

GC28-1150-3
File No. S370-36

Program Product

**MVS/Extended Architecture
System Programming
Library: System Macros
and Facilities
Volume 1**

MVS/System Product:

JES3 Version 2	5665-291
JES2 Version 2	5740-XC6

IBM

Fourth Edition (June, 1987)

This is a major revision of, and obsoletes GC28-1150-2. See the Summary of Amendments following the Contents for a summary of the changes made to this manual. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

This edition applies to Version 2 Release 2.2.0 and all subsequent releases of MVS/System Product 5665-291 or 5740-XC6 until otherwise indicated in new editions or Technical Newsletters. Changes are made periodically to the information herein; before using this publication, in connection with the operation of IBM systems, consult the latest *IBM System/370 Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products or services do not state or imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product in this publication is not intended to state or imply that only IBM's product may be used. Any functionally equivalent product may be used instead.

Publications are not stocked at the address given below. Requests for 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, Information Development, Department D58, Building 921-2, PO Box 390, Poughkeepsie, New York 12602. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Preface

This two-volume publication describes supervisor and scheduler facilities that the system programmer can use. In this publication, a system programmer is defined as a programmer whose programs run in supervisor state, system key 0-7 or access APF-authorized libraries. The publication included the macro instructions and parameters used to obtain the functions.

Volume 1 contains descriptions of the supervisor and scheduler services available to a system programmer. Most of the services described are supervisor services; however, the scheduler functions available through the use of the DYNALLOC macro instruction are also described. Volume 1 includes a description of the DYNALLOC macro instruction. Some of the topics discussed in Volume 1 are also discussed in *Supervisor Services and Macro Instructions*; however in Volume 1, these topics are extended to include functions that are restricted to system programmers or used primarily by system programmers.

Volume 2, GC28-1151, contains the formats and descriptions of the supervisor macro instructions. Volume 2 provides system programmers with the information necessary to code the macro instructions. Each macro instruction is completely described, in Volume 2, but restrictions, requirements, and environmental considerations for the effective use of each macro is explained in Volume 1.

Publications referenced:

A Structured Approach to Describing and Searching Problems, SC34-2129

The Considerations of Physical Security in a Computer Environment, G520-2700

Data Security Controls and Procedures - A Philosophy for DP Installations, G320-5649

MVS/Extended Architecture Checkpoint/Restart User's Guide, GC26-4012

MVS/Extended Architecture Data Administration: Macro Instruction Reference, GC26-4014

MVS/Extended Architecture Debugging Handbook Volume 1, LC28-1164

MVS/Extended Architecture Debugging Handbook Volume 2, LC28-1165

MVS/Extended Architecture Debugging Handbook Volume 3, LC28-1166

MVS/Extended Architecture Debugging Handbook Volume 4, LC28-1167

MVS/Extended Architecture Debugging Handbook Volume 5, LC28-1168

MVS/Extended Architecture Diagnostic Techniques, LY28-1199

MVS/Extended Architecture Job Control Language, GC28-1148

MVS/Extended Architecture Message Library: System Messages, GC28-1156

MVS/Extended Architecture Operations: System Commands, GC28-1206

MVS/Extended Architecture Planning: Global Resource Serialization, GC28-1062

MVS/Extended Architecture Interactive Problem Control System User's Guide and Reference, GC28-1297

MVS/Extended Architecture Supervisor Services and Macro Instructions, GC28-1154

MVS/Extended Architecture System Initialization Logic, LY28-1200

MVS/Extended Architecture System Programming Data Administration, GC26-4010

MVS/Extended Architecture System Programming Library: Initialization and Tuning, GC28-1149

MVS/Extended Architecture System Programming Library: Service Aids, GC28-1159

MVS/Extended Architecture System Programming Library: System Modifications, GC28-1152

MVS/Extended Architecture System Programming Library: User Exits, GC28-1147

MVS/Extended Architecture System Programming Library: 31-Bit Addressing, GC28-1158

MVS/Extended Architecture Virtual Storage Access Method (VSAM) Administration Guide, GC26-4015

OS/VS Mass Storage System Extensions Messages, SH35-0041

Resource Access Control Facility (RACF): General Information Manual, GC28-0722

Security Assessment Questionnaire, GX20-2381

System Programming Library: RACF, SC28-1343

370-Extended Architecture: Principles of Operation, GA22-7085

Notes:

