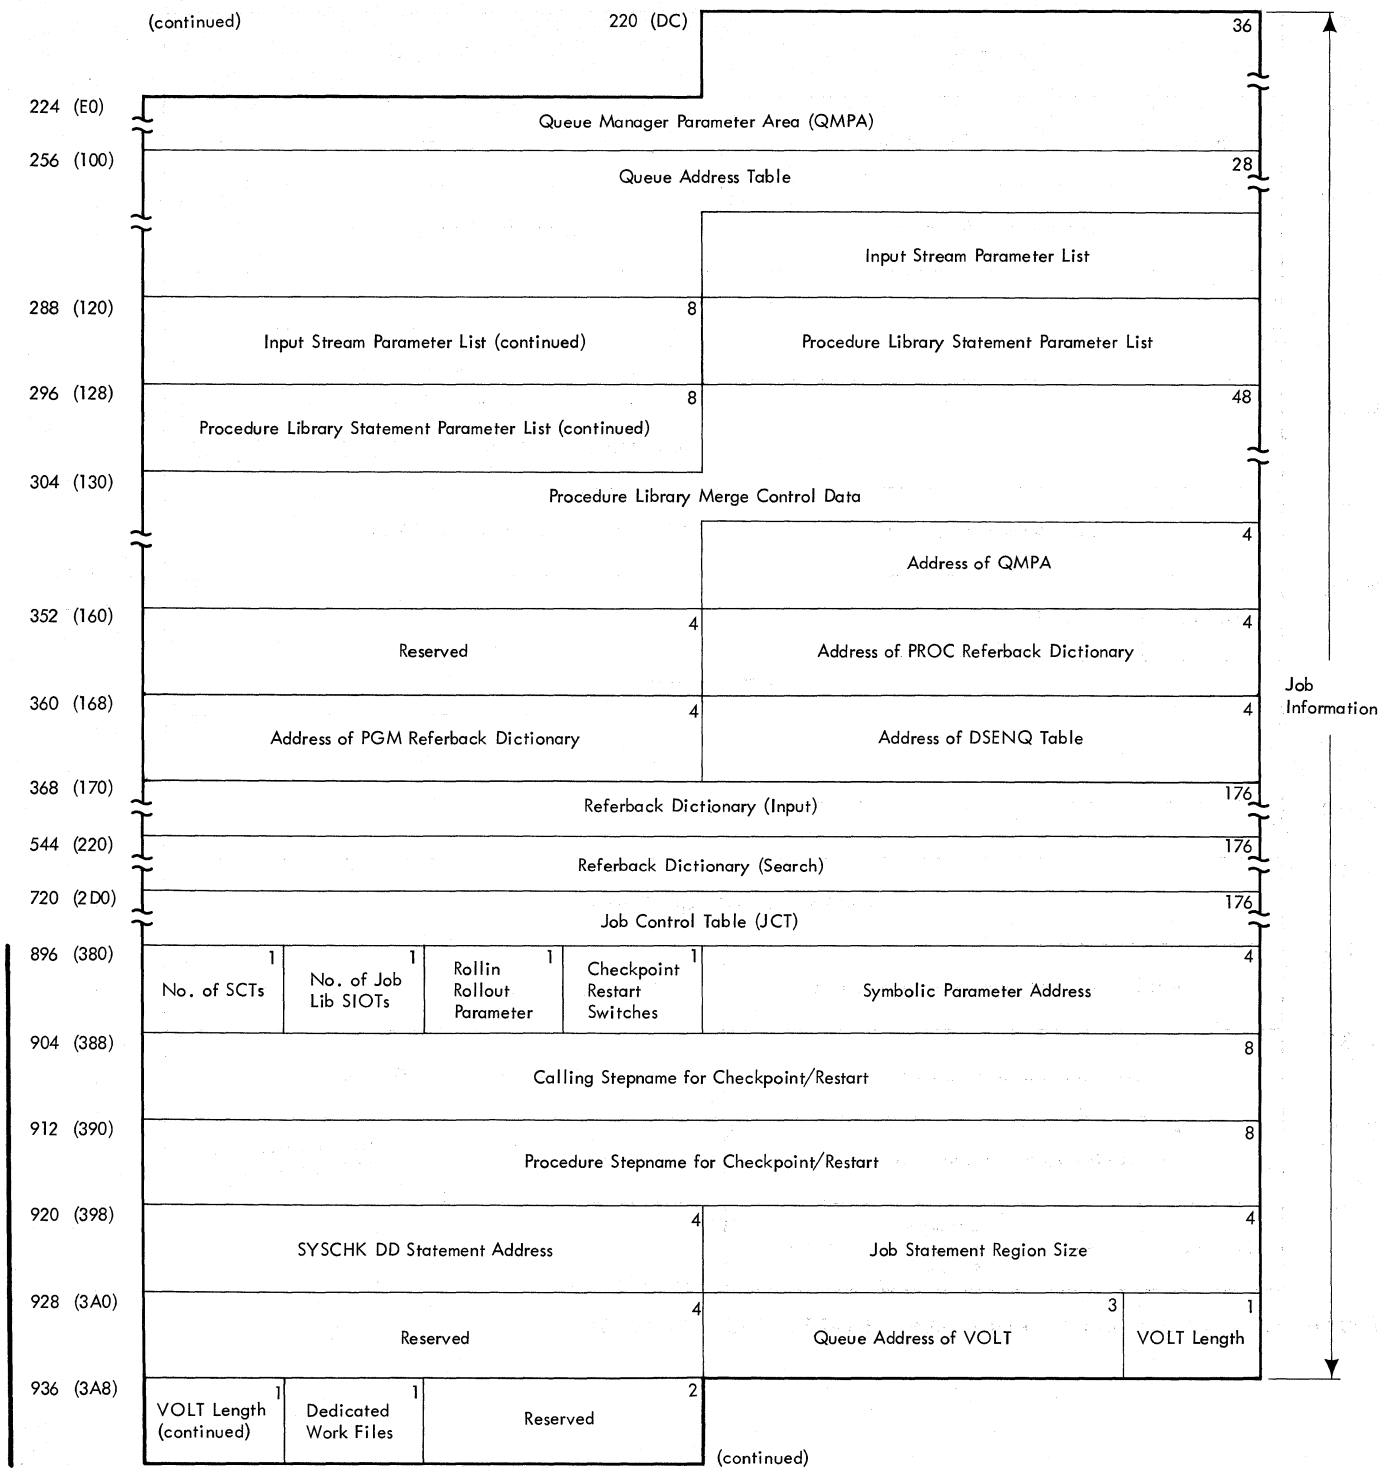



Figure 33. Interpreter Work Area (IWA) (Part 2 of 3)



• Figure 33. Interpreter Work Area (IWA) (Part 3 of 3)

JOB CONTROL TABLE (JCT)

Description: The job control table (JCT) (Figure 34) is created in the interpreter work area by the job statement processor routine of the interpreter. It contains information from the JOB statement, job status information, and pointers to other tables in the job's input queue entry. When the interpreter has processed all steps of a job, the JCT is written into the appropriate input queue according to priority; it is read back into main storage by the initiator job selection and job delete routines.

Although most of the fields in the job control table are self-explanatory, the following require further description:

- **Job Status Indicators:** The sixth byte of the JCT indicates the status of the job as shown below:

<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	1	A JOBLIB DD statement is included with the job
1	1	Job flush
2	1	Job step canceled by condition codes
3	1	Step flush
4	1	JCT ABEND
5	1	Job failed
6	1	Job includes a cataloged procedure
7	1	Job is a "no setup" job

- **Checkpoint/Rescart Indicators:** This two byte field indicates the checkpoint/restart status as shown below:

Byte 1		
<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	1	Warm start
1		Not used by MFT
2		Not used
3	1	Checkpoint taken for this step
4	1	Intra-step checkpoint/restart to be done
5	1	Step restart to be done
6-7	0	Must be set to zero

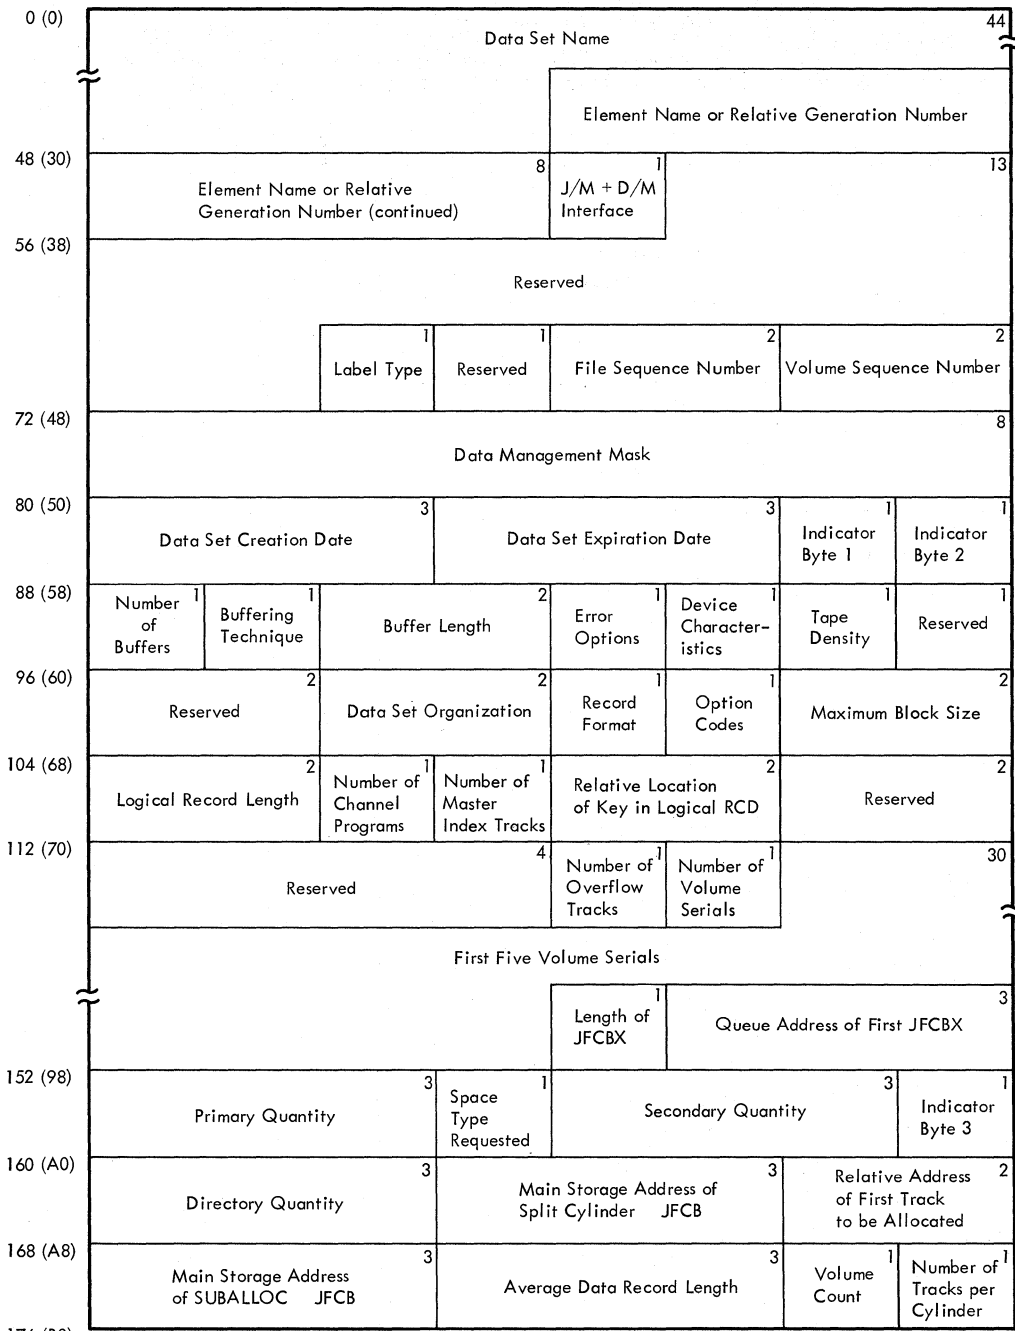
Byte 2		
<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	1	SYSCHK DD statement is included with the job
1	1	RD keyword parameter is not NC
2	1	No restart is to be done
3	1	No checkpoints are to be taken
4	1	Do restart if necessary
5-6		Not used
7	1	DSDR processing has not successfully ended

- **SYSOUT Classes:** The first 36 bits of the five-byte field are used to indicate the system output classes that contain data. The four remaining bits are reserved.

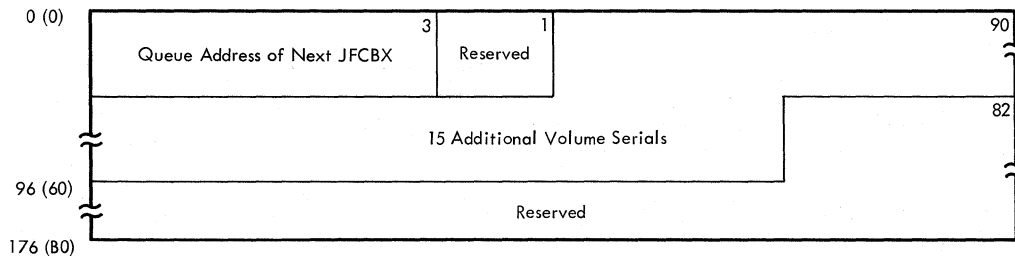
Mapping Macro Instruction: IEFAJCTB

0 (0)	Address in Queue of JCT	3	Table ID = 00	1	Internal Job Serial Number	1	Job Status Indicators	1	Message Class	1	Message Level and Job Priority	1
8 (8)	Job Name											8
16 (10)	Teleprocessing Terminal Name											8
24 (18)	Address in Queue of PDQ	3	Reserved	1	Address in Queue of GDG Bias Count Table	3	Reserved	1				
32 (20)	Address in Queue of First SCT	3	Reserved	1	Address in Queue of First SMB	3	Reserved	1				
40 (28)	Address in Queue of Job ACT	3	Reserved	1	Address in Queue of First DSB	3	Reserved	1				
48 (30)	Address in Queue of Last DSB	3	Reserved	1	Key of SMB Track	2	First Job Condition Code	2				
56 (38)	First job Condition Operator	1	Reserved	1								28
Reserved for Seven Additional Job Condition Codes and Operators												
											Checkpoint/Restart Indicators	2
88 (58)	TTR of DSENG Table (MVT Only)	3	Zeros	1	Region Parameter (MVT Only)	2	Queue Ident. (MVT Only)	1	No. of Steps	1		
96 (60)	TTR of Compressed TIOT (MVT Only)	3	No. of Job Tracks on SYS1.JOBQE (MVT only)	1	Checkpoint Data Set Device Type							4
104 (68)	TTR of JFCB for Checkpoint Data Set	3	Zeros	1	Number of Checkpoints	2	Vol. of Checkpoint Data Set	1	Reserved	1		
112 (70)	TTR of SCT for First Step to Run	4	Reserved	1	Length of Checkpoint ID	1						16
Checkpoint Identification											17	
136 (88)	Reserved											
											SYSOUT Classes	1
152 (98)	SYSOUT Classes (Continued)											4

•Figure 34. Job Control Table (JCT)



Job File Control Block



Job File Control Block Extension

Figure 35. Job File Control Block (JFCB) and Extension (JFCBX)

JOB FILE CONTROL BLOCK (JFCB) AND EXTENSION (JFCBX)

Description: A job file control block (JFCB) (Figure 35) is constructed in subpool zero (from information in a DD statement) by the interpreter DD statement processor routine. The JFCB is written into the job's input queue entry, and retrieved when a DCB with the corresponding name is opened. The information in the JFCB, which describes the characteristics of a data set, may be modified by the open routine.

A JFCB contains enough space to record five volume serials. If more than five volume serials are specified, enough job file control block extensions (JFCBXs) to contain the additional volume serials are constructed; each JFCBX can contain up to fifteen additional volume serials.

Additional information on the contents of the JFCB and JFCBX may be found in the publication, IBM System/360 Operating System: System Control Blocks, Form C28-6628.

Mapping Macro Instruction: IEFJFCBN

LIFE-OF-TASK (LOT) BLOCK

Description: The 348-byte life-of-task (LOT) block (Figure 36) is built in a main storage area obtained from subpool 253. It stores information for scheduling functions, and is used by system task control and initiators. It is created by the Job Select module for initiating problem programs, and by system task control for initiating readers and writers.

The LOT block contains the linkage control table (LCT), a two-level register save area (REGSAVE), an input queue manager parameter area (QMGR1), an output queue manager parameter area (QMGR2), the address of the ECB list, the address of the PIB, and the address of the SPIL.

LINKAGE CONTROL TABLE (LCT)

Description: The linkage control table (LCT) (Figure 37) is built in a main storage area obtained from subpool 253 by the initiator initialization routine. It is a communications area used by the routines of the initiator.

Most of the fields in the LCT are self-explanatory; it should be noted, however, that the job termination status bit is the low-order bit of the one-byte device features field.

Mapping Macro Instruction: IEFALLCT

MASTER SCHEDULER RESIDENT DATA AREA

Description: The master scheduler resident data area (Figure 38), which is in the nucleus area of main storage, contains information used by the queue initialization, command scheduling, initiator, and I/O device allocation routines. Its location is stored in the CVTMSER field of the communication vector table.

Most of the fields in the master scheduler resident data area are self-explanatory; those fields that require further explanation are described below:

- Queue Formatting Switch: If the high-order bit of this field is on, it indicates that the queue data set must be formatted.
- Transient Reader TTR: This field is used by the transient reader suspend routine to store the address of the work queue data set where the reader information was placed when the reader was suspended.
- DEFINE Control Information: If the high-order bit of this field is on, it is a DEFINE operation; if off, it is IPL time. The second bit indicates that a list of the partitions' sizes and job class(es) has been requested; the third bit indicates that there is an adjacent partition check; the fourth bit is set when initialization is complete to allow DEFINE commands to be accepted; the fifth bit is set on when the

operator has requested partition changes at IPL; the sixth bit indicates that a small partition cannot terminate because of the DEFINE operation; the seventh bit indicates that a DEFINE command has been issued during operation; the eighth bit indicates that the system has storage protection.

- **Status Flags:** When set on, status flags indicate:

<u>Bit</u>	<u>Meaning</u>
0	System Initialization in progress
1	DISPLAY JOBNAMEs
2	Reserved
3	VARY/UNLOAD summary
4	Queue hold-release
5	DISPLAY ACTIVE processing
6-7	Reserved




Figure 36. Life-of-Task (LOT) Block

0 (0)	1	Address of Job Step CSCB		3	Address of I/O Supervisor UCB Lookup Table		4		
8 (8)	TCB Address			4	1	Linkor's Register Save Area Address	3		
16 (10)	JCT Address			4	SCT Address			4	
24 (18)	Queue Address of SCT			4	Allocate/IEFVPOST Communication Block Address			4	
32 (20)	Error Code			4				16	
Communications Area									
				Address of Register Save Area and QMPA				4	
56 (38)	1	JFCB Housekeeping Indicators	1	Current Step Number	1	Action Code	4	Address of Current SMB	4
64 (40)	Counter for Assigning Unique Volume Serials to Passed Data Set Volumes			4	Address of Message Class QMPA			4	
72 (48)	Return Address to System Task Control Routine			4	Address of Initiator CSCB			4	
Timer Work Area									
96 (60)	JOB LIB DCB Address			4	Allocate/Terminate Parameter List Address			4	

•Figure 37. Linkage Control Table (LCT)

• Log Status Flags:

<u>Bit</u>	<u>Meaning</u>
0	Log Data Set Sysout Scheduling
1	Log Threshold Reached

• MFT Switches: When set on, flags indicate:

<u>Bit</u>	<u>Meaning</u>
0	Transient Reader Active
1	Transient Reader in Core
2	Pending START command for transient reader
3	MFT Environment switch
4	System Assigned Reader is Running
5	Core storage is in System

- Initialization Switches: When set on, flags indicate:

<u>Bit</u>	<u>Meaning</u>
0	IPL switch
1	SYSOUT IPL
2	SYSOUT job start
3-4	Reserved
5	34 Security
6	Queue initialized
7	Procedure catalog initialized

- Pending Flags: When set on, flags indicate:

<u>Bit</u>	<u>Meaning</u>
0	IPL Date
1	Region busy
2	Command move completed
3	Interpreter command return
4	System Input control purge request
5	System output control purge request
6	Blank start pending (REQ=1, START BLANK=0)
7	Console command suppressed by WTO/WTOR Exit Routine

- ECB Flags: When set on, flags indicate:

<u>Bit</u>	<u>Meaning</u>
0	External interrupt
1	WTO or WTOR
2	WTL
3	Console Attention key hit
4	System Input
5	System Output
6	Master command routine
7	Summary bit, Vary UCB scan required

- Resident Switches: When set on, switches indicated:

<u>Bit</u>	<u>Meaning</u>
0	IPL has been completed
1	WTO or WTOR pending
2	Console usage, Primary or alternate
3	Log purge request
4	Reader has reached end of file, or Start reader
5	New reader pending
6	New writer pending
	New writer pending (Modify)
7	Job notification (1=yes)

- Fetch Flags: When set on, flags indicate.

<u>Bit</u>	<u>Meaning</u>
0	Named Fetch
1	Defer current command execution sequence
2	TCB Tree Trace Fetch (Locate)
3	Auxiliary FETCH given
4	Reply bit to Request attention
5	Pseudo-SYSOUT flag
6	DISPLAY STATUS
7	Queue hold-release

- Mapping Macro Instruction: IEEBASEB.

0	(0)	Address of CSCB Chain		4	Group Queue Pointer (MVT only)		4	
8	(8)	Master Scheduler ECB			4	Communications Task IPL ECB		4
16	(10)	Address of Job Queue UCB			4	Address of PROCLIB UCB		4
24	(18)	Queue Formatting Switch	Address of Set Auto Command Parameter List		3	Address of System Log Control Table		4
32	(20)	Status Flags	Number of Tracks in Initiator Stack	Interpreter Counter	2	Initiator Protection Key Mask	Minimum Initiator Partition Size	2
40	(28)	Minimum Problem Program Partition Size		2	Log Status Flags	Reserved	System Log ECB	4
48	(30)	Reserved						46
						ID of console that entered DEFINE	Reserved	1
96	(60)	Core Storage Low Boundary			4	Subpool 255 Boundary Box		4
						First FQE Pointer		4
104	(68)	Low Boundary Pointer			4	High Boundary Pointer		4
112	(70)	Transient Reader, Pending CSCB Pointer			4	MFT Switches	Transient Reader CSCB Pointer	3
120	(78)	Transient Reader TTR			4	DEFINE Control Information		4
128	(80)	Size of Scheduler			4	Address of ECB Chain for Readers		4

MFT Area

•Figure 38. Master Scheduler Resident Data Area (Part 1 of 2)

136	(88)	Initialization Switch	1	PCP System Switches	1	Pending Flags	1	ECB Flags	1	Resident Switches Status Flags	1	Fetch Flags	1	Command Verb	Common Area	
144	(90)	Command Verb (cont.)										8	Variable Communication Field			
152	(98)	Variable Communication Field (cont.)										8	Msg. Generation Control	2		
160	(A0)	Pointer to Character Before List						4	Master ECB							4
168	(A8)	Pointer to ECB in SJQ Entry of Job Using Console						4	ECB for Allocation							4
176	(B0)	Pointer to Primary UCB						4	Pointer to Alternate UCB							4
184	(B8)	Pointer to Pseudo-Disable Switch						4	Pointer to Second Highest Priority Problem Program TCB (MFT)							4
192	(C0)	Pointer to Highest Priority Problem Program TCB (MFT)						4								4

• Figure 38. Master Scheduler Resident Data Area (Part 2 of 2)

PARTITION INFORMATION BLOCK

The 40-byte partition information block (PIB) (Figure 39) contains information used by the command processing and scheduler routines. Its location is stored in the TCBPIB field at displacement 124 (decimal) of the task control block (TCB).

Although most of the fields in the partition information block are self-explanatory, the following require further description:

- **ECB Address:** Contains the address of ECB to be posted by job selection when the partition is made quiescent for partition redefinition.
- **"No Work" ECB for the Initiator:** This ECB is posted by small partitions requesting service, the queue manager when a job has been enqueued, and by the DEFINE and START command routines.
- **Status A Information:**

<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	0	Stop initiator
	1	START INIT issued
1	1	Partition active
2	1	Pending command
3	1	Transient reader operating
4	1	Partition is to be terminated by IEFSD599 when it next gets control
5	1	Partition is involved in redefinition
6	1	System-assigned transient reader operating in this partition
7	1	Problem program is running

• Status B Information:

<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	1	Logical tracks added for initiator
1	1	LOT block exits
2	1	SPIL has been created
3	1	Reserved
4	1	Unending task present in partition

- SPIL Address: The small partition information list (SPIL) is applicable to large partitions only.
- Job Class Codes: Contains one to three codes for the partition, arranged in descending numerical order, i.e., GRP3 is in the second byte of the field, followed by GRP2 and GRP1. The first byte contains the protection key for the partition, if the system has the storage protection feature.

• Internal Queue Status Bits:

<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	1	A large partition in which the DSDR processing step for a small partition (less than 12K) is to be executed
1	1	A restart reader has been started in place of a user assigned reader
2	1	A DEFINE command has been received and the partition is processing jobs on its internal queue.
3-7		Reserved

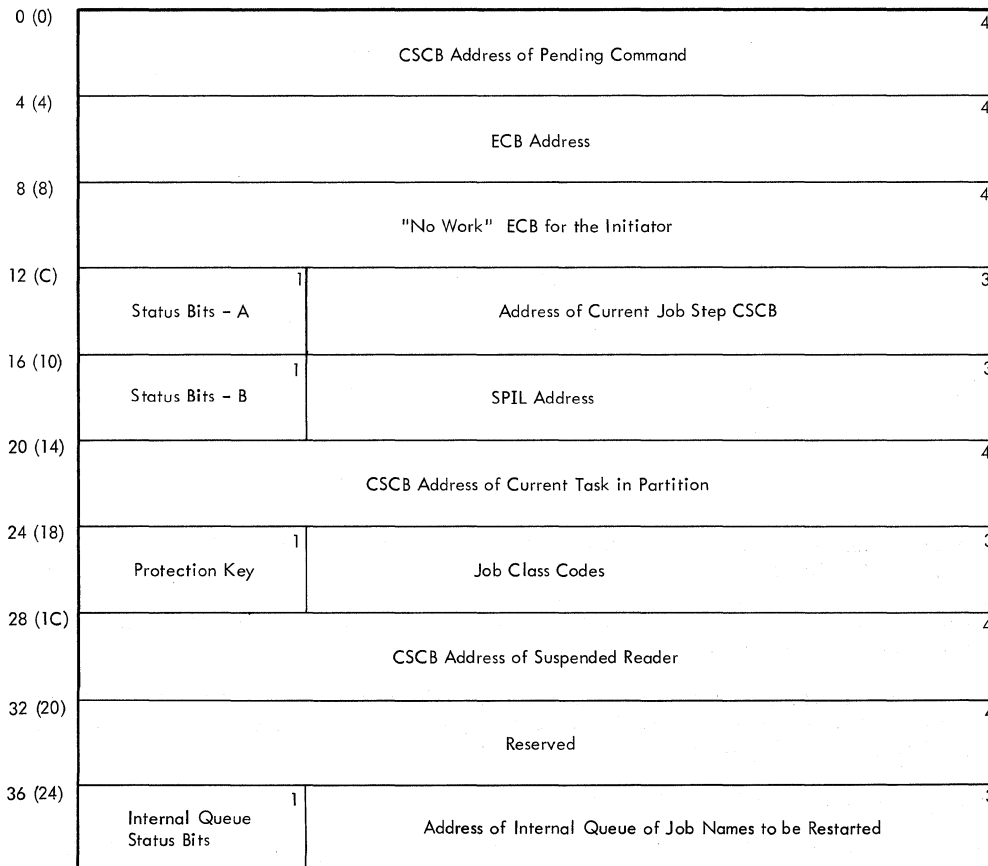


Figure 39. Partition Information Block (PIB)

SMALL PARTITION INFORMATION LIST (SPIL)

Description: The 32-byte small partition information list (SPIL) (Figure 40) is a storage area for information pertaining to small partition scheduling. It is built in main storage obtained from subpool 0. The address of the ECBs provides for information to be passed between the small partition and the large partition that is performing initiation, allocation, or termination functions for the small partition.

Most of the fields in the small partition information block are self explanatory; however, the status bits field is described below.

Bits 0 and 1 contain ones if a START writer command has been entered.

Bit 2 contains a one if a SPIL pointer has been stored in the PIB.

Bit 3 contains a one if a problem program has requested termination.

Bits 0-7 contain zeros if a START INIT command was entered.

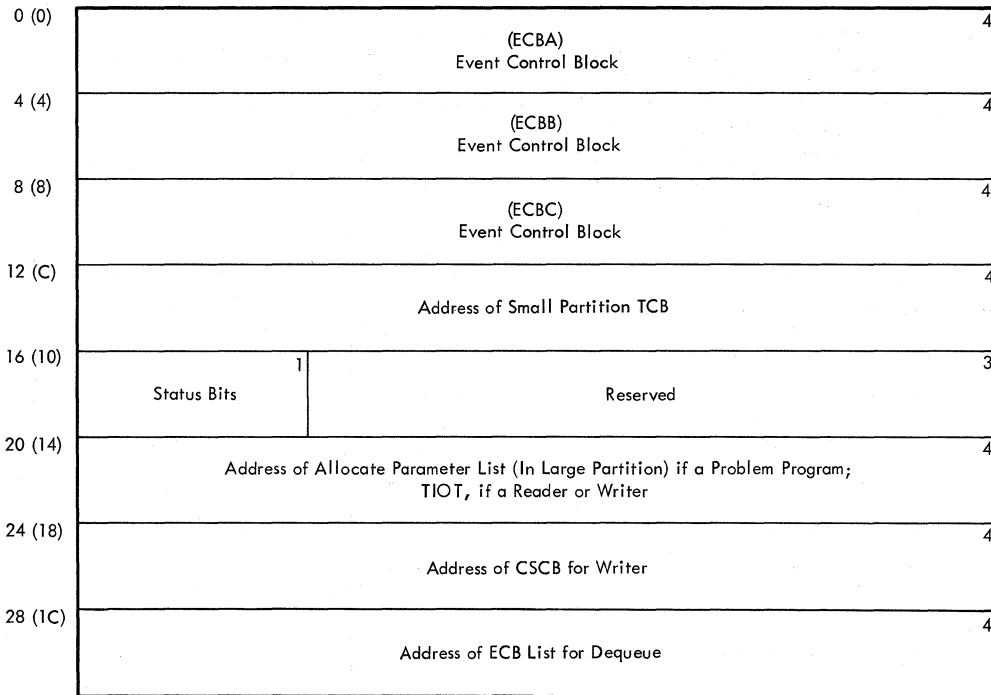


Figure 40. Small Partition Information List (SPIL)

STEP CONTROL TABLE (SCT)

Description: The step control table (SCT) (Figure 41), is used to pass control information to the DD routine of the interpreter and to the initiator routines, which also contribute information to the table. This table is created and initialized by the execute statement processor routine of the interpreter when an EXEC statement is read. One SCT is created for each step of a job.

If the step is part of a previously cataloged procedure, the name of the step that called the procedure, if any, is entered. The following variable-content and indicator fields are included in the table:

BYTE 4: Internal Step Status Indicators:

Bit	Setting	Meaning
0	1	Step can be rolled out
1	1	Roll step out if necessary
2	1	Do not restart step
3	1	Do not take a checkpoint

4	1	Restart if necessary
5	1	Graphics - alter protect key
6	1	Graphics - ABEND exit
7	1	Step failed

PARM Count or Step Status Code:

- a. Interpreter: The number of characters specified in the PARM parameter of the EXEC statement is placed in this entry.
- b. Initiator: This table entry contains the condition code returned by the processing program.

BYTE 67: Step Type Indicators:

<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0	1	EXEC statement contains PGM=*.stepname.ddname
1	1	SYSIN is specified as DD*
2	1	SYSOUT is specified
3	1	JFCB housekeeping is complete
4-7		Reserved

BYTE 104: Step Status

<u>Bit</u>	<u>Setting</u>	<u>Meaning</u>
0-1		Reserved
2	1	SCTMCVOL
3		Reserved
4	1	SCTSTPLB
5-7		Reserved

Mapping Macro Instruction: IEFASCTB

STEP INPUT/OUTPUT TABLE (SIOT)

Description: The Step Input/Output Table (SIOT) (Figure 42), makes DD statement available to the initiator for use as a source of information for the TIOT and for providing DD information to allocation and disposition routines. When a DD statement is read, the interpreter creates a new SIOT and places the DD information into it. The individual bits of the disposition byte and of indicator bytes 56 through 59 in the SIOT are set to one to indicate the following conditions:

BYTE 55: Scheduler Disposition

<u>Bit</u>	<u>Meaning</u>
0	Reserved
1	Retain volume
2	Private volume
3	Pass data set
4	Keep data set
5	Delete data set
6	Catalog data set
7	Uncatalog data set

BYTE 56: Indicator Byte Number 1

<u>Bit</u>	<u>Meaning</u>
0	Dummy data set
1	SYSIN data set
2	Split (primary)
3	Split (secondary)
4	Suballocate
5	Parallel mount
6	Unit affinity
7	Unit separation

0	(0)	Queue Address of SCT			3	Table ID (02)	1	Internal Step Status Indicators	1	Maximum Step Running Time		3		
8	(8)	PARM Count or Step Status Code at Termination		2	Length of Allocate Work Area, or Number of SIOTs			2	Queue Address of First SIOT Entry		3	Reserved	1	
16	(10)	Queue Address of Allocate Work Area			3	Reserved			1	Queue Address of Next SCT		3	Reserved	1
24	(18)	Queue Address of First SMB for Next Step			3	Reserved			1	Queue Address of Last SMB for This Step		3	Reserved	1
32	(20)	Queue Address of First ACT Entry for This Step			3	Reserved			1	Queue Address of VOLT		3	Reserved	1
40	(28)	Queue Address of Dsname Table for This Step			3	Reserved			1	Name of Step That Called Procedure				
48	(30)	Name of Step That Called Procedure (Continued)						8	Step Name					
56	(38)	Step Name (Continued)						8	Relative Pointer to Step Entry in ACT		2	Length of VOLT		2
64	(40)	Number of SIOTs in This Step	1	Number of Setup Messages	1	Number of JFCBs to Allocate	1	Step Type Indicators	1	Queue Address of SCTX				4
72	(48)	X'00'	1	Hierarchy 0 Region Address				3	X'01'	1	Hierarchy 1 Region Address			3
80	(50)	Reserved											8	
88	(58)	Hierarchy 0 Region Size		2	Hierarchy 1 Region Size			2	Reserved		2	Step Dispatching Priority (MVT only)		2
96	(60)	Step SYSIN count for SMF						4	Queue Address of PGM = *, stepname, dname SIOT					4
104	(68)	Extension of Internal Step Status Indicators	1	Queue Address of the Step TIOT				3	Program Name					4
112	(70)	Program Name (Continued)						8	Length (in Bytes) of Dsname Table for This Step		2	First Step Condition Code		2
120	(78)	First Step Condition Operator	1	Queue Address of First Condition SCT				3						36
Second Through Seventh Step Condition Entries														
160	(A0)	Eighth Step Condition Code		2	Eighth Step Condition Operator		1	Queue Address of Eighth Condition SCT			3	Reserved		2
168	(A8)	Queue Address of the First DSB in Message Class			3	Number of Message Class DSBs for this Step	1	Step Status	1	Queue Address of Last Legitimate SMB				3
176	(B0)													

•Figure 41. Step Control Table (SCT)

BYTE 57: Indicator Byte Number 2

<u>Bit</u>	<u>Meaning</u>
0	Channel affinity
1	Channel separation
2	Volume affinity
3	JOELIB DD statement
4	Unlabeled (no labels)
5	Pool DD statement
6	Defer mounting
7	Received data set

BYTE 58: Indicator Byte Number 3

<u>Bit</u>	<u>Meaning</u>
0	Volume reference
1	SYSIN expected (procedures only)
2	Allocate work table volume block indicator
3	Volume reference in step
4	SYSOUT was specified
5	NEW data set
6	MOD data set
7	OLD or SHR data set

BYTE 59: Indicator Byte Number 4

<u>Bit</u>	<u>Meaning</u>
0	Set by reader to indicate GDG single
4	Step processed
5	Intra-step volume affinity
6	Data set is in passed data set queue (PDQ)
7	1 = old or modified data set 0 = new data set

BYTE 92: Conditional Disposition

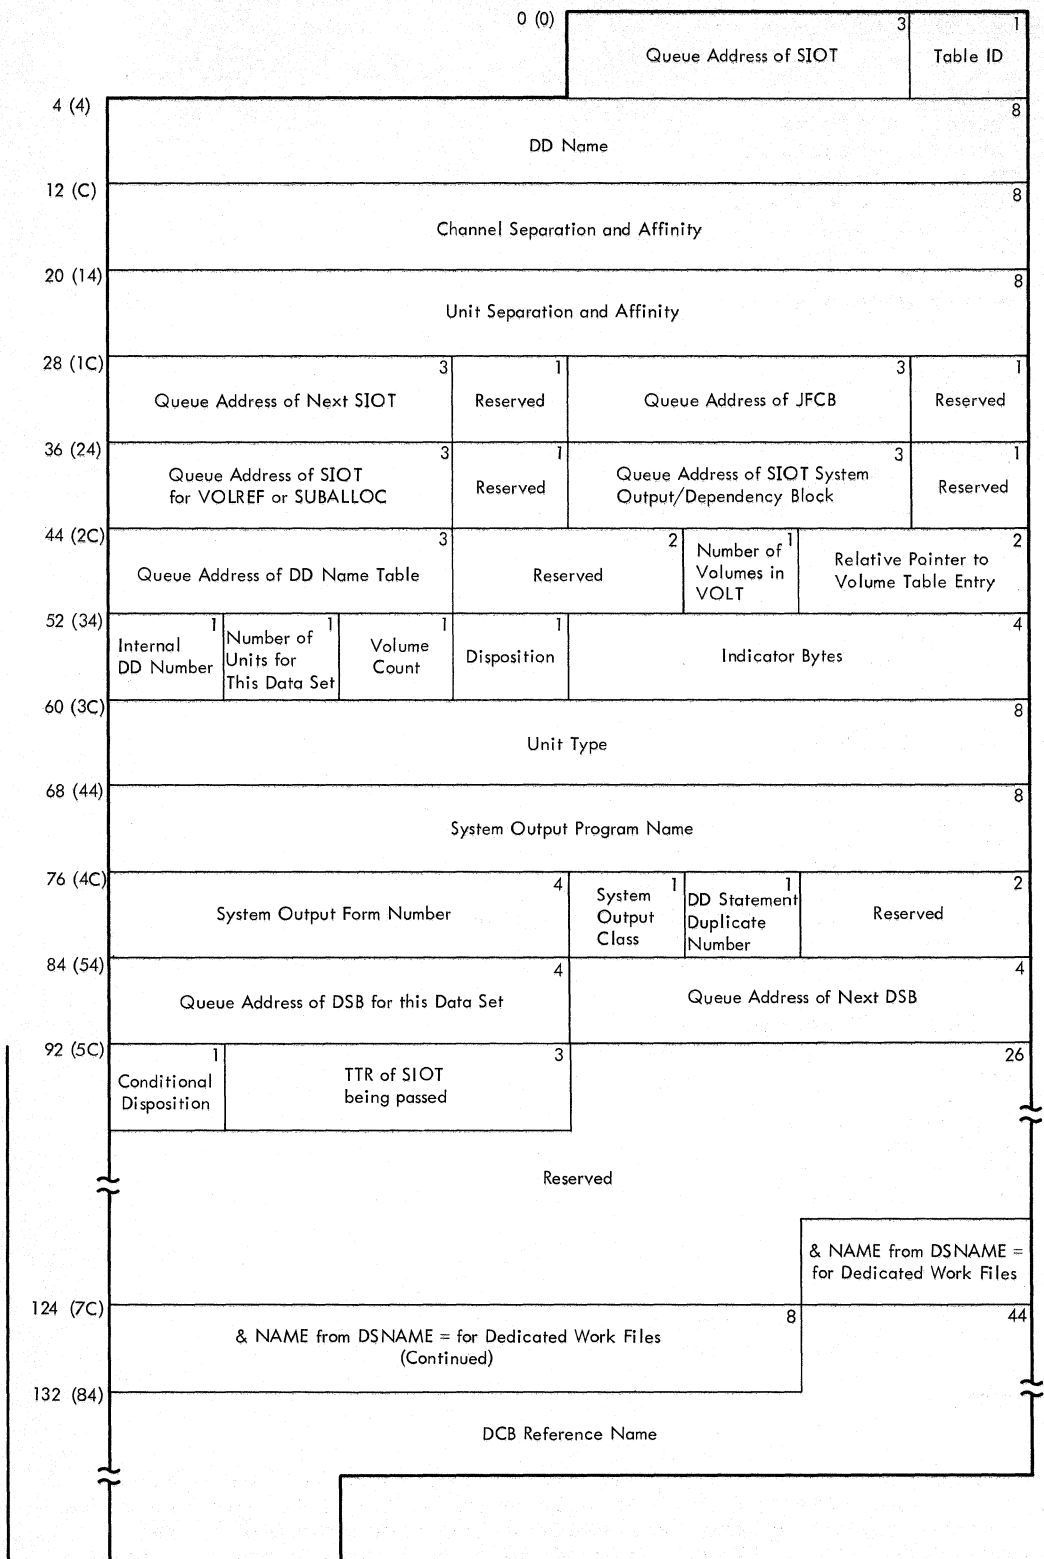
<u>Bit</u>	<u>Meaning</u>
0-3	Reserved
4	Keep data set
5	Delete data set
6	Catalog data set
7	Uncatalog data set

Mapping Macro Instruction: IEFASIOT

TASK INPUT/OUTPUT TABLE (TIOT)

Description: The Task Input/Output Table (TIOT) (Figure 43) provides data management routines with the addresses of the JFCBs and devices allocated to the data sets in a job step or system task. It is constructed by the I/O device allocation routine in main storage obtained from subpool zero. The allocation routine also places a copy of the TIOT on the appropriate job class queue with the other tables for the job step. After the step completes processing, the TIOT is brought in from the job queue and placed in the upper portion of the partition. The step is then terminated, and the TIOT is deleted.

For further information on the TIOT, see IBM System/360 Operating System: System Control Blocks, Form C28-6628.



•Figure 42. Step Input/Output Table (SIOT)

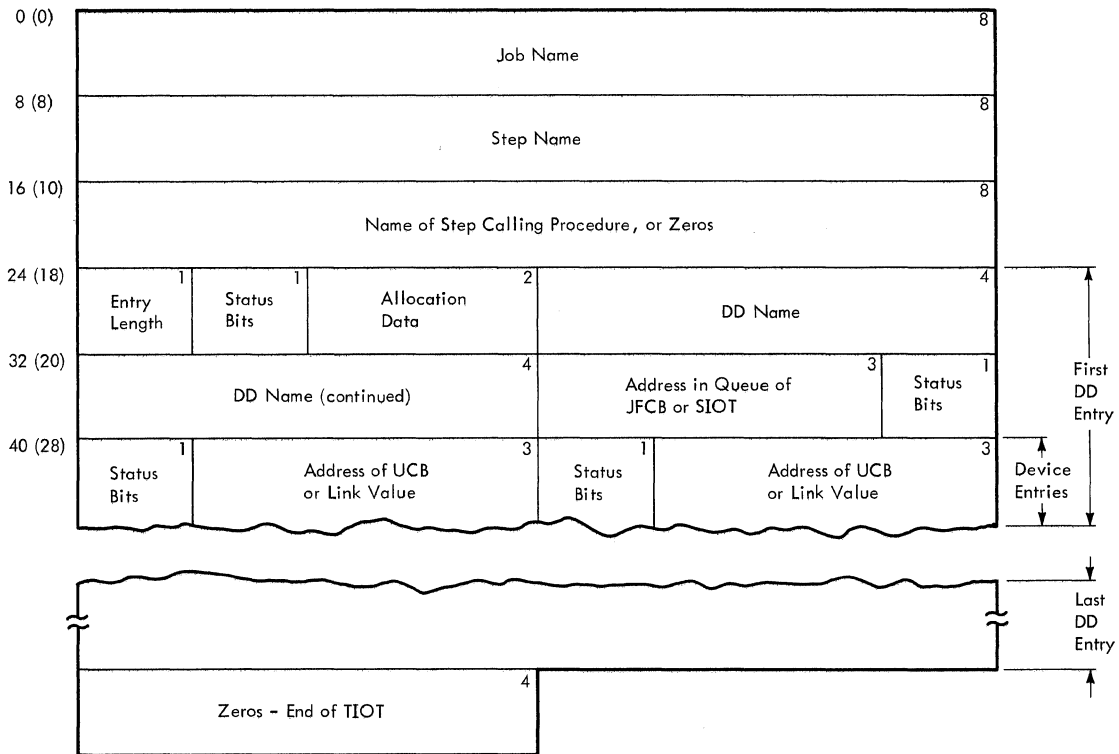


Figure 43. Task Input/Output Table (TIOT)

Appendix B: MFT Modules

This appendix contains a table of unique MFT modules, a group of tables showing the modules of each major component, a list matching entry point and control section names with source module names, and a brief description of each of the modules used by MFT. If you are looking for a specific module and know only the major component and routine name, use Tables 3-14 which give a cross-reference to the source module. The source modules are in turn listed alphanumerically for easy access. If you know the source module name, go directly to the module descriptions.

Unique MFT Modules

Table 2 lists all modules that are unique to MFT. This table is organized alphabetically by major component.

Table 2. MFT Modules

<u>ABEND:</u>	<u>I/O Device Allocation:</u>	<u>Queue Management:</u>
IEAGTM05	IEFSD551	IEFSD514
IEAGTM06	IEFSD552	IEFSD572
IEAGTM08	IEFSD557	
IEAGTM09		
<u>Communications Task:</u>	<u>Master Scheduler Task:</u>	
IEECIR45	IEECIR50	
IEEVWTOR	IEEDFIN1	
	IEEDFIN2	
<u>Initiator:</u>	IEEDFIN3	<u>Reader/Interpreter:</u>
IEFSD510	IEEDFIN4	IEFSD530
IEFSD511	IEEDFIN5	IEFSD531
IEFSD512	IEEDFIN6	IEFSD532
IEFSD513	IEEDFIN7	IEFSD533
IEFSD515	IEEDFIN8	IEFSD536
IEFSD516	IEEDFIN9	IEFSD537
IEFSD517	IEESD561	
IEFSD518	IEESD562	
IEFSD519	IEESD563	<u>System Task Control:</u>
IEFSD540	IEESD564	IEESD590
IEFSD541	IEESD565	IEESD591
IEFSD553	IEESD566	IEESD592
IEFSD554	IEESD571	IEFSD534
IEFSD555	IEFSD569	IEFSD535
IEFSD556		IEFSD587
IEFSD558	<u>Nucleus:</u>	IEFSD588
IEFSD559	IEESD568	
IEFSD589	IEFSD567	
IEFSD598		
IEFSD599		

Major Component Modules

Tables 3 through 14 list all MFT modules according to major component. The tables appear in alphabetical order by component name. Within each component, routine names are listed alphabetically with a cross-reference to the module name.

Table 3. ABEND Modules

Routine	Source Module
ABDUMP	IEAATM04
Indicative Dump	IEAATM03
Initialization	IEAGTM00
Input/Output Purge	IEAGTM06
Linkage	IEAATM01
Main Storage Allocation	IEAATM02
Termination	IEAGTM05
DAR Core Image Dump	IEAGTM08
DAR Task Reinstatement	IEAGTM09

Table 4. Communication Task Modules

Console Device Processor	IEECVPM
Console Interrupt	IEECVCRA
External Interrupt	IEECVCRX
Initialization Routine	IEECVCTI
Purge RQE	IEECVED2
Router	IEECVCTR
Wait	IEECVCTW
Write-to-Operator	IEECVWTO
Write-to-Operator-With-Reply	IEEVWTOR
EXCP OPEN/CLOSE	IEECVOC
MCS Comm Task Router	IEECMAWR
MCS Console Switch	IEECMCSW
MCS Device Interface	IEECMDSV
MCS 1052 Device Support	IEECMPMX
MCS 1403/1443 Device Support	IEECMPMP
MCS 2540 Device Support	IEECMPMC
MCS 2740 Device Support	IEEC2740
MCS Delete Operator Message	IEECMDOM
MCS WTCR Processor (SVC 35)	IEECMWSV
WTCR Purge (End of Job)	IEAGTM07

Table 5. Initiator Modules

Alternate Step Deletion	IEFSD516
Data Set Integrity	IEFSD541
Dequeue by Jobname Interface	IEFSD519
ENQ/DEQ Purge	IEFSD598
Job Deletion	IEFSD517
Job Initiation	IEFSD511
Job Selection	IEFSD510
Job Suspension	IEFSD168
Linkage to IEFSD510	IEFSD555
Linkage to IEFSD511	IEFSD558
Linkage to IEFSD512	IEFSD553
Linkage to IEFSD515	IEFSD559
Linkage to IEFSD516	IEFSD554
Linkage to IEFSD534	IEFSD589
Linkage to IEFSD541	IEFSD540
Partition Recovery	IEFSD518
Problem Program Interface	IEFSD513
Set Problem Program State	IEFSD556
Small Partition Module	IEFSD599
Step Deletion	IEFSD515
Step Initiation	IEFSD512

Table 6. I/O Device Allocation Modules

Routine	Source Module
Allocation Control	IEFXCSSS
Allocation Entry	IEFSD21Q
Allocation Exit	IEFSD41Q
Allocation Recovery Messages	IEFSJMSG
Allocation Recovery	IEFXJIMP
Automatic Volume Recognition	IEFXV001
Automatic Volume Recognition Messages	IEFVMSG
Automatic Volume Recognition Non-standard Label Routine	IEFXVNSL
DADSM Error Recovery	IEFXT003
Decision Allocation	IEFS5000
Demand Allocation	IEFWA000
Device Strikeout	IEFX300A
EXEC Statement Condition Code Processor	IEFVKIMP
EXEC Statement Condition Code Processor Messages	IEFVKMSG
External Action Messages	IEFWD001
External Action Interface	IEFWD000
JFCB Housekeeping Control and Allocate Processing	IEFVMLS1
JFCB Housekeeping Error Message Processing	IEFVMLS6
JFCB Housekeeping Error Messages	IEFVMLS7
JFCB Housekeeping Fetch DCB Processing	IEFVM2LS
JFCB Housekeeping GDG All Processing	IEFVM4LS
JFCB Housekeeping GDG Single Processing	IEFVM3LS
JFCB Housekeeping Patterning DSCB	IEFVM5LS
JFCB Housekeeping Unique Volume ID	IEFVM76
Mount Control-Volume Routine Linkage Module	IEFMCVCL
Linkage Module	IEFWCFAK
Linkage Module	IEFWDFA
Linkage Module	IEFWSWIN
Linkage Module	IEFXJFAK
Linkage to JFCB Housekeeping	IEFVMMS1
Linkage to JFCB Housekeeping	IEFVMFAK
Linkage to IEFXJIMP	IEFSD551
Linkage to IEFXJIMP	IEFSD552
Linkage to IEFXV001	IEFAVFAK
Linkage to Mount Control Volume	IEFCVFAK
Message Module	IEFWSTRT
Message Module	IEFXAMSG
Non-Recovery Error	IEFXKIMP
Non-Recovery Error Messages	IEFXKMSG
Separation Strikeout	IEFXH000
Space Request	IEFXT00D
VARY Interface and TICT Compression	IEFXT002
TIOT Construction	IEFWCIMP
Unsolicited Device Interrupt Handler	IEFVPOST
Wait for Space Decision	IEFSD097
Wait for Unallocation	IEFSD195

Table 7. Interpreter Modules

Routine	Source Module
Command Statement	IEFVHM
CPO Allocation Subroutine	IEFVSD12
CPC	IEFVHG
Continuation Statement	IEFVBC
DD* Statement Generator	IEFVHB
DD Statement Processor	IEEFVDA
Data Set Name Table Construction	IEFVDBSD
Dictionary Entry	IEFVGI
Dictionary Search	IEFVGS
End-of-File	IEFVHAA
EXEC Statement Processor	IEFVEA
Get Parameter	IEFVGK
Get	IEFVHA
Housekeeping	IEFVHHB
Initialization	IEFVH1
Initialization	IEFVH2
Interface	IEFSD533
Job and Step Enqueue	IEFVHH
Job Statement Processor	IEFVHA
Job Validity Check	IEFVHEC
Linkage Module	IEFSD537
Message Module	IEFVGM1
Message Module	IEFVGM2
Message Module	IEFVGM3
Message Module	IEFVGM4
Message Module	IEFVGM5
Message Module	IEFVGM6
Message Module	IEFVGM7
Message Module	IEFVGM8
Message Module	IEFVGM9
Message Module	IEFVGM10
Message Module	IEFVGM11
Message Module	IEFVGM12
Message Module	IEFVGM13
Message Module	IEFVGM14
Message Module	IEFVGM15
Message Module	IEFVGM16
Message Module	IEFVGM17
Message Module	IEFVGM18
Message Module	IEFVGM70
Message Module	IEFVGM78
Message Processing	IEFVGM
Null Statement	IEFVHL
Operator Message	IEFSD536
Post-Scan	IEFVHF
Pre-Scan Preparation	IEFVHEB
Queue Management Interface	IEFVHQ
Router	IEFVHE
Scan	IEFVFA
SCD Construction	IEFVSD13
Symbolic Parameter Processing	IEFVFB
Termination	IEFVHN
Test and Store	IEFVGT
Transient Reader Restore	IEFSD531
Transient Reader Suspend	IEFSD530
Transient Reader Suspend Tests	IEFSD532
Vary Identification	IEFVHCB

Table 8. Master Scheduler Modules

Routine	Source Module
DEFINE Command Final Processor	IEEDFIN9
DEFINE Final Processor	IEEDFIN3
DEFINE Initialization	IEEDFIN1
DEFINE Keyword Scan	IEEDFIN7
DEFINE Listing	IEEDFIN4
DEFINE Message	IEEDFIN5
DEFINE Syntax Check and Router	IEEDFIN2
DEFINE System Reinitialization	IEEDFIN8
DEFINE Time-Slice Syntax Check	IEEDFIN6
DISPLAY A	IEESD566
Look-up Routine	IEEVFRX
Queue Search	IEESD564
Queue Search Setup	IEESD563
Service	IEESD565
Syntax Check	IEESD562
Time-Slice Syntax Check	IEEDFIN6
Wait/Router	IEECIR50
Resident Volume Initialization	IEFPRES
Message Module	IEFK1MSG
System Log Initialization	IEEVLIN
System Log Open Initializer	IEEVLIN2
System Log Output Writer	IEEVIOUT
System Log Dispatcher	IEEVIDSP
System Log Wait Routine	IEELWAIT
System Log SVC (SVC 36)	IEE0303F
System Log SVC (SVC 36 - second load)	IEE0403F
DISPLAY CONSOLES	IEEXEDNA
Master Scheduler Resident Data Area	IEESD568
Master Scheduler Initialization	IEESD569
User Dummy WIO/WTOR Exit	IEECVCTE
Console Initialization	IEECVCTI

Table 9. Queue Management Modules

Assign	IEFQASGQ
Assign/Start	IEFQAGST
Branch	IEFQMLK1
Control	IEFQBVMMS
Delete	IEFQDELQ
Dequeue	IEFQMDQQ
Dequeue by Jobname	IEFQLOCDQ
Dummy	IEFQMDUM
Enqueue	IEFQMNQQ
Interpreter/Queue Manager Interlock	IEFSD572
Message Module	IEFSD311
Queue Formatting	IEFORMAT
Queue Initialization	IEFSD055
Queue Manager Table Breakup	IEFSD514
Read/Write	IEFQMRAW
Resident Main Storage Reservation	IEFPRESD
Unchain	IEFQMUNQ

Table 10. SVC 34 Modules

Routine	Source Module
CSCB Creation	IEE0803D
CSCB Marking	IEE0703D
DEFINE, MOUNT, CANCEL	IEE0571
HALT	IEE1403D
Message Assembly	IEE0503D
Message Assembly	IEE2103D
Reply Processor	IEE1203D
Router	IEE0403D
SET Command Handler	IEE0903D
SET Command	IEE0603D
START and STOP INIT	IEE0561
Translator/Chain Manipulator	IEE0303D
VARY and UNLOAD	IEE1103D
MCS Reply Processor	IEE1A03D
MCS Reply Messages	IEE1B03D
RJE Commands	IEE1503D
LOG and WRITELOG Routine	IEE1603D
VARY ONGFX/OFFGFX Handler	IEE1703D
System Management Facilities	IEE2303D
MODE Command Handler	IEE2603D
DISPLAY R Handler	IEE2903D
VARY and UNLOAD Processor II	IEE3103D
HARDCPY Message routine	IEE4103D
VARY Scan and Router	IEE4203D
VARY MSTCONS Handler	IEE4303D
VARY Keyword Scan	IEE4403D
STOP Command Handler	IEE4503D
VARY ON/OFFLINE of Consoles and Message Handler	IEE4603D
VARY HARDCPY Handler	IEE4703D
VARY CONSOLE Message routine	IEE4803D
VARY CONSOLE Handler	IEE4903D

Table 11. System Output Writer Modules

Class Name Setup	IEFSD081
Command Processing	IEFSD083
Data Set Delete	IEFSD171
Data Set Writer Interface	IEFSD070
DSB Handler	IEFSD085
Initialization	IEFSD080
Job Separator	IEFSD094
Linkage Module	IEF078SD
Linkage Module	IEF079SD
Linkage Module	IEF082SD
Linkage Module	IEF083SD
Linker	IEFSD078
Linkage to Queue Manager Delete	IEFSD079
Main Logic	IEFSD082
Message Module	IEFSD096
Print Line	IEFSD095
Put	IEFSD089
SMB Handler	IEFSD086
Standard Writer	IEFSD087
Transition	IEFSD088
Wait	IEFSD084

Table 12. System Restart Modules

Routine	Source Module
Delete	IEFSD303
Initialization	IEFSD300
Jobnames Table	IEFSD302
Linkage Module	IEF300SD
Linkage Module	IEF304SD
Message Module	IEFSD312
Purge Queue Construction	IEFSD301
Scratch Data Sets	IEFSD304
Scratch Data Sets	IEFSD308
TTR and NN to MBCCCHR Conversion	IEFSD310

Table 13. System Task Control Modules

Allocation Interface	IEEVACTL
Internal JCL Reader	IEEVICLR
Interpreter Control	IEEVRCTL
JCL Edit	IEEVJCL
Linkage to IEFSD535	IEFSD587
Linkage to IEE534SD	IEFSD588
Linker	IEESD591
Link-Table	IEEVINKT
LPSW	IEFSD534
Message Writer	IEEVMSG1
Message Writer	IEEVMSG
Message Writing	IEEVOMSG
POST	IEESD592
Problem Program Mode	IEFSD535
QMPA Builder	IEEVMRA
START Syntax Check	IEEVSTAR
Termination Interface	IEEVTCTL
Write TIOT on Disk	IEESD590

Table 14. Termination Modules

Routine	Source Module
Disposition and Unallocation Messages	IEFZGMSG
VARY Interface and Disposition and Unallocation Messages	IEFZHMSG
Disposition and Unallocation	IEFZGJB1
Disposition and Unallocation	IEFZGST1
DSB Processing	IEFYTVMS
Job Statement Condition Code Processor	IEFVJIMP
Job Statement Condition Code Processor Messages	IEFVJMSG
Job Termination Control	IEFZAJB3
Job Termination Exit	IEFSD31Q
Message Blocking	IEFYSVMS
Message Module	IEFWTERM
Message	IEFIDMPM
Restart Preparation	IEFRPREP
Step Termination Control	IEFYNIMP
Step Termination Control Routine Messages	IEFYNMSG
Step Termination Data Set Driver	IEFYPIB3
Step Terminate Exit	IEFSD22Q
Step Termination Messages	IEFYPMMSG
System Output Interface	IEFSD017
Termination Entry	IEFSD42Q
User Accounting Routine Linkage	IEFACTLK
User Dummy Accounting	IEFACTFK

Module Descriptions

This section contains a brief description of each of the modules used by MFT. An alphabetic list of the entry point and control section names, together with the name of the module that contains them, is provided to allow cross-referencing between modules. Modules are listed alphabetically by module name; associated with each module is a descriptive name, which indicates the major component of the system to which the module belongs. Each module contains a brief statement of the purpose of the module. Where applicable, the description includes the names of the module's entry points, the names of the modules to which it passes control, the major tables and work areas to which it refers, its attributes, the names of the control sections it contains, and a page reference to the detailed writeup in the Job Management section.

Entry Point or Control Section Name	Module Name	Entry Point or Control Section Name	Module Name
GO	IEFSD515	IEESD592	IEESD592
IEAGENQ1	IEAGENQ1	IEEVACTL	IEEVACTL
IEAGENQ2	IEAGENQ2	IEEVICLR	IEEVICLR
IEAQCT00	IEE0903D	IEEVJCL	IEEVJCL
IEAOTI01	IEAOTI01	IEEVLIN	IEEVLIN
IEEBA1	IEECVCRA	IEEVIDSP	IEEVLDSP
IEEBC1PE	IEECVCRX	IEEVLNKT	IEEVLNKT
IEECIR45	IEECVCTW	IEEVLOUT	IEEVLOUT
IEECIR50	IEECIR50	IEEVMSG1	IEEVMSG1
IEECMDOM	IEECMDOM	IEEVOMSG	IEEVOMSG
IEECMDSV	IEECMDSV	IEEVRCTL	IEEVRCTL
IEECMWRT	IEECMAWR	IEEVRFRX	IEEVRFRX
IEECMWSV	IEECMWSV	IEEVSMBA	IEEVSMBA
IEECMWTL	IEECMWTL	IEEVSMG	IEEVSMG
IEECVCTI	IEECVCTI	IEEVSTRT	IEEVSTAR
IEECVCTR	IEECVCTR	IEEVICTL	IEEVTCTL
IEECVCTW	IEECMAWR	IEEXEDNA	IEEXEDNA
IEECVPM	IEECVPM	IEE0303D	IEE0303D
IEECVPRG	IEECVED2	IEE0303F	IEE0303F
IEECVXIT	IEECVCTE	IEE0403D	IEE0403D
IEEDFIN1	IEEFIN1	IEE0403F	IEE0403F
IEEDFIN2	IEEDFIN2	IEE0503D	IEE0503D
IEEDFIN3	IEEDFIN3	IEE0603D	IEE0603D
IEEDFIN4	IEEDFIN4	IEE0703D	IEE0703D
IEEDFIN5	IEEDFIN5	IEE0803D	IEE0803D
IEEDFIN6	IEEDFIN6	IEE1103D	IEE1103D
IEEDFIN7	IEEDFIN7	IEE1203D	IEE1203D
IEEDFIN8	IEEDFIN8	IEE1403D	IEE1403D
IEEDFIN9	IEEDFIN9	IEE1603D	IEE1603D
IEEDPART	IEEDFIN2	IEE1703D	IEE1703D
IEELG02	IEELOG02	IEE1A03D	IEE1A03D
IEELWAIT	IEELWAIT	IEE1B03D	IEE1B03D
IEEMSER	IEESD568	IEE2103D	IEE2103D
IEEPDISC	IEEPDISC	IEE2303D	IEE2303D
IEESD562	IEESD562	IEE2603D	IEE2603D
IEESD563	IEESD563	IEE2903D	IEE2903D
IEESD564	IEESD564	IEE3103D	IEE3103D
IEESD565	IEESD565	IEE4103D	IEE4103D
IEESD566	IEESD566	IEE4203D	IEE4203D
IEESD567	IEESD567	IEE4303D	IEE4303D
IEESD590	IEESD590	IEE4403D	IEE4403D
IEESD591	IEESD591	IEE4503D	IEE4503D

Entry Point or Control Section Name	Module Name
IEE4603D	IEE4603D
IEE4703D	IEE4703D
IEE4803D	IEE4803D
IEE4903D	IEE4903D
IEE591SD	IEESD591
IEFACTLK	IEFACTFK
IEFACTLK	IEFACTLK
IEFACTRT	IEFACTRT
IEFCVCL1	IEFCVFAK
IEFCVCL1	IEFMCVOL
IEFCVCL2	IEFCVFAK
IEFCVOL2	IEFMCVOL
IEFCVCL3	IEFCVFAK
IEFCVCL3	IEFMCVOL
IEFDPCST	IEFVPOST
IEFSDSRP	IEFSDSRP
IEFICR	IEEVICIR
IEFIDMPM	IEFIDMPM
IEFIRC	IEFSD533
IEFJOB	IEFQRES
IEFKG	IEFSD532
IEFORMAT	IEFORMAT
IEFPH2	IEFSD531
IEFQAGST	IEFQAGST
IEFQASGN	IEFQASGQ
IEFQASNM	IEFQASGQ
IEFQDELE	IEFQDELQ
IEFQMDQ2	IEFQMDQQ
IEFQMDUM	IEFQMDUM
IEFQMNQ2	IEFQMNQQ
IEFQMSSS	IEFQBVMS
IEFQMSSS	IEFQMDUM
IEFQMSSS	IEFQMLK1
IEFQMRW	IEFQMRW
IEFQMUNC	IEFQMUNQ
IEFRCLN1	IEFRCLN1
IEFRCLN2	IEFRCLN2
IEFRPREP	IEFRPREP
IEFRSTRT	IEFRSTRT
IEFSD012	IEFVSD12
IEFSD017	IEFSD017
IEFSD055	IEFSD055
IEFSD068	IEFSD168
IEFSD070	IEFSD070
IEFSD071	IEFSD171
IEFSD078	IEFSD078
IEFSD078	IEF078SD
IEFSD079	IEFSD079
IEFSD079	IEF079SD
IEFSD080	IEFSD080
IEFSD081	IEFSD081
IEFSD082	IEF082SD

Entry Point or Control Section Name	Module Name
IEFSD082	IEFSD082
IEFSD083	IEFSD083
IEFSD083	IEF083SD
IEFSD084	IEFSD084
IEFSD085	IEFSD085
IEFSD086	IEFSD086
IEFSD087	IEFSD087
IEFSD088	IEFSD088
IEFSD089	IEFSD089
IEFSD090	IEFVSD13
IEFSD094	IEFSD094
IEFSD095	IEFSD095
IEFSD095	IEFSD195
IEFSD096	IEFSD096
IEFSD097	IEFSD097
IEFSD300	IEF300SD
IEFSD300	IEFSD300
IEFSD301	IEFSD301
IEFSD302	IEFSD302
IEFSD303	IEFSD303
IEFSD304	IEFSD304
IEFSD304	IEF304SD
IEFSD305	IEFSD305
IEFSD308	IEFSD308
IEFSD310	IEFSD310
IEFSD311	IEFSD311
IEFSD312	IEFSD312
IEFSD510	IEFSD510
IEFSD511	IEFSD511
IEFSD512	IEFSD512
IEFSD512	IEFSD553
IEFSD513	IEFSD513
IEFSD514	IEFSD514
IEFSD515	IEFSD515
IEFSD516	IEFSD516
IEFSD517	IEFSD517
IEFSD518	IEFSD518
IEFSD519	IEFSD519
IEFSD530	IEFSD530
IEFSD531	IEFSD531
IEFSD534	IEFSD534
IEFSD535	IEFSD535
IEFSD537	IEFSD537
IEFSD540	IEFSD540
IEFSD541	IEFSD541
IEFSD554	IEFSD554
IEFSD555	IEFSD555
IEFSD556	IEFSD556
IEFSD557	IEFSD557
IEFSD558	IEFSD558
IEFSD559	IEFSD559
IEFSD567	IEFSD567

Entry Point or Control Section Name	Module Name
IEFSD569	IEFSD569
IEFSD572	IEFSD572
IEFSD573	IEFSD572
IEFSD587	IEFSD587
IEFSD588	IEFSD588
IEFSD589	IEFSD589
IEFSD597	IEFSD597
IEFSD598	IEFSD598
IEFSD599	IEFSD599
IEFSD71M	IEFSD171
IEFSD83M	IEFSD083
IEFSD85M	IEFSD085
IEFSD86M	IEFSD086
IEFSD877	IEFSD087
IEFSD897	IEFSD089
IEFSMR	IEFRSTRT
IEFUCBL	IEFWA000
IEFVAWAT	IEFSD195
IEFVDA	IEFVDA
IEFVDBSD	IEFVDBSD
IEFVEA	IEFVEA
IEFVFA	IEFVFA
IEFVFB	IEFVFB
IEFVGI	IEFVGI
IEFVGR	IEFVGR
IEFVGM	IEFVGM
IEFVGM1	IEFVGM1
IEFVGM2	IEFVGM2
IEFVGM3	IEFVGM3
IEFVGM4	IEFVGM4
IEFVGM5	IEFVGM5
IEFVGM6	IEFVGM6
IEFVGM7	IEFVGM7
IEFVGM8	IEFVGM8
IEFVGM9	IEFVGM9
IEFVGM10	IEFVGM10
IEFVGM11	IEFVGM11
IEFVGM12	IEFVGM12
IEFVGM13	IEFVGM13
IEFVGM14	IEFVGM14
IEFVGM15	IEFVGM15
IEFVGM16	IEFVGM16
IEFVGM17	IEFVGM1M
IEFVGM18	IEFVGM18
IEFVGM19	IEFVGM19
IEFVGM70	IEFVGM70
IEFVGM78	IEFVGM78
IEFVGS	IEFVGS
IEFVGT	IEFVGT
IEFVHA	IEFVHA
IEFVHAA	IEFVHAA
IEFVHB	IEFVHB

Entry Point or Control Section Name	Module Name
IEFVHC	IEFVHC
IEFVHCB	IEFVHCB
IEFVHE	IEFVHE
IEFVHEB	IEFVHEB
IEFVHEC	IEFVHEC
IEFVHF	IEFVHF
IEFVHG	IEFVHG
IEFVHH	IEFVHH
IEFVHHB	IEFVHHB
IEFVHL	IEFVHL
IEFVHM	IEFVHM
IEFVHN	IEFVHN
IEFVHQ	IEFVHQ
IEFVHR	IEFSD536
IEFVH1	IEFVH1
IEFVH2	IEFVH2
IEFVJ	IEFVJIMP
IEFVJA	IEFVJA
IEFVJMSG	IEFVJMSG
IEFVK	IEFVKIMP
IEFVKMJ1	IEFVKMSG
IEFVKMSG	IEFVKMSG
IEFVM	IEFVMLS1
IEFVMCVL	IEFVMFAK
IEFVMCVL	IEFVMLS1
IEFVMQMI	IEFVMLS1
IEFVMSSGR	IEFVMLS6
IEFVM1	IEFVMLS1
IEFVM1	IEFVMMS1
IEFVM2	IEFVM2LS
IEFVM3	IEFVM3LS
IEFVM4	IEFVM4LS
IEFVM5	IEFVM5LS
IEFVM6	IEFVMLS6
IEFVM7	IEFVMLS7
IEFVM76	IEFVM76
IEFVRR	IEFVRR
IEFVRRCA	IEFVRR
IEFVRRCB	IEFVRR
IEFVRR1	IEFVRR1
IEFVRR2	IEFVRR2
IEFVRR3	IEFVRR3
IEFVSDRA	IEFVSDRA
IEFVSDRD	IEFVSDRD
IEFVSMBR	IEFVSMBR
IEFV15XL	IEFXJIMP
IEFV15XL	IEFSD551
IEFVR2AE	IEFVRR2
IEFVR3AE	IEFVRR3
IEFWA000	IEFWA000
IEFWA002	IEFWA000
IEFWA7	IEFWA000

Entry Point or Control Section Name	Module Name	Entry Point or Control Section Name	Module Name
IEFWC000	IEFWCFAK	IEFZH	IEFZHMSG
IEFWC000	IEFWCIMP	IEF085SD	IEFSD085
IEFWC002	IEFWCIMP	IEF086SD	IEFSD086
IEFWD000	IEFWDFAK	IEF850SD	IEFSD085
IEFWD000	IEFWD000	IEG056	IEAGENQ1
IEFWD001	IEFWD001	IEG056	IEAGENQ2
IEFWDMSG	IEFWD000	IGCXL07B	IEECMCSW
IEFWLISD	IEFSD21Q	IGCXM07B	IEECMCSW
IEFWSTRT	IEFWSVRT	IGCXN07B	IEECMCSW
IEFWSWIT	IEFWSWIN	IGCX007B	IEECMCSW
IEFWTERM	IEFWTERM	IGC0001C	IEAGTMOA
IEFW1FAK	IEFSD41Q	IGC0003E	IEECVWTC
IEFW1FAK	IEF41FAK	IGC0005B	IEFVSMBR
IEFW2FAK	IEFSD41Q	IGC00060	IEAAS00
IEFW2FAK	IEF41FAK	IGC0101C	IEAATM01
IEFW21SD	IEFSD21Q	IGC0103D	IGC0103D
IEFW21SD	IEFSD557	IGC0103E	IEEVWTOR
IEFW22SD	IEFSD22Q	IGC0111C	IEAATMOA
IEFW31SD	IEFSD31Q	IGC0201C	IEAATM02
IEFW41SD	IEFSD41Q	IGC0211C	IEAATM2A
IEFW41SD	IEF41FAK	IGC0221C	IEAATM2B
IEFW42SD	IEFSD42Q	IGC0301C	IEAATM03
IEFXA	IEFXCSSS	IGC0401C	IEAATM04
IEFXAMSG	IEFXAMSG	IGC048	IEAGENQ1
IEFXH000	IEFXH000	IGC048	IEAGENQ2
IEFXJMSG	IEFXJMSG	IGC0501C	IEAGTM05
IEFXJX5A	IEFSD552	IGC0601C	IEAGTM06
IEFXJX5A	IEFXJIMP	IGC0701C	IEAGTM00
IEFXJ000	IEFXJFAK	IGC0801C	IEAGTM08
IEFXJ000	IEFXJIMP	IGC0901C	IEAGTM09
IEFXKMSG	IEFXKMSG	IGC0907B	IEECMWT1
IEFXK000	IEFXKIMP	IGC0B01C	IEAATM0E
IEFXT000	IEFXT00D	IGC0C01C	IEAATM0C
IEFXT002	IEFXT002	IGC0D01C	IEAATM0D
IEFXT003	IEFXT003	IGC0E01C	IEAATM0E
IEFXVMSG	IEFXVMSG	IGC1803D	IEESD571
IEFXVNSL	IEFXVNSL	IGC1903D	IEESD561
IEFXV001	IEFAVFAK	IGF2603D	IGF2603D
IEFXV001	IEFXV001	IGF2703D	IGF2703D
IEFXV002	IEFXV002	LOC	IEFLOCDQ
IEFX3000	IEFX300A	LOCCAN	IEFLOCDQ
IEFX5000	IEFX5000	LOCDQ	IEFLOCDQ
IEFYN	IEFYNIMP	MSOFF	IEFXJMSG
IEFYNMSG	IEFYNMSG	MSRCV	IEFXJMSG
IEFYF	IEFYFJB3	MSSYS	IEFXJMSG
IEFYPMMSG	IEFYPMMSG	SD304MG1	IEFSD312
IEFYS	IEFYSVMS	SD304MG2	IEFSD312
IEFYT	IEFYTVMS	SD305MG1	IEFSD312
IEFZA	IEFZAJB3	SD55MSG1	IEFSD311
IEFZG	IEFZGST1	SD55MSG2	IEFSD311
IEFZGJ	IEFZGJB1	SD55MSG3	IEFSD311
IEFZGMSG	IEFZGMSG	SMALLGO	IEFSD599

Entry Point or Control Section Name	Module Name
SMALTERM	IEFSD515
SMALTERM	IEFSD559
STRMSG01	IEFYNMSG
VM7000	IEFVMLS1
VM7055	IEFVMLS1
VM7055AA	IEFVMLS1
VM7060	IEFVMLS1
VM7065	IEFVMLS1
VM7070	IEFVMLS1
VM7090	IEFVMLS1
VM7100	IEFVM2LS
VM7130	IEFVMLS1
VM7150	IEFVM3LS
VM7200	IEFVM4LS
VM7300	IEFVM5LS
VM7370	IEFVMLS1
VM7600	IEFVM76
VM7700	IEFVMLS1
VM7742	IEFJMLS1
VM7750	IEFVMLS1
VM7850	IEFVMLS1

Entry Point or Control Section Name	Module Name
VM7900	IEFVMLS1
VM7950	IEFVMLS1
XIIB32	IEFX5000
XTEA0	IEFXT002
XTEA1	IEFXT002
XTEB3	IEFXT002
XTP00	IEFXT00D
XTRDJ	IEFXT002
XUUB00	IEFXT003
XUUH06	IEFXT003
X33B42	IEFX300A
X55C86	IEFX5000
X55D3G	IEFX5000
YPPMSG1	IEFYPMMSG
YPPMSG2	IEFYPMMSG
XPS631	IEFZHMSG
ZG0E60	IEFZHMSG
ZK0D1	IEFZHMSG
ZK0E1	IEFZHMSG
ZPOQM	IEFZGJB1
ZPOQMGR1	IEFZGST1

Module Descriptions

IEAAST00: Supervisor - STAE Service Routine

This routine is entered when the STAE macro instruction (SVC 60) is issued. It creates, cancels, or overlays a STAE control block according to the options specified. It prepares the task to intercept the scheduled abnormal termination (ABEND) processing.

- Entry: IGC00060
- Attributes: Non-resident, reentrant
- Control Section: IGC00060

IEAATMOA: Supervisor -- ABEND Linkage Routine

This routine checks for valid and invalid recursions. For an invalid recursion, control is passed to the Damage Assessment Routine. For a valid recursion, a bit is set in the TCFLGS field of the TCB to prevent an ABDUMP from being attempted. IEAATMOA determines the amount of main storage required by ABEND, and transfers control to the appropriate ABEND load module.

- Entry: IGC0111C
- Exits: XCTL to IEAATM02 if main storage must be 'stolen'
 - to IEAATM03 if main storage is available and an indicative dump is requested
 - to IEAATM04 if main storage is available and ABDUMP is requested
 - to IEAATM05 if main storage is available and no dump is requested
 - to IEAATM08 if the scheduler is failing and no ABDUMP is requested
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0111C

IEAATMOB: Supervisor - ABEND/STAE Interface Routine

This routine is the first ABEND/STAE Interface load module to receive control from the ABEND routine when abnormal termination has been scheduled for a task operating in a STAE environment.

- Entry: IGC0B01C

- Attributes: Non-resident, reentrant
- Control Section: IGC0B01C

IEAATMOC: Supervisor - ABEND/STAE Interface Routine

- Entry: IGC0C01C
- Attributes: Non-resident, reentrant
- Control Section: IGC0C01C

IEAATMOD: Supervisor - ABEND/STAE Interface Routine

- Entry: IGC0D01C
- Attributes: Non-resident, reentrant
- Control Section: IGC0D01C

IEAATMOE: Supervisor - ABEND/STAE Interface Routine

- Entry: IGC0E01C
- Attributes: Non-resident, reentrant
- Control Section: IGC0E01C

IEAATM01: Supervisor -- ABEND Validity Check Routine

This routine performs validity checking of the MSS (main storage supervisor) queue, the load list, the active RB list, and the DEB queue. It dequeues invalid control blocks, or terminates the queue at the point of error, and sets bits in the ABEND appendage to the boundary box to indicate invalid control blocks found on one or more lists.

- Entry: IGC0101C
- Exit: XCTL to IEAATMOA
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0101C

IEAATM02: Supervisor -- ABEND Steal LRB Main Storage Routine

This routine 'steals' main storage needed for ABEND functions that cannot be obtained via a GETMAIN macro instruction. The main storage is stolen from programs represented by LRBs on the loaded program list.

- Entry: IGC0201C
- Exits: XCTL to IEAATM2A if there is no loaded program list or if enough main storage is not available from the LRBS
to IEAATM2B if IEAATM02 has acquired the necessary main storage
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0201C

IEAATM03: Supervisor -- ABEND Indicative Dump Routine

This routine accumulates the information for an indicative dump and stores it in main storage.

- Entry: IGC0301C
- Exit: XCTL to IEAGTM05
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0301C

IEAATM04: Supervisor -- ABEND/ABDUMP Routine

This routine determines if the user wants a full or partial ABDUMP, initiates the ABDUMP output, and calls an SVC 51 (ABDUMP).

- Entry: IGC0401C
- Exits: XCTL to IEAATM03 for an indicative dump if the DCB has failed to open
to IEAGTM05 for initialization of the next task
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0401C

IEAATM2A: Supervisor -- ABEND Steal Problem Program Main Storage Routine

This routine 'steals' main storage needed for ABEND functions from the lower end of the partition when it cannot be acquired either by a GETMAIN macro instruction or by the steal routine provided by IEAATM02.

- Entry: IGC0211C

- Exits: XCTL to IEAATM03 if indicative dump is requested
to IEAATM04 if ABDUMP is requested
to IEAGTM05 if no dump is requested or if a dump was previously attempted and failed
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0211C

IEAATM2B: Supervisor -- ABEND LRB Stack Routine

This routine moves the LRBS whose main storage was stolen by IEAATM02 to contiguous locations in the low end of the freed area and resets the chain pointers in the LRBS.

- Entry: IGC0221C
- Exits: XCTL to IEAATM03 if indicative dump is requested
to IEAATM04 if ABDUMP is requested
to IEAGTM05 if no dump is requested or if a dump was previously attempted and failed
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0221C

IEAGENQ1: Supervisor -- Enqueue Service Routine

This routine constructs and processes control blocks to serialize the use of resources in a multiprogramming environment.

- Entry: IEAGENQ1
- Exit: EXIT routine or to the dispatcher
- Tables/Work Areas: Minor QCB, Major QCB, Queue element
- Attributes: Reenterable
- Control Sections: IGC048 and IEG056

IEAGENQ2: Supervisor -- Shared DASD Enqueue Service Routine

This routine is the enqueue service routine for systems that include the Shared DASD option. It is identical to IEAGENQ1 except that additional processing is performed

when a shared direct-access device is requested through the RESERVE macro instruction.

- Entry: IEAGENQ2
- Exits: EXIT routine or to the dispatcher
- Tables/Work Areas: Minor QCB, Major QCB, Queue element
- Attributes: Reentrantable
- Control Sections: IGC048 and IEG056

IEAGTMOA: Supervisor -- ABEND STAE Test Routine

This routine prevents asynchronous exits and stores the completion code (if not previously stored). It determines if control should be returned to STAE after a purge error, if the Graphics Abend Exit routine should be entered, and if a valid STAE is in effect.

- Entry: IGC0001C
- Exits: XCTL to IEAGTM00 to continue ABEND processing
to IEAATMOB if a valid STAE is in effect
to IEAGTM08 if an invalid ABEND recursion or a primary DAR recursion has occurred
to IEAGTM09 if a secondary DAR recursion has occurred
to IEAATMOB if an ABEND was issued by the Purge routine during STAE processing
EXIT to caller if graphics program

- Attributes: Reentrant, disabled, privileged

- Control Section: IGC0001C

IEAGTM00: Supervisor -- ABEND Initialization Routine

This routine provides purging for IQEs and WTOR requests, and cancels the task timer element.

- Entry: IGC0701C
- Exit: XCTL to IEAATMOD if IEAGTM00 was entered from STAE
to IEAGTM05 if this is a normal end
to IEAGTM06 if this is an abnormal end

- Attributes: Reentrant, disabled for SVC interruptions, privileged

- Control Section: IGC0701C

IEAGTM05: Supervisor -- ABEND Termination Routine

This routine closes all data sets, purges the timer queue, resets the TCB fields, frees main storage, and transfers control to the job scheduler.

- Entry: IGC0501C
- Exits: XCTL to IEFSD51K for scheduler-size partitions
to IEFSD599 for small partitions
- Attributes: Reentrant, disabled for external and I/O interruptions, privileged
- Control Section: IGC0501C

IEAGTM06: Supervisor -- ABEND Input/Output Purge Routine

This routine purges I/O operations in process and outstanding I/O requests.