1. All references to RACF in this publication indicate the program product Resource Access Control Facility Version 1 Release 6 (5740-XXH).
2. All references to Assembler H in this publication indicate the program product Assembler H Version 2 (5668-962).
3. All references to RMF in this publication indicate the program product Resources Measurement Facility (5665-274)

Contents

Introduction	1-1
Subtask Creation and Control	1-3
Creating a New Task (ATTACH)	1-3
Changing the Defaults of ATTACH	1-4
Issuing an Internal START or REPLY Command (MGCR)	1-6
Communicating with a Problem Program (EXTRACT, QEDIT)	1-7
Providing an EXTRACT Answer Area	1-11
Changing the Priority of a Task (CHAP)	1-12
Program Management	1-13
Residency and Addressing Mode of Programs	1-13
Placement of Modules in Storage	1-14
Addressing Mode	1-14
Specifying Where the Module is to be Loaded (LOAD)	1-14
Synchronous Exits (SYNCH)	1-15
Using Checkpoint/Restart	1-15
Using Re-entrant Modules	1-16
Serialization	1-17
When Resource Serialization Is Needed	1-17
Serialization Requirements	1-17
Locking	1-18
Categories of Locks	1-18
Types of Locks	1-19
Classes of Locks	1-21
Locking Hierarchy	1-23
CML Lock Considerations	1-24
Obtaining, Releasing, and Testing Locks (SETLOCK)	1-24
Altering the Dispatching Queue (INTSECT)	1-25
Using the Must-Complete Function (ENQ/DEQ)	1-26
Characteristics of the Must-Complete Function	1-26
Programming Notes	1-26
Limiting Global Resource Serialization Requests	1-27
Shared Direct Access Storage Devices (Shared DASD)	1-28
Devices that Can be Shared	1-28
Volume/Device Status	1-29
System Configuration	1-29
Volume Handling	1-30
Macro Instructions Used with Shared DASD (RESERVE, EXTRACT)	1-30
Indicating Event Completion (POST)	1-37
Cross Memory POST	1-37
Bypassing the POST Routine	1-37
Waiting for Event Completion (WAIT)	1-38

Writing POST Exit Routines	1-38	
Identifying and Deleting Exit Routines	1-39	
Initializing Extended ECBs and ECB Extensions	1-39	
POST Interface with Exit Routines	1-40	
Re-entry to POST from a POST Exit	1-41	
Example of Using a POST Exit Function	1-41	
Branch Entry to the POST Service Routine	1-42	
Branch Entry to the WAIT Service Routine	1-44	
Suspension and Resumption of Request Blocks	1-45	
Waiting for an Event to Complete (SUSPEND)	1-45	
Resuming Execution of a Suspended Request Block (RESUME)	1-47	
Transferring Control for SRB Processing (TCTL)	1-49	
Using the BRANCH= YES Option of CALLDISP (CALLDISP)	1-50	
Reporting System Characteristics	1-51	
Collecting Information About Resources and Their Requestors (GQSCAN)	1-51	
Using the SRM Reporting Interface to Measure Subsystem Activity	1-55	
Reporting Software Error Symptoms (SYMREC)	1-56	
Writing Applications That Use SYMREC	1-56	
The Format of the Symptom Record	1-58	
Symptom Strings — SDB Format	1-59	
Using EREP and IPCS to Format Symptom Record Reports	1-59	
Programming Notes for SYMREC Applications	1-60	
Communication	1-67	
Interprocessor Communication	1-67	
Service Classes	1-67	
Status Indicators	1-69	
Writing and Deleting Messages (WTO, WTOR, DOM, and WTL)	1-70	
Routing the Message	1-70	
Writing a Multiple-Line Message	1-72	
Embedding Label Lines in a Multiline Message	1-72	
Using the Authorized Parameters of WTO and WTOR	1-72	
Deleting Messages Already Written	1-73	
Identifying Messages to be Deleted	1-73	
Limiting the Extent of Message Deletion	1-74	
Writing to the System Log	1-74	
Inter-Address Space Communication	1-74	
Asynchronous Inter-Address Space Communication	1-75	
Synchronous Inter-Address Space Communication	1-80	
Designing a PC Routine	1-98	
Recovery Considerations	1-100	
Virtual Storage Management	1-103	
Allocating and Freeing Virtual Storage (GETMAIN, FREEMAIN)	1-104	
The BRANCH Parameter	1-104	
The KEY Parameter	1-105	
Using Cell Pool Services (CPOOL)	1-105	
Using Storage Subpools	1-106	
Obtaining Information about the Allocation of Virtual Storage	1-109	
Using the VSMLIST Work Area	1-109	
Managing SWA Control Blocks	1-117	
Assigning SWA Control Blocks	1-117	
SWA Control Blocks	1-118	