- Entry: IGC0601C
- Exit: XCTL to IEAATM01 for normal exit
to IEAGTM09 if a system task or "must complete" task
- Attributes: Reentrant, disabled for I/O and external interruptions, privileged
- Control Section: IGC0601C

IEAGTM08: Supervisor -- DAR Core Image Dump Routine

This routine attempts to write a core image dump to a preallocated data set. It also processes primary DAR recursions.

- Entry: IGC0801C
- Exits: XCTL to IEAGTM05 if called for a dump only
to IEAGTM09 to continue DAR processing
- Attributes: Refreshable, disabled, privileged
- Control Section: IGC0801C

IEAGTM09: Supervisor -- DAR Task Reinstatement Routine

This routine attempts reinstatement of failing system tasks. It also processes secondary DAR recursions and failing tasks which are in "Must Complete" status.

- Entry: IGC0901C
- Exits: XCTL to IEAGTM05 for non-system task with or without non-critical resources
Branch to dispatcher for secondary DAR recursion, critical resources, or system task failure.
- Attributes: Refreshable, disabled, privileged
- Control Section: IGC0901C

IEA0TI01: Supervisor -- Timer Second Level Interruption Handler

This routine maintains the timer queue when the timer option is not specified during system generation. It handles only the normal six hour interruptions.

- Entry: IEA0TI01
- Exit: To Timer/External FLIH
- Tables/Work Areas: SHPC, T4PC, LTPC
- Attributes: Reenterable, disabled for system interruptions, resident, supervisor mode
- Control Section: IEA0TI01

IEECIR50: Master Scheduler -- Wait/Router Routine

This routine waits until a command is issued, analyzes the command and passes control to the appropriate processing module.

- Entry: IEECIR50
- Exits: IEESD562, IEEDFIN1
- Attributes: Read-only, reenterable, resident in nucleus.
- Control Section: IEECIR50
- Page Reference: 49

IEECMAWR: Communications Task -- Router Module

This module waits for the posting of a communications task ECB, determines the type of interruption service required (external,

attention, I/O, WTO, or DCM), and passes control to other communications task modules for further processing.

- Entry: IEECMWRT
- Exit: IEECMCSW, IEECMDSV, IEECMWSV, IEECMWTL, IEECMDOM, Dispatcher
- Tables/Work Areas: CVT, EIL, UCM, WQE
- Attributes: Reentrant, refreshable
- Control Section: IEECVCTW

IEECMSW: Communications Task -- Console Switch Module

This routine provides console switching as a result of an unrecoverable I/O error on a console device, as a result of an external interruption, or as a result of a VARY command, and provides hard copy switching from a console device of SYSLOG.

- Entry: IGCXL07B
- Exit: IEECMAWR, IEECMDSV
- Tables/Work Areas: CVT, CXSA, RQE, UCM, WQE
- Attributes: Reentrant, refreshable
- Control Sections: IGCXL07B, IGXCM07B, IGXCN07B, IGXC007B

IEECMDOM: Communications Task -- DOM Service Module

This module marks for deletion specified WQEs on the system output queue.

- Entry: IEECMDOM
- Exit: IEECMAWR, IEECVDT1
- Tables/Work Areas: CVT, DCM, UCM, WQE
- Attributes: Reentrant, refreshable
- Control Section: IEECMDOM

IEECMDSV: Communications Task -- Device Service Module

This module provides the interface with device support processors and provides console and system output queue management.

- Entry: IEECMDSV
- Exit: IEECMAWR, IEECMWSV, IEECMCSW, Device Support Processors
- Tables/Work Areas: IEEBASEB, CVT, EIL, UCM, WQE

- Attributes: Reentrant, refreshable
- Control Section: IEECMDSV

- Tables/Work Areas: UCM
- Attributes: Reenterable
- Control Section: IEIBC1PE
- Page Reference: 45

IEECMWSV: Communications Task -- WTO(R) Service Module

This module puts unprocessed WQES on appropriate console output queues

- Entry: IEECMWSV
- Exit: IEECMDSV, IEECMAWR
- Tables/Work Areas: UCM, WQE
- Attributes: Reentrant, refreshable
- Control Section: IEECMWSV

IEECMWTI: Communications Task -- NIP Message Buffer Writer Module

This module issues SVC 36 to write NIP messages to SYSLOG. If SYSLOG has not been initialized or not specified as the hard copy log, it issues SVC 35 to write the NIP messages to the operator.

- Entry: IEECMWTI
- Exit: Return to caller
- Control Section: IGC0907B

IEECVCRA: Communications Task -- Console Interruption Routine

This routine notifies the wait routine that a console read has been requested.

- Entry: IEEBAl
- Exit: Return to IOS
- Tables/Work Areas: ECB, UCM, UCB
- Attributes: Reenterable
- Control Section: IEFPAl
- Page Reference: 44

IEECVCRX: Communications Task -- External Interruption Routine

This routine switches control from the primary console device to an alternate console device when an external interruption occurs.

- Entry: IEIBC1PE
- Exit: Return to IOS

IEECVCTE: Communications Task -- User Dummy WTO/WTOR Exit Routine

This routine takes the place of the user's WTO/WTOR exit routine when an exit routine was specified at system generation, but none was supplied.

- Entry: IEECVXIT, from IEECMWSV
- Exit: Return to caller
- Control Section: IEECVXIT

IEECVCTI: Console Initialization Routine

This routine prints out the NIP message buffer in systems with the MCS option, and initializes the console configuration.

- Entry: IEECVCTI, from IEESD569
- Exit: To IEESD569
- Tables/Work Areas: CVT, EIL, UCB, and UCM
- Attributes:
- Control Section: IEECVCTI

IEECVCTR: Communications Task -- Router Routine

This routine determines the type of request or interruption that occurred, and passes control to the appropriate processing routine.

- Entry: IEECVCTR
- Exits: XCTL to IEECVPMX (IGC0107B), IEECVPMC (IGC1107B), or IEECVPMP (IGC2107B)
- Tables/Work Areas: UCM, SVRB, UCB
- Attributes: Reenterable
- Control Section: IEECVCTR
- Page Reference: 44

IEECVCTW: Communications Task -- Wait Routine

This routine waits on all communications task ECBs associated with WTC/WTOR macro instructions.

- Entry: IEECIR45
- Exit: None
- Tables/Work Areas: TCB, ECB, UCM
- Attributes: Reenterable
- Control Section: IEECIR45
- Page Reference: 44

IEECVED2: Communications Task -- Purge RQE Routine

This routine scans and purges all outstanding request queue elements (RQEs) pertaining to the terminating task.

- Entry: IEECVPRG
- Exits: End-of-task, and ABEND
- Tables/Work Areas: RQE, WQE, JCM, CVT
- Attributes: Reenterable
- Control Section: IEECVPRG

IEECVPM: Communications Task -- Console Device Processor Routine

This routine performs console read and write operations and checks for errors.

- Entry: IEECVPM
- Exit: XCTL to IEECVCTR (IGC0007B)
- Tables/Work Areas: DCB, UCB, UCM
- Attributes: Reenterable
- Control Section: IEECVPM
- Page Reference: 44

IEECVWTO: Communications Task -- Write-to-Operator Routine

This routine processes all WTO macro instructions.

- Entry: IGC0003E
- Exit: Return to calling program
- Tables/Work Areas: WQE, UCM, CVT, RQE
- Attributes: Reenterable
- Control Section: IGC0003E
- Page Reference: 45

IEEDFIN1: Master Scheduler -- DEFINE Command Initialization Routine

This routine sets up data areas for partition definition, issues a DEFINE CCMAND BEING PROCESSED message to all active consoles, and passes control to the appropriate processing module.

- Entry: IEEDFIN1
- Exits: IEEDFIN3, IEEDFIN4, IEEDFIN5
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN1
- Page Reference: 50

IEEDFIN2: Master Scheduler -- DEFINE Command Syntax Check and Router Routine

This routine checks the syntax of DEFINE command statements. If a syntax error is discovered, the statement is ignored and an error message is issued. If the syntax is correct, the information is stored and control is passed to the appropriate routine.

- Entry: IEEDFIN2, IEEDPART
- Exits: IEEDFIN5, IEEDFIN6, IEEDFIN7
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN2
- Page Reference: 50

IEEDFIN3: Master Scheduler -- DEFINE Command Validity Check Routine

This routine determines that all information for the partition redefinition is correct, and passes control to the appropriate processing routine.

- Entry: IEEDFIN3
- Exits: IEEDFIN8, IEEREXIT
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN3
- Page Reference: 52

IEEDFIN4: Master Scheduler -- DEFINE Command Listing Routine

This routine lists partition definitions.

- Entry: IEEDFIN4
- Exits: IEEDFIN3, IEEDFIN5
- Attributes: Read-only, reenterable

- Control Section: IEEDFIN4

- Page Reference: 52

IEEDFIN5: Master Scheduler -- DEFINE Command Message Routine

This routine contains texts for operator messages required for DEFINE command processing. The message is constructed according to a code passed by the calling routine. IEEDFIN5 issues the requested message and passes control to IEEDFIN2 or the dispatcher.

- Entry: IEEDFIN5
- Exits: IEEDFIN1, IEEDFIN2 or return to calling program
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN5
- Page Reference: 52

IEEDFIN6: Master Scheduler -- Time-Slice Syntax Check Routine

This routine checks the TMSL subparameters for proper syntax.

- Entry: IEEDFIN6
- Exits: IEEDFIN2, IEEDFIN5, IEEDPART
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN6
- Page Reference: 52

IEEDFIN7: Master Scheduler -- Keyword Scan Routine

This routine checks keyword parameters for syntax errors. If a syntax error is discovered, the erroneous entry and all following entries are ignored, and an error message is generated. If the syntax is correct, the information is stored.

- Entry: IEEDFIN7
- Exits: IEEDFIN2, IEEDFIN3, IEEDFIN4, IEEDFIN5, IEEDPART
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN7
- Page Reference: 52

IEEDFIN8: Master Scheduler -- System Reinitialization Routine

This routine checks the redefinition information to assure that the request is valid. If no error is found, this routine constructs a list of ECBs (one for each of the affected partitions) to be posted when the jobs in the partitions have terminated. It enqueues on the partition boundary boxes, updates system control blocks and the affected boundary boxes, and then dequeues.

- Entry: IEEDFIN8
- Exits: IEEDFIN9, IEEREXIT
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN8
- Page Reference: 53

IEEDFIN9: Master Scheduler -- Command Final Processor Routine

This routine issues a message to all active consoles that processing is complete and updates the task control blocks affected by time-slicing if time-slicing is specified.

- Entry: IEEDFIN9
- Exits: IEEDFIN5
- Attributes: Read-only, reenterable
- Control Section: IEEDFIN9
- Page Reference: 53

IEELOG02: Master Scheduler -- Log Open Initialization Module

This routine opens the system log at IPL time.

- Entry: IEELOG02
- Exit: IEESD569
- Tables/Work Areas: CVT, UCB, UCM, TIOT, M/S resident data area, JFCB, IEELCA, DCB.
- Attributes: Refreshable
- Control Section: IEELOG02

IEELWAIT: Master Scheduler -- Log Wait and Writer Module

This module writes data from the log buffer to the system log.

- Entry: IEELWAIT

- Exit: To Dispatcher
- Tables/Work Areas: CVT, LCA, MRC
- Attributes: Resident
- Control Section: IEELWAIT

- Control Section: IEESD562
- Page Reference: 49

IEEPDISC: Display Consoles Get Region Routine

This routine obtains a region of main storage, and sets up an environment for the execution of the DISPLAY CONSOLES command, and then frees the region when control is returned.

- Entry: IEEPDISC, from IEEVATT1
- Exit: To IEEXEDNA, Return to Master Task (SVC 3)
- Attributes: Read-only, reentrant, resident
- Control Sections: IEEPDISC

IEESD563: Master Scheduler -- Queue Search Setup Routine

This routine determines which queue is to be searched, reads and scans the queue control record, establishes parameters for the search, and transfers control to the queue search module. IEESD563 will write out updated queue control records.

- Entry: IEESD563
- Exits: XCTL to IEESD564 to search queue; XCTL to IEESD565 at completion
- Tables/Work Areas: QCR, QMPA, CVT, CSCB
- Attributes: Read-only, reenterable
- Control Section: IEESD563
- Page Reference: 49

IEESD561: SVC 34 -- START and STOP INIT Routine

This routine processes the START and STOP INIT commands.

- Entry: IGC1903D
- Exit: Return to caller, or to IEE0503D for an error message
- Tables/Work Areas: CSCB, PIB, M/S resident data area, CVT, XSA
- Attributes: Reenterable, Transient read-only
- Control Section: IGC1903D
- Page Reference: 47

IEESD564: Master Scheduler -- Queue Search Module

This routine searches the work queues for the execution of the queue manipulation commands.

- Entry: IEESD564
- Exit: XCTL to IEESD563
- Tables/Work Areas: QCR, CSCB, CVT, QMPA, XSA
- Attributes: Read-only, reenterable
- Control Section: IEESD564
- Page Reference: 49

IEESD562: Master Scheduler -- Syntax Check Routine

This routine checks syntax of the command and sets internal codes for queue search, if required.

- Entry: IEESD562
- Exits: XCTL to IEESD563 for queue search, to IEESD566 for DISPLAY active, or to IEEXEDNA for DISPLAY CONSOLES
- Attributes: Read-only, reenterable
- External References: None

IEESD565: Master Scheduler -- Service Routine

This routine frees storage obtained by IEESD563, links to the queue manager to enqueue an entry or queue control record on SYS1.SYSJOBQE, or links to write a message.

- Entry: IEESD565
- Exit: Return to caller
- Tables/Work Areas: QMPA, CSCB, QCR, CVT
- Attributes: Read-only, reenterable

- External References: IEFQMNQ2, IEE0503D
- Control Section: IEESD565
- Page Reference: 49

IEESD566: Master Scheduler -- DISPLAY A Routine

This routine builds a table and constructs operator messages according to the processing required by a DISPLAY A command.

- Entry: IEESD566
- Exit: Return to caller (IEECIR50)
- Tables/Work Areas: QMPA, CSCB, XSA, QCR, CVT
- Attributes: Read-only, reenterable
- Control Section: IEESD566
- Page Reference: 50

IEESD568: Nucleus -- Master Scheduler Resident Data Area

This routine contains the master scheduler resident data area.

- Entry: IEEMSER
- Exit: None
- Attributes: Not reusable
- Control Section: IEEMSER
- Page Reference: 87

IEESD571: SVC 34 -- DEFINE, MOUNT, CANCEL Routine

This routine schedules the execution of the DEFINE, MOUNT, and CANCEL (for active jobs only) commands.

- Entry: IGC1803D
- Exits:
 - MOUNT - XCTL to IGC0103D
 - DEFINE - Return to caller
 - CANCEL - Active and cancelable -- Enter ABTERM to force cancel
 - Active and not cancelable -- POST and mark CSCB inactive; XCTL to IEE0803D
 - XCTL to IEE0503D and IEE2103D due to error.
- Tables/Work Areas: CSCB, PIB, M/S resident data area, CVT

- Attributes: Reenterable
- Control Section: IGC1803D
- Page Reference: 47

IEESD590: System Task Control -- Write TIOT on Disk

This routine writes the TIOT which is used by Job Selection (IEESD510) and checks for a small partition writer.

- Entry: IEESD590
- Exits: XCTL to IEFSD510 (small partition writer) or XCTL to IEFSD591
- Tables/Work Areas: TIOT, SPIL
- Attributes: Reenterable
- Control Section: IEESD590
- Page Reference: 75

IEESD591: System Task Control -- Linker Routine

This routine transfers control between system task control and an interpreter or system output writer.

- Entry: IEESD591, IEE591SD
- Exit: XCTL to IEEVTCTL
- Tables/Work Areas: CSCB, CVT, PIB, IWA, QMPA
- Attributes: Reenterable
- Control Section: IEESD591
- Page Reference: 75

IEESD592: System Task Control -- POST Routine

This routine checks for an error indication in the CSCB. It posts the error condition or a valid condition.

- Entry: IEESD592
- Exit: XCTL to IEFSD510
- Tables/Work Areas: None
- Attributes: Reenterable
- Control Section: IEESD592
- Page Reference: 75

IEEVACTL: System Task Control -- Allocation Interface Routine

This routine sets up the interface between system task control and the I/O device allocation routine.

- Entry: IEEVACTL
- Exits: To IEFSD21Q, IEEVMSG1, IEEVSMMSG, IEEVTCTL, or IEEVRWTC
- Attributes: Reentrantable
- Control Section: IEEVACTL
- Page Reference: 75

IEEVICLR: System Task Control -- Internal JCL Reader

This routine reads the internal job control language used in starting a reader or writer.

- Entry: IEEVICLR, IEFICR
- Exit: Return to caller
- Tables/Work Areas: DCBD
- Attributes: Read-only, reentrantable
- Control Section: IEEVICLR

IEEVJCL: System Task Control -- JCL Edit Routine

This routine constructs the internal job control language used in the START reader and START writer command execution routines.

- Entry: IEEVJCL, from IEEVSTRT
- Exit: XCTL to IEERCTL
- Tables/Work Areas: SDT, CSCD
- Attributes: Reentrantable
- Control Section: IEEVJCL

IEEVLDSP: Master Scheduler -- Log Dispatcher Routine

This routine puts the log data set on the system output queue.

- Entry: IEEVLDSP
- Exit: Master Scheduler
- Tables/Work Areas: IEEBASEA, CT, IEELCA, UCB, JFCB.

- Attributes: Reentrant
- Control Section: IEEVLDSP

IEEVLIN: Master Scheduler -- Log Initialization Routine

This routine initializes the system log.

- Entry: IEEVLIN
- Exit: IEESD569, IEEVLIN2
- Tables/Work Areas: UCM, CVT, UCB, TIOT, M/S resident data area, IEELCA.
- Attributes: Refreshable
- Control Section: IEEVLIN

IEEVLNKT: System Task Control -- Link-Table Module

This routine contains the table of routines that is scanned by IEEVACTL as a validity check for program linking.

- Entry: IEEVLNKT
- Attributes: Non-executable
- Control Section: IEEVLNKT

IEEVLOUT: Log Data Set Reinitialization Routine

This routine opens and closes the log data set to reinitialize the DS1LSTAR and DS1TRBAL fields of the DSCB associated with the log data set.

- Entry: IEEVLOUT, from IEFSD171
- Exit: IEFSD171
- Tables/Work Areas: CVT, DSCB, ICA, M/S Resident Data Area
- Attributes: Reentrantable
- Control Section: IEEVLOUT

IEEVMSG1: System Task Control -- Message Writer Routine

This routine writes messages to the operator as required by system task control.

- Entry: from IEEVRCTL, IEEVACTL, or IEEVTCTL
- Exit: Return to caller
- Control Section: IEEVMSG1

IEEVOMSG: System Task Control -- Message Writing Routine

This routine assembles and writes messages to the operator.

- Entry: IEEVOMSG
- Exit: Return to caller
- Control Section: IEEVOMSG

IEEVRCTL: System Task Control -- Interpreter Control Routine

This routine provides an interface between system task control and an interpreter.

- Entry: IEEVRCTL
- Exits: To IEFVH1 and IEEVACTL
- Tables/Work Areas: CVT, CSCB
- Control Section: IEEVRCTL
- Page Reference: 74

IEEVRFRX: Master Scheduler -- Table Lookup Routine

This routine can be used to obtain the following information; the CVT address, the contents of a CVT entry, or the contents at the CVT pointer address, a pointer to the TCB or the RB, the TIOT pointer, the TIOT entry, the TIOT TTR, or the TIOT UCB pointer. The routine can also be used to insert a TIOT pointer, a TIOT TTR, or a TIOT UCB pointer in the CVT.

- Entry: IEEVRFRX
- Exit: Return to calling program
- Tables/Work Areas: CVT, TCB, RB, TIOT, UCB
- Attributes: Reenterable
- Control Section: IEEVRFRX

IEEVSMBA: System Task Control -- QMPA Builder

This routine constructs a queue manager parameter area (QMPA) referring to the message class queue for the use of the I/O Device Allocation routine.

- Entry: IEEVSMBA
- Exit: To IEEVACTL

- Tables/Work Areas: QMPA, LCT, SMB, IOB

- Control Section: IEEVSMBA

IEEVSMMSG: System Task Control -- Message Writer Routine

This routine writes messages to the operator as required by the master scheduling task and system task control.

- Entry: IEEVSMMSG, from IEEVMSG1, IEFSD533, IEFUH1, or IEEVACTL
- Exit: Return to caller
- Control Section: IEEVSMMSG

IEEVSTAR: System Task Control -- Start Command Syntax Check Routine

This routine checks the syntax of a START command, and builds a start descriptor table (SDT) containing the parameters of the command.

- Entry: IEEVSTRT
- Exits: To IEEVJCL, IEFSD533, or IEE0503D
- Tables/Work Areas: SDT, M/S Resident Data Area, CVT, M/S TIOT, UCB XSA, and CSCB.
- Attributes: Reenterable
- Control Section: IEEVSTRT
- Page Reference: 74

IEEVTCTL: System Task Control -- Termination Interface Routine

This routine initializes the necessary tables for terminating a task that was established via a START or MCUNT command.

- Entry: IEEVTCTL, from IEEVWILK, IEEVACTL or IEFW31SD
- Exit: To IEFW42SD or IEEVOMSG, , then return to supervisor
- Tables/Work Areas: TCB, JCT, SCT, LCT, and CSCB
- Attributes: Reenterable. Character Dependence Type C
- Control Section: IEEVTCTL

IEEVWTOR: Communications Task --
Write-to-Operator With Reply Routine

This routine processes all WTOR macro instructions.

- Entry: IGC0103E
- Exit: Return to calling program
- Tables/Work Areas: WQE, RQE, UCM, CVT
- Attributes: Reenterable
- Control Section: IGC0103E
- Page Reference: 45

IEEXEDNA: DISPLAY CONSOLES Processor

This routine processes the DISPLAY command with the CONSCLES operand and displays the system console configuration on the requesting console.

- Entry: IEEXEDNA to IEESD562
- Exit: To IEECIR50
- Attributes: Reentrant
- Control Sections: IEEXEDNA

IEE0303D: SVC 34 -- Translator/Chain Manipulator

This routine translates lowercase letters into uppercase, and manipulates the CSCB chain as requested by the caller of SVC 34.

- Entry: IEE0303D
- Exit: To IEE0403D, or return to caller
- Tables/Work Areas: CVT, M/S resident data area, CSCB, XSA
- Control Section: IEE0303D

IEE0303F: SVC 36 -- WRITE-TO-LOG

This module copies text records from an input area to the log buffer and posts the log ECB when the buffer is full.

- Entry: IEE0303F
- Exit: Returns to Master Scheduler, IEE0403F.
- Tables/Work Areas: IEEBASEA, IEELCA, CVT
- Attributes:

- Control Section: IEE0303F

IEE0403D: SVC 34 -- Router Routines

This routine identifies the command verb, ensures that the console has authority to enter the command, and passes control to the appropriate routine.

- Entry: IEE0403D
- Exit: Depending on command verb, via XCTL to another SVC 34 module
- Tables/Work Areas: M/S resident data area, XSA, CSCB
- Control Section: IEE0403D
- Page Reference: 47

IEE0403F: SVC 36 (Load 2) -- Log Buffer Management Module

This module opens, closes, and switches system log buffers.

- Entry: IEE0403F
- Exit: IEE0303F
- Tables/Work Areas: IEEBASEA, IEELCA, UCB, JFCB, DCB, CVT, TICT.
- Attributes: Reentrant
- Control Section: IEE0403F

IEE0503D: SVC34 -- Message Assembly Routine

This routine assembles and edits messages for the command scheduling routine, and writes the messages to the operator.

- Entry: IEE0503D
- Exit: Branch on register 14
- Attributes: Reenterable, read-only
- Control Section: IEE0503D

IEE0603D: SVC 34 -- SET Command Routine

This routine processes the SET command.

- Entry: IEE0603D
- Exits: To IEE0903D, IEE0503D, or return to caller

- Tables/Work Areas: XSA, CVT, M/S resident data area
- Attributes: Reenterable, self-relocating, read only transient
- Control Section: IEE0603D

IEE0703D: SVC 34 -- CSCB Marking Routine

This routine schedules the execution of the STOP and MCDIFY commands by finding and updating the appropriate CSCB and by issuing a POST macro instruction to the master scheduling task.

- Entry: IEE0703D
- Exits: Branch on register 14, or XCTL to IEE0803D, IEE0503D or IEE2103D.
- Tables/Work Areas: M/S Resident Data Area, XSA, CVT, CSCB, SPL
- Attributes: Reenterable, self-relocating, read-only, transient
- Control Section: IEE0703D

IEE0803D: SVC34 -- CSCB Creation Routine

This routine schedules the execution of commands that cannot be completely processed by the command scheduling routines. It performs this function by adding a CSCB to the CSCB chain and issuing a POST macro instruction to the master scheduling task. It also processes the DISPLAY T command.

- Entry: IEE0803D
- Exit: IEE0503D, IEE2103D, IEE2903D, or return to caller
- Tables/Work Areas: XSA, M/S resident data area, CVT, CSCB, and UCM
- Attributes: Reenterable, transient, partially disabled.
- Control Section: IEE0803D

IEE0903D: SET Command Handler

This routine processes the date and time operands of the SET command.

- Entry: IEAQOT00
- Exit: SVC 3
- Tables/Work Areas: CVT

- Attributes: Reenterable, supervisor state, disabled for system interrupts, transient
- Control Section: IEAQOT00

IEE1103D: SVC 34 -- VARY and UNLOAD Scan and Router Routine

This routine examines the command and its operand and routes the command to the appropriate processing module.

- Entry: IEE1103D
- Exit: IEE2303D for VARY ONLINE, CNGFX, and CONSOLES operands when SMF is present, to IEE3103D for all other VARY operands and UNLOAD, and to IEE0503D for errors.
- Tables/Work Areas: XSA, CVT, M/S resident data area, and UCM.
- Attributes: Reenterable, self-relocating, read-only, and transient.
- Control Sections: IEE1103D

IEE1203D: SVC 34 -- Reply Processor

This routine checks the validity of the operator's reply command, and moves the operator's reply (if valid) to the buffer of the user that issued the respective WIOR.

- Entry: IEE1203D
- Exit: Return to caller
- Tables/Work Areas: CVT, UCM, WQE, RQE, CXSA
- Attributes: Reenterable
- Control Section: IEE1203D

IEE1403D: SVC 34 -- HALT Routine

This routine schedules the execution of the HALT command by adding a CSCB to the CSCB chain and by issuing a POST macro instruction to the master scheduling task.

- Entry: IEE1403D
- Exit: IFBSTAT
- Tables/Work Areas: XSA, M/S resident data area, CVT, and CSCB
- Attributes: Reenterable
- Control Section: IEE1403D

IEE1603D: SVC 34 -- Log and Writelog Processor Routine

This routine issues a WTL macro instruction when a LOG command is issued, and stores the WRITELCG command and posts the Log ECB, for WRITELCG processing.

- Entry: IEE1603D, from IEE0403D
- Exit: IEE0503D for errors, and return to caller of SVC 34.
- Tables/Work Areas: XSA, CVT, LCA, and M/S resident data area.
- Attributes: Reentrant, self-relocating, read-only, and transient.
- Control Sections: IEE1603D

IEE1703D: SVC 34 -- VARY ONGFX/OFFGFX

This routine processes the GVARY command. It checks the parameters for validity and if an error is found, it passes control to IEE0503D via an XCTL macro instruction. If the parameters are valid, the routine sets appropriate bits in the Overall Control Table (OCT) of the GFX reader task. It then issues a POST macro instruction on the ECB in the OCT for each graphics device (2250) placed in the online status.

- Entry: IEE1703D
- Exit: IEE0503D, return to issuer of SVC 34
- Tables/Work Areas: CVT, OCT, XSA
- Attributes: Reenterable, read-only, self-relocating
- Control Section: IEE1703D

IEE1A03D: SVC 34 -- MCS Reply Processor Routine

The purpose of this routine is to process valid operator replies to WTOR macro instructions.

- Entry: IEE1A03D
- Exit: To IEE1B03D to issue error messages or return to the caller of SVC 34.
- Control Sections: IEE1A03D

IEE1B03D: SVC 34 -- MCS Reply Message Routine

This routine assembles, edits, and broadcasts the accepted reply to a WTOR macro instruction for the MCS Reply Processor

routine (module IEE1A03D) of the Command Scheduling routine, and to write error messages to the operator whose command is in error.

- Entry: IEE1B03D, from IEE1A03D
- Exit: Return to the caller of SVC 34
- Control Sections: IEE1B03D

IEE2103D: SVC 34 -- Message Assembly Routine

This routine assembles and edits messages for the command scheduling routine, and writes the messages to the operator.

- Entry: IEE2103D
- Exit: Branch on register 14
- Attributes: Reenterable, self-relocatory, read-only, transient
- Control Section: IEE2103D

IEE2303D: SVC 34 -- SMF Processor

This routine initially processes the ONLINE, ONGFX and CONSCLES operand of the VARY command when the system has the SMF option. It builds and issues an SMF record for each device placed in online status.

- Entry: IEE2303D
- Exit: IEE3103D
- Tables/Work Areas: CVT, SMCA, XSA
- Attributes: Reentrant, read-only, self-relocating
- Control Sections: IEE2303D

IEE2903D: SVC 34 -- Display Requests Routine

This routine displays to the requesting operator the ID of all outstanding WTCRs, the unit name of each device for outstanding MOUNT messages, and an indication as to whether any AVR mount messages are pending.

- Entry: IEE2903D, from IEE0803D
- Exit: Return to caller of SVC 34
- Tables/Work Areas: Message work area
- Attributes: Reentrant, Refreshable, transient
- Control Sections: IEE2903D

IEE3103D: SVC 34 -- Vary and Unload Processor Routine

This routine processes the UNLOAD command, all VARY command operands in a system without the MCS option, and VARY ONLINE and OFFLINE operands for non-console devices in a system with the MCS option. In addition, it passes control to the appropriate MCS processors for processing of console devices.

- Entry: IEE3103D, initially from IEE1103D or IEE2303D and returns from IEE4203D and IEE4603D.
- Exit: IEE1703D for VARY ONGFX/OFFGFX, IEE4203D for VARY ONLINE/OFFLINE/CONSOLES with no keywords after CONSOLES, IEE4303D for VARY MSTCONS, IEE4403D for VARY HARDCPY/CONSOLES with keywords, IEE4703D for VARY HARDCPY without keywords, IEE0503D for errors, Return to Caller of SVC 34 for UNLOAD.
- Tables/Work Areas: XSA, UCM, CVT, M/S resident data area, and UCB.
- Attributes: Reentrant, self-relocating, read-only, transient.
- Control Sections: IEE3103D

IEE4103D: SVC 34 -- Hardcopy Message Issuing Routine

This routine issues messages concerning the status of the hard copy log and frees storage obtained for those messages.

- Entry: IEE4103D, from IEE4703D and IEE4803D
- Exit: IEE4203D if multiple units specified in VARY command and units remain to be processed, IEE0503D for errors, or return to caller of SVC 34 if all units have been processed
- Tables/Work Areas: XSA, message area, UCB, CVT, XSA, and UCM
- Attributes: Reentrant, transient
- Control Sections: IEE4103D

IEE4203D: SVC 34 -- Vary Unit Field Scan and Router Routine

This module performs authority and operand validity checking, and passes control to the routine that will process the command.

- Entry: IEE4203D, from IEE3103D, IEE4103D, IEE4403D, and IEE4603D
- Exit: To IEE3103D for processing VARY ONLINE/OFFLINE of non-console units, to IEE4603D for processing of VARY ONLINE/OFFLINE CONSOLES of console units, and for processing when errors in syntax are found or when multiple units were specified and units remain to be processed, IEE4903D for processing of VARY CONSOLES, and to IEE0503D and IEE2103D when other errors occur.
- Tables/Work Areas: XSA, CVT, UCM, and UCB.
- Attributes: Reentrant, self-relocating, read-only, transient
- Control Sections: IEE4203D

IEE4303D: SVC 34 -- VARY MSTCCNS Routine

This routine processes the VARY MSTCONS command.

- Entry: IEE4303D, from IEE3103D
- Exit: To IEE0503D or IEE2103D on errors, SVC 72 to Console Switch Routine (module IEECMCSW) and upon return to caller of SVC 34
- Tables/Work Areas: UCB, CVT, XSA, and UCM
- Attributes: Reentrant, self-relocating, read-only, transient
- Control Sections: IEE4303D

IEE4403D: SVC 34 -- Vary Keyword Scan Routine

This routine determines the validity of VARY CONSOLE-HARDCPY keywords, and to set appropriate bits in the XSA.

- Entry: IEE4403D, from IEE3103D
- Exit: To IEE4203D if VARY CONSOLES, to IEE4703D if VARY HARDCPY, to IEE0503D if errors.
- Tables/Work Areas: XSA, UCM, CVT, and UCB
- Attributes: Reentrant, transient
- Control Sections: IEE4403D

IEE4503D: SVC 34 -- Periodic STOP Command Handler Routine

This routine processes the commands STOP JOBNAME/STATUS/SPACE/DSNAME.

- Entry: IEE4503D, from IEE0403D
- Exit: IEE0503D for errors, and return to caller of SVC 34
- Tables/Work Areas: XSA, M/S resident data area, CVT, and UCM
- Attributes: Reentrant, self-relocating, read-only, transient
- Control Section: IEE4503D

IEE4603D: SVC 34 -- VARY ONLINE/OFFLINE of Consoles and Message Routing Routine

This routine processes VARY ONLINE/OFFLINE for all MCS consoles and dispatches error messages for syntax errors.

- Entry: IEE4603D, from IEE4203D to process VARY ONLINE/OFFLINE or to dispatch WTOs, IEE4903D to dispatch WTOs.
- Exit: IEE3103D when multiple units specified and only non-console units remain to be processed, IEE4203D if unit field indicates more console (and possible non-console) units remain to be processed, IEE0503D and IEE2103D for errors, return to caller of SVC 34 if all units have been processed.
- Tables/Work Areas: XSA, CVT, UCB, UCM, and M/S resident data area.
- Attributes: Reentrant, self-relocating, read-only, transient
- Control Sections: IEE4603D

IEE4703D: SVC 34 -- VARY HARDCPY Processor Routine

This routine processes VARY HARDCPY commands.

- Entry: IEE4703D, from IEE3103D if command has no keyword operands, IEE4203D if command has keyword operands.
- Exit: To IEE4103D to issue a hardcopy message, IEE0503D or IEE2103D on errors.
- Tables/Work Areas: XSA, UCM, M/S resident data area, CVT and UCB.
- Attributes: Reentrant, transient

- Control Sections: IEE4703D

IEE4803D: SVC 34 -- VARY CONSOLE Information Message Routine

This routine constructs a message which shows the current status of the varied console.

- Entry: IEE4803D, from IEE4903D
- Exit: To IEE4103D to issue the message, IEE0503D for errors
- Tables/Work Areas: XSA, message area, UCB, CVT, and UCM
- Attributes: Reentrant, transient
- Control Sections: IEE4803D

IEE4903D: SVC 34 -- VARY CONSOLE Processor Routine

This module processes the VARY CONSOLE command.

- Entry: IEE4903D, from IEE4203D
- Exit: To IEE4803D to construct console message, IEE4603D to dispatch error messages
- Table/Work Areas: XSA, CVT, UCB, and UCM
- Attributes: Reentrant, self-relocating, read-only, transient
- Control Sections: IEE4903D

IEFACTFK: Termination -- User Dummy Accounting Routine

This routine takes the place of the user's accounting routine when a user accounting routine was specified at system generation, but none was supplied.

- Entry: IEFACTLK
- Exit: Return to caller
- Control Section: IEFACTLK

IEFACTLK: Termination -- User Accounting Routine Linkage Routine

This routine provides linkage between the termination routine and the user's accounting routine. It also sets up the required parameter list -- including the execution time of the job step -- and reads the first record of the account control table.

- Entry: IEFACTLK

- Exits: To user's accounting routine, return to caller.
- Tables/Work Areas: LCT, JCT, SCT, JACT, SACT, QMPA
- Control Section: IEFACTLK

IEFACTRT: Termination -- Dummy Accounting Routine

This routine takes the place of the user-supplied accounting routine.

- Entry: IEFACTRT
- Exit: Return to caller
- Control Section: IEFACTRT

IEFAVFAK: I/C Device Allocation -- Linkage to IEFXV001

This routine passes control to the AVR routine (IEFXV001) via and XCTL macro instruction.

- Entry: IEFXV001
- Exit: XCTL to IEFXV001
- Control Section: IEFXV001

IEFCVFAK: I/C Device Allocation -- Linkage to IEFMCVOL

This routine passes control to Mount Control-Volume Routine IEFMCVOL via an XCTL macro instruction to one of three entry points, IEFCVOL1, IEFCVOL2, or IEFCVOL3.

- Entries: IEFCVOL1, IEFCVOL2, IEFCVOL3
- Exits: XCTL to IEFCVOL1, IEFCVOL2, IEFCVOL3
- Control Section: IEFCVOL1

IEFSDSRP: Data Set Descriptor Record Processing Routine

This routine processes the job queue information in the DSDR record to make a restarting job's queue entry reflect the environment when the checkpoint was taken.

- Entry Point: IEFSDSRP
- Exit: Return to caller
- Table/Work Areas: JCT, SCT, SIOT, JFCB, TIOT, UCB, CVT, VOLT, TCB, QMPA, CSCB, DCBD, DCB, JFCBX, SCTX, LCT
- Attributes: Reenterable
- Control Section: IEFSDSRP

IEFIDMPM: Termination -- Message Module

This routine contains the messages used by the Indicative Dump routine.

- Entry: IEFIDMPM
- Attributes: Non-executable
- Control Section: IEFIDMPM

IEFLOCDQ: Queue Management -- Dequeue by Jobname Routine

This routine searches a queue for a named job or list of named jobs, and can return information, or dequeue or cancel the job.

- Entry: LOCDQ, LOCCAN, LCC
- Exit: Return to caller
- Tables/Work Areas: QCR, LTH
- Attributes: Reenterable
- External References: IEFCNVRT, IEFRDWRT
- Page Reference: 71

IEFMCVOL: I/O Device Allocation -- Mount Control-Volume Routine

This routine will have a control volume mounted when a data set called for in a job step cannot be located on any currently mounted control volume.

- Entries: IEFCVOL1, IEFCVOL2, IEFCVOL3
- Exits: IEFVM1, IEFVMCVL, IEFVM6, IEFYN (IEFW41SD)
- Tables/Work Areas: LCT, JCT, SCT, SIOT, JFCB, VOLT, QMPA, UCB
- Attributes: Reusable
- Control Sections: IEFCVOL1, IEFCVOL2, IEFCVOL3

IEFORMAT: Queue Management -- Queue Formatting Routine

This routine places the work queue data set in the format required by the MFT queue management routines.

- Entry: IEFORMAT, from IEFSD055
- Exit: Return to IEFSD055
- Tables/Work Areas: DCB, DEB
- Attributes: Reusable

- Control Section: IEFORMAT
- Page Reference: 49

IEFQAGST: Queue Management -- Assign/Start Routine

This routine sets up an ECB/IOB and prepares the queue manager parameter area for the assign routine.

- Entry: IEFQAGST
- Exit: Return to caller
- Tables/Work Areas: Q/M resident data area, QMPA, CVT
- Attributes: Reenterable
- Control Section: IEFQAGST
- Page Reference: 55

IEFQASGQ: Queue Management -- Assign Routine

This routine assigns records to a queue entry and assigns logical tracks as required.

- Entry: IEFQASGN
- Exit: Return to caller
- Tables/Work Areas: Q/M resident data area, QMPA, CVT
- Attributes: Reenterable
- Control Sections: IEFQASGN, IEFQASNM
- Page Reference: 55

IEFQBVMS: Queue Management -- Control Routine

This routine inspects the function code in the queue manager parameter area and, on the basis of this code, branches to the appropriate queue management routines.

- Entry: IEFQMSSS
- Exits: To IEFQAGST, IEFQMRAW, IEFQMNQO, or IEFQASGQ, return to caller
- Tables/Work Areas: QMPA
- Attributes: Reenterable
- Control Section: IEFQMSSS

IEFQDELE: Queue Management -- Delete Routine

This routine makes those logical tracks assigned to a queue entry available for assignment to other queue entries.

- Entry: IEFQDELE
- Exit: Return to caller
- Tables/Work Areas: LTH, QMPA, QCR, Q/M resident data area, CVT
- Attributes: Reenterable
- Control Section: IEFQDELE
- Page Reference: 59

IEFQMDQO: Queue Management -- Dequeue Routine

This routine removes the highest priority entry from an input queue or a system output queue.

- Entry: IEFQMDQ2
- Exit: Return to caller
- Tables/Work Areas: CVT, Q/M resident data area, QCR, LTH
- Attributes: Reenterable
- Control Section: IEFQMDQ2
- Page Reference: 59

IEFQMDUM: Queue Management -- Dummy Module

This routine prevents the occurrence of an unresolved external reference to module IEFQMSSS during system generation.

- Entry: IEFQMDUM
- Attributes: Non-Executable
- Control Section: IEFQMSSS

IEFQMLK1: Queue Management -- Branch Routine

This routine branches to the appropriate queue management routine on the basis of an assign or read/write function code issued by an initiator.

- Entry: IEFQMSSS
- Exits: To IEFQASGQ or IEFQMRAW

- Tables/Work Areas: QMPA
- Attributes: Reenterable
- Control Section: IEFQMSSS

IEFQMNOQ: Queue Management -- Enqueue Routine

This routine places an entry in an input queue or an output queue at the requested priority.

- Entry: IEFQMNO2
- Exit: Return to caller
- Tables/Work Areas: CVT, Q/M resident data area, QMPA, QCR, LTH
- Attributes: Reenterable
- Control Section: IEFQMNO2

IEFQMRAW: Queue Management -- Read/Write Routine

This routine performs the conversion of a TTR into a MBBCCHHR and reads or writes up to 15 records of the work queue data set.

- Entry: IEFQMRAW
- Exit: Return to caller
- Tables/Work Areas: Q/M resident data area, QMPA, CVT, IOB/ECB
- Attributes: Reenterable
- Control Section: IEFQMRAW

IEFQMUNC: Queue Management -- Unchain Routine

This routine removes a task from the queue management no-work chain.