SWA Manager	1-120
SWA Manager Move Mode	1-121
SWA Manager Locate Mode	1-124
SWA Manager Error Processing	1-125
Real Storage Management	1-127
Fixing/Freeing Virtual Storage Contents	1-128
PGFIX/PGFREE Completion Considerations	1-129
Input to Page Services	1-130
Virtual Subarea List (VSL)	1-130
Page Service List (PSL)	1-131
Short Page Service List (SSL)	1-131
Branch Entry to the PGSER Routine	1-132
Branch Entry to MVS/370 Page Services	1-133
Cross Memory Mode	1-133
Non-Cross Memory Mode	1-134
The Nucleus	1-135
Linking to Routines in the DAT-OFF Nucleus (DATOFF)	1-135
Using System Provided DAT-OFF Routines (DATOFF)	1-136
Writing User DAT-OFF Routines	1-138
Obtaining Information about CSECTs in the DAT-ON Nucleus (NUCLKUP)	1-139
Normal and Abnormal Program Termination	1-141
Recovery Termination Manager	1-141
Invoking the Recovery Termination Manager	1-142
Processing Program Interruptions (SPIE, ESPIE)	1-144
Interruption Types	1-145
Intercepting System Errors	1-146
Using the SLIP Command	1-147
Obtaining an SVC Dump During Slip Processing	1-147
Bypassing Dump Suppression	1-147
System Trace Facilities	1-148
Performing Branch Tracing	1-148
Performing Address Space Tracing	1-148
Performing Explicit Tracing (PTRACE)	1-148
Dumping Virtual Storage	1-149
Using the IPCS Macro Instructions	1-149
Using the SDUMP Macro Instruction	1-150
Obtaining an SVC Dump	1-153
Obtaining a Summary Dump	1-154
Suppressing SDUMPs and SYSMDUMPs	1-156
Using Dump Data Sets	1-157
Using the Dumping Services Commands	1-157
Canceling and Restarting the DUMPSRV Address Space	1-157
Getting More Than One SYSMDUMP	1-158
Providing Recovery Routines	1-158
Providing Information for Dump Analysis and Elimination	1-159
Selecting a Recovery Routine	1-160
System Environment	1-161
ESTAE-Type Recovery Routines	1-168
Using the FESTAE Macro Instruction	1-170
Special Considerations	1-170
Recovery Routine Guidelines	1-183

Uses of Resource Managers	1-191
Protecting the System	1-193
System Integrity	1-193
Documentation on System Integrity	1-193
Installation Responsibility	1-193
Elimination of Potential Integrity Exposures	1-194
Using the Authorized Program Facility (APF)	1-198
APF Authorization	1-198
Using APF	1-200
Authorization Results Under Various Conditions	1-202
Guidelines for Using APF	1-203
Resource Access Control Facility (RACF)	1-203
Defining a Resource to RACF (RACDEF)	1-204
Identifying a RACF-Defined User (RACINIT)	1-204
Checking RACF Authorization (RACHECK)	1-204
Retrieving and Encrypting Data (RACXTRT)	1-204
Building In-Storage Profiles (RACLIST)	1-204
System Authorization Facility (SAF)	1-205
MVS Router	1-205
Interface to the MVS Router (RACROUTE)	1-207
Changing System Status (MODESET)	1-208
Generating an SVC	1-208
Generating Inline Code	1-208
Protecting Low Storage (PROTPSA)	1-209
Exit Routines	1-211
Using Asynchronous Exit Routines	1-211
Stage 1 Initialization	1-211
Stage 2 Scheduling	1-213
Stage 3 Execution	1-213
Establishing a Timer Disabled Interrupt Exit	1-214
DIE Characteristics	1-216
Timer Queue Element Control	1-218
User-Written SVC Routines	1-221
Writing SVC Routines	1-221
Programming Conventions for SVC Routines	1-223
Inserting SVC Routines Into the Control Program	1-225
Modifying the SVC Table at Execution Time (SVCUPDTE)	1-226
Subsystem SVC Screening	1-227
UCB Scan Services	1-229
Invoking IOSVSUCB	1-229
Input to IOSVSUCB	1-229
Limiting the UCB Scan	1-230
Output from IOSVSUCB	1-231
Example Using IOSVSUCB	1-232
Obtaining Information from the Input/Output Supervisor (IOS)	1-234
Dynamic Allocation	1-235
Introduction to SVC 99 Functions	1-236
Concepts Needed to Understand SVC 99 Processing	1-236
Processing Control Features	1-237