- Entry: IEFQMUNC
- Exit: Return to caller
- Tables/Work Areas: CVT, Q/M resident data area, QCR
- Attributes: Reenterable
- Control Section: IEFQMUNC

IEFQRESM: Queue Management -- Resident Main Storage Reservation Module

This routine reserves 140 bytes of resident main storage for the queue-management-

opened DCB/DEB and the master queue control record at nucleus initialization time.

- Attributes: Non-executable
- Control Section: IEFJOB

IEFRCLN1: Restart Reader Linkage

This routine receives control from IEFVRRRC and LINKS to interpreter initialization routine IEFVH1.

- Entry: IEFRCLN1
- Exit: XCTL to IEFVRRRC at entry IEFVRRCA
- Attributes: Reenterable
- Control Section: IEFRCLN1

IEFRCLN2: Restart Reader Linkage

This routine receives control from IEFVRRRC and LINKS to interpreter initialization routine IEFVH1.

- Entry: IEFRCLN2
- Exit: XCTL to IEFVRRRC at entry IEFVRRCB
- Attributes: Reenterable
- Control Section: IEFRCLN2

IEFRPREP: Termination -- Restart Preparation Routing

This routine determines whether a job step that has been abnormally terminated can be restarted.

- Entry: IEFRPREP from IEFYNIMP
- Exit: Return to caller
- Attributes: Reenterable
- Tables/Work Areas: LCT, JCT, SCT, PDQ, QMPA
- Control Section: IEFRPREP

IEFRSTRT: Restart SVC Issuing Routine

This routine issues the Restart SVC. When called by its alias, IEFSMR, it issues the Restart SVC and then returns to the caller.

- Entry: IEFRSTRT, IEFSMR

- Exit: SVC 52 (RESTART), return to caller
- Attributes: Reenterable
- Control Sections: IEFRSTRT

IEFSD017: Termination -- System Output Interface Routine

This routine provides an interface between the termination entry routine and system output processing.

- Entry: IEFSD017
- Exit: To IEFSD42Q
- Control Section: IEFSD017

IEFSD055: Queue Management -- Queue Initialization Routine

This routine constructs a resident DEB/DCB, passes control to the queue formatting routine or the first phase of system restart, initializes the queue manager resident data area, and (if required) passes control to the second phase of the system restart routine.

- Entry: IEFSD055, from IEFQINTZ
- Exits: To IEFORMAT, IEF300SD, or IEF304SD
- Attributes: Reusable
- Control Section: IEFSD055
- Page Reference: 54

IEFSD070: System Output Writer -- Data Set Writer Interface Routine

This routine passes control to the standard data set writer or to the user-supplied data set writer routine.

- Entry: IEFSD070
- Exits: To IEFSD087 or user-supplied routine via LINK, or to IEFSD171 via XCTL
- Attributes: Reenterable
- Control Section: IEFSD070
- Page Reference: 72

IEFSD078: System Output Writer -- Linker Routine

This routine determines whether the record obtained from the output queue entry is a DSB or SMB, and passes control, accordingly, to the DSB or SMB processor.

- Entry: IEFSD078
- Exits: To IEFSD085, IEFSD086, or IEFSD079
- Attributes: Reenterable
- Control Section: IEFSD078

IEFSD079: System Output Writer -- Link to Queue Manager Delete Routine

This routine passes control to the delete routine to delete the current output queue entry.

- Entry: IEFSD079
- Exits: To IEFQDELQ and IEFSD082
- Tables/Work Areas: QMPA
- Attributes: Reenterable
- Control Section: IEFSD079
- Page Reference: 79

IEFSD080: System Output Writer -- Initialization Routine

This routine initializes the system output writer by obtaining main storage for a parameter list and the output DCB, and opening the output DCB.

- Entry: IEFSD080
- Exit: To IEFSD081
- Tables/Work Areas: DCB, CSCB, TIOT, JFCB
- Attributes: Reenterable
- Control Section: IEFSD080

IEFSD081: System Output Writer -- Class Name Setup Routine

This routine obtains main storage for, and initializes, a list of ECB pointers, ECBs, and queue management communication elements, depending on the system output classes specified for the writer.

- Entry: IEFSD081
- Exit: To IEFSD082
- Tables/Work Areas: CSCB, ECB
- Attributes: Reenterable
- Control Section: IEFSD081

IEFSD082: System Output Writer -- Main Logic Routine

This routine obtains main storage for QMPAs and internal work areas, dequeues output queue entries, checks for operator commands, and passes control to the appropriate routine.

- Entry: IEFSD082
- Exits: IEFSD083, IEFSD084, IEFSD078
- Tables/Work Areas: CSCB, ECB
- Attributes: Reenterable
- Control Section: IEFSD082

IEFSD083: System Output Writer -- Command Processing Routine

This routine processes MODIFY and STOP commands that apply to the writer.

- Entry: IEFSD083
- Exits: To IEFSD081 or IEEVTCTL
- Tables/Work Areas: CSCB, DCB, QMPA, ECB
- Attributes: Reenterable
- Control Sections: IEFSD083, IEFSD83M

IEFSD084: System Output Writer -- Wait Routine

This routine waits for an entry to be enqueued in an output queue corresponding to a class available to the writer.

- Entry: IEFSD084
- Exit: To IEFSD082
- Attributes: Reenterable
- Control Section: IEFSD084
- Page Reference: 72

IEFSD085: System Output Writer -- DSB Handler Routine

This routine initializes for data set processing, and informs the operator of the pause option in effect.

- Entry: IEFSD085, IEF085SD, or IEF850SD
- Exit: To IEFSD171
- Attributes: Reenterable

- Control Sections: IEFSD085, IEFSD85M
- Page Reference: 73

IEFSD086: System Output Writer -- SMB Handler

This routine initializes for message processing, and extracts each message from the current SMB.

- Entry: IEFSD086, IEF086SD
- Exits: To IEFSD088, IEFSD089, IEFQMNQQ, IEFQMRW, IEFSD085, IEFSD078
- Tables/Work Areas: SMB, UCB, QMPA, TIOT, CSCB, TCB
- Attributes: Reenterable
- Control Sections: IEFSD086, IEFSD86M

IEFSD087: System Output Writer -- Standard Writer Routine

This routine gets records from a data set.

- Entry: IEFSD087
- Exits: To IEFSD088, IEFSD089, IEFSD078
- Tables/Work Areas: DCB
- Attributes: Reenterable
- Control Sections: IEFSD087, IEFSD87M
- Page Reference: 73

IEFSD088: System Output Writer -- Transition Routine

This routine handles the transition between messages and data sets, and between data sets.

- Entry: IEFSD088
- Exit: To IEFSD089
- Tables/Work Areas: DCB
- Attributes: Reenterable
- Control Section: IEFSD088

IEFSD089: System Output Writer -- Put Routine

This routine formats records as required and issues PUT macro instructions to write them on the output unit.

- Entry: IEFSD089
- Exit: To IEFSD088

- Tables/Work Areas: DCB
- Attributes: Reenterable
- Control Sections: IEFSD089, IEFSD89M

IEFSD094: System Output Writer -- Job Separator Routine

This routine prints or punches a job name and system output class designation on the writer's output device.

- Entry: IEFSD094
- Exits: To IEFSD088, IEFSD089, IEFSD095, IEFSD078
- Attributes: Reenterable
- Control Section: IEFSD094

IEFSD095: System Output Writer -- Print Line Routine

This routine constructs the block letters used to separate jobs processed by a system output writer when the output data set is to be printed.

- Entry: IEFSD095
- Exit: Return to caller
- Attributes: Reenterable
- Control Section: IEFSD095

IEFSD096: System Output Writer -- Message Module

This routine contains message headers and texts for messages to the operator.

- Entry: IEFSD096
- Attributes: Non-executable
- Control Section: IEFSD096

IEFSD097: I/O Device Allocation -- Wait for Space Decision Routine

This routine makes the decision whether to wait for direct access space, and provides an interface with the I/O device allocation space request routine so that retry and additional recovery passes may be made.

- Entry: IEFSD097
- Exit: Branch on register 14
- Tables/Work Areas: LCT, TIOT, UCB

- Attributes: Read-only, reenterable
- Control Section: IEFSD097

IEFSD168: Initiator -- Job Suspension

This routine causes a terminated job to be reenqueued so that the job can be reactivated.

- Entry: IEFSD068
- Exit: Branch to IEFSD598 to purge resources, branch to IEFSD510 to reinitiate job
- Tables/Work Areas: QMPA, LCT, JCT, SCD, SCT
- Attributes: Reenterable
- Control Section: IEFSD068
- External Reference: IEFQMRAW, IEFQMNQ2, IEFVSDRA

IEFSD171: System Output Writer -- Data Set Delete Routine

This routine obtains records from an output queue entry, and deletes system output data sets.

- Entry: IEFSD071
- Exits: To IEEVLOUT, IEFQMNQ2, IEF850SD, IEF086SD, IEFSD078, or IEFQMRAW
- Tables/Work Areas: DCB, SMB, UCB, CVT, QMPA, TIOT, CSCB, TCB
- Attributes: Reenterable
- Control Sections: IEFSD071, IEFSD71M

IEFSD195: I/O Device Allocation -- Wait for Deallocation Routine

This routine provides the I/O device allocation routine with the ability to wait for deallocation to occur during the execution of another task, when allocation cannot be completed because of current allocations.

- Entry: IEFVAWAT
- Exit: Return to caller
- Tables/Work Areas: JCT, SCT, SICT, LCT, ECG, CSCB
- Attributes: Read-only, reenterable
- Control Section: IEFSD095

IEFSD21Q: I/O Device Allocation -- Allocation Entry Routine

This routine provides an interface for entry to the I/O device allocation routine operating in a multiprogramming environment.

- Entry: IEFW21SD
- Exits: To IEFVK, IEFVM or IEFWD000
- Tables/Work Areas: JCT, LCT, SCT, SMB, QMPA, CVT
- Attributes: Read-only, reenterable
- Control Section: IEFWLISD

IEFSD22Q: Termination Routine -- Step Terminate Exit Routine

This routine provides an interface between the termination routine and the step delete or alternate step delete routine when a step has been terminated.

- Entry: IEFW22SD
- Exit: Return to caller of termination routine
- Tables/Work Areas: JCT, SCT, SMB, LCT, QMPA, ECB
- Attributes: Read-only, reenterable
- Control Section: IEFW22SD
- Page Reference: 71

IEFSD300: System Restart -- Initialization Routine

This routine reads all QCRs and logical track header records into main storage, builds tables A, B, and C, and removes from Table A all the LTH entries corresponding to logical tracks in the free-track queue or in one of the other queues.

- Entry: IEFSD300
- Exit: To IEFSD301
- Tables/Work Areas: System restart work area, Table A, Table B, Table C
- Attributes: Reenterable
- Control Section: IEFSD300

IEFSD301: System Restart -- Purge Queue Construction Routine

This routine searches Table A for the last LTH corresponding to each queue entry,

determines the type of entry, and constructs the purge queue.

- Entry: IEFSD301
- Exit: To IEFSD302
- Tables/Work Areas: System restart work area, Table A, Table C purge queue, interpreter jobnames table
- Attributes: Reenterable
- Control Section: IEFSD301

IEFSD302: System Restart -- Jobnames Table Routine

This routine removes from Table A all logical tracks assigned to dequeued input or RJE queue entries, and builds a table of jobnames for incomplete input and RJE queue entries and dequeued input queue entries.

- Entry: IEFSD302
- Exit: To IEFSD303
- Tables/Work Areas: System restart work area, Table A, Table C, and the interpreter/initiator jobnames table
- Attributes: Reenterable
- Control Section: IEFSD302

IEFSD303: System Restart -- Delete Routine

This routine creates a queue entry of the remaining logical tracks and deletes that entry, thus assigning those tracks to the free-track queue.

- Entry: IEFSD303
- Exit: Return to caller
- Tables/Work Areas: System restart work area, QMPA, Table A
- Attributes: Reenterable
- Control Section: IEFSD303

IEFSD304: System Restart -- Scratch Data Sets Routine

This routine informs the operator of the names of jobs being processed by an interpreter, and scratches temporary data sets generated for incomplete input queue entries.

- Entry: IEFSD304
- Exits: To IEFSD055, IEFSD308

- Tables/Work Areas: CVT, UCB address look-up table
- Attributes: Reenterable
- Control Section: IEFSD304

IEFSD305: System Restart -- Reenqueue Routine

This routine dequeues the entries in the purge queue and reenqueuees them in the appropriate input or output queue and informs the operator of the names of jobs in the process of initiation.

- Entry: IEFSD305
- Exit: IEFSD304
- Tables/Work Areas: System restart work area, purge queue, JCT, SCT, JFCB, DSB, SCD, SIOT.
- Attributes: Reenterable
- Control Section: IEFSD305

IEFSD308: System Restart -- Scratch Data Sets Routine

This routine scratches the temporary data sets generated for incomplete input queue entries.

- Entry: IEFSD308
- Exit: Return to caller
- Tables/Work Areas: DSCB, DCB, UCB, CVT, VTOC, DEB
- Attributes: Reenterable
- Control Section: IEFSD308

IEFSD31Q: Termination Routine -- Job Termination Exit Routine

This routine provides an interface between the termination routine and the step delete or alternate step delete routine when the last step of a job has been terminated.

- Entry: IEFW31SD
- Exit: Return to caller of termination routine
- Tables/Work Areas: JCT, SCT, SMB, QMPA, ECB, CVT, M/S resident data area
- Attributes: Read-only, reenterable
- Control Section: IEFW31SD

IEFSD310: System Restart -- TTR and NN to MBBCCHHR Conversion Routine

This routine converts a relative record address (NN) or a relative track and record address (TTR) to an actual disk address (MBBCCHHR).

- Entry: IEFSD310
- Exit: Return to caller
- Tables/Work Areas: CVT
- Attributes: Reenterable
- Control Section: IEFSD310

IEFSD311: Queue Management -- Message Module

This routine contains the messages required by the queue initialization routine (IEFSD055).

- Entry: IEFSD311, SD55MSG1, SD55MSG2, SD55MSG3
- Attributes: Non-executable
- Control Section: IEFSD311

IEFSD312: System Restart -- Message Module

This routine contains the messages required by the system restart routines.

- Entry: IEFSD312, SD304MG1, SD304MG2, SD305MG1
- Attributes: Non-executable
- Control Section: IEFSD312

IEFSD41Q: I/O Device Allocation -- Allocation Exit Routine

This routine provides an interface for exit from the I/O device allocation routine operating in a multiprogramming environment.

- Entry: IEFW41SD, IEFW1FAK, IEFW2FAK
- Exits: To IEFVM, or return to caller
- Tables/Work Areas: JCT, ICT, SCT, SMB, QMPA
- Attributes: Read-only, reenterable
- Control Section: IEFW41SD

IEFSD42Q: Termination Routine -- Termination Entry Routine

This routine provides an interface for entry to the termination routine operating

in a multiprogramming environment. it also provides an interface for entry to the LOG function if a LOG data set is scheduled to be added to the SYSOUT queue.

- Entry: IEFW42SD
- Exit: To IEFYN
- Tables/Work Areas: JCT, SCT, SMB, LCT, TIOT
- Attributes: Read-only, reenterable
- Control Section: IEFW42SD

IEFSD510: Initiator -- Job Selection Routine

This routine selects a system or problem program job. This module executes only in a large (scheduler-size) partition.

- Entry: IEFSD510
- Exits: Branch to IEFSD511 or IEFSD515, LINK to IEFSD519, XCTL to IEFSD589, SMALTERM, or IEEVSTAR
- Tables/Work Areas: LOT block, CSCB, SPIL, CVT, TCB, PIB
- Attributes: Read-only, reenterable
- Control Section: IEFSD510
- External References: IEFQMDQQ, IEFQMUNC
- Page Reference: 63

IEFSD511: Initiator -- Job Initiation Routine

This routine initializes information pertaining to a job.

- Entry: IEFSD511
- Exit: Branch to IEFSD541
- Tables/Work Areas: Life-of-Task Block, CSCB, JCT, SCT, SCD, PIB, IOB2
- Attributes: Read-only, Reenterable
- Control Section: IEFSD511
- External References: IEFQMRAW
- Page Reference: 68

IEFSD512: Initiator -- Step Initiation Routine

This routine passes control to allocation as a closed subroutine via a LINK macro instruction. If an allocation error occurs, it passes control to the Alternate Step Deletion routine. Otherwise, it continues normally and schedules a job step.

- Entry: IEFSD512
- Exits: Branch to IEFSD513, IEFSD516, or IEFSD518, XCTL to IEFSD510
- Tables/Work Areas: LOT Block, JCT, SCT, APL, TIOT, CSCB, IOB1, IOB2
- Attributes: Read-only, reenterable
- Control Section: IEFSD512
- External References: IEFQMRAW, IEFSD556, IEFSD514
- Page Reference: 68

IEFSD513: Initiator -- Problem Program Interface

This routine prepares the partition for problem program execution by moving the TIOT to the highest available storage area.

The routine also opens JOBLIB and FETCH DCBS, if required. A final check is made to determine if a CANCEL command has been received for the job before the problem program is brought into the partition and given control. If scheduling was performed for a small partition, IEFSD513 communicates with the small partition.

- Entry: IEFSD513
- Exits: XCTL to problem program, ABEND, or to IEFSD510
- Tables/Work Areas: LOT Block, Transfer Parameter List, TIOT, User's Parameter List, TCB, CVT, PIB, CSCB, SPIL, APL, JCT, SCT, DCB
- Attributes: Read-only, reenterable
- Control Section: IEFSD513
- Page Reference: 69

IEFSD514: Queue Management -- Table Breakup Routine

This routine reads and writes tables which may be required by the job scheduler. The routine breaks the tables into 176-byte records, writes the records on disk, and

retrieves the records from disk to reconstruct the tables in main storage.

- Entry: IEFSD514
- Exit: Return to caller
- Tables/Work Areas: QMPA, TBR Parameter List
- Attributes: Read-only, reenterable
- Control Section: IEFSD514
- External References: IEFQASGN, IEFQMRAW
- Page Reference: 59

IEFSD515: Initiator -- Step Deletion Routine

This routine retrieves the TIOT and Life-of-Task Block from disk, reads in the JCT and SCT, and branches to termination, which is used as a closed subroutine. It also reads in the SCT for the next step to be scheduled, if required.

- Entry: IEFSD515, SMALTERM, or GO
- Exits: XCTL to IEFSD512 or Branch to IEFSD517 or IEFSD510
- Tables/Work Areas: Life-of-Task Block, Terminate Parameter List, CVT, TCB, PIB, IOB, CSCB, DCB, JCT, SCT, SPIL
- Attributes: Read-only, reenterable
- Control Section: IEFSD515
- External References: IEFQMRAW, IEFSD514, IEFSD42Q, IEFSD598
- Page Reference: 70

IEFSD516: Initiator -- Alternate Step Deletion Routine

This routine provides an interface with termination when an allocation error occurs during step initiation. Termination is used as a closed subroutine. If required, this routine reads the SCT of the next step to support job flushing.

- Entry: IEFSD516
- Exits: Branch to IEFSD512 or IEFSD517
- Tables/Work Areas: Life-of-Task block, CSCB, Terminate Parameter List, SCT
- Attributes: Read-only, reenterable

- Control Section: IEFSD516
- External References: IEFQMRAW, IEFSD42Q
- Page Reference: 71

IEFSD517: Initiator -- Job Deletion Routine

This routine deletes the disk queue entry for a terminated job and unchains and deletes the CSCB for the job.

- Entry: IEFSD517
- Exit: Branch to IEFSD510
- Tables/Work Areas: CSCB, Life-of-Task block, SPIL
- Attributes: Read-only, reenterable
- Control Section: IEFSD517
- External References: IEFQDELE, IEFSD598
- Page Reference: 71

IEFSD518: Initiator -- Partition Recovery Routine

This routine determines the status of main storage required for a checkpoint/restart.

- Entry: IEFSD518
- Exits: Return to IEFSD512
- Tables/Work Areas: SPIL, CVT, TCB, JCT, PIB, LOT, QMPA, CSCB
- Attributes: Reenterable
- Control Section: IEFSD518
- External Reference: IEFQMRAW, IEFQMNQ2, IEFSD598
- Page Reference: 71

IEFSD519: Queue Management -- Dequeue by Jobname Interface Routine

This routine builds a seven-word parameter list used by IEFLOCDQ to locate a job by jobname on the checkpoint/restart internal queue.

- Entry: IEFSD519
- Exit: Return to IEFSD510
- Tables/Work Areas: IOT, PIB

- Attributes: Reenterable
- Control Section: IEFSD519
- External Reference: IEFLOCDQ, IEFQMRAW
- Page Reference: 71

IEFSD530: Interpreter -- Transient Reader Suspend Routine

This routine closes the reader input data set and procedure library, and saves data required to restore the reader.

- Entry: IEFSD530
- Exit: Return to caller
- Tables/Work Areas: IWA, TIOT, LWA, QMPA, CVT, UCB, MSRC, PIB, CSCB
- Attributes: Read-only, reenterable
- Control Section: IEFSD530
- External References: IEFSD514, IEF-QMRAW, IEFQASNM, IEFQASGN
- Page Reference: 62

IEFSD531: Interpreter -- Transient Reader Restore Routine

This routine restores the information required to "restart" a transient reader after it has been suspended. It reopens the reader input data set and procedure library.

- Entry: IEFSD531
- Exit: XCTL to IEFVHCB
- Tables/Work Areas: IWA, TIOT, QMPA, CVT, UCB, MSRC, PIB, CSCB
- Attributes: Read-only, reenterable
- Control Sections: IEFSD531, IEFPH2
- External References: IEFSD514, IEF-QMRAW, IEFQASNM, IEFQASGN
- Page Reference: 62

IEFSD532: Interpreter -- Transient Reader Suspend Tests

This routine determines the status of a transient reader. IEFSD532 receives control from IEFVHH after a job has been enqueued.

- Entry: IEFKG

- Exits: XCTL to IEFVHN or IEFSD530, or branch to IEFVHHB
- Tables/Work Areas: IWA, LWA, QMPA, PIB, CVT
- Attributes: Read-only, reenterable
- Control Section: IEFKG
- Page Reference: 62

IEFSD533: Interpreter -- Interface Routine

This routine provides an interface between the reader/interpreter and system task control.

- Entry: IEFIRC
- Exits: XCTL to IEFSD537. RETURN to STC if error.
- Tables/Work Areas: CSCB, CVT, QMPA
- Attributes: Reenterable, read-only
- Control Section: IEFIRC

IEFSD534: System Task Control -- IPSW Routine

This routine places system task control in problem program mode by loading a PSW.

- Entry: IEFSD534
- Exit: IEFVSTRT
- Tables/Work Areas: None
- Attributes: Reenterable
- Control Section: IEFSD534

IEFSD535: System Task Control -- Problem Program Mode Routine

This routine puts system task control in problem program mode for ABEND.

- Entry: IEFSD535
- Exit: IEEVICTL
- Tables/Work Areas: None
- Attributes: Reenterable
- Control Section: IEFSD535

IEFSD536: Interpreter -- Operator Message Routine

This routine writes a message to the operator when an I/O error or CPO full condition

has occurred. The routine also sets proper indicators to cause a cleanup of the current job.

- Entry: IEFVHR
- Exits: Return to caller, XCTL to IEFVHN, or LINK to IEFSD308
- Tables/Work Areas: IWA, JCT, LWA, UCB, CVT, PIB, CSCB, Master Scheduler resident data area
- Attributes: Read-only, reenterable
- Control Section: IEFVHR
- Page Reference: 61

IEFSD537: Interpreter -- Linkage Module

This routine provides an interface between system task control and a reader. It also frees the interpreter entrance list (NEL) and associated areas if a reader is being terminated or suspended.

- Entry: IEFSD537
- Exits: LINK to IEFVH1, or IEFSD531, or Return to system task control
- Tables/Work Areas: NEL
- Attributes: Read-only, reenterable
- Control Section: IEFSD537

IEFSD540: Initiator -- Linkage to IEFSD541

This routine provides an interface linkage to IEFSD541 via an XCTL macro instruction.

- Entry: IEFSD540
- Exit: XCTL to IEFSD541
- Tables/Work Areas: Same as caller
- Attributes: Read-only, reenterable
- Control Section: IEFSD540

IEFSD541: Initiator -- Data Set Integrity

This routine enqueues on explicit data sets and thus prevents concurrent, and impairing, access between tasks.

- Entry: IEFSD541
- Exit: Branch to IEFSD512
- Tables/Work Areas: LOT Flock, IOB1, IOB2, JCT, SCT, CSCB, SPIL, DSENG Table, Minor Name List, ENQ supervisor list.

- Attributes: Read-only
- Control Section: IEFSD541
- External References: IEFQMRAW
- Page Reference: 68

IEFSD551: I/O Device Allocation -- Linkage to IEFXJIMP

This routine provides an interface linkage to IEFXJIMP via an XCTL macro instruction in the 30K design package.

- Entry: IEFV15XL
- Exit: XCTL to IEFXJIMP
- Tables/Work Areas: Same as caller
- Attributes: Read-only, reenterable
- Control Section: IEFV15XL

IEFSD552: I/O Device Allocation -- Linkage to IEFXJIMP

This routine provides an interface linkage to IEFXJIMP via an XCTL macro instruction in the 30K design package.

- Entry: IEFXJX5A
- Exit: XCTL to IEFXJIMP
- Tables/Work Areas: Same as caller
- Attributes: Read-only, reenterable
- Control Section: IEFXJX5A

IEFSD553: Initiator -- Linkage to IEFSD512

This routine provides a linkage to IEFSD512 via an XCTL macro instruction.

- Entry: IEFSD512
- Exit: XCTL to IEFSD512
- Tables/Work Areas: Same as caller
- Attributes: Read-only, reenterable
- Control Section: IEFSD512

IEFSD554: Initiator -- Linkage to IEFSD516

This routine provides a linkage to IEFSD516 via an XCTL macro instruction.

- Entry: IEFSD554
- Exit: XCTL to IEFSD516
- Tables/Work Areas: Same as caller

- Attributes: Read-only, reenterable
- Control Section: IEFSD554

IEFSD555: Initiator -- Linkage to IEFSD510

This routine provides linkage to IEFSD510 via an XCTL macro instruction.

- Entry: IEFSD555
- Exit: XCTL to IEFSD510
- Tables/Work Areas: Same as caller.
- Attributes: Read-only, reenterable
- Control Section: IEFSD555

IEFSD556: Initiator -- Set Problem Program State Return

This routine establishes the allocation routine in a problem program state, upon entry.

- Entry: IEFSD556
- Exit: LPSW to IEFW21SD
- Tables/Work Areas: Same as caller.
- Attributes: Read-only, reenterable
- Control Section: IEFSD556

IEFSD557: I/O Device Allocation -- Interface Routine

This routine provides an interface between system task control and allocation.

- Entry: IEFW21SD
- Exit: IEFWSD21
- Tables/Work Areas: ECB, IOB
- Attributes: Reenterable
- Control Section: IEFSD557

IEFSD558: Initiator -- Linkage to IEFSD511

This routine provides a linkage to IEFSD511 via an XCTL macro instruction.

- Entry: IEFSD558
- Exit: IEFSD511
- Attributes: Read-only, reenterable
- Control Section: IEFSD558

IEFSD559: Initiator -- Linkage to IEFSD515

This routine provides a linkage to IEFSD515 via an XCTL macro instruction.

- Entry: SMALTERM
- Exit: IEFSD515
- Attributes: Read-only, reenterable
- Control Section: IEFSD559

IEFSD567: Nucleus -- Device-End Interrupt Handler Routine

This routine handles unsolicited device-end interrupts from a disk storage unit.

- Entry: IEFSD567
- Exit: Return to caller
- Tables/Work Areas: None
- Attributes: Reenterable
- Control Section: IEFSD567
- External Reference: Communications Task TCB

IEFSD569: Master Scheduler -- Initialization Routine

This routine initializes the communications task and the system log. It issues the READY message and formats the job queue, as well as typing out the automatic commands and invoking processing of the automatic commands. This routine establishes partitioning of main storage at system initialization and readies the partitions for the START command. This routine is called out at system generation by the macro, SGIEE0VV.

- Entry: IEFSD569
- Exit: IEE0503D, Branch to dispatcher
- Attributes: Read-only, non-reenterable
- Control Section: IEFSD569
- Page Reference: 48

IEFSD572: Queue Management -- Interpreter/Queue Manager Interlock Routine

This routine determines if a possible interlock condition exists between the queue manager and the reader. The routine issues a message requesting the operator to reply with either WAIT, to wait for space to be freed, or CANCEL, to cancel the job.

- Entry: IEFSD572
- Exits: ABEND, or return to caller
- Attributes: Read-only, reenterable
- Control Section: IEFSD572, IEFSD573
- External Reference: IEFQDELQ
- Page Reference: 57

IEFSD587: System Task Control -- Linkage to IEFSD535

This routine provides a linkage to IEFSD535 via a LINK macro instruction.

- Entry: IEFSD587
- Exit: IEFSD535
- Attributes: Read-only, reenterable
- Control Section: IEFSD587

IEFSD588: System Task Control -- Linkage to IEE534SD

This routine links to IEE534SD to bring the suspended reader into the assigned partition so that upon return, the initiator will be in supervisor state.

- Entry: IEFSD588
- Exit: LINK to IEE534SD
- Tables/Work Areas: Same as caller
- Attributes: Read-only, reenterable
- Control Section: IEFSD588

IEFSD589: Initiator -- Linkage to IEFSD534

This routine links to system task control so that upon return, the initiator will be in supervisor state.

- Entry: IEFSD589
- Exit: LINK to IEFSD534
- Tables/Work Areas: Same as caller
- Attributes: Read-only, reenterable
- Control Section: IEFSD589

IEFSD597: Initiator -- Shared DASD ENQ/DEQ Purge Routine

This routine is the purge routine for systems that include the shared DASD feature. In addition to purging all resources enqueued by a job step, but not dequeued, IEFSD597 also releases reserved devices.

- Entry: IEFSD597
- Exit: Return to caller
- Tables/Work Areas: Major QCB, Minor QCB, QEL, TCB, SVRB, CVT, ABTERM
- Attributes: Read-only, reenterable, disabled
- Control Section: IEFSD597

IEFSD598: Initiator -- ENQ/DEQ Purge Routine

This routine purges all resources enqueued by a job step, but not dequeued.

- Entry: IEFSD598
- Exit: Return to caller
- Tables/Work Areas: Major QCB, Minor QCB, QEL, TCB, SVRB, CVT, ABTERM
- Attributes: Read-only, Reenterable, disabled
- Control Section: IEFSD598
- Page Reference: 70

IEFSD599: Initiator -- Small Partition Module

This routine provides an interface with the scheduler in a large partition to initiate and terminate small partitions.

- Entry: IEFSD599, SMALLGO
- Exits: ABEND, or XCTL to problem program or writer
- Tables/Work Areas: SPII, allocate parameter list (API)
- Attributes: Read-only, reenterable
- Control Section: IEFSD599
- External Reference: IEFQMUNC
- Page Reference: 66

IEFVDA: Interpreter -- DD Statement Processor

This routine constructs and adds entries to a JFCB and SIOT from the complete logical DD statement in the internal text buffer.

- Entry: IEFVDA
- Exit: To IEFVHF
- Tables/Work Areas: IWA, LWA, SIOT, JFCB, JCB, SCT

- Attributes: Read-only, reenterable
- Control Section: IEFVDA

IEFVDBSD: Interpreter -- Data Set Name Table Construction Routine

This routine creates a data set name table.

- Entry: IEFVDBSD
- Exit: To IEFVDA
- Attributes: Reenterable
- Control Section: IEFVDBSD

IEFVEA: Interpreter -- EXEC Statement Processor

This routine constructs or updates an SCT, and, if necessary, a joblib JFCB and SIOT from the complete logical EXEC statement in the internal text buffer.

- Entry: IEFVEA, from IEFVFA
- Exit: To IEFVHF
- Tables/Work Areas: IWA, EXEC work area, interpreter key table, JCT, SCT, SIOT, QMPA, procedure override table.
- Attributes: Read-only, reenterable
- Control Section: IEFVEA

IEFVFA: Interpreter -- Scan Routine

This routine scans the card image of a JOB, EXEC, or DD statement, performs error checking of JCL syntax, builds internal text, and, when a complete logical statement (including continuations and overrides) has been scanned, passes control to the appropriate statement processor.

- Entry: IEFVFA
- Exits: To IEFVGM, IEFVHQ, IEFVHF, IEFVJA, IEFVDA, IEFVEA
- Tables/Work Areas: IWA, scan routine work area, interpreter key table, QMPA, internal text buffer, scan dictionary.
- Attributes: Read-only, reenterable
- Control Section: IEFVFA

IEFVFB: Interpreter -- Symbolic Parameter Processing Routine

This routine processes symbolic parameters by creating symbolic parameter table buffer entries to assign values to symbolic parameters, and extracts those values and places

them in the intermediate text buffer when a symbolic parameter is used.

- Entry: IEFVFB
- Exit: Return to caller
- Tables/Work Areas: IWA, IWA SYMBUF, Intermediate Text Buffer, QMPA
- Attributes: Read-only, reenterable
- Control Section: IEFVFB

IEFVGI: Interpreter -- Dictionary Entry Routine

This routine constructs entries for the refer-back dictionary.

- Entry: IEFVGI
- Exit: Return to caller
- Tables/Work Areas: Refer-back dictionary, auxiliary work area, IWA, QMPA
- Control Section: IEFVGI

IEFVGK: Interpreter -- Get Parameter Routine

This routine searches the internal text buffer for the next parameter, performs basic error checking, and passes control to the appropriate keyword routine.

- Entry: IEFVGK
- Exit: Return to caller
- Tables/Work Areas: Local work area, IWA, internal text buffer, KBT, PDT.
- Control Section: IEFVGK

IEFVGM: Interpreter -- Message Processing Routine

This routine constructs SMBs containing interpreter error messages and JCL statement images, assigns space for these SMBs in the message class output queue entry, and writes the SMBs into the entry.

- Entry: IEFVGM
- Exit: Return to caller
- Tables/Work Areas: QMPA, SMB, SCD, IWA, JCT
- Attributes: Reenterable, character dependence type C
- Control Section: IEFVGM

IEFVGM1: Interpreter -- Message Module

This routine contains interpreter messages 01-07.

- Attributes: Non-executable
- Control Section: IEFVGM1

IEFVGM2: Interpreter -- Message Module

This routine contains interpreter messages 08-0F.

- Attributes: Non-executable
- Control Section: IEFVGM2

IEFVGM3: Interpreter -- Message Module

This routine contains interpreter messages 10-17.

- Attributes: Non-executable
- Control Section: IEFVGM3

IEFVGM4: Interpreter -- Message Module

This routine contains interpreter messages 18-1F.

- Attributes: Non-executable
- Control Section: IEFVGM4

IEFVGM5: Interpreter -- Message Module

This routine contains interpreter messages 20-27.

- Attributes: Non-executable
- Control Section: IEFVGM5

IEFVGM6: Interpreter -- Message Module

This routine contains interpreter messages 28-2F.

- Attributes: Non-executable
- Control Section: IEFVGM6

IEFVGM7: Interpreter -- Message Module

This routine contains interpreter messages 30-37.

- Attributes: Non-executable
- Control Section: IEFVGM7

IEFVGM8: Interpreter -- Message Module

This routine contains interpreter messages 50-57.

- Attributes: Non-executable
- Control Section: IEFVGM8

IEFVGM9: Interpreter -- Message Module

This routine contains interpreter messages 58-5F.

- Attributes: Non-executable
- Control Section: IEFVGM9

IEFVGM10: Interpreter -- Message Module

This routine contains interpreter messages 60-67.

- Attributes: Non-executable
- Control Section: IEFVGM10

IEFVGM11: Interpreter -- Message Module

This routine contains interpreter messages 68-6F.

- Attributes: Non-executable
- Control Section: IEFVGM11

IEFVGM12: Interpreter -- Message Module

This routine contains interpreter messages 70-77.

- Attributes: Non-executable
- Control Section: IEFVGM12

IEFVGM13: Interpreter -- Message Module

This routine contains interpreter messages 78-7F.

- Attributes: Non-executable
- Control Section: IEFVGM13

IEFVGM14: Interpreter -- Message Module

This routine contains interpreter messages 88-8F.

- Attributes: Non-executable
- Control Section: IEFVGM14

IEFVGM15: Interpreter Message -- Module

This routine contains interpreter messages 90-97.

- Attributes: Non-executable
- Control Section: IEFVGM15

IEFVGM16: Interpreter -- Message Module

This routine contains interpreter messages A0-A7.

- Attributes: Non-executable
- Control Section: IEFVGM16

IEFVGM17: Interpreter -- Message Module

This routine contains interpreter messages 56-5D.

- Attributes: Non-executable
- Control Section: IEFVGM17

IEFVGM18: Interpreter -- Message Module

This routine contains interpreter messages 80-87.

- Attributes: Non-executable
- Control Section: IEFVGM18

IEFVGM19: Interpreter -- Message Module

This routine contains interpreter messages 3E-45.

- Attributes: Non-executable
- Control Section: IEFVGM19

IEFVGM70: Interpreter -- Message Module

This routine contains interpreter messages 38-3F.

- Attributes: Non-executable
- Control Section: IEFVGM70

IEFVGM78: Interpreter -- Message Module

This routine contains interpreter messages 08-0D.

- Attributes: Non-executable
- Control Section: IEFVGM78

IEFVGS: Interpreter -- Dictionary Search Routine

This routine searches the refer-back dictionary for the address of a previously-defined SCT, SIOT, or JFCB.

- Entry: IEFVGS

- Exit: Return to caller
- Tables/Work Areas: Auxiliary work area, IWA, QMPA, refer-back dictionary
- Control Section: IEFVGS

IEFVGT: Interpreter -- Test and Store Routine

This routine performs operations on a parameter as indicated in the appropriate parameter descriptor table entry.

- Entry: IEFVGT
- Exit: Return to keyword routine
- Tables/Work Areas: Internal text buffer, PDT, local work area, IWA
- Control Section: IEFVGT

IEFVHA: Interpreter -- Get Routine

This routine reads statements from the input stream and the procedure library.

- Entry: IEFVHA
- Exits: IEFVHC, IEFVHB, IEFVHAA, IEFSD536, IEFVGM
- Tables/Work Areas: IWA, JCT, DCB
- Attributes: Read-only, reenterable
- Control Section: IEFVHA

IEFVHAA: Interpreter -- End-of-File Routine

This routine determines the conditions under which an end-of-file condition has occurred, and sets switches and passes control accordingly.

- Entry: IEFVHAA
- Exits: IEFVHC or IEFVHN
- Tables/Work Areas: IWA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHAA

IEFVHB: Interpreter -- DD * Statement Generator Routine

This routine generates a "SYSIN DD *" statement for data in the input stream, when no such statement was included.

- Entry: IEFVHB

- Exits: To IEFVHC, IEFVHA, IEFVGM
- Tables/Work Areas: IWA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHE

IEFVHC: Interpreter -- Continuation Statement Routine

This routine determines whether the current statement should be a continuation, and, if so, determines whether it is a valid continuation statement.

- Entry: IEFVHC
- Exits: To IEFVHEB, IEFVHCB, IEFVGM
- Tables/Work Areas: IWA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHC

IEFVHCB: Interpreter -- Verb Identification Routine

This routine identifies the verb in a control statement.

- Entry: IEFVHCB
- Exits: To IEFVHE, IEFVHM, IEFVHA, IEFVGM, IEFVHL
- Tables/Work Areas: IWA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHCB

IEFVHE: Interpreter -- Router

This routine determines the conditions under which it was entered, and passes control to the appropriate routine.

- Entry: IEFVHE
- Exits: To IEFVHEB, IEFVHH, IEFVHEC
- Tables/Work Areas: IWA
- Attributes: Read-only, reenterable
- Control Section: IEFVHE

IEFVHEB: Interpreter -- Pre-Scan Preparation Routine

This routine determines whether a message is required or additional work queue space is required before a statement is scanned. If so, it causes the message to be written or the work queue space to be assigned.

- Entry: IEFVHEB
- Exits: To IEFVHQ, IEFVGM, IEFVHG, IEFVFA
- Tables/Work Areas: IWA, JCT, SCT, QMPA
- Attributes: Read-only, reenterable
- Control Section: IEFVHEB

IEFVHEC: Interpreter -- Job Validity Check Routine

This routine determines whether an SCT has been built for the current job; if not, the routine constructs an SCT.

- Entry: IEFVHEC
- Exits: To IEFVGM, IEFVHH
- Tables/Work Areas: IWA, JCT, SCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHEC

IEFVHF: Interpreter -- Post-Scan Routine

This routine determines the conditions under which it was entered, and passes control accordingly.

- Entry: IEFVHF
- Exits: To IEFVHG, IEFVHEB, IEFVHCB, IEFVHA
- Tables/Work Areas: IWA, CWA
- Attributes: Read-only, reenterable
- Control Section: IEFVHF

IEFVHG: Interpreter -- CPC Routine

This routine writes system input data sets on a direct-access device. If IEFVHG is unable to obtain enough space to complete writing a data set, control passes to IEFVHR. If the input reaches end-of-file, control passes to IEFVHAA. If a /* is found following DD DATA, control passes to IEFVHA to read the next record. If a // is found, control passes to IEFVHC to identify the verb.

- Entry: IEFVHG
- Exits: To IEFSD536, IEFVGM, IEFVHQ, IEFVHAA, IEFVHA, IEFVHC, or IEFVHEB
- Tables/Work Areas: IWA, JCT, SIOT, VOLT, TIOT, LWA, SCT, JFCB, UCB, QMPA, CWA

- Attributes: Read-only, reenterable
- Control Section: IEFVHG
- Page Reference: 61

IEFVHH: Interpreter -- Job and Step Enqueue Routine

This routine places the SCT, DSNT, VOLT, and JCT in the job's queue entry, and determines whether the interpreter is to enqueue jobs.

- Entry: IEFVHH
- Exits: To IEFKFG, IEFVHQ, IEFSD532, IEFVHHB, IEFVHN
- Tables/Work Areas: IWA, JCT, SCT, QMPA, NEL
- Attributes: Read-only, reenterable
- Control Section: IEFVHH

IEFVHHB: Interpreter -- Housekeeping Routine

This routine initializes for merging a cataloged procedure.