Functions Available Through SVC 99	1-239
Dynamic Allocation	1-239
Dynamic Unallocation	1-241
Dynamic Concatenation	1-243
Dynamic Deconcatenation	1-244
Dynamic Information Retrieval	1-244
Installation Options For SVC 99 Functions	1-245
Space and Unit Defaults	1-245
Mounting Volumes and Bringing Devices Online	1-247
Installation Input Validation Routine for SVC 99	1-248
Requesting SVC 99 Functions	1-249
Programming Considerations When Using SVC 99	1-249
SVC 99 Parameter List	1-251
Request Block Pointer	1-253
Request Block	1-253
Text Pointers	1-255
Text Units	1-255
Detailed Review of Dsname Allocation Processing	1-256
Checking for Environmental Conflicts	1-256
Using an Existing Allocation	1-257
Using a New Allocation	1-259
Considerations When Requesting Dsname Allocation	1-260
SVC 99 Return Codes	1-262
Information Reason Codes	1-262
Error Reason Codes	1-264
SVC 99 Text Units, by Function	1-270
Dsname Allocation Text Units	1-273
DCB Attribute Text Units	1-291
Non-JCL Dynamic Allocation Functions	1-302
Dynamic Unallocation Text Units	1-306
Dynamic Concatenation Text Units	1-309
Dynamic Deconcatenation Text Unit	1-310
Text Units for Removing the In-Use Attribute Based on Task-ID	1-310
Ddname Allocation Text Units	1-311
Dynamic Information Retrieval Text Units	1-312
Example of a Dynamic Allocation Request	1-316
Index	X-1

Figures

1. Setting Up the Buffer for MGCR 1-6
2. EXTRACT ECB, CIB Pointers, and Token 1-7
3. Command Input Buffer Contents 1-7
4. Example Using the EXTRACT and QEDIT Macros 1-9
5. EXTRACT Answer Area Fields 1-11
6. Assembler Definition of AMODE/RMODE 1-13
7. Summary of Locking Characteristics 1-19
8. Requests for Shared/Exclusive Locks 1-20
9. Valid Volume Characteristic and Device Status Combinations 1-29
10. Example of an Interlock Environment 1-32
11. Example of Subroutine Issuing RESERVE and DEQ 1-36
12. Bypassing the POST Routine 1-38
13. ECB Extension (ECBE) 1-39
14. Extended ECB 1-40
15. Data Areas Post Exit Example 1-41
16. POST Function and Branch Entry Points 1-43
17. POST Branch Entry Input 1-43
18. POST Branch Entry Output 1-44
19. GQSCAN Results with STEP, SYSTEM, SYSTEMS, or ALL 1-53
20. GQSCAN Results with LOCAL or GLOBAL 1-54
21. EBCDIC Characters Printed or Displayed on an MCS Console 1-70
22. PC Number Indexing Linkage and Entry Tables 1-87
23. Authorization and Linkage Macro Instructions 1-89
24. PC/PT Linkage Conventions 1-91
25. Declared Storage For Cross Memory Examples 1-92
26. Entry Table Descriptions for Examples 1-93
27. Linkage Table and Entry Table Connection 1-95
28. Linkage and Entry Tables for a Global Service 1-97
29. Characteristics of a Non-Space Switch PC Routine 1-99
30. Characteristics of a Space Switch PC Routine 1-100
31. Characteristics of the Valid Storage Subpools 1-107
32. MVS/XA Virtual Storage Map 1-108
33. Format of the VSMLIST Work Area 1-109
34. Description of VSMLIST Work Area 1-110
35. Allocated Storage Information for Subpools in a Specified Area 1-112
36. Format of Subpool Descriptor 1-113
37. Format of Allocated Block Descriptor 1-113
38. Allocated Storage Information for the Private Area 1-114
39. Allocated Storage Information for a Subpool List 1-114
40. Format of Free Space Descriptor 1-115
41. Unallocated Storage Information for CSA and PVT Subpools 1-116
42. Format of Region Descriptor 1-116
43. Format of Unallocated Block Descriptor 1-117
44. Format of a SWA Control Block 1-118