- Entry: IEFVHHB
- Exits: IEFVHA, IEFVHEB
- Tables/Work Areas: IWA
- Attributes: Read-only, reenterable
- Control Section: IEFVHHB

IEFVHL: Interpreter -- Null Statement Routine

This routine determines the conditions under which the null statement was encountered, and passes control to the proper routine.

- Entry: IEFVHL
- Exits: To IEFVHCB, IEFHEC, IEFVHE, IEFVHA
- Tables/Work Areas: IWA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHL

IEFVHM: Interpreter -- Command Statement Routine

This routine tests for valid command verbs, and, if the verb is valid, issues SVC 34 to schedule execution of the command.

- Entry: IEFVHM
- Exits: To IEFVHA, IEFVGM
- Tables/Work Areas: IWA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVHM

IEFVHN: Interpreter -- Termination Routine

This routine closes the input stream and procedure library data sets, frees main storage used by the interpreter, and builds the interpreter exit list.

- Entry: IEFVHN
- Exit: Return to caller
- Tables/Work Areas: IWA, JCT, CSCB, QMPA
- Attributes: Read-only, reenterable
- Control Section: IEFVHN
- Page Reference: 62

IEFVHQ: Interpreter -- Queue Management Interface Routine

This routine is a common interface between the queue management routines and the interpreter. If an I/C error occurs, IEFVHR receives control. Queue management may be unable to allocate space for a job's input data. If, in this case, the operator replies CANCEL to the message which is issued, IEFVHG receives control.

- Entry: IEFVHQ
- Exits: Return to caller, IEFSD536, or IEFVHG
- Tables/Work Areas: IWA, JCT, QMPA, CSCB
- Attributes: Read-only, reenterable
- Control Section: IEFVHQ

IEFVH1: Interpreter -- Initialization Routine

This routine initializes the interpreter; it obtains main storage for and initializes the IWA, local work areas, and DCBs.

- Entry: IEFVH1
- Exit: To IEFVH2
- Tables/Work Areas: UCB, CSCB, IWA, DCB, local work area

- Attributes: Not reusable
- Control Section: IEFVH1

IEFVH2: Interpreter -- Initialization Routine

This routine opens the input stream data set and the procedure library data set, and obtains main storage for a buffer for procedure library records.

- Entry: IEFVH2
- Exit: To IEFVHA
- Tables/Work Areas: IWA, UCB, TIOT
- Control Section: IEFVH2
- Attributes: Not reusable

IEFVJA: Interpreter -- Job Statement Processor

This routine initializes a JCT and job ACT from the complete logical job statement in the internal text buffer.

- Entry: IEFVJA
- Exit: To IEFVHF
- Tables/Work Areas: IWA, job work area, interpreter key table, JCT, ACT, QMPA
- Attributes: Read-only, reenterable
- Control Section: IEFVJA

IEFVJIMP: Termination -- JOB Statement Condition Code Processor

This routine tests the condition codes specified in the JOB statement to determine whether the remaining steps in the job are to be run.

- Entry: IEFVJ
- Exits: To IEFVK or IEFZA
- Tables/Work Areas: LCT, JCT, SCT
- Control Section: IEFVJ

IEFVJMSG: Termination -- JOB Statement Condition Code Processor Messages

This routine contains the messages issued to the programmer by the JOB statement condition code processor.

- Entry: IEFVJMSG
- Attributes: Non-executable
- Control Section: IEFVJMSG

IEFVKIMP: I/O Device Allocation -- EXEC Statement Condition Code Processor

This routine tests the condition codes specified in the EXEC statement to determine whether the next step of the job is to be run.

- Entry: IEFVK
- Exits: IEFVS, IEFLB
- Tables/Work Areas: JCT, LCT, SCT
- Control Section: IEFVK

IEFVKMSG: I/O Device Allocation -- EXEC Statement Condition Code Processor Messages

This routine contains the messages issued to the programmer by the EXEC -- statement condition -- code processor.

- Entry: IEFVKMJ1
- Attributes: Non-executable
- Control Section: IEFVKMSG

IEFVMFAK: I/O Device Allocation -- Linkage to IEFVMLS1

This routine passes control to entry point IEFVMCVL of the JFCB housekeeping module IEFVMLS1 via the XCTL macro instruction.

- Entry: IEFVMCVL
- Exit: To IEFVMCVL
- Control Section: IEFVMCVL

IEFVMLS1: I/O Device Allocation -- JFCB Housekeeping Control Routine and Allocate Processing Routines

The control routine obtains the required SIOTs, determines the processing required for each, and passes control to the appropriate routine. The allocate processing routine performs the processing required in certain refer-back situations, when the data set is cataloged or passed, and when unit name is specified.

- Entry: IEFVM, IEFVMCVL, IEFVMQMI, VM7000, VM7055, MV7055AA, VM7060, MV7070, MV7090, MV7130, VM7370, MV7700, MV7742, MV7750, MV7850, VM7900, MV7950
- Exits: To IEFVM2LS, IEFVM3LS, IEFVM4LS, IEFVM5LS, IEFVM6LS, and IEFXCSSS
- Tables/Work Areas: LCT, JCT, PDQ, SIOT, JFCB, QMPA

- Control Section: IEFVM1

IEFVMS6: I/O Device Allocation -- JFCB Housekeeping Error Message Processing Routine

This routine prepares error messages for the JFCB housekeeping routines.

- Entry: IEFVMSGR
- Exit: Return to caller
- Tables/Work Areas: JCT, LCT
- Control Section: IEFVM6

IEFVMS7: I/O Device Allocation -- JFCB Housekeeping Error Messages

This routine contains the messages issued by the JFCB housekeeping routines.

- Entry: IEFVM7
- Attributes: Non-executable
- Control Section: IEFVM7

IEFVMS1: I/O Device Allocation -- Linkage to JFCB Housekeeping

This routine provides a linkage to the JFCB housekeeping routines for the step flush function.

- Entry: IEFVM1
- Exit: XCTL to IEFVMS1
- Attributes: Read-only, reenterable
- Control Section: IEFVM1

IEFVPOST: I/O Device Allocation -- Unsolicited Device Interrupt Handler

This routine handles the posting of unsolicited device interruptions for I/O device allocation operating in a multiprogramming environment.

- Entry: IEFDPOST
- Exits: To IEAOPT01 or Return to caller
- Tables/Work Areas: CSCB, ECB, TCB
- Attributes: Read-only, reenterable, disabled, resident
- Control Section: IEFDPOST

IEFVM2LS: I/O Device Allocation -- JFCB Housekeeping Fetch DCB Processing Routine

This routine updates the SIOT, SCT, JFCB and VOLT with information required for the allocation of devices for the fetch DCB.

- Entry: VM7100
- Exit: To IEFVMS1
- Tables/Work Areas: LCT, SCT, SIOT, JFCB, VOLT
- Control Section: IEFVM2

IEFVM3LS: I/O Device Allocation -- JFCB Housekeeping GDG Single Processing Routine

This routine obtains the fully qualified name of a member of a generation data group (GDG), and completes the required information in the JFCB, VOLT, and SIOT for that member.

- Entry: VM7150
- Exit: To IEFVMS1
- Tables/Work Areas: LCT, SIOT, GDG Bias Count table, JFCB
- Control Section: IEFVM3

IEFVM4LS: I/O Device Allocation -- JFCB Housekeeping GDG All Processing Routine

This routine builds an SIOT, JFCB, and VOLT, and PDQ entries for each member of the GDG.

- Entry: VM7200
- Exit: To IEFVMS1
- Tables/Work Areas: LCT, SCT, VOLT, PDQ, SIOT, JFCB
- Control Section: IEFVM4

IEFVM5LS: I/O Device Allocation -- JFCB Housekeeping Patterning DSCB Routine

This routine establishes DCB control information within a JFCB.

- Entry: VM7300
- Exit: To IEFVMS1
- Tables/Work Areas: LCT, SCT, SIOT, DSCB, JFCB
- Control Section: IEFVM5

IEFVM76: I/C Device Allocation -- JFCB Housekeeping Unique Volume ID Routine

This routine creates unique volume serials for unlabeled tape data sets, when the disposition is "PASS".

- Entry: VM7600
- Exit: Return to caller
- Tables/Work Areas: SIOT, JFCB, JFCBX
- Control Section: IEFVM76

IEFVRR3: Reinterpretation Control Routine

This routine passes control among the routines that modify the queue entry of a restart step so that they appear as they were prior to the initiation of the step.

- Entry: IEFVRR3, IEFVRRCA, IEFVRRCB
- Exit: Return to caller
- Attributes: Read-only reenterable
- Tables/Work Areas: NEL, JCT, SCT, SIOT, JFCB, JFCBX, VOLT, SMB, DSENQ, SCD, DSB, QMPA
- Control Section: IEFVRR3

IEFVRR1: Dequeue Interface Routine

This routine interfaces with queue management to cause a specific job to be dequeued and the JCT for that job to be read into main storage.

- Entry: IEFVRR1
- Exit: Return to caller
- Tables/Work Areas: QMPA, JCT
- Attributes: Read-only, reenterable
- Control Section: IEFVRR1

IEFVRR2: Table Merge Routine

This routine merges the reinterpreted queue entry tables of a restart step with the original queue tables for that step.

- Entry: IEFVRR2, IEFVRR2AE
- Exit: Return to caller
- Tables/Work Areas: QMPA, JCT, ACT, SMB, SCT, SIOT, JFCB, DSENQ, VOLT, JFCBX, NEL
- Attributes: Reenterable
- Control Section: IEFVRR2

IEFVRR3: Reinterpretation Delete/Enqueue Routine

This routine deletes the reinterpreted input and output queue entries of a restart step, constructs the internal JCL necessary for processing a checkpoint restart, and reenqueues the job's queue entry.

- Entry: IEFVRR3, IEFVRR3AE
- Exit: Return to caller
- Tables/Work Areas: QMPA, JCT, SCT, SIOT, JFCB
- Attributes: Reenterable
- Control Section: IEFVRR3

IEFVSDRA: Restart Activation Routine

This routine issues a START Restart Reader command for one or more jobnames. This routine is entered from IEFSD168 during a problem program restart or IEFSD305 during a warm start.

- Entry: IEFVSDRA
- Exit: Return to caller
- Tables/Work Areas: CSCB, CVT, TCB
- Attributes: Reenterable
- Control Section: IEFVSDRA

IEFVSDRD: Restart Determination Routine

This routine initiates automatic restarts.

- Entry: IEFVSDRD
- Exit: To IEFSD305
- Tables/Work Areas: JCT, SCT, QMPA, CVT, SIOT, ICT
- Attributes: Reenterable
- Control Section: IEFVSDRD

IEFVSD12: Interpreter -- CPO Allocation Subroutine

This routine sets up a JFCB and allocates space on a direct-access device for a system input data set.

- Entry: IEFSD012
- Exit: Return to caller
- Attributes: Reenterable
- Tables/Work Areas: IWA, QMPA, LWA, SIOT, TIOT, UCB, JFCB, JCT, CSCB
- Control Section: IEFSD012
- External References: IEFVHQ

IEFVSD13: Interpreter -- SCD Construction Routine

This routine constructs an SCD entry for each system output class defined for a job, and assigns space for all DSBs that will be required.

- Entry: IEFSD090
- Exit: Return to caller
- Tables/Work Areas: IWA, QMPA, DD work area, SCD, SCT, SIOT, JCT, JFCB
- Control Section: IEFSD090

IEFVSMBR: SMB Reader Routine

This routine reads the SMBs associated with a restarting job and converts the JCL statements to their original format.

- Entry: IGC0005B
- Exits: If called during restart reader processing, return to caller; if called during restart, XCTL to the first load of restart housekeeping.
- Tables/Work Areas: QMPA, DCB, JCT, SMB, RRCWKAR, SCT
- Attributes: Reenterable
- Control Section: IEFVSMBR

IEFWA000: I/O Device Allocation -- Demand Allocation Routine

This routine establishes data set device requirements, and allocates in response to specific unit requests.

- Entry: IEFWA000, IEFUCBL

- Exits: To IEFWD000, IEFX3000, IEFX5000
- Tables/Work Areas: UCB Address List, DMT, UCB, LCT, SCT, SIOT, VOLT, AWT
- Control Sections: IEFWA7, IEFWA002

IEFWCFAK: I/O Device Allocation -- Linkage Module

This routine passes control to the TIOT construction routine.

- Entry: IEFWC000, IEFWC002
- Exit: To IEFWCIMP
- Control Section: IEFWC000, IEFWC002

IEFWCIMP: I/O Device Allocation -- TIOT Construction Routine

This routine calculates the main storage required for the TIOT, builds the TIOT, and processes requests for direct-access space.

- Entry: IEFWC000
- Exits: To IEFXJIMP, IEFWDIMP
- Tables/Work Areas: JCT, SCT, LCT, SIOT, VOLT, AWT, TIOT
- Control Section: IEFWC000

IEFWDFAK: I/O Device Allocation -- Linkage Module

This routine passes control to the external action routine.

- Entry: IEFWD000
- Exit: To IEFWD000
- Control Section: IEFWD000

IEFWD000: I/O Device Allocation -- External Action Routine

This routine causes the correct volumes for the step to be mounted on the appropriate units.

- Entry: IEFWD000, IEFWDMSG
- Exits: To IEFXT000, IEFW41SD, IEFXK000
- Tables/Work Areas: SCT, LCT, TIOT, UCB
- Control Section: IEFWD000, IEFWDMSG

IEFWD001: I/O Device Allocation -- External Action Messages

This routine contains a directory and the messages used in the external action routine.

- Entry: IEFWD001
- Attributes: Non-executable
- Control Section: IEFWD001

IEFWSTRT: I/O Device Allocation -- Message Module

This routine contains the message issued to the operator when a job is started and the messages issued to the operator when a job is terminated due to ABEND, condition codes, or JCL errors.

- Entry: IEFWSTRT
- Attributes: Non-executable
- Control Section: IEFWSTRT

IEFWSWIN: I/O Device Allocation -- Linkage Module

This routine passes control to the decision allocation routine.

- Entry: IEFWSWIT
- Exit: To IEFX5000
- Control Section: IEFWSWIT

IEFWTERM: Termination -- Message Module

This routine contains the message issued to the operator when a job is terminated normally, or when it is terminated because of a JCL error found in the interpreter or initiator.

- Entry: IEFWTERM
- Attributes: Non-executable
- Control Section: IEFWTERM

IEFXAMSG: I/O Device Allocation -- Message Module

This routine contains the messages issued by the allocation control routine.

- Entry: IEFXAMSG
- Attributes: Non-executable
- Control Section: IEFXAMSG

IEFXCSSS: I/O Device Allocation -- Allocation Control Routine

This routine calculates table space requirements and obtains the main storage for the tables used or built during allocation.

- Entry: IEFXA
- Exits: To IEFXJ, IEFWA, IEFWC
- Tables/Work Areas: JCT, SCT, ICT, UCB, SIOT, VOLT, AWT
- Control Section: IEFXA

IEFXH000: I/O Device Allocation -- Separation Strikeout Routine

This routine strikes from AWT entries, the bits corresponding to devices that would violate separation or affinity requests.

- Entry: IEFXH000
- Exit: Return to caller
- Tables/Work Areas: ICT, AWT, AVT, UCB
- Control Section: IEFXH000

IEFXJFAK: I/O Device Allocation -- Linkage Module

This routine passes control to the allocation recovery routine.

- Entry: IEFXJ000
- Exit: To IEFXJIMP
- Control Section: IEFXJ000

IEFXJIMP: I/O Device Allocation -- Allocation Recovery Routine

This routine informs the operator of the allocation recovery options available, and passes control to the proper routine to comply with his request.

- Entry: IEFXJ000, IEFV15XL, IEFXJX5A
- Exits: To IEFXCSSS, IEFSD095, IEFW41SD
- Tables/Work Areas: ICT, AWT, JCT, CVT, UCB, SCT, SIOT
- Control Section: IEFXJ000

IEFXJMSG: I/O Device Allocation -- Allocation Recovery Messages

This routine contains the messages used by the allocation recovery routine.

- Entry: MSRCV, MSSYS, MSCFF
- Attributes: Non-executable
- Control Section: IEFXJMSG

IEFXKIMP: I/O Device Allocation --
Non-Recovery Error Routine

This routine cancels the step when a lack of available devices has been discovered after the TIOT is constructed.

- Entry: IEFXK000
- Exit: To IEFW41SD
- Tables/Work Areas: LCT, SCT, UCB, TIOT
- Control Section: IEFXK000

IEFXKMSG: I/O Device Allocation --
Non-Recovery Error Routine Messages

This routine contains the messages used by the non-recovery error routine.

- Entry: IEFXKMSG
- Attributes: Non-executable
- Control Section: IEFXKMSG

IEFXT00D: I/O Device Allocation -- Space
Request Routine

This routine obtains space on direct-access devices for requesting data sets.

- Entry: IEFXT000
- Exits: To IEFW41SD, IEFXK000, IEFWD000
- Tables/Work Areas: LCT, TIOT, UCB, JCT, SIOT, JFCB, PDQ
- Control Section: XTTP00, IEFXT000

IEFXT002: I/O Device Allocation -- VARY
Command Interface TIOT Compression Routine

This routine reduces the TIOT to its final size and provides an interface with the VARY command.

- Entry: IEFXT002, XTTRDJ, XTTEB3, XTTEA1, XTTEA0
- Exits: to IEFXKIMP, IEFXT003, IEF41FAK
- Tables/Work Areas: LCT, TIOT, UCB, JCT, SIOT, JFCB
- Control Section: IEFXT002

IEFXT003: I/O Device Allocation -- DADSM
Error Recovery Routine

This routine determines what action should be taken when the request for space on a particular volume fails.

- Entry: IEFXT003, XUUH06, XUUB00
- Exits: To IEFXT00D, IEFXT002
- Tables/Work Areas: LCT, TIOT, UCB, JCT, SIOT, JFCB
- Control Section: IEFXT003

IEFXVMSG: I/O Device Allocation --
Automatic Volume Recognition Messages

This routine contains the messages used by the automatic volume recognition (AVR) routine.

- Entry: IEFXVMSG
- Attributes: Non-executable
- Control Section: IEFXVMSG

IEFXVNSL: I/O Device Allocation --
Automatic Volume Recognition --
Non-Standard Label Routine

This routine processes non-standard labels for the AVR routine.

- Entry: IEFXVNSL
- Exit: Return to caller
- Control Section: IEFXVNSL

IEFXV001: I/O Device Allocation --
Automatic Volume Recognition Routine

This routine finds and allocates volumes pre-mounted by the operator.

- Entry: IEFXV001
- Exits: IEFWC000, IEFX5000, IEFXJ000
- Tables/Work Areas: JCT, SCT, AWT, AVT, VOLT, SIOT, LCT, UCB
- Control Section: IEFXV001

IEFXV002: I/O Device Allocation --
Automatic Volume Recognition, Label
Processing

This routine reads the label of a newly mounted volume, extracts the serial number, and places it into the UCB for the corresponding device.

- Entry: IEFXV002
- Exits: To IEFXVNSL via CALL, return to caller.
- Tables/Work Areas: LUT, UCB, CVT, DEB, IOB

- Attributes: Reusable
- Control Section: IEFXV002

IEFX300A: I/O Device Allocation -- Device Strikeout Routine

This routine modifies the primary and secondary bit patterns in AWT entries to complete the allocation to a data set.

- Entry: IEFX3000, X33P42
- Exit: Return to caller
- Tables/Work Areas: AWT, AVT, UCB, LCT
- Control Section: IEFX3000

IEFX5000: I/O Device Allocation -- Decision Allocation Routine

This routine selects devices for data sets with multiple unit possibilities.

- Entry: IEFX5000, XIIB32, X55C86, X55D3G
- Exits: To IEFWC000, IEFXJ000
- Tables/Work Areas: LCT, AWT, AVT, UCB
- Control Section: IEFX5000

IEFYNIMP: Termination -- Step Termination Control Routine

This routine passes control among the modules of the step termination routine and, when required, passes control to the job termination routine.

- Entry: IEFYN
- Exits: To IEFW22SD, IEFYPJB3, IEFVJIMP, IEFZAJB3, IEFPRP
- Tables/Work Areas: JCT, SCT, LCT
- Control Section: IEFYN

IEFYNMSG: Termination -- Step Termination Control Routine Messages

This routine contains the messages required for the step termination control routine.

- Entry: IEFYNMSG, STRMSG01
- Attributes: Non-executable
- Control Section: IEFYNMSG

IEFYPJB3: Termination -- Step Termination Data Set Driver Routine

This routine obtains SIOTs and to pass control to the disposition and unallocation routine.

- Entry: IEFYP
- Exits: To IEFZG, IEFYNIMP
- Tables/Work Areas: ICT, TIOT, UCB, QMPA, SIOT, TCB
- Control Section: IEFYP

IEFYPMSG: Termination -- Step Termination Messages

This routine contains the messages required by the step termination routine.

- Entry: IEFYPMSG, YPPMSG1, YPPMSG2
- Attributes: Non-executable
- Control Section: IEFYPMSG

IEFYSVMS: Termination -- Message Blocking Routine

This routine blocks system messages into SMBs, and places SMBs into the message class queue entry.

- Entry: IEFYS
- Exit: Return to caller
- Tables/Work Areas: LCT, SCT, SMB
- Attributes: Reenterable
- Control Section: IEFYS

IEFYTVMS: Termination -- DSB Processing Routine

This routine places data set blocks in the space reserved for them in the output queue entries.

- Entry: IEFYT
- Exit: Return to caller
- Tables/Work Areas: JCT, SCT, TIOT, SIOT, QMPA, DSCB, LCT, CVT, JFCB
- Attributes: Reenterable
- Control Section: IEFYT

IEFZAJB3: Termination -- Job Termination Control Routine

This routine provides entry to the job termination routine, obtains PDQ blocks, and passes control to the disposition and unallocation routine.

- Entry: IEFZA
- Exits: To IEFZGJ, IEFW31SD
- Tables/Work Areas: LCT, JCT, PDQ, UCB, QMPA
- Control Section: IEFZA

IEFZGJB1: Termination -- Disposition and Deallocation Routine

This routine directs the disposition and deallocation of those data sets that remain to be processed at job termination: passed data sets that were not received, and retained data sets that were not referred to.

- Entry: IEFZGJ, ZP0QM
- Exit: Return to caller
- Tables/Work Areas: JCT, PDQ, JFCB, LCT, QMPA, UCB
- Control Section: IEFZGJ

IEFZGMSG: Termination -- Disposition and Deallocation Messages

This routine contains the messages required for the disposition and deallocation routine (IEFZGJB1).

- Entry: IEFZGMSG
- Attributes: Non-executable
- Control Section: IEFZGMSG

IEFZGST1: Termination -- Disposition and Deallocation Routine

This routine directs the disposition of data sets as specified in the DISP field of the DD statement, and makes the associated units available for allocation to other data sets.

- Entry: IEFZG, ZP0QMGR1
- Exit: Return to caller

- Tables/Work Areas: LCT, PDQ, SIOT, TIOT, UCB, JFCB, QMPA

- Control Section: IEFZG

IEFZHMSG: Termination -- VARY Command Interface and Disposition and Deallocation Message Routine

This routine prepares messages to the programmer and to the operator for the disposition and allocation routines. It also provides an interface with the VARY command.

- Entry: IEFZH, ZG0E60, ZK0D1, ZK0E1, XPS631
- Exit: Return to caller
- Tables/Work Areas: LCT, QMPA, SMB
- Control Section: IEFZH

IEF078SD: System Output Writer -- Linkage Module

This routine transfers control to module IEFSD078.

- Entry: IEFSD078
- Exit: To IEFSD078
- Attributes: Reenterable

IEF079SD: System Output Writer -- Linkage Module

This routine transfers control to IEFSD079.

- Entry: IEFSD079
- Exit: To IEFSD079
- Attributes: Reenterable

IEF082SD: System Output Writer -- Linkage Module

This routine passes control to the system output writer main processing routine.

- Entry: IEFSD082
- Exit: To IEFSD082
- Control Section: IEFSD082

IEF083SD: System Output Writer -- Linkage Module

This routine passes control to the system output writer command processing routine.

- Entry: IEFSD083
- Exit: IEFSD083
- Control Section: IEFSD083

IEF300SD: System Restart -- Linkage Module

This routine provides a linkage to the system restart initialization routine.

- Entry: IEFSD300
- Exits: To IEFSD300, IEFSD055
- Attributes: Reenterable

IEF304SD: System Restart -- Linkage Module

This routine provides a linkage to the system restart scratch data sets routine.

- Entry: IEFSD304
- Exits: To IEFSD304, IEFSD055
- Attributes: Reenterable
- Control Section: IEFSD304

IEF41FAK: I/O Device Allocation -- Linkage Module

This routine provides a linkage to the allocation exit routine during step flush.

- Entry: IEFW41SD, IEFW1FAK, IEFW2FAK
- Exit: To IEFW41SD
- Attributes: Read-only, reenterable
- Control Section: IEFW41SD

IGC0103D: SVC -- Master Command EXCP Routine

This routine processes the MOUNT Command.

- Entry: IGC0103D
- Attributes: Reenterable, transient
- Control Section: IGC0103D.
- Page Reference:

IGF2603D: SVC 34 -- Machine Status Control Routine

This routine is available only for the model 85. It processes the status parameter of the MODE command.

- Entry: IGF2603D
- Exit: IGF2703D
- Tables/Work Areas: CVT, XSA
- Attributes: Reenterable, read-only, self-relocating
- Control Section: IGF2603D

IGF2703D: SVC34 - Machine Status Control Routine

This routine is available only for the model 85. It processes all parameters of the MODE command but the status parameter.

- Entry: IGF2703D
- Exit: Return to issuer of SVC 34
- Tables/Work Areas: CVT, XSA
- Attributes: Reenterable, read-only, self-relocating
- Control Sections: IGF2703D

Appendix C: Flowcharts

This appendix includes the MFT flowcharts that are different from MVT. For the flowcharts on allocation, termination, and sys-

tem restart, see IBM System/360 Operating System: MVT Job Management, Program Logic Manual, Form Y28-6660.

Chart 01. Task Dispatcher (Without Time Slicing)

Note - 'Old' Is the TCB Address of the Task Currently in Control. 'New' Is the TCB Address of the Task to be Given Control.

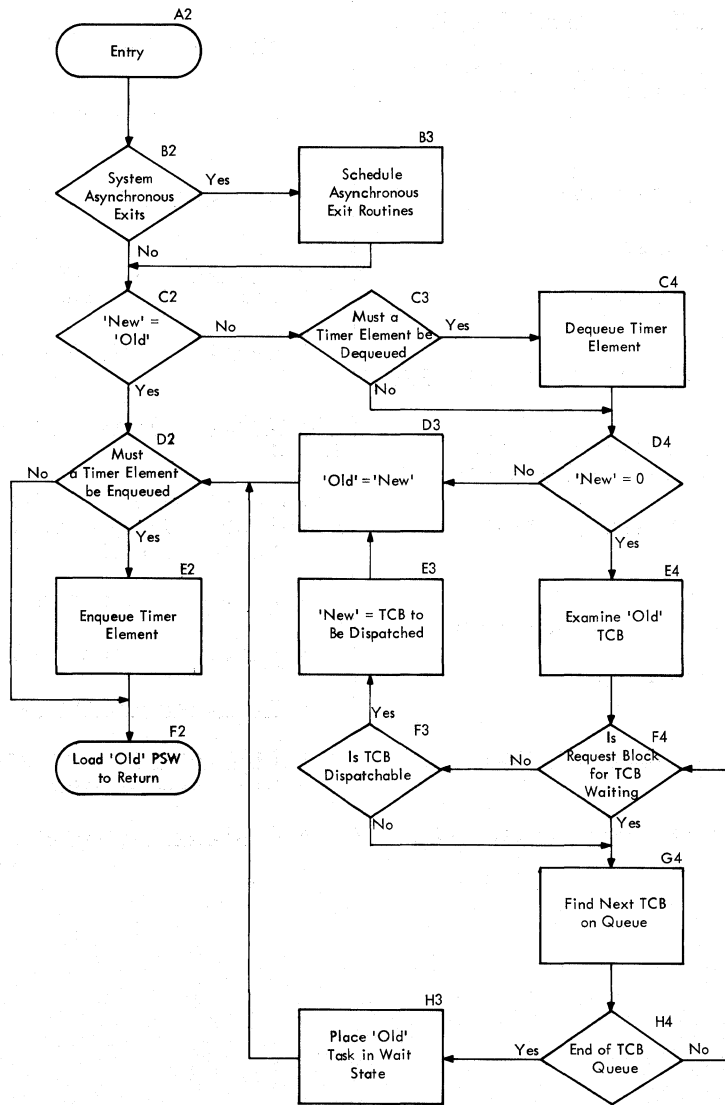


Chart 02. Task Dispatcher (With Time Slicing)

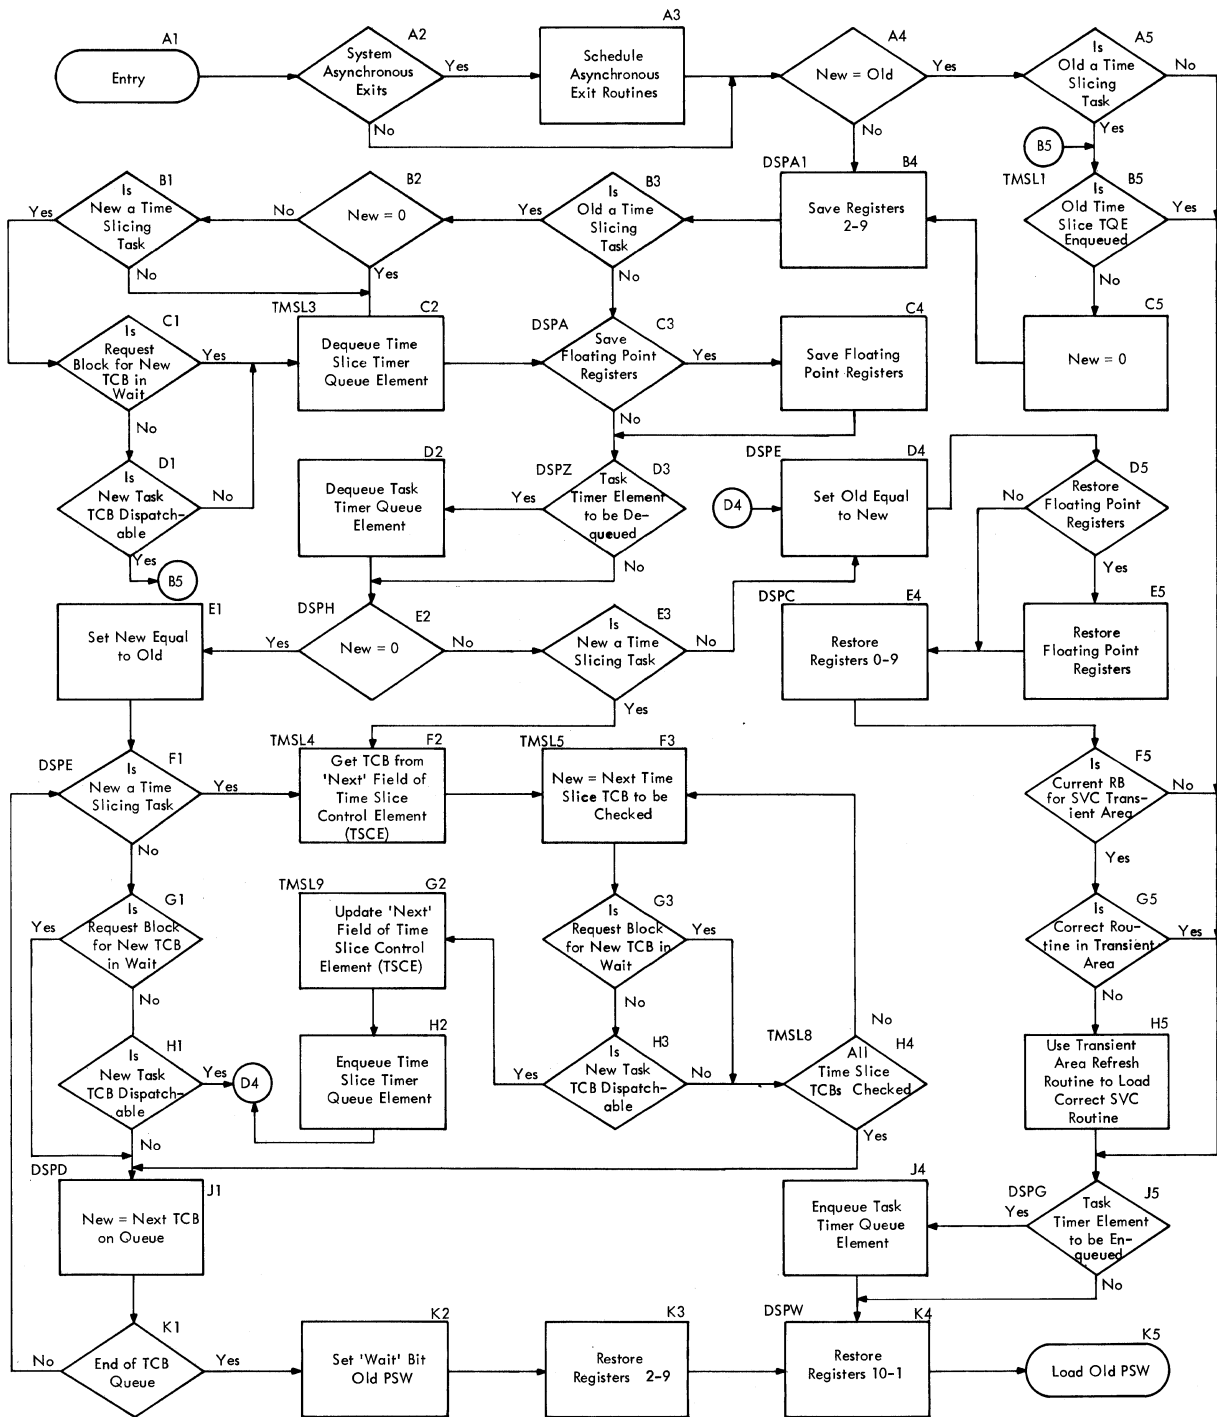


Chart 03. ABEND and DAR Control Flow (Part 1 of 2)

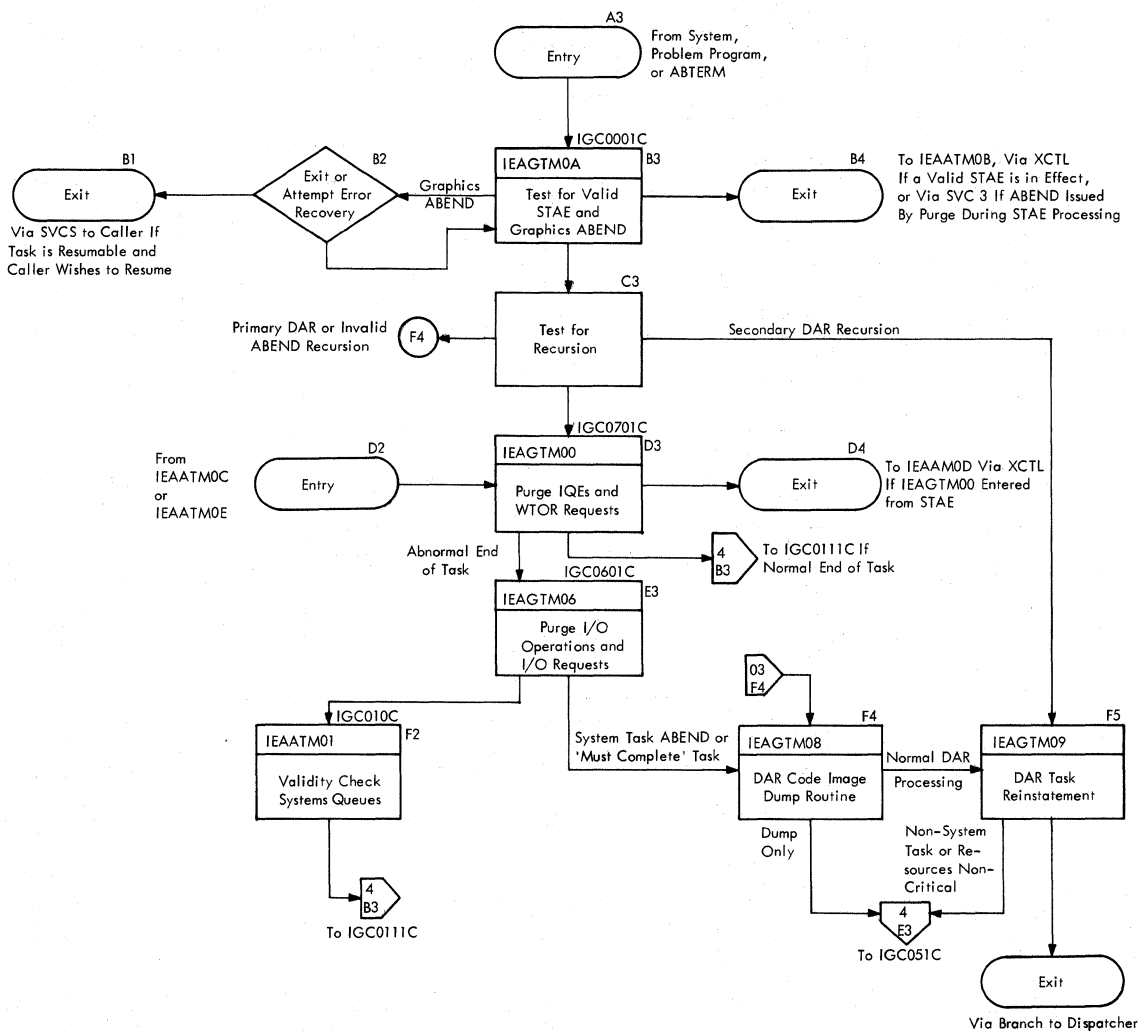


Chart 04. ABEND and DAR Control Flow (Part 2 of 2)

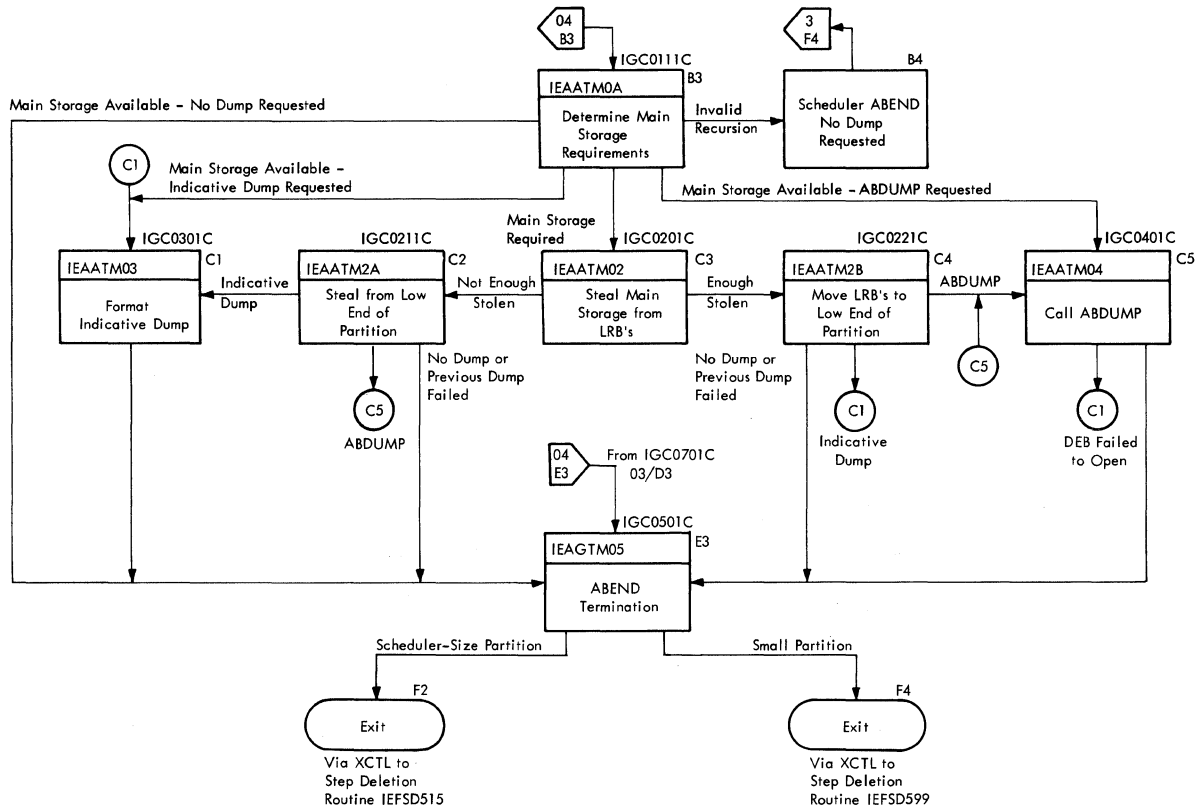


Chart 05. Small Partition Routine (Part 1 of 4)

Note -
 At Entry,
 Small Partition
 Has Zero
 Protection
 in TCB, PSW
 and Hardware.
 Also, PSW is
 Supervisor
 State.

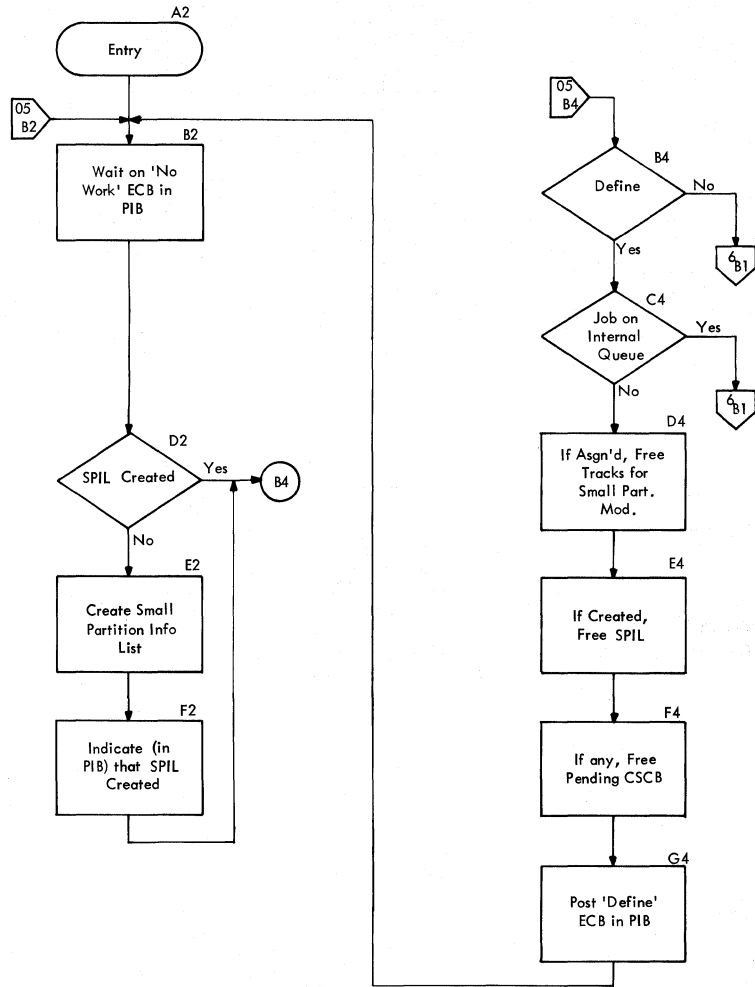


Chart 06. Small Partition Routine (Part 2 of 4)

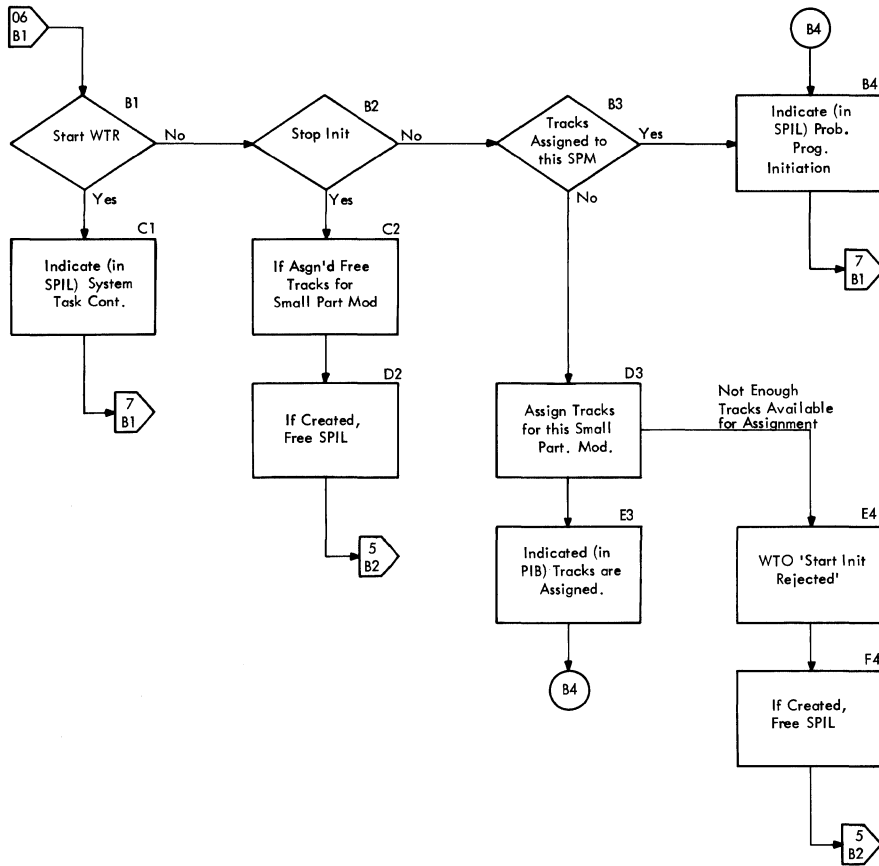


Chart 07. Small Partition Routine (Part 3 of 4)

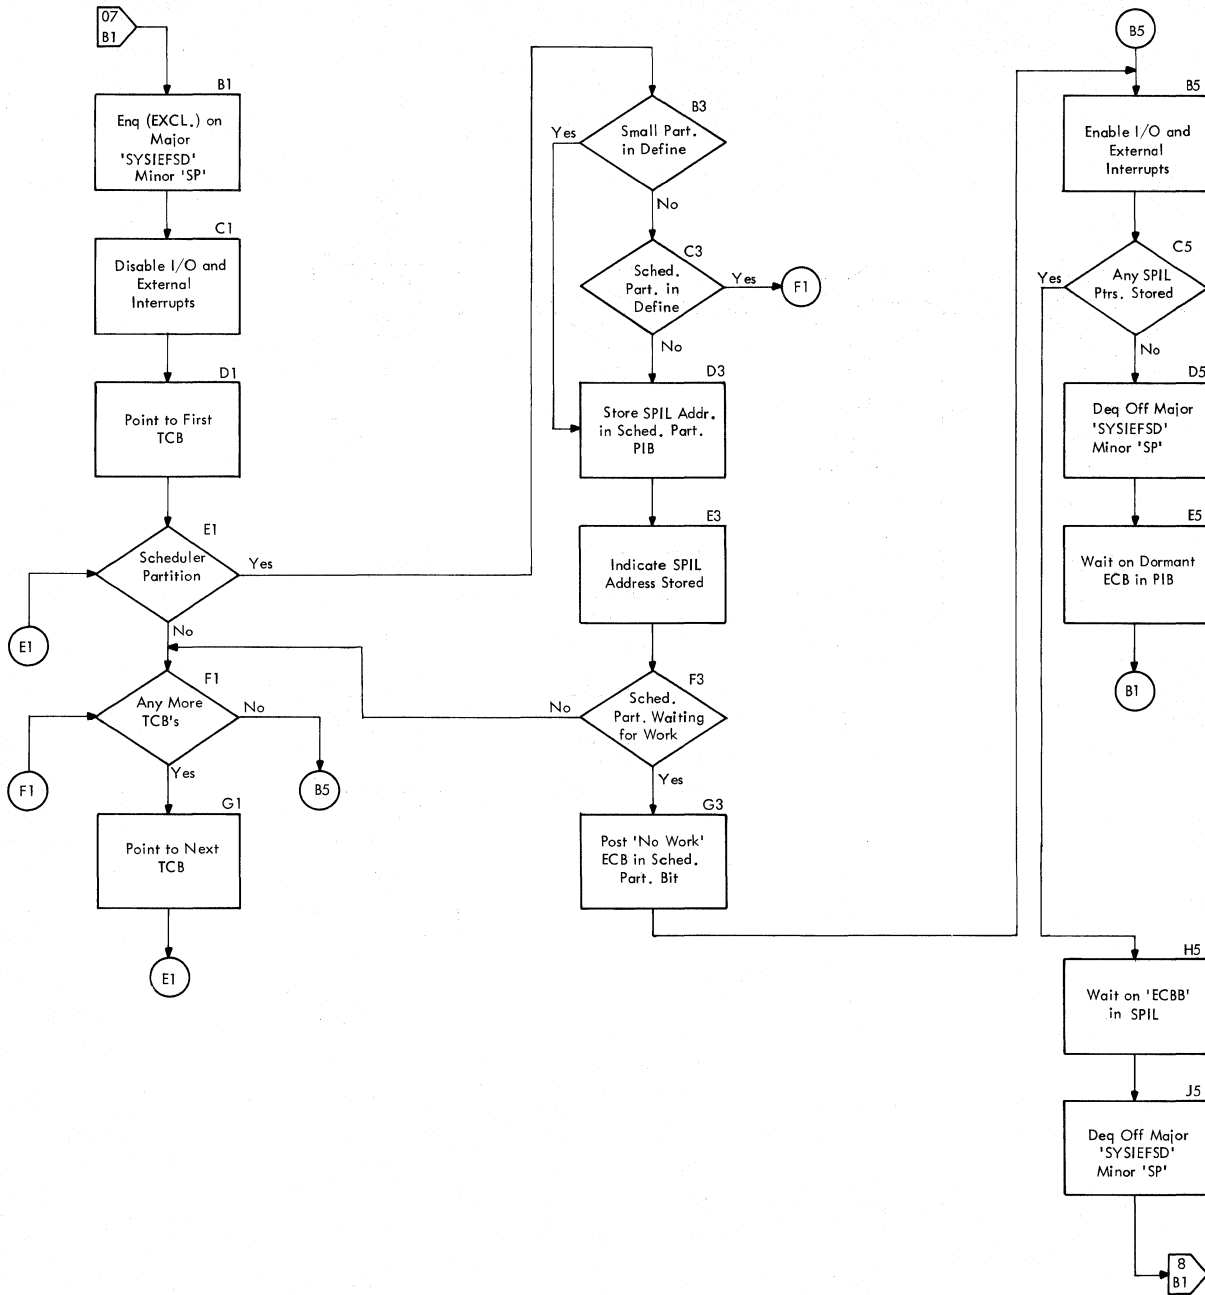


Chart 08. Small Partition Routine (Part 4 of 4)

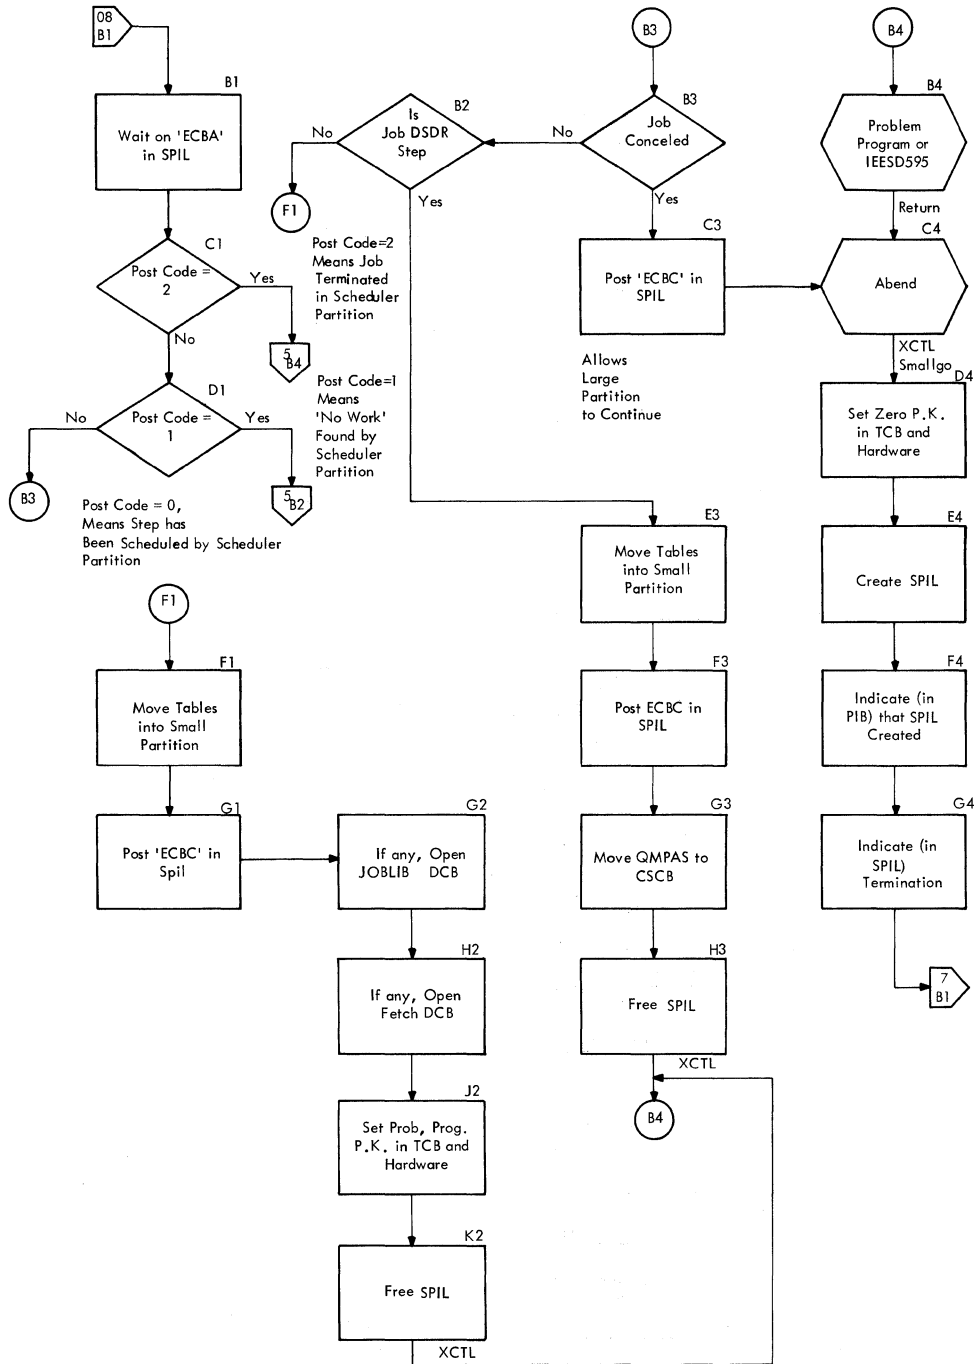


Chart 09. Master Scheduler Task

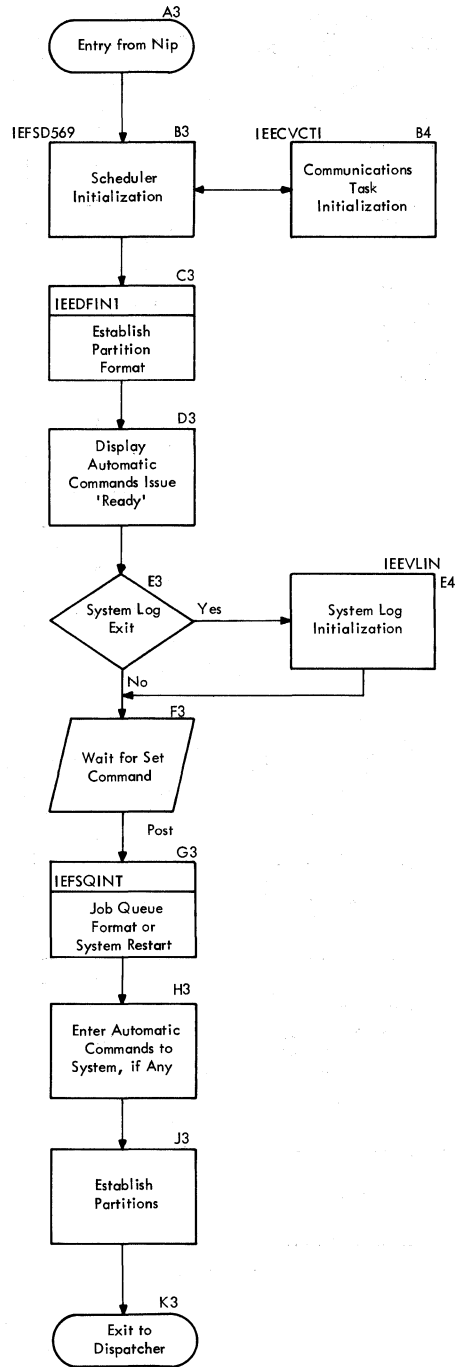


Chart 10. Queue Search

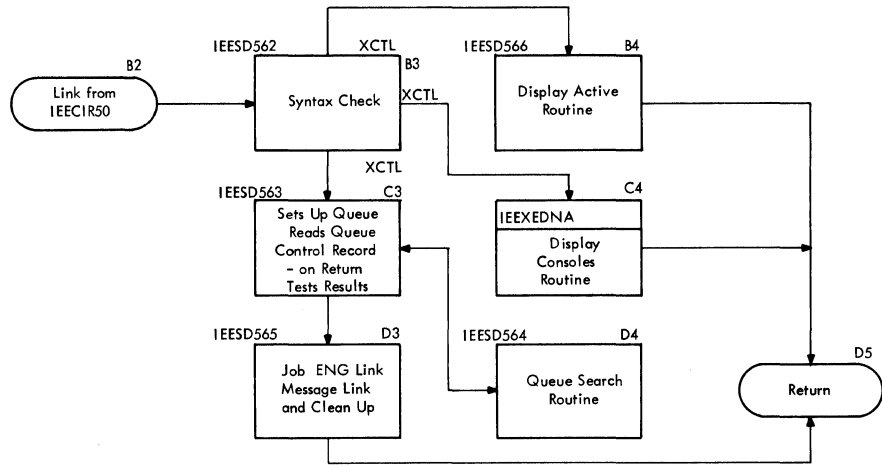


Chart 11. Queue Manager Table Breakup Routine

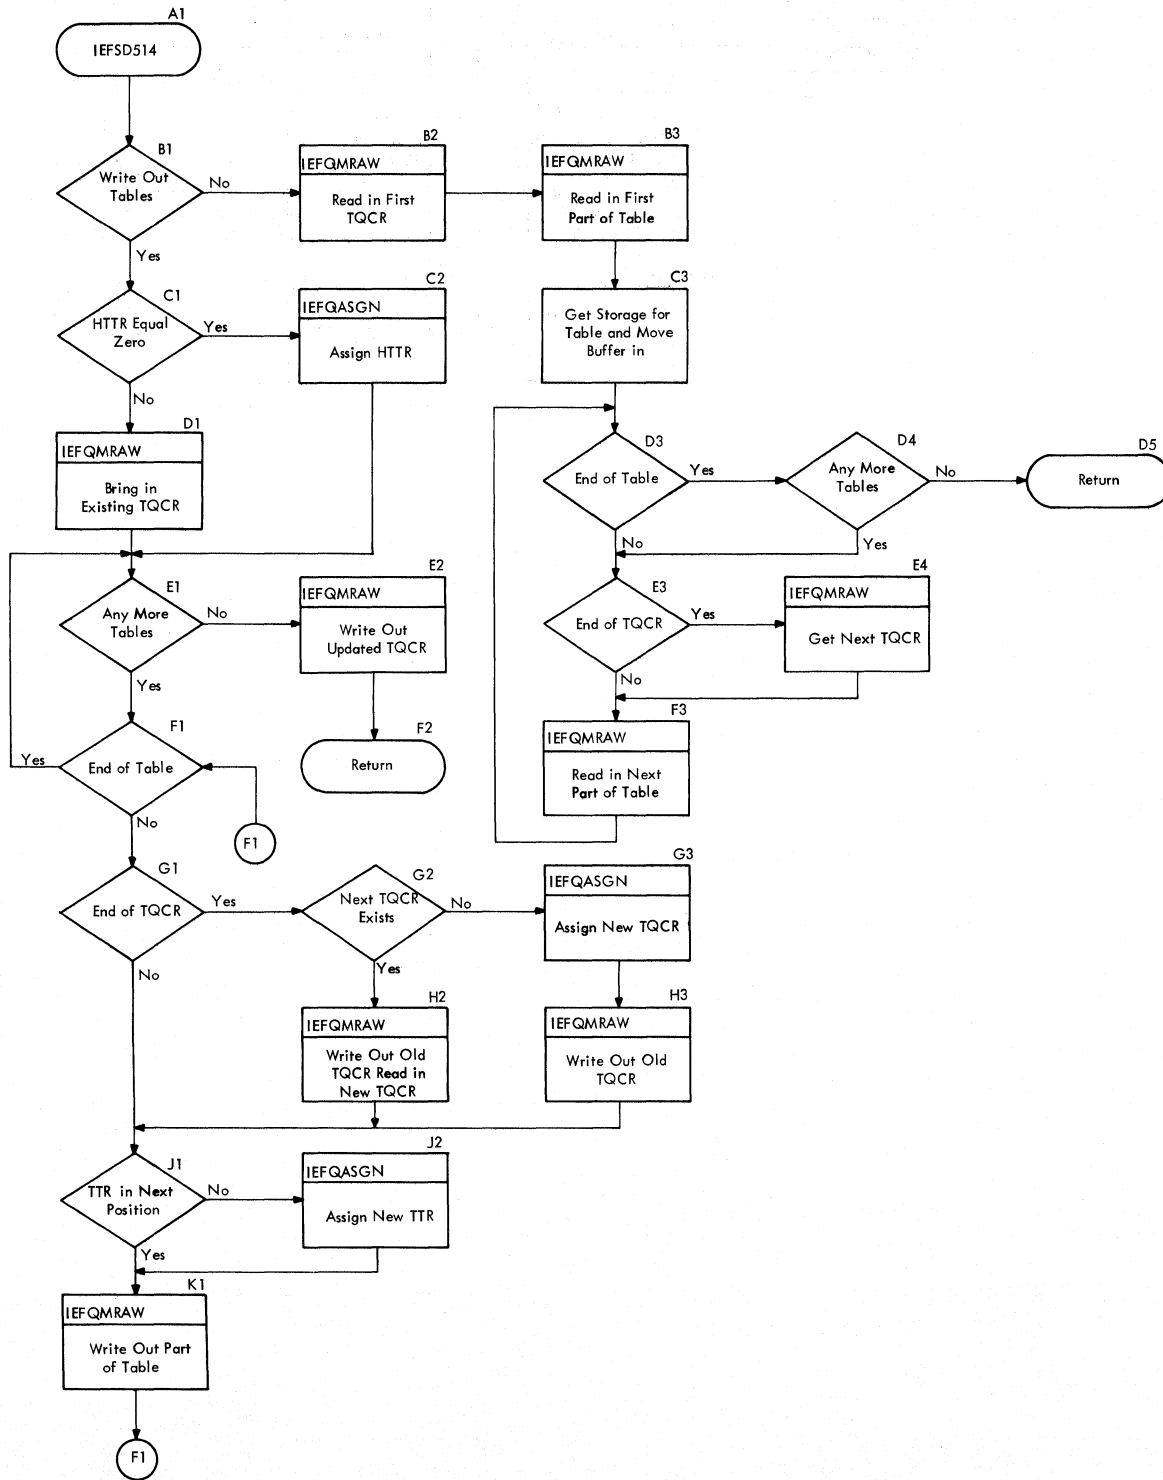


Chart 12. Master Scheduler Resident Command Processor

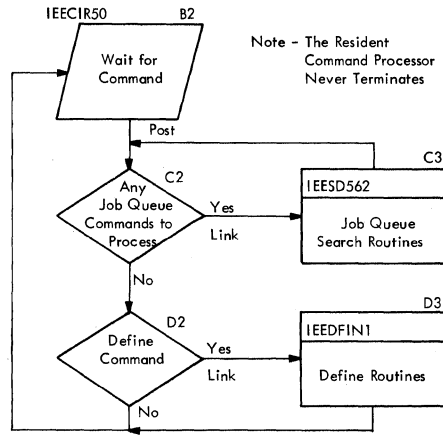


Chart 13. SVC 34 Command Processing (Part 1 of 3)

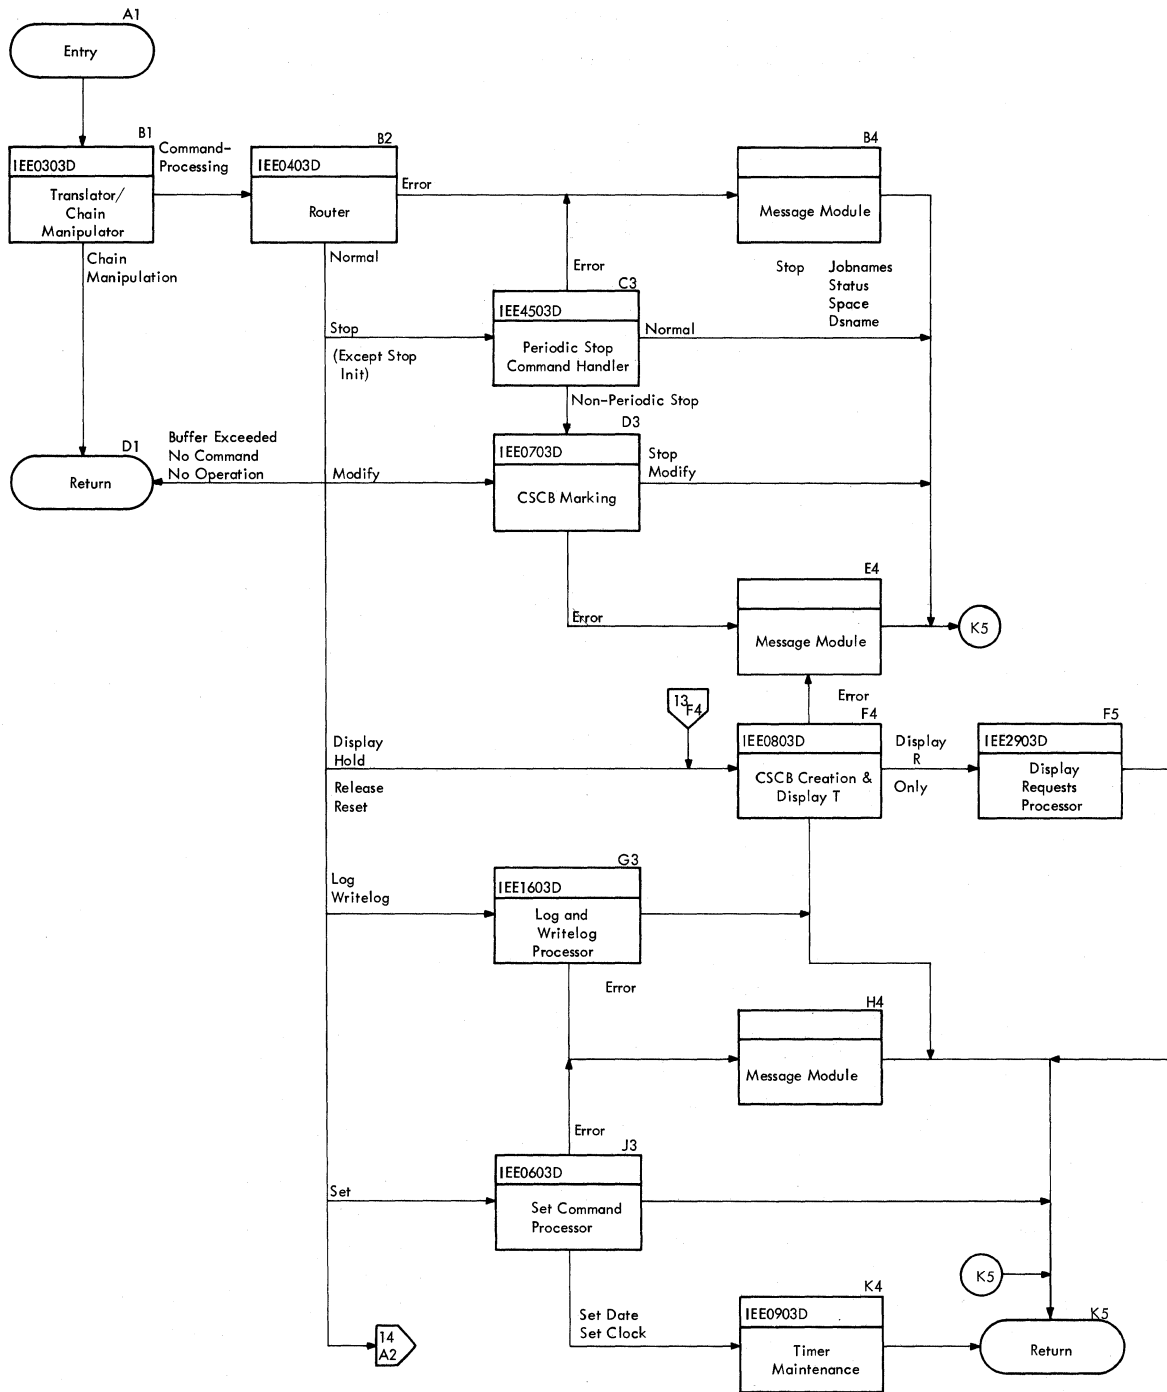


Chart 14. SVC 34 Command Processing (Part 2 of 3)

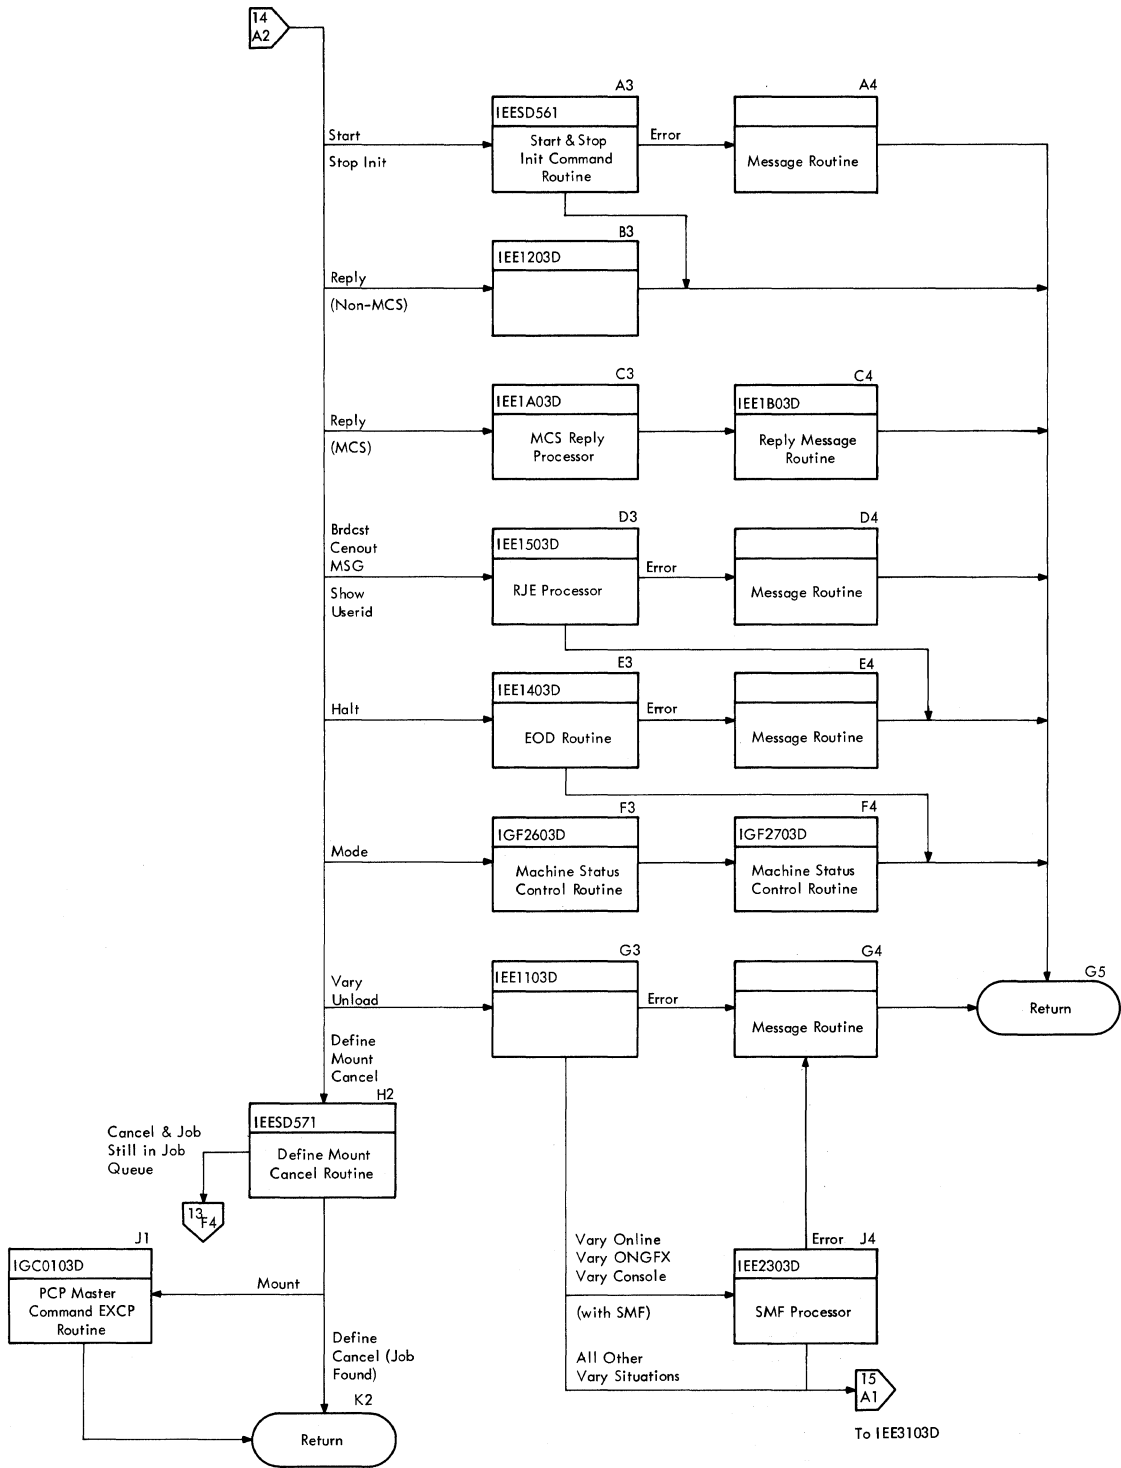


Chart 15. SVC 34 Command Processing (Part 3 of 3)