45. Format of the JFCB 1-118
46. SWA Control Blocks That Can Reside Above 16 Megabytes 1-119
47. SWA Control Blocks That Must Reside Below 16 Megabytes 1-119
48. Valid Function Codes for a Move Mode Request 1-121
49. EPA Size and Contents for Move Mode Requests 1-122
50. Valid Function Codes for a Locate Mode Request 1-124
51. EPA Size and Contents for Locate Mode Requests 1-125
52. DAT-OFF Routines Available to Users 1-135
53. Virtual Storage Map of DAT-ON Nucleus 1-139
54. Key Fields in the SDWA 1-166
55. ESTAE Environment 1-169
56. Routing Control to Recovery Routines 1-178
57. Assigning Authorization via SETCODE 1-202
58. Authorization Rules 1-202
59. Asynchronous Exit Data Area Configuration 1-212
60. Programming Conventions for SVC Routines 1-223
61. Parameter List for the UCB Scan Routine (IOSVSUCB) 1-230
62. Device Classes 1-230
63. Example of the UCB Scan Routine (IOSVSUCB) 1-232
64. JCL DD Statement Facilities not Supported by Dynamic Allocation 1-240
65. Structure of the SVC 99 Parameter List 1-252
66. SVC 99 Return Codes 1-262
67. Class 2 Error Reason Codes (Unavailable System Resource) 1-265
68. Class 3 Error Reason Codes (Invalid Parameter List) 1-266
69. Class 4 Error Reason Codes (Environmental Error) 1-267
70. Class 7 Error Reason Codes (System Routine Error) 1-269
71. Verb Code 01 (Dname Allocation) - Text Unit Keys, Mnemonics, and functions 1-271
72. Verb Code 01 (DCB Attributes) - Text Unit Keys, Mnemonics, and Functions 1-290
73. Verb Code 01 (Non-JCL Dname Functions) - Text Unit Keys, Mnemonics, and Functions 1-302
74. Verb Code 02 (Dynamic Unallocation) - Text Unit Keys, Mnemonics, and Functions 1-306
75. Verb Code 03 (Dynamic Concatenation) - Text Unit Keys, Mnemonics, and Functions 1-309
76. Verb Code 04 (Dynamic Deconcatenation) - Text Unit Key, Mnemonic, and Function 1-309
77. Verb Code 05 (Remove-In-Use Processing Based on Task-ID) - Text Unit Keys, Mnemonics, and Functions 1-310
78. Verb Code 06 (Dname Allocation) - Text Unit Keys, Mnemonics, and Functions 1-311
79. Verb Code 07 (Dynamic Information Retrieval) - Text Unit Keys, Mnemonics, and Functions 1-312
80. Example of a Dynamic Allocation Request 1-317
81. Parameter List Resulting From Dynamic Allocation Example 1-318

Summary of Amendments

Summary of Amendments for GC28-1150-3 MVS/System Product Version 2 Release 2

This major revision describes how to use the SYMREC macro, which is new. It also describes changes to the DATOFF, DOM, VSMLOC, WTO, and WTOR macros, and changes that affect:

- The use of address space by PC/AUTH.
- User-written SVC routines.
- The allocation default module (IEFAB445).
- Dumps obtained via SDUMP.
- The jobstep DD limit.

Summary of Amendments for GC28-1150-2 MVS/System Product Version 2 Release 1.3

This major revision contains information about the new macro, IOSINFO, in support of System Product Version 2 Release 1.3 and minor technical and editorial changes.

Summary of Amendments for GC28-1150-1 as Updated January 31, 1984 by TNL GN28-0910

This technical newsletter incorporates three categories of changes:

- Changes in support of System Product Version 2, Release 1.2.
- Technical updates in support of the Resource Access Control Facility (RACF) Version 1, Release 6, which include the entirely new RACXTRT macro.
- Maintenance updates and editorial changes.