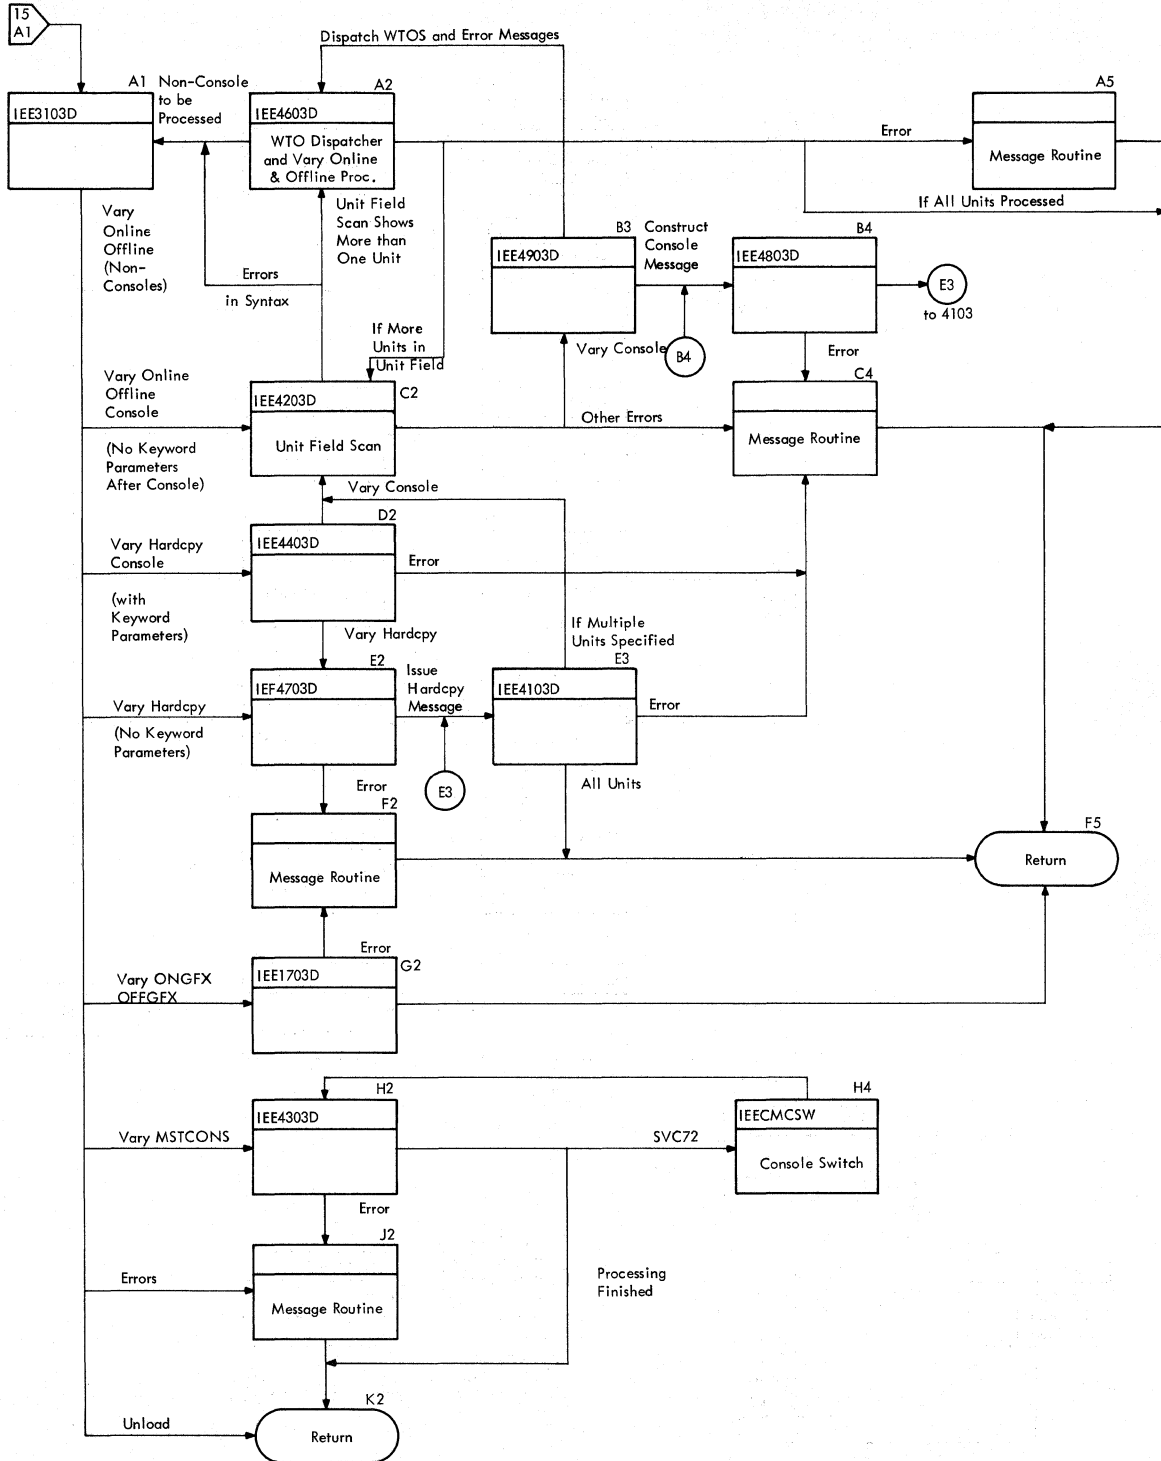


Chart 16. Communications Task

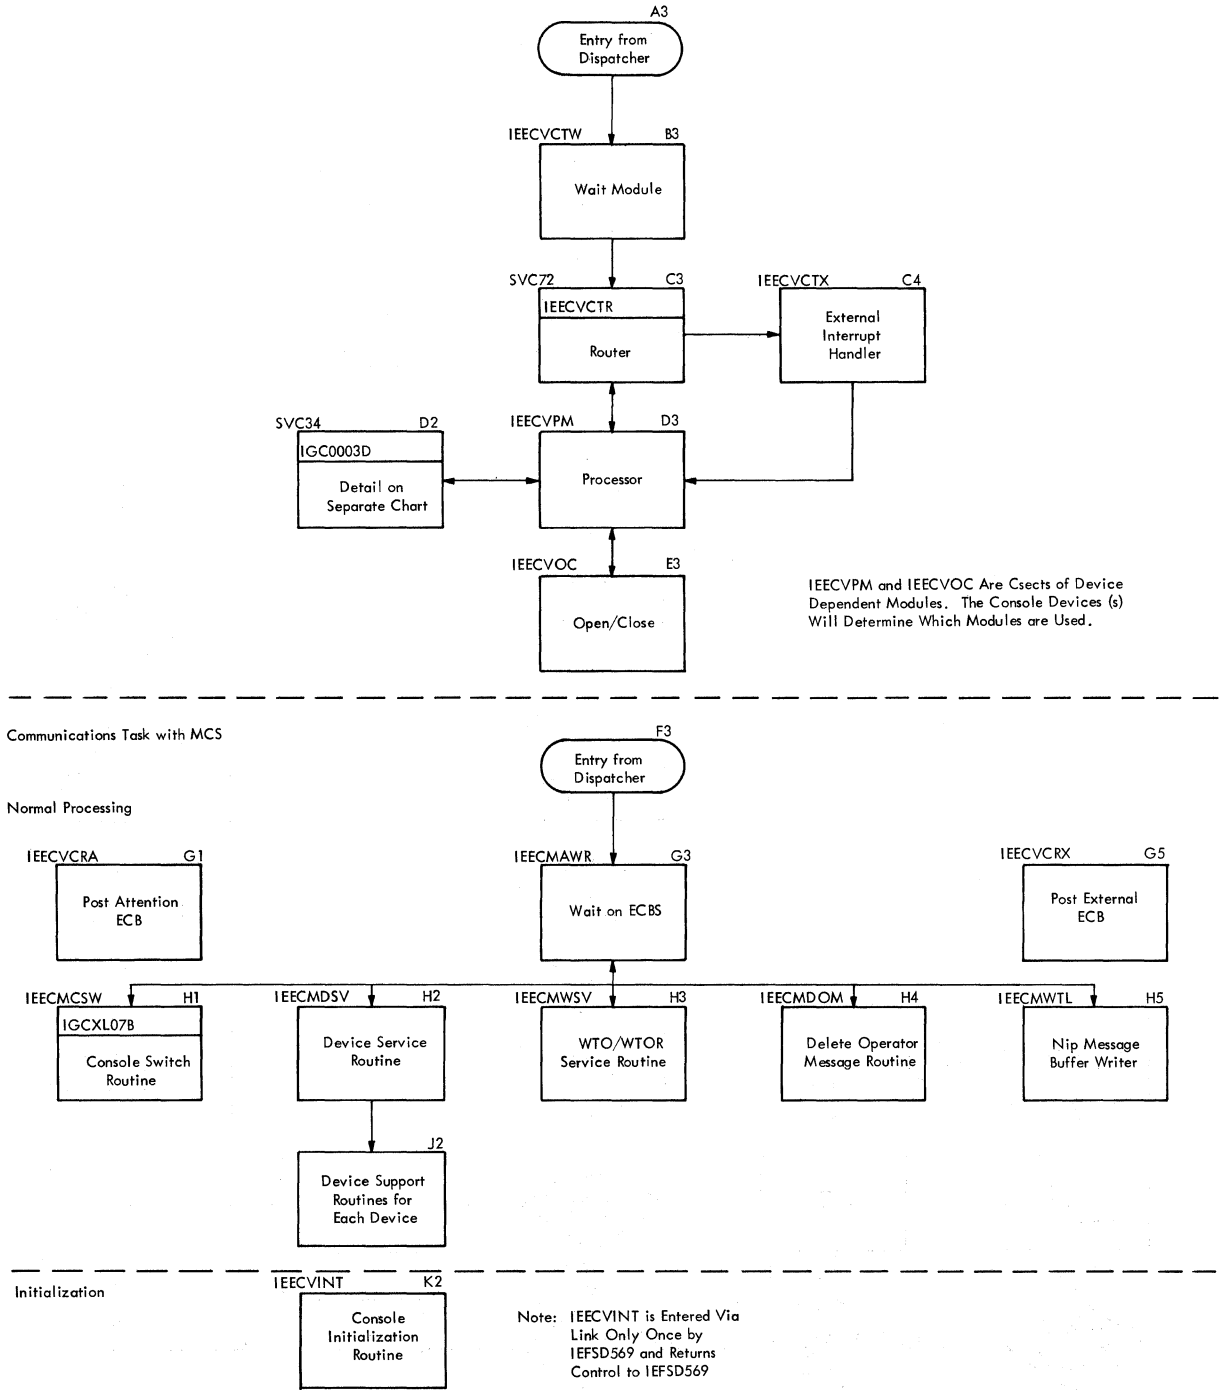


Chart 17. IEFSD518 - Partition Recovery Routine

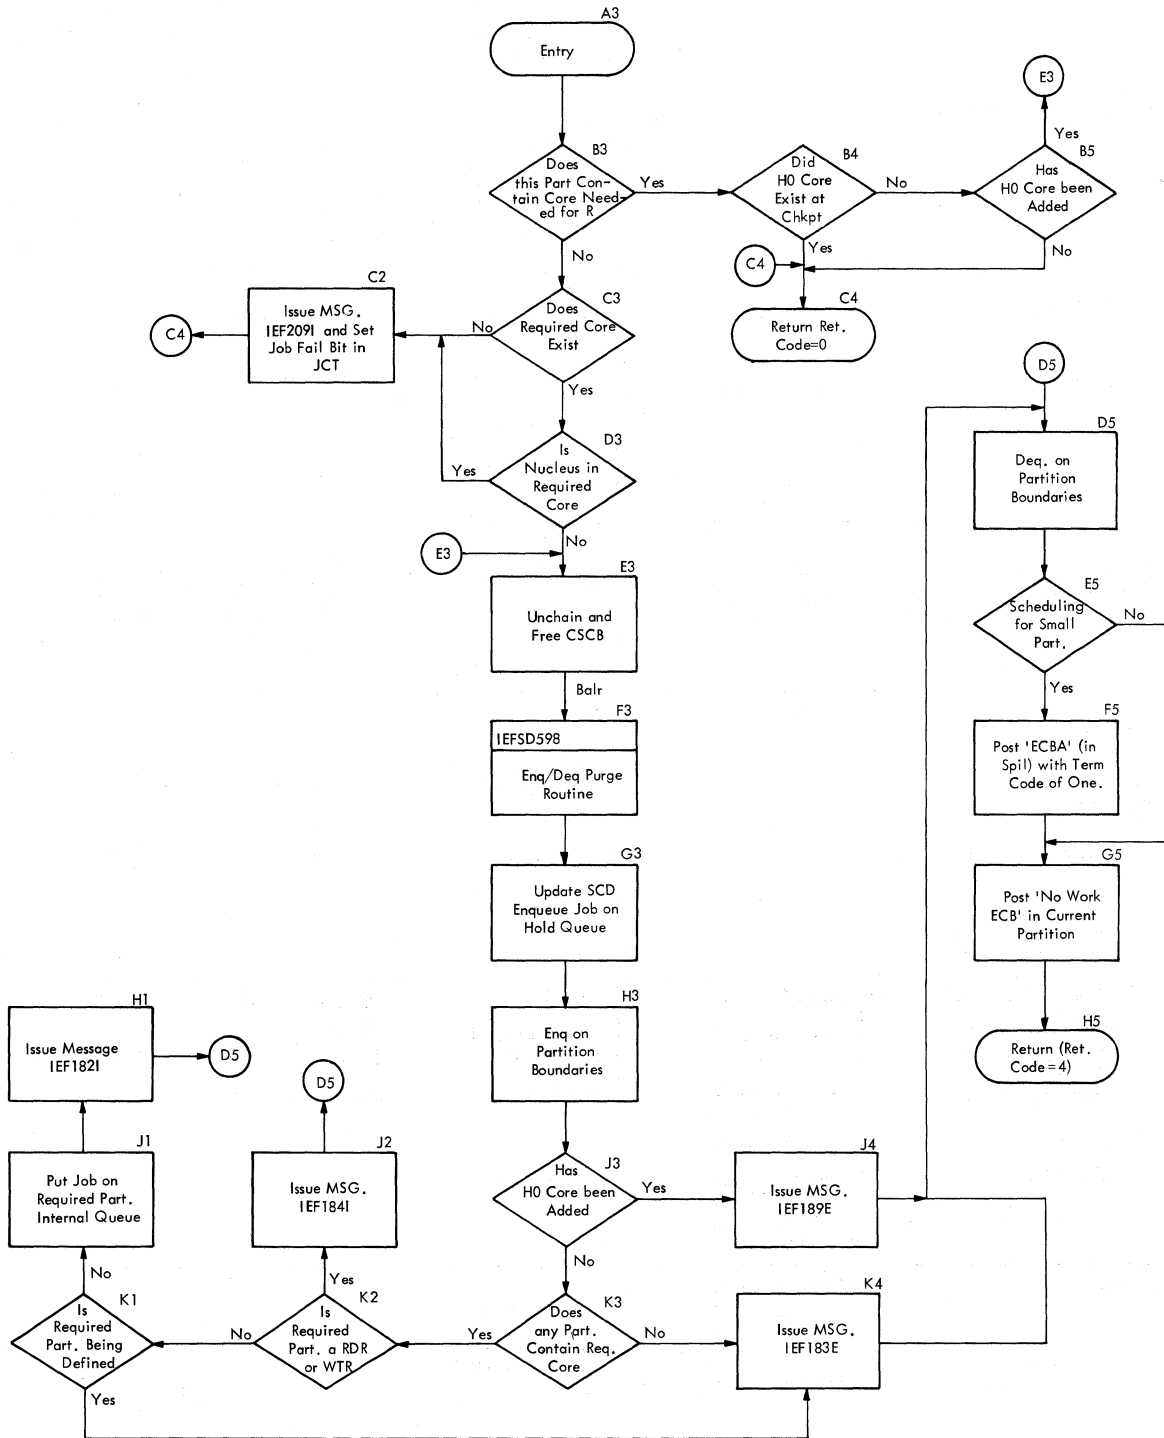


Chart 19. Job Selection Routine (Sheet 1 of 5)

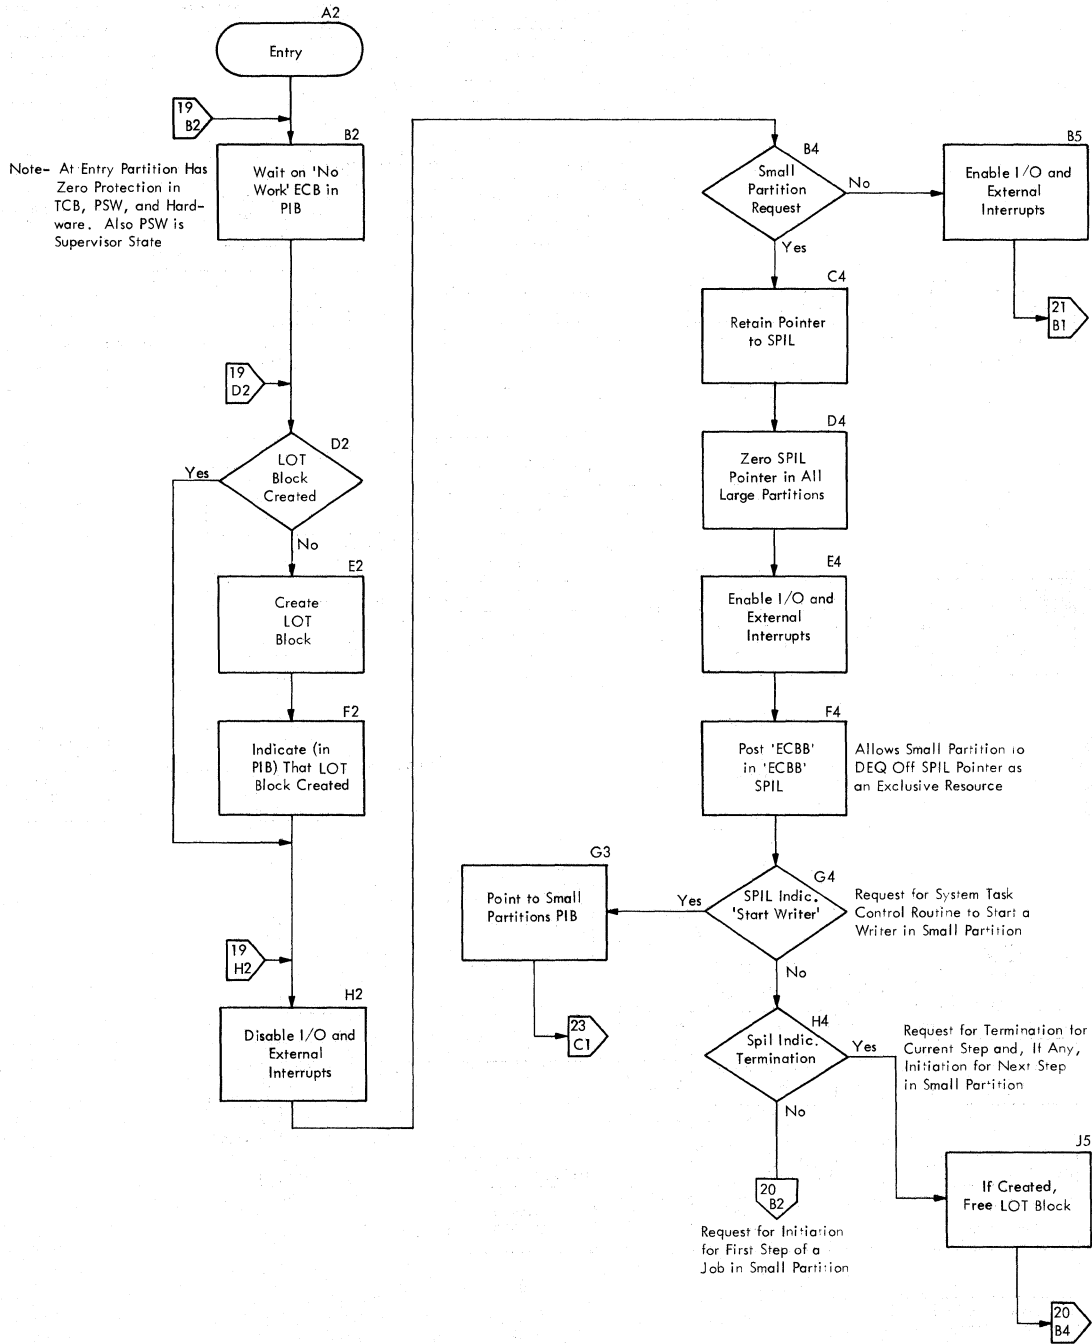


Chart 21. Job Selection Routine (Sheet 3 of 5)

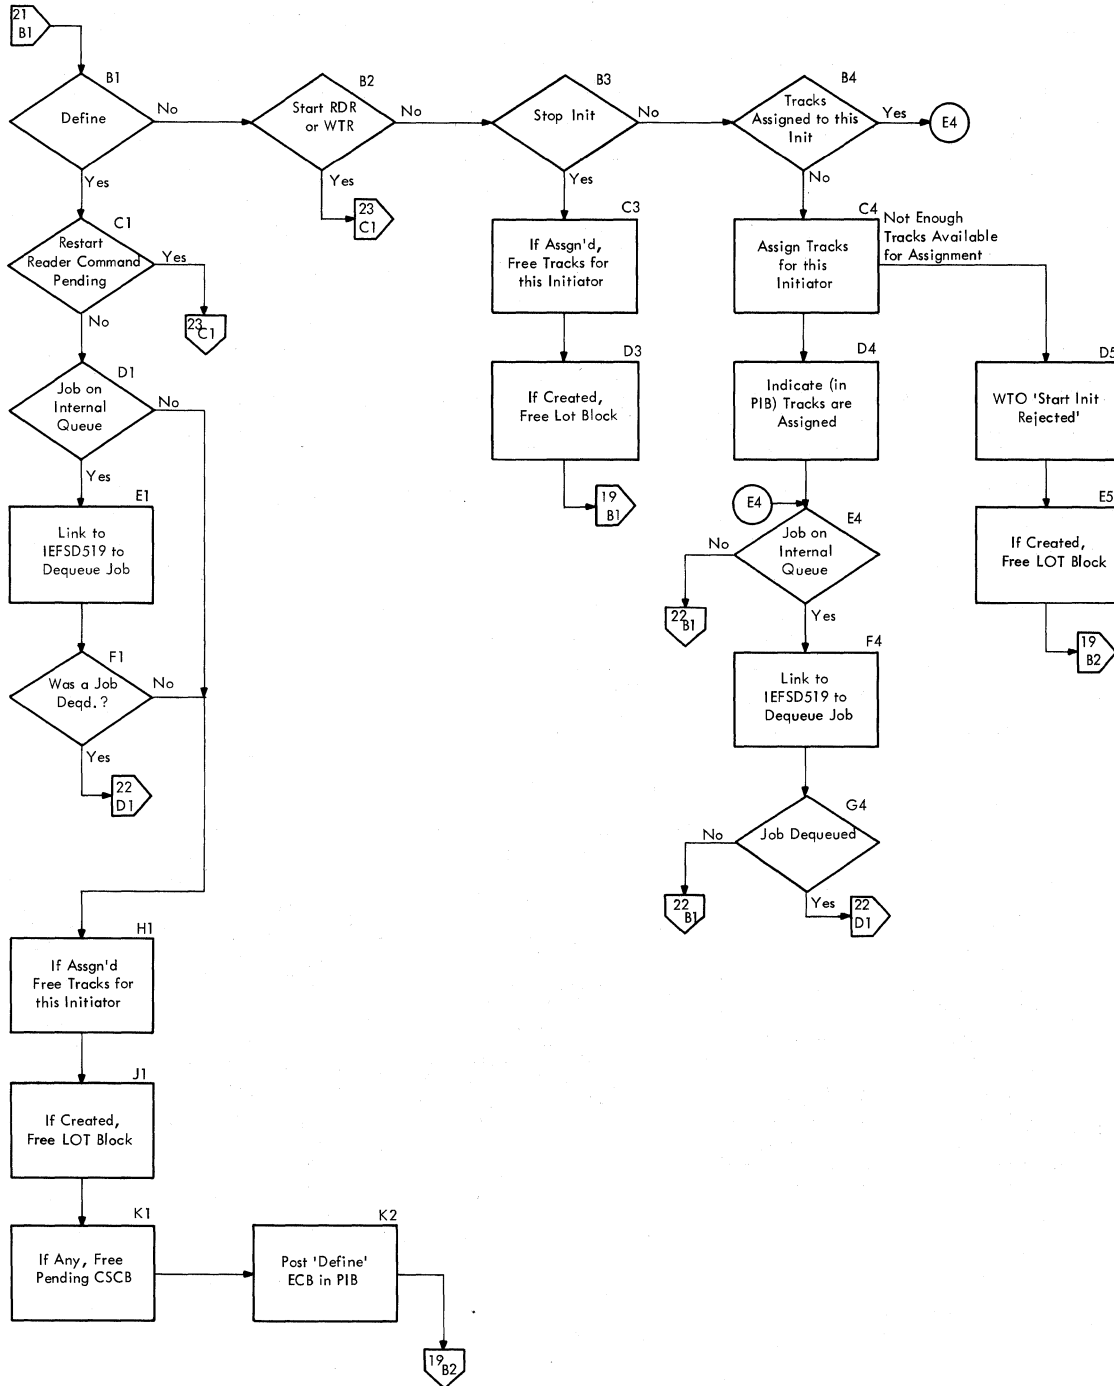


Chart 22. Job Selection Routine (Sheet 4 of 5)

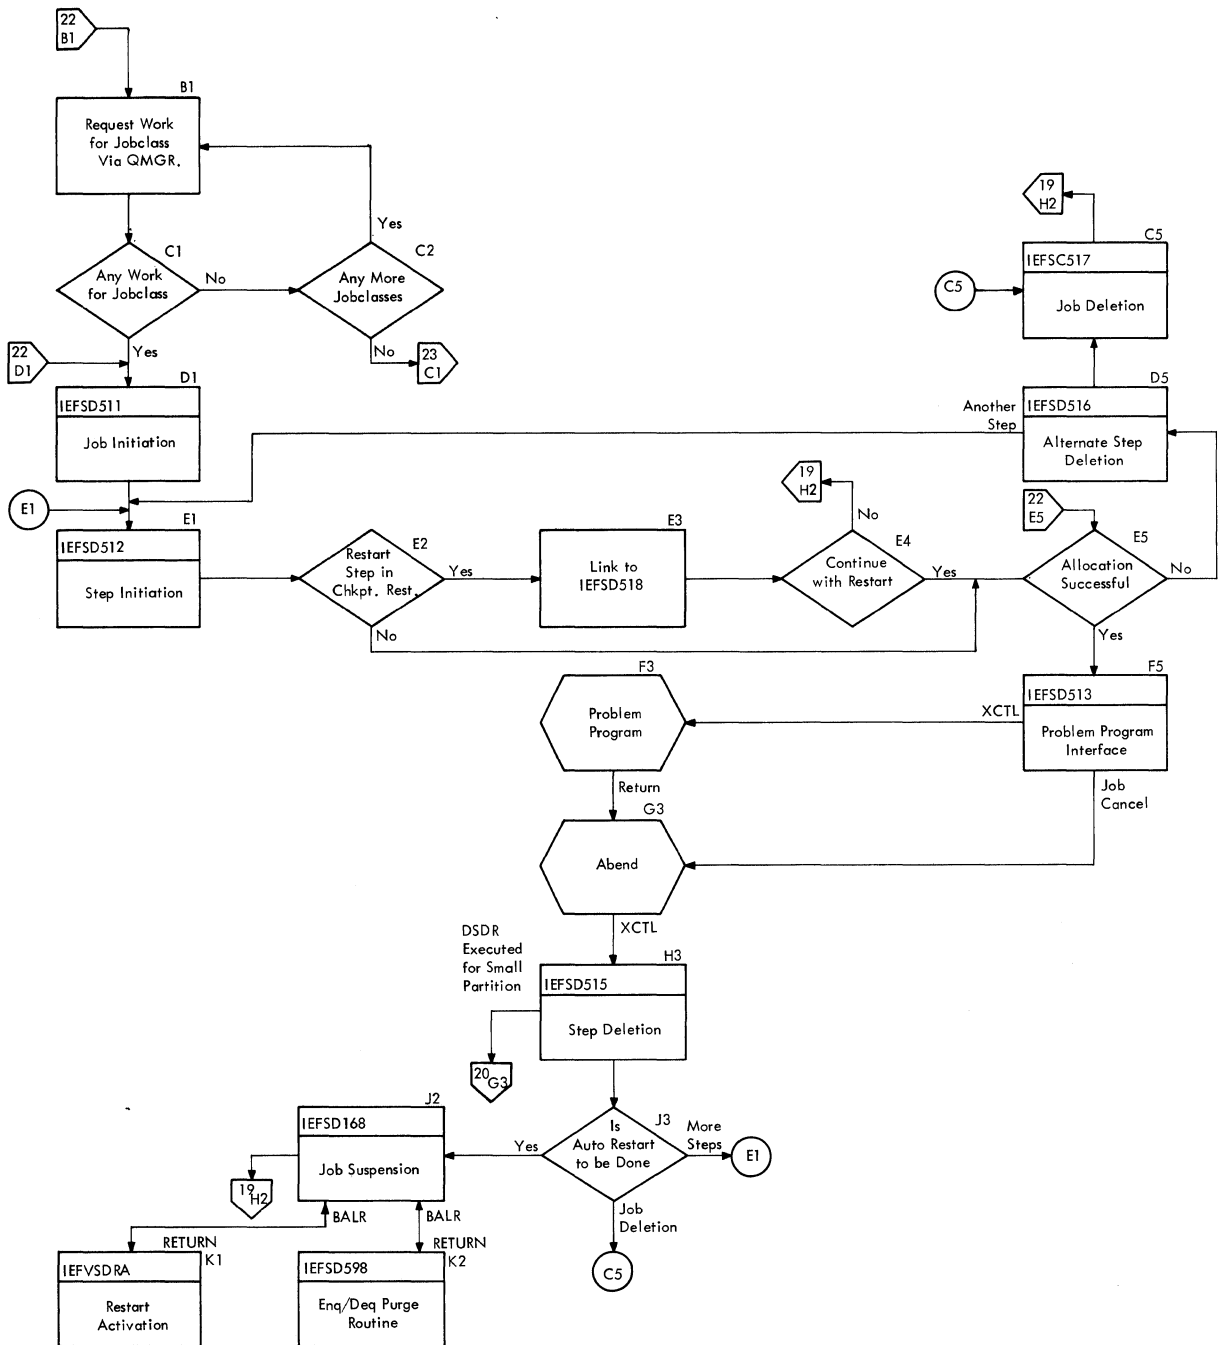
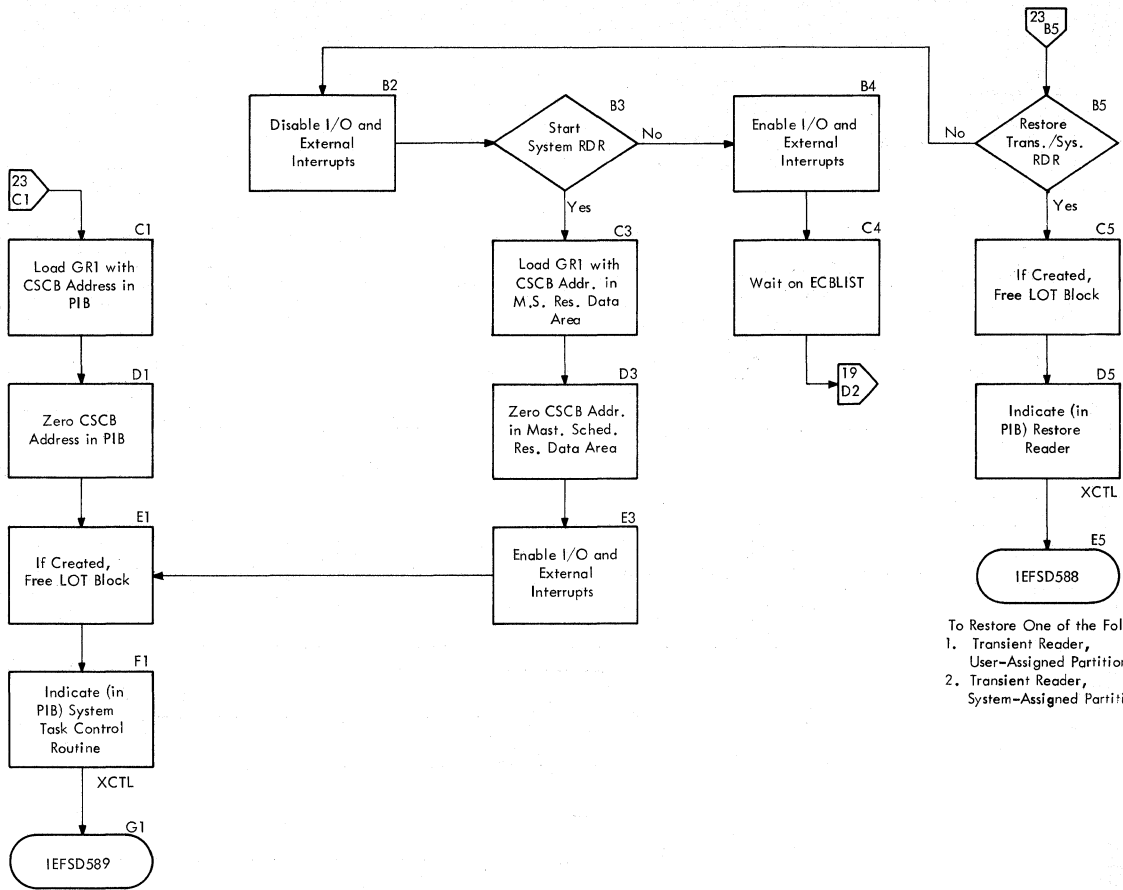


Chart 23. Job Selection Routine (Sheet 5 of 5)



To Restore One of the Following
 1. Transient Reader, User-Assigned Partition
 2. Transient Reader, System-Assigned Partition

Allows System Task Control Routine to Initially Start One of the Following

1. Resident Reader
2. Transient Reader, User-Assigned Partition
3. Transient Reader, System-Assigned Partition
4. Writer, This Partition
5. Writer, Small Partition

Chart 24. Reader/Interpreter (Sheet 1 of 3)

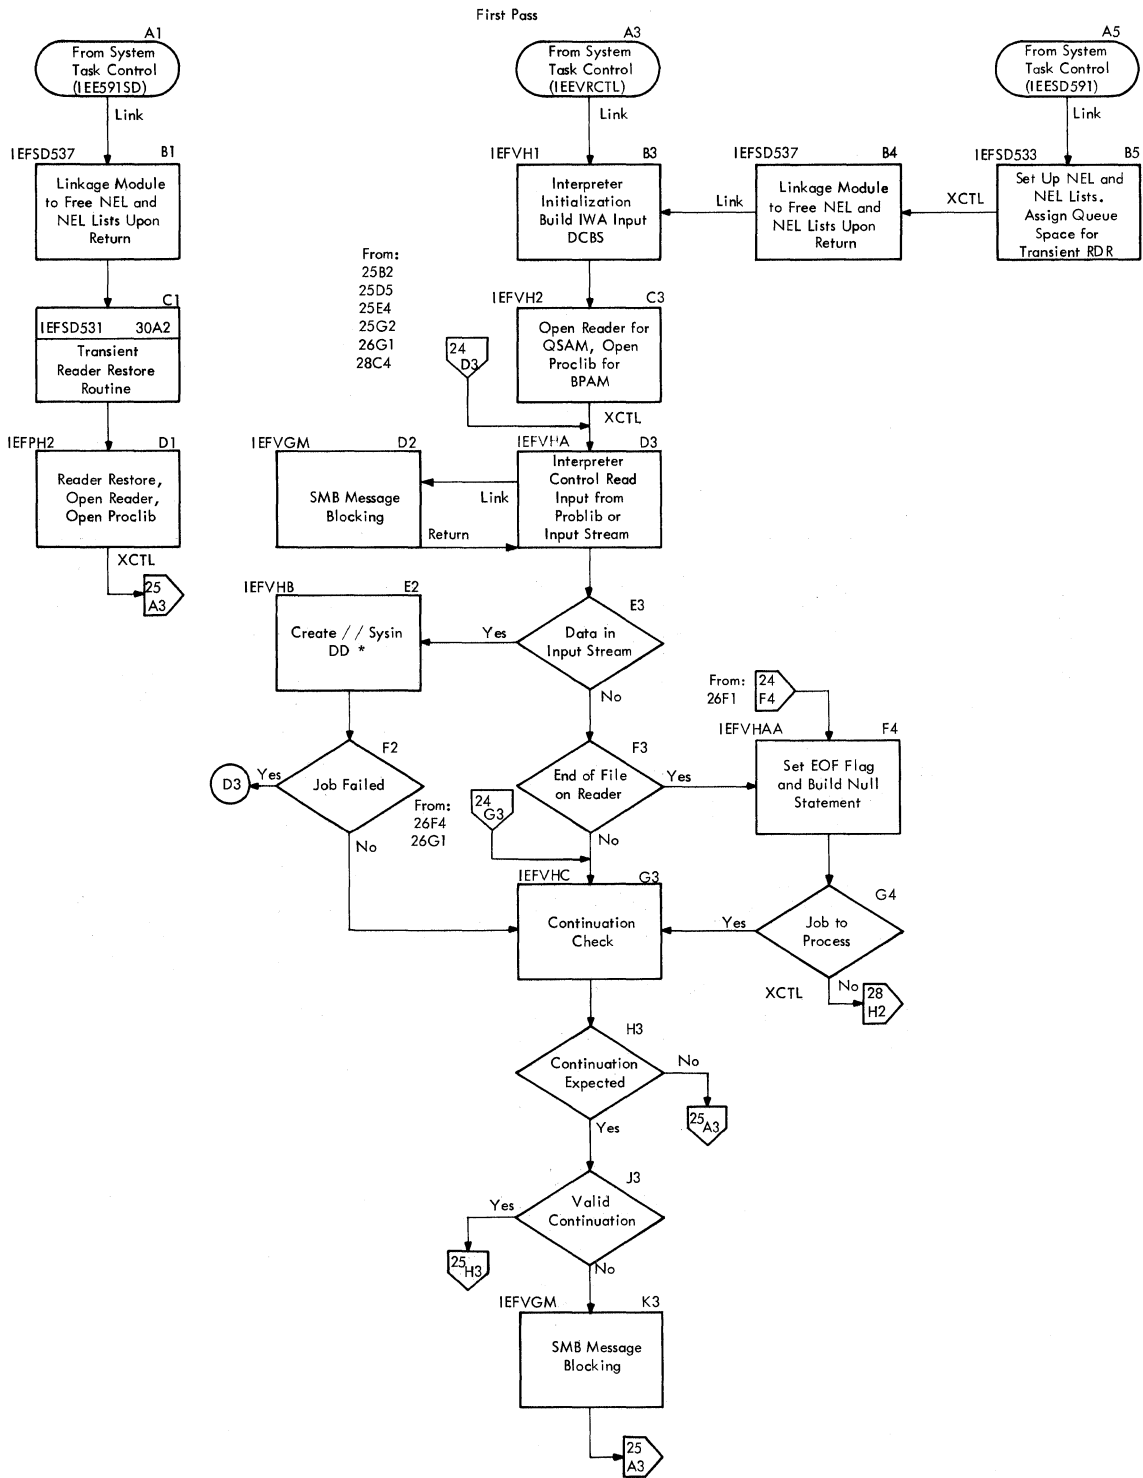


Chart 25. Reader/Interpreter (Sheet 2 of 3)

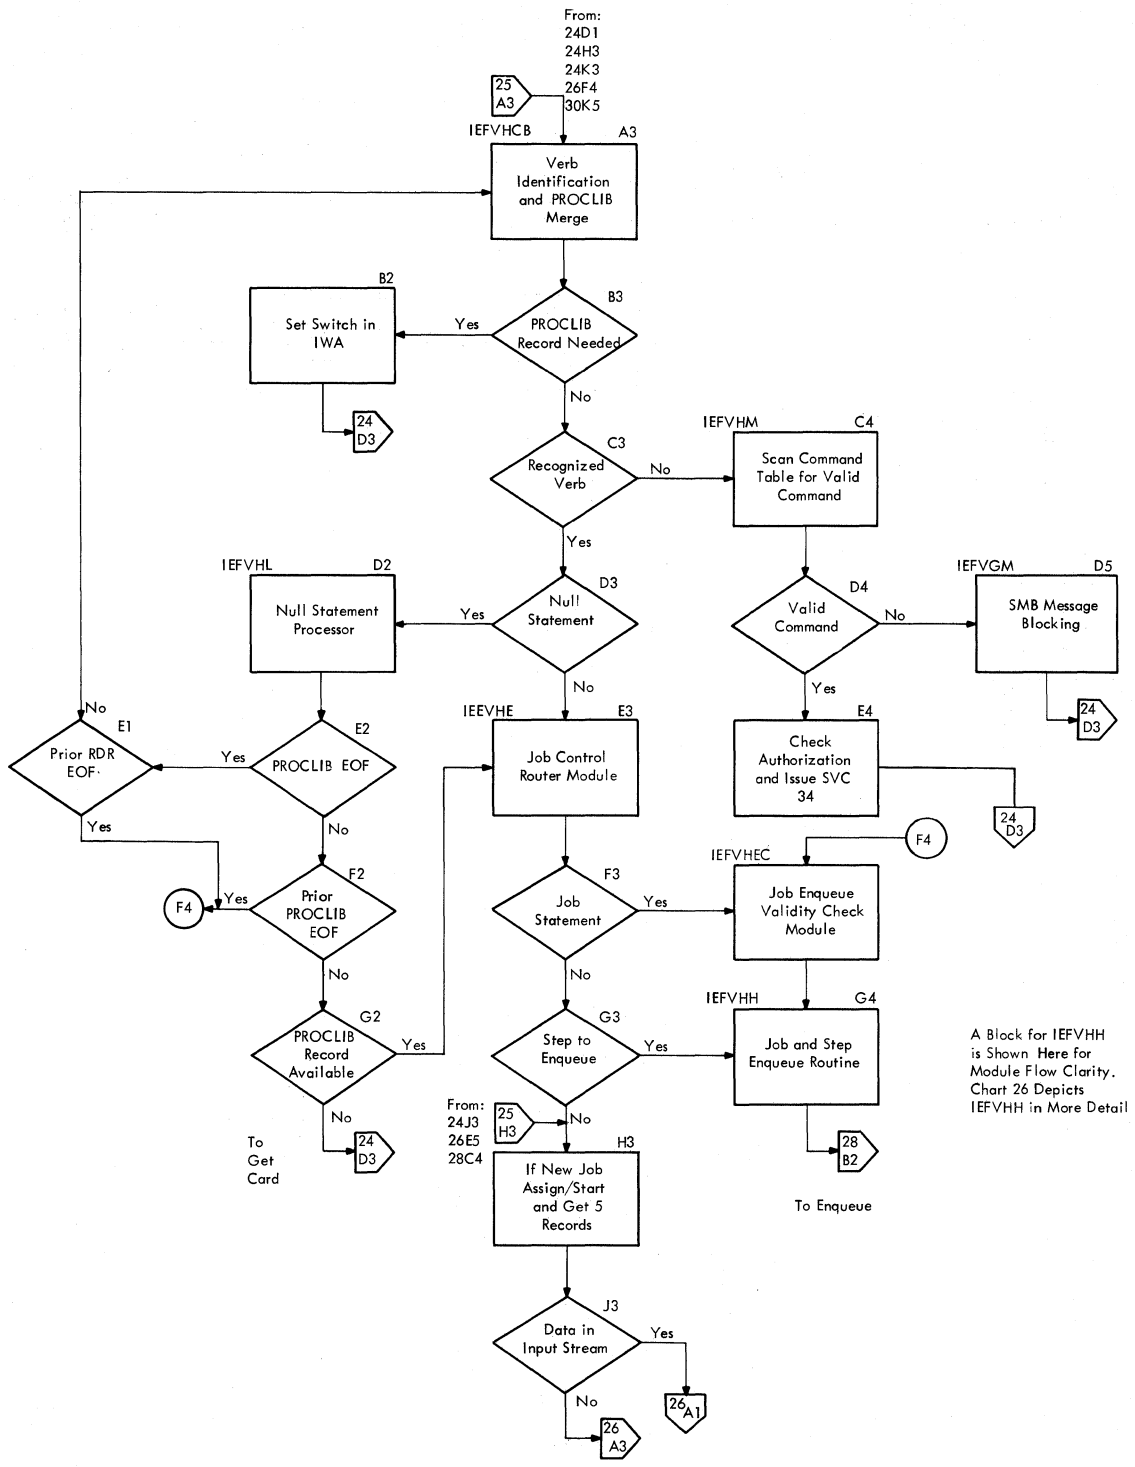


Chart 26. Reader Interpreter (Sheet 3 of 3)

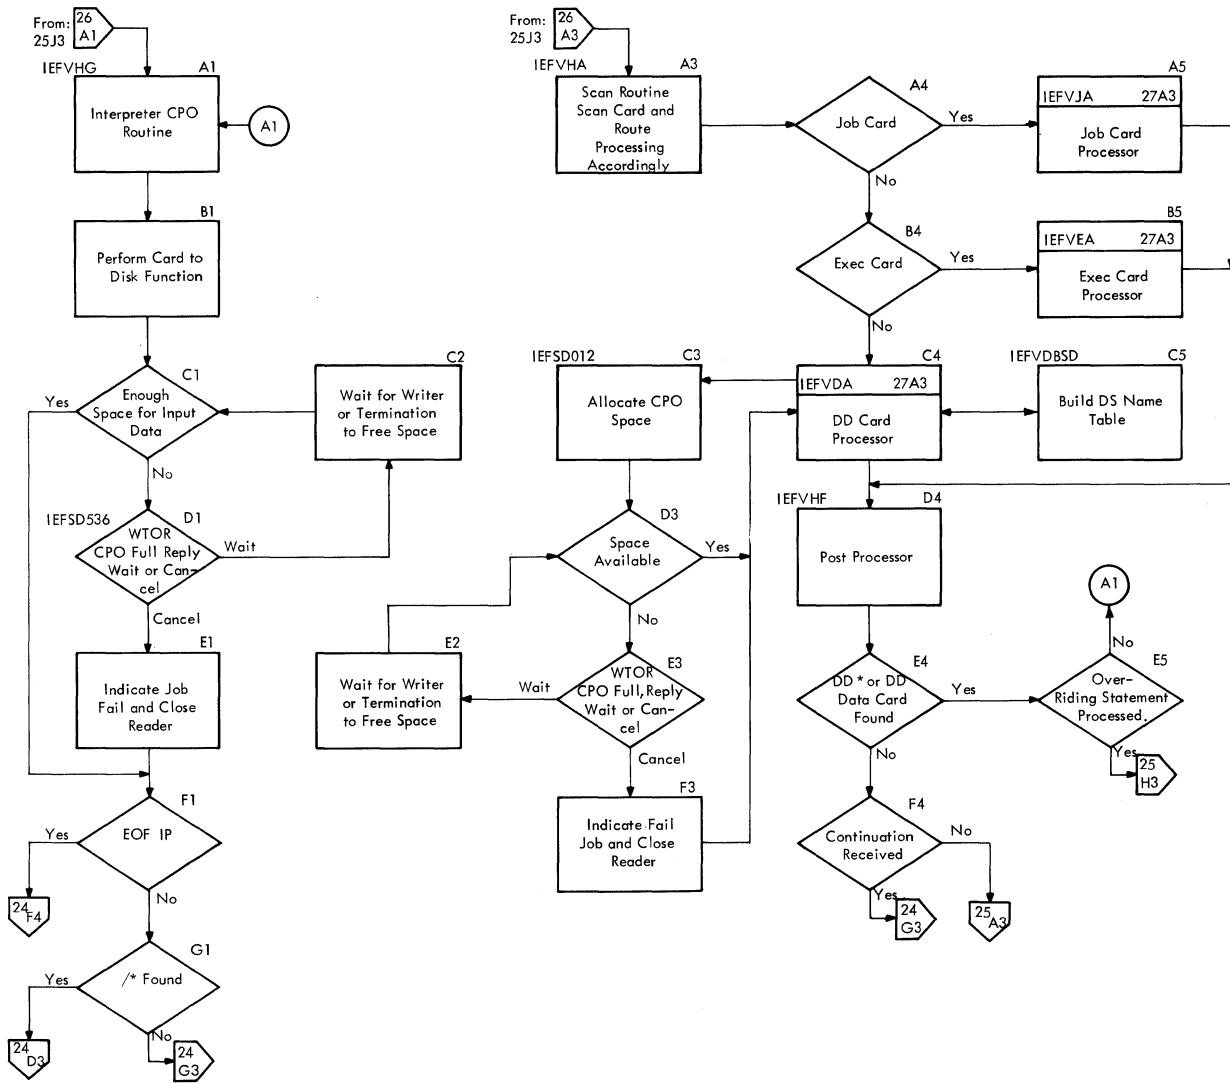


Chart 27. JCL Statement Processor

JCL Processing Modules
 IEFVJA, IEFVEA and IEFVDA
 Function by Driving
 Subroutines. Their
 General Flow is Described
 on this Chart

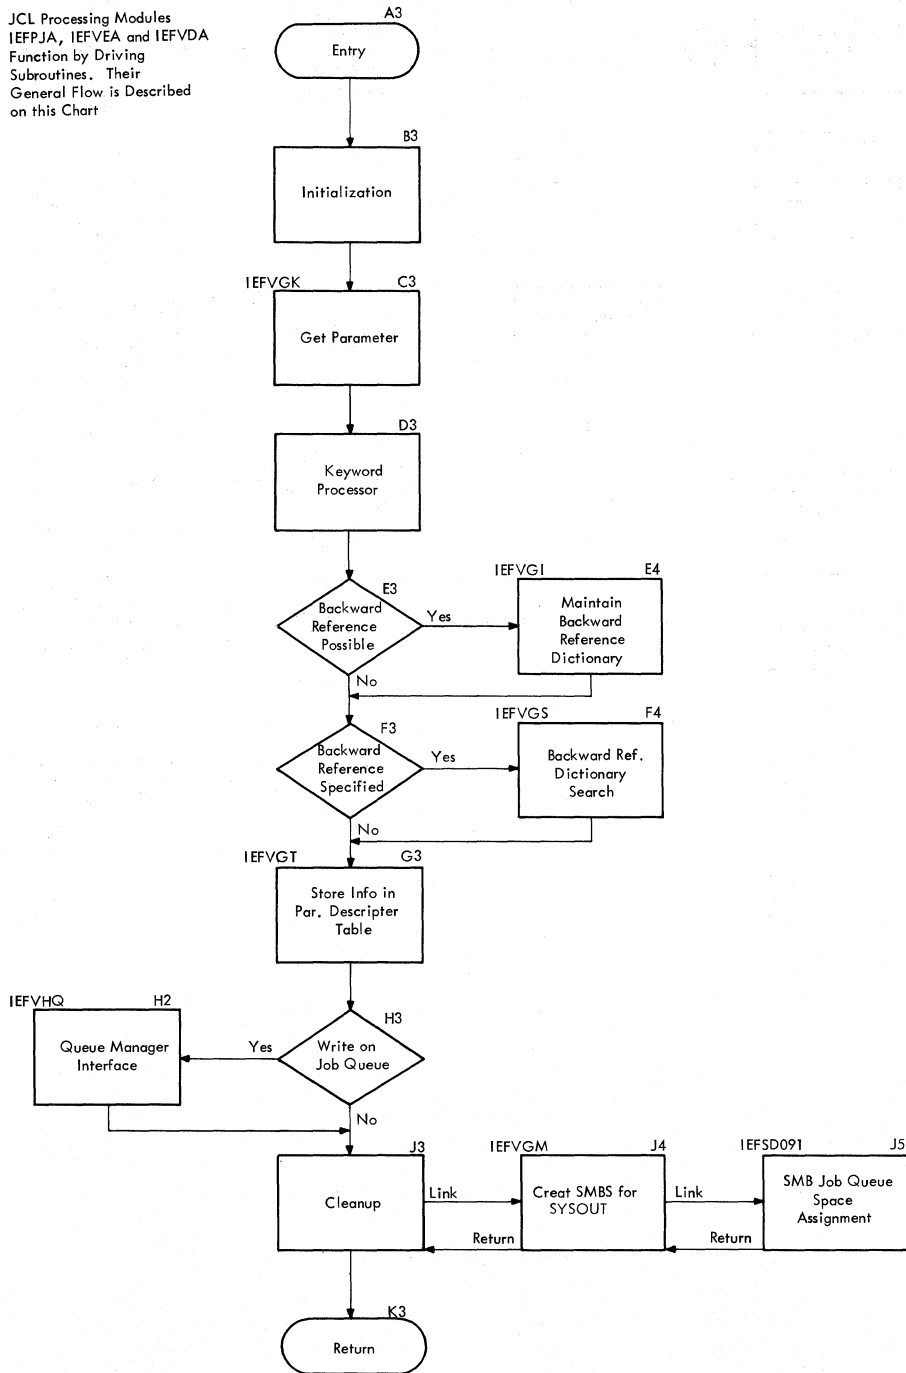


Chart 28. Job and Step Enqueue Routine

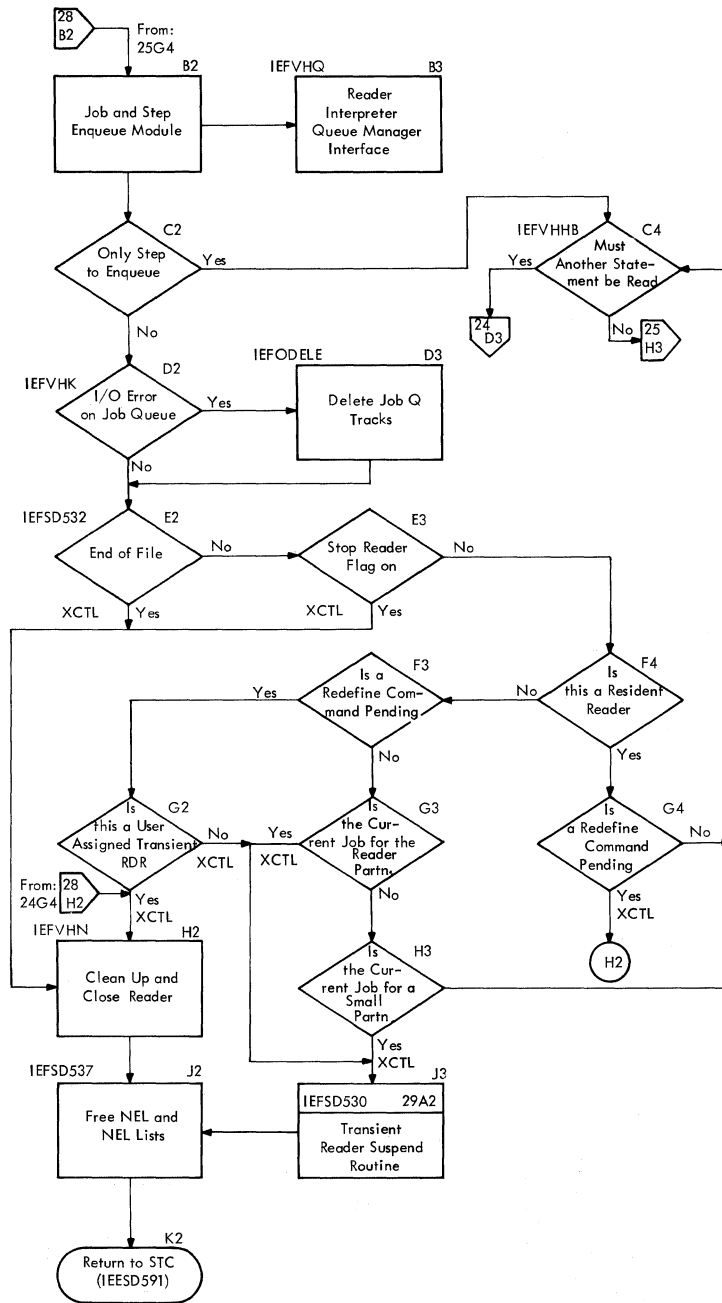


Chart 29. Transient Reader Suspend Routine

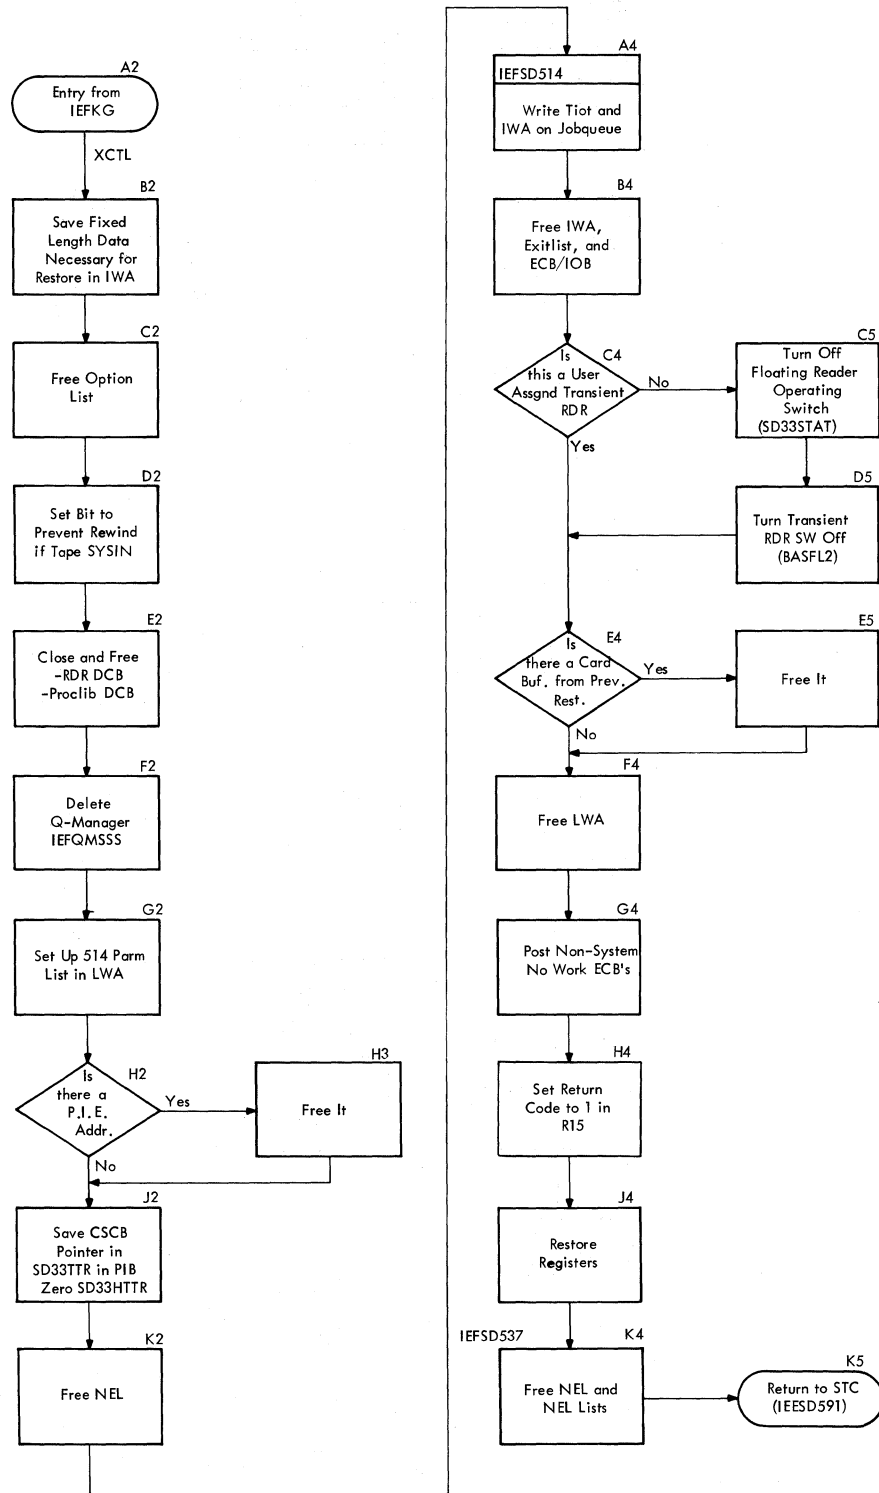


Chart 30. Transient Reader Restore Routine

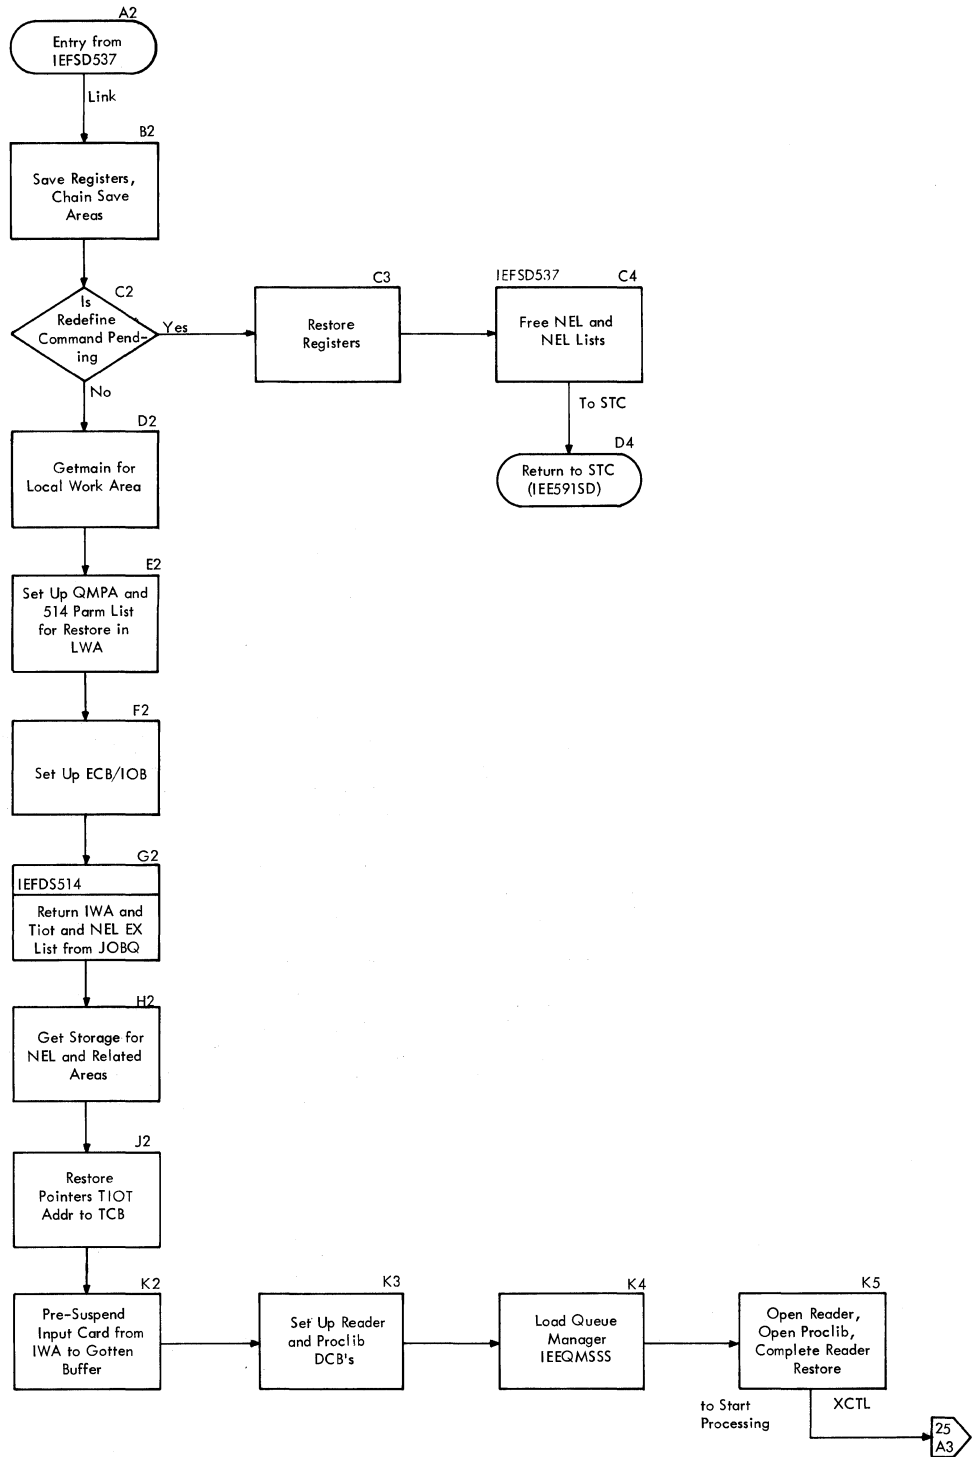


Chart 31. System Output Writer Control Flow

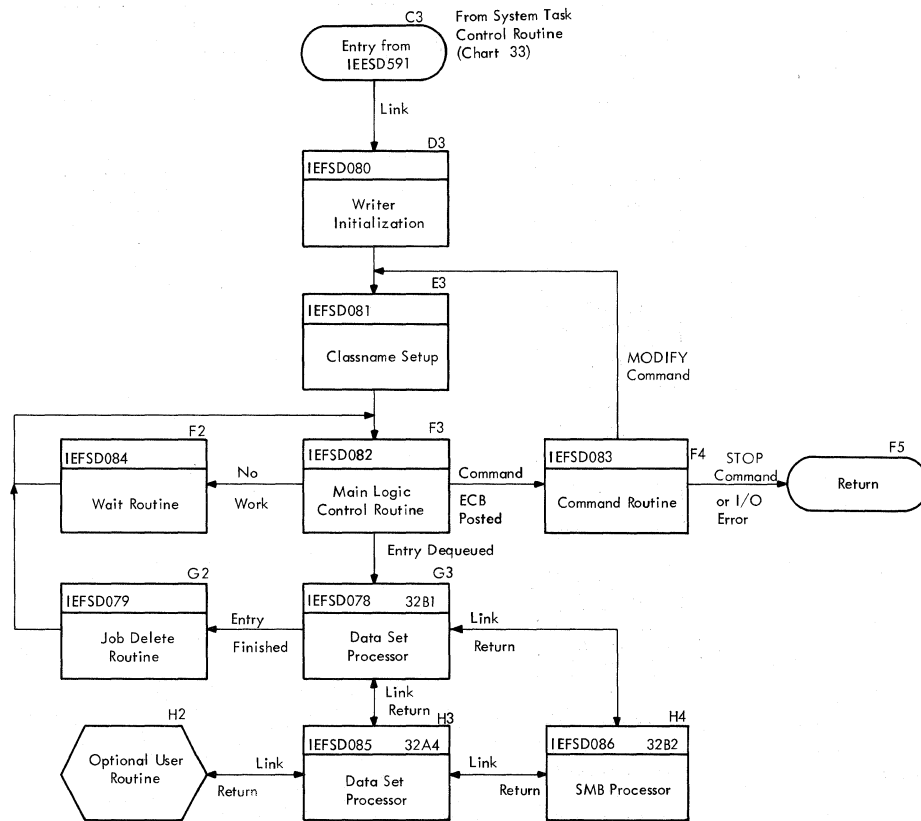


Chart 32. System Output Writer

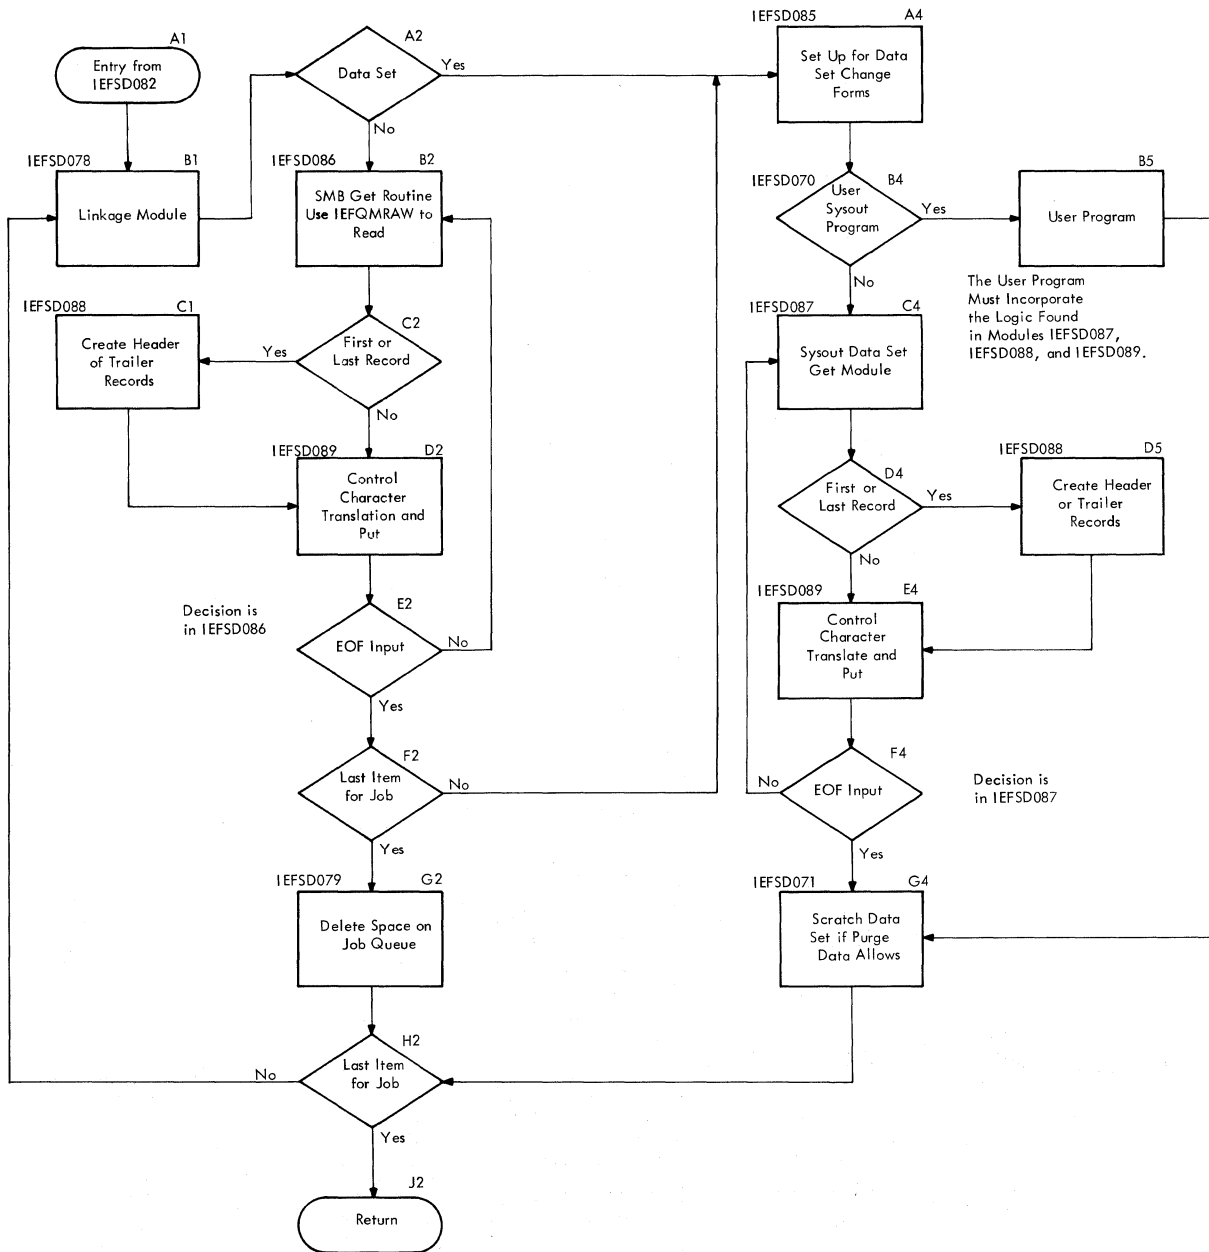
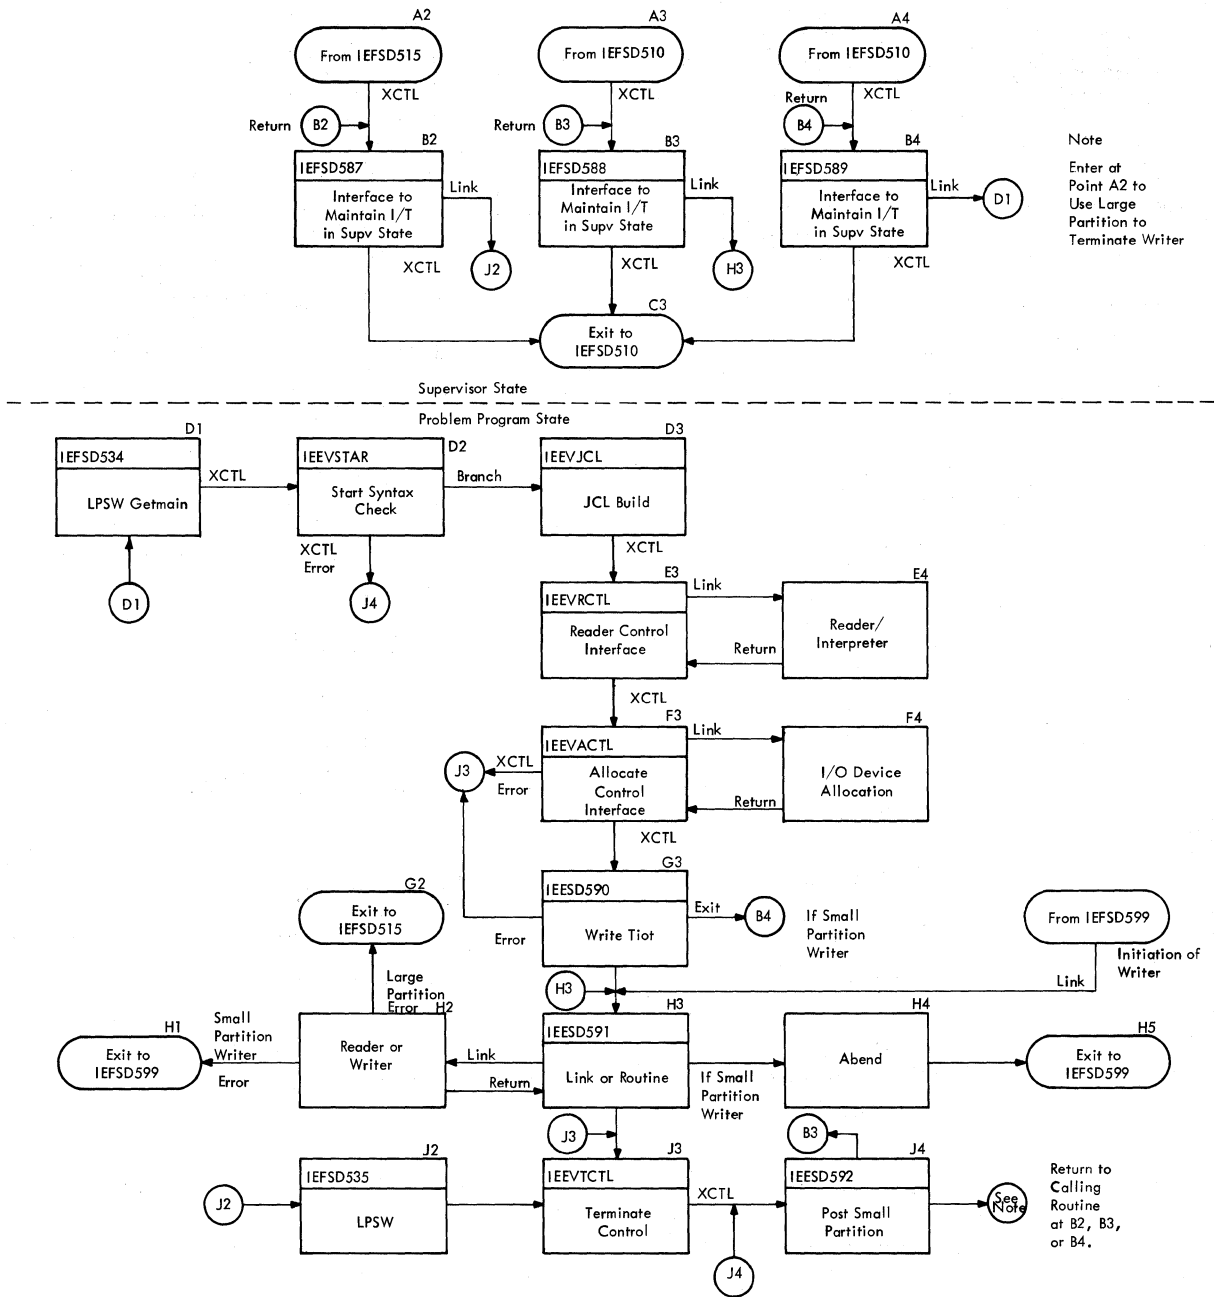


Chart 33. System Task Control



Index

Indexes to program logic manuals are consolidated in the publication IBM System/360 Operating System: Program Logic Manual Master Index, Form Y28-6717. For additional information about any subject listed below, refer to other publications listed for the same subject in the Master Index.

Where more than one page reference is given, the major reference is first.

- ABDUMP 30
- ABEND service routine 29-30
- ABEND/STAE interface routine 28-29
- Abnormal termination processing 28-31
- ABTERM 29
- Access methods 10-11
- Accounting routine 63
- Active request block queue 33
- Allocate parameter list (APL) 68-69,71
- Allocate register save area (ARSA) 68-69,71
- Alternate console 42-43
- APL (see allocate parameter list)
- ARSA (see allocate register save area)
- ASB (see automatic SYSIN batching)
- ASIR (see ABEND/STAE interface routine)
- Assign/Start routines 55-57
- ATTACH macro instruction 31,33
- Automatic commands 18,40-41
- Automatic SYSIN batching 55

- BBX (see boundary box)
- BLDL routines 11
- Boundary box (BBX) 35

- CANCEL command 47,56-57
- Catalog management 10,16
- CCH (see Channel-check routine)
- Channel-check routine 37
- Checkpoint/Restart 37
- Command processing 39-40,42,47-48
- Command scheduling control block (CSCB) 76-77,63
 - chain 49
 - creation routine 47,35,18
- Communication task 42-46,35,39
 - control flow 45
 - dispatching 24-27
 - SVC 72 43-44
- Communication vector table (CVT) . 21-23,14
- Contents supervision 33-34
- Control program functions 10-11
 - (also see data management, job management, and task management)
- Control program organization 11
 - non-resident portion 11
 - resident portion 11
- Core storage
 - (see main storage hierarchy support)

- CSCB (see command scheduling control block)
- CVT (see communication vector table)

- DADSM (see direct access device space management)
- Damage Assessment routine 29,30-31
- DAR (see Damage Assessment routine)
- Data control block (DCB) 15,54
- Data event block (DEB) 54
- Data set 15-16
- Data set block (DSB) 72-73
- Data set control block (DSCB) 15
- Data set descriptor record (DSDR) 68
 - (also see checkpoint/restart)
- Data set enqueue table (DSENQ) 68,79
- Data set input stream 15
- Data set integrity 68
- DCB (see data control block)
- DEB (see data event block)
- DEFINE command processing 47,50-53
- Defining control program areas 18
- Definition routines 50-53
- DELETE macro instruction 34
- Delete operator-message (DOM) macro instruction 39,42,46
- DEQ macro instruction 32-33,59
- Dequeue
 - queue manager dequeue routine 54,59
 - supervisory routine 32-33
- Device allocation 16
- Direct access device space management (DADSM) 10,16
- Dispatcher 21-30
 - with time-slicing 27-30
 - without time-slicing 21-27
 - (also see communication vector table, task control block, and task dispatching)
- Dispatching priority 21
- DISPLAY command 47,50
 - (also see command processing)
- DOM macro instruction (see delete operator-message macro instruction)
- DSB (see data set block)
- DSCB (see data set control block)
- DSDR (see data set descriptor record)
- DSENQ (see data set enqueue table)
- DSNAME parameter 68
- Dynamic area..... 12,9
 - partition organization 9

- ECB (see event control block)
- ECB/IOB 49
 - (also see event control block and input/output block)
- EIL (see event indication list)
- End-of-volume
 - (see open/close/end-of-volume)
- ENQ macro instruction 32-33
- ENQ/DEQ Purge routine 70-71

ENQ/DEQ routine	32-33	Interruption supervision	21,45-46
Enqueue		IOS (see input/output supervisor)	
queue manager enqueue routine	59	IPL (see initial program loading)	
supervisory routine	68	IQE (see interruption queue element)	
Entering commands	43	IWA (see interpreter work area)	
Entry to job management		JCL (see job control language)	
after IPL	40	JCLS (see job control language set)	
following step execution	41	JCT (see job control table)	
EOV (end-of-volume)		JFCB (see job file control block)	
(see open/close/end-of-volume)		JFCBX	
Error handling	36-37	(see job file control block extension)	
Event control block (ECB)	31,32,43	Job class	54-55
Event indication list (EIL)	44	Job control language (JCL)	73
Extended save area (XSA)		Job control language set (JCLS)	74
		Job control table (JCT)	84-85,15,59
FINCH routine	34	Job deletion	71,67-68
Fixed area	9,13	Job file control block (JFCB)	86-87,15
(also see input/output error handling,		Job file control block extension	
SVC transient area, SVCLIB partitioned		(JFCBX)	86-87
data set, and system queue area)		Job initiation	68
Free track queue	55	(also see step control table)	
FREEMAIN macro instruction	35	Job management	39,15
		control flow	40-42
GDG (see generation data group)		job scheduler function	39
General system initialization	18,48-49	(also see command processing,	
Generation data group (GDG)	148	communication task, and master	
GETMAIN macro instruction	35	scheduler)	
Graphic console	46	Job processing	53
		(also see initiator/terminator, input	
HALT command	39-40,47	stream, reader/interpreter, START, and	
Hierarchy support		system output writers)	
(see main storage hierarchy support)		Job queue	54-59,49
HOLD command	39-40	initialization	54-55
H0 (see main storage hierarchy support)		(also see queue control record)	
H1 (see main storage hierarchy support)		Job scheduler	39
		Job selection	63,49
I/O supervisor		(also see command processing, life of	
(see input/output supervisor)		task block, and partition information	
IDENTIFY macro instruction	34	block)	
Inactive partition	22	Job step timer	28
Initial program loading (IPL)	11,12,13	Job stream (see input stream)	
Initiating system tasks		Job termination (see job deletion)	
(see system task control)			
Initiator/Terminator	62-72	Large partition	62
(also see data set integrity, ENQ/DEQ		LCS (see main storage hierarchy support)	
purge routines, job deletion, job		LCT (see linkage control table)	
initiation, job selection, small		Life of task block (LOT)	87,88
partition scheduling, and step		Link library option	
deletion)		(see resident reenterable routine area)	
Input job queue	55-56,63	LINK macro instruction	33
Input stream		(also see ATTACH macro instruction)	
data sets	15	Link pack area (LPA)	
Input/output		(see resident reenterable routine area)	
device allocation	16	Link parameter list (LPL)	65
error handling	36-37	Linkage control table (LCT)	87,89,71
supervisor	10-11	LINKLIB partitioned data set	11
Integrated operator's console	46	LOAD macro instruction	33-34
Interlocks, system	57	Loaded program list	33
(also see queue manager)		Loaded program request block (LPRB)	33
Interpreter entrance list (NEL) ..	61,74-75	Loaded request block (LRB)	33
Interpreter work area (IWA)	79-83,62	Local work area (LWA)	62
Interruption queue element	30	Log task	21,46-47
(also see task dispatching and task		Logical track	54,55-60
switching)		Logical track header (LTH)	55
		LOT (see life of task block)	

LPA (link pack area)
 (see resident reenterable routine area)
 LPL (see link parameter list)
 LPRB (see loaded program request block)
 LRB (see loaded request block)
 LTH (see logical track header)
 LWA (see local work area)

M/S resident data area
 (see master scheduler resident data area)
 Machine check handler (MCH) 36
 Main storage hierarchy support 12
 Main storage initialization ... 11-12,48-49
 (also see job queue initialization,
 master scheduler initialization,
 nucleus initialization program, and
 READY message)
 Main storage organization 11-12
 Main storage supervision 21,35
 Master scheduler task (MST) 46-53,35
 dispatching 21,24-25
 initialization 14,48-49
 resident data area 87-92
 (also see SVC 34 and task control block)
 MCH (see machine check handler)
 MCS (see multiple console support)
 MODE command 47,40
 MOUNT command 47,40
 MST (see master scheduler task)
 MSTCCN (master console) (see VARY command)
 Multiple console support (MCS) 46
 Must complete 33

NEL (see interpreter entrance list)
 NIP (see nucleus initialization program)
 No work ECB 59
 Nondispatchable tasks 24
 Nonresident
 readers (see transient reader)
 SVC routines (see SVC transient area)
 writers (see system output writers)
 Nucleus 11,18
 Nucleus initialization program (NIP) ... 18
 (also see general system initialization)

ONGFX/OFFGFX (see VARY command)
 ONLINE/OFFLINE (see VARY command)
 OPEN macro instruction 73-74
 Open/close/end-of-volume 10-11
 Output work queue 54,72-73
 Output writer 72-73
 Overlay supervisor 36,21

Partition 9,35
 definition 50-53
 organization 12-13
 recovery 71
 task control block 21-23
 Partition information block
 (PIB) 92-93,18,63
 location 31
 Passed data set queue (PDQ) 97
 PDQ (see passed data set queue)
 PIB (see partition information block)

POST macro instruction 32
 PRB (see program request block)
 Priority
 dispatching 9,21-27
 job 54-55
 Program request block (PRB) 33
 Program status word (PSW) 37
 Protection keys, storage 9,12,21
 PSW (see program status word)
 Purge routine 70-71

QCB (see queue control block)
 QCR (see queue control record)
 QEL (see queue element)
 QMPA (see queue manager parameter area)
 Qname 32
 Queue control block (QCB) 32
 Queue control record (QCR) 49,54-60
 Queue element (QEL) 32
 Queue manager 54-60
 functions 54
 job queue initialization 55
 parameter area (QMPA) 55
 (also see input work queue and output
 work queue)

Queues
 (see free track queue, input work queue,
 job queue, output work queue, and task
 control block)

RAM (see resident access method)
 RB (see request block)
 Reader/Interpreter 60-62
 resident reader 61
 transient reader 61
 (also see input stream, input work
 queue, and system task control)
 READY message 48
 Recording/Recovery routines 36-37
 (also see Damage Recovery routines)
 RELEASE command 46
 Remote job entry (RJE) 45
 Reply queue element (RPQE) 44
 Request block (RB) 33
 Request queue element (RQE) 30
 RESET command 46
 Resident access method (RAM) 11
 Resident reenterable load module option 11
 Resident reenterable routines 11
 Resident SVC (RSVC) area 11,34
 Restart reader 63
 RJE (see remote job entry)
 RMS/85 48
 Rname 32
 RPQE (see reply queue element)
 RQE (see request queue element)
 RSVC (see resident SVC)

SCD (see system output class directory)
 Scheduler (see initiator/terminator)
 SCT (see step control table)
 SDT (see start descriptor table)
 SER (see system environment recording)
 SEREP (see system environment recording)
 SET command 15,41
 SIOT (see step input/output table)

SIRB (see system interruption request block)	Task supervision	31
Small partition	Task switching	24
information list (SPIL)	TER (see Table Breakup routine)	
module	TCB (see task control block)	
scheduling	Terminator routines	67
SPIL (see small partition information list)	(also see initiator/terminator)	
SQA (see system queue area)	TIME macro instruction	35
STAE (specify task asynchronous exit)	Time slice control element (TSCE)	27-28, 23
service routine	Time-slicing	50, 23
START	CVTISCE field of CVT	23
Start descriptor table (SDT)	dispatcher	27-28
STC (see system task control)	Timer supervision	35-36
Step control table (SCT)	timer pseudo clock	36
Step deletion	timer queue element (TQE)	35
Step initiation	TIOT (see task input/output table)	
Step input/output table (SIOT) ...	TQCR (see table queue control record)	
Step termination (see step deletion)	TQE (see timer supervision timer queue element)	
STIMER macro instruction	Track stacking	55
STOP command	Transient area	
Storage protection (see protection keys)	input/output	12
Subpools	SVC	12, 44
Supervisor request block (SVRB)	Transient reader	61
SVC transient area (see transient area)	system assigned	61
SVC 34	user assigned	61
SVC 35	TSCE (see time slice control element)	
SVC 72	TTIMER macro instruction	35
SVCLIB partitioned data set		
SVRB (see supervisor request block)		
SYNCH macro instruction	UCB (see unit control block)	
Syntax check	UCM (see unit control module)	
DEFINE command	Unit control block (UCB)	43-44
master scheduler	Unit control module (UCM)	43-44
System area (see fixed area)	UNLOAD command	47
System environment recording	UPL (see user's parameter list)	
System initialization	User options	11
System input readers	(also see BLDL list, resident access method, resident SVC, and system queue area)	
(see reader/interpreter)	User's parameter list (UPL)	69
System interruption request block (SIRB)		
System output class directory (SCD)	Validity check	47, 52
System output writers	Volume table (VOLT)	15
nonresident	Volume table of contents (VTCC)	14
resident	VTOC (see volume table of contents)	
System queue area (SQA)		
boundary box	WAIT macro instruction	32
..	WQE (see WTO queue element)	
System restart	Write-to-operator (WTO)	
System task control (STC)	macro instruction	42-45
	queue element (WQE)	45
Table Breakup routine (TER)	reply queue element (WTCR)	42
Table queue control record (TQCR) ...	Writer (see system output writers)	
Task control block (TCB)	WTO (see write-to-operator)	
TCB queue	WTOR	
TCBRRP field	(see write-to-operator reply queue element)	
TCBFLGS field		
TCBTCB field	XCTL macro instruction	34
Task creation	XSA (see extended save area)	
Task dispatching		
Task input/output table (TIOT)		
Task management		
(also see contents supervision, interruption supervision, main storage supervision, overlay supervision, task supervision, and timer supervision)		

Indexes to program logic manuals are consolidated in the publication IBM System/360 Operating System: Program Logic Manual Master Index, Form Y28-6717. For additional information about any subject listed above, refer to other publications listed for the same subject in the Master Index.

READER'S COMMENT FORM

IBM System/360 Operating System
Control Program With MFT
Program Logic Manual

Form Y27-7128-4

- Is the material:

	Yes	No
Easy to read?	<input type="checkbox"/>	<input type="checkbox"/>
Well organized?	<input type="checkbox"/>	<input type="checkbox"/>
Complete?	<input type="checkbox"/>	<input type="checkbox"/>
Well illustrated?	<input type="checkbox"/>	<input type="checkbox"/>
Accurate?	<input type="checkbox"/>	<input type="checkbox"/>
Suitable for its intended audience?	<input type="checkbox"/>	<input type="checkbox"/>

- How did you use this publication?

<input type="checkbox"/> As an introduction to the subject	Other
<input type="checkbox"/> For additional knowledge	

- Please check the items that describe your position:

<input type="checkbox"/> Customer personnel	<input type="checkbox"/> Operator	<input type="checkbox"/> Sales Representative
<input type="checkbox"/> IBM personnel	<input type="checkbox"/> Programmer	<input type="checkbox"/> Systems Engineer
<input type="checkbox"/> Manager	<input type="checkbox"/> Customer Engineer	<input type="checkbox"/> Trainee
<input type="checkbox"/> Systems Analyst	<input type="checkbox"/> Instructor	Other

- Please check specific criticism(s), give page number(s), and explain below:

<input type="checkbox"/> Clarification on page(s)	<input type="checkbox"/> Deletion on page(s)
<input type="checkbox"/> Addition on page(s)	<input type="checkbox"/> Error on page(s)

Explanation:

• Thank you for your cooperation. No postage necessary if mailed in the U.S.A.

YOUR COMMENTS, PLEASE . . .

This manual is part of a library that serves as a reference source for systems analysts, programmers and operators of IBM systems. Your answers to the questions on the back of this form, together with your comments, will help us produce better publications for your use. Each reply will be carefully reviewed by the persons responsible for writing and publishing this material. All comments and suggestions become the property of IBM.

Note: Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.


Cut Along Line

Fold

Fold

FIRST CLASS
PERMIT NO. 81
POUGHKEEPSIE, N.Y.

BUSINESS REPLY MAIL
NO POSTAGE STAMP NECESSARY IF MAILED IN THE UNITED STATES



POSTAGE WILL BE PAID BY ...

IBM Corporation
P.O. Box 390
Poughkeepsie, N.Y. 12602

Attention: Programming Systems Publications
Department D58

Fold

Fold

IBM System/360 Printed in U.S.A. Y27-7128-4



International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, N.Y. 10601
[USA Only]

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
[International]



International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, N.Y. 10601
[USA Only]

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
[International]