Introduction

The system facilities described in this publication include both supervisor and scheduler services. The supervisor services provide the resources that your programs need while assuring that as many of these resources as possible are being used at a given time. The scheduler services described in this publication are the scheduler functions that are available through the use of the dynamic allocation macro instruction (DYNALLOC). Knowing the conventions and characteristics of the system facilities will help you to design more efficient programs.

Volume 1 describes those supervisor services that should be restricted in use to systems programmers and installation-approved personnel. If a particular topic includes a description of a macro instruction, the macro instruction is given in parentheses after the topic heading. *Volume 1* includes a description of the DYNALLOC macro instruction. The supervisor macros and parameters are described in *Volume 2*. The topics described in *Volume 1* are:

Subtask Creation and Control: Occasionally, you can have your program executed faster and more efficiently by dividing parts of it into subtasks that compete with each other and with other tasks for execution time. This topic includes information about task creation, using an internal START, and communication with a problem program.

Program Management: You can use the supervisor to aid communication between segments of a program. This topic includes information about the residency and addressing mode of a module, loading a module, synchronous exits, checkpoint/restart, and re-entrant modules.

Serialization: Portions of some tasks depend on the completion of events in other tasks, which requires planned task synchronization. Planning is also required when more than one program uses a serially reusable resource. Locking, the must-complete function, shared direct access storage devices, waiting for an event to complete, and indicating event completion are discussed in this topic.

Reporting System Characteristics: Collecting information about resources and their requestors and using the SRM and SYMREC reporting interfaces are described in this topic.

Communication: This topic is divided into four distinct and different types of communication. These are:

- Interprocessor communication available through the use of the SIGP instruction
- Communication with the operator available through the use of the WTO, WTOR, and DOM macro instructions
- Asynchronous inter-address space communication available through the use of the SCHEDULE macro instruction
- Synchronous inter-address space communication available through the use of cross memory facilities

Virtual Storage Management: Virtual storage allows you to write large programs without the need for complex overlay structures. This topic describes how to allocate and free virtual storage. It also includes descriptions of the VSM functions, available through the use of the VSMLIST, VSMLOC, and VSMREGN macro instructions, and a description of managing SWA control blocks.

Real Storage Management: The supervisor administers the use of real storage and directs the movement of virtual pages between auxiliary storage and real storage in page size blocks. The services provided allow you to release virtual storage contents, load virtual storage areas into real storage, and page out virtual storage areas from real storage.

The Nucleus: This topic includes descriptions of the functions available through the use of the DATOFF and NUCLKUP macro instructions.

Normal and Abnormal Program Termination: The supervisor provides facilities for writing exit routines to handle specific types of interruptions. It is not likely, however, that you will be able to write routines to handle all types of abnormal conditions. The supervisor therefore provides for termination of your program when you request it by issuing an ABEND macro instruction or when the control program detects a condition that will degrade the system or destroy data. This topic describes the recovery termination manager, system trace facilities, recovery routines, the use of the SPIE/ESPIE macro instructions to process program interruptions, the use of the SLIP command to intercept errors, and the use of the SDUMP macro to obtain a dump of virtual storage.

Protecting the System: This topic includes the maintenance of system integrity, the use of the authorized programming facility, the use of the resource access control facility, changing system status, and protecting low storage.

Exit Routines: Two types of exit routines are described in this topic. They are asynchronous exit routines and timer disabled interrupt exits.

User-Written SVC Routines: This topic contains information needed to write SVC routines. It includes the characteristics of the SVC routines, program conventions for SVC routines, and ways to insert SVC routines into the control program.

UCB Scan Services: This topic describes the function of the UCB scan routine (IOSVSUCB). This routine allows you to scan each unit control block (UCB) in the system or in a specified device class.

Dynamic allocation (SVC 99) Services: This topic describes the functions provided by dynamic allocation (SVC 99). A description of the parameter list used to request SVC 99 functions, the SVC 99 return codes, error codes, and information codes are included.

Subtask Creation and Control

The control program creates a task when it initiates execution of the job step; this task is the job step task. You can create additional tasks in your program. If you do not, however, the job step task is the only task in a job being executed. The benefits of a multiprogramming environment are still available even with only one task in the job step; work is still being performed for other jobs when your task is waiting for an event, such as an input operation, to occur.

The advantage in creating additional tasks within the job step is that more of your tasks are competing for control than the single job step task. When a wait condition occurs in one of your tasks, it is not necessarily a task from some other job that gets control; it might be one of your tasks, a portion of your job.

The general rule is that you should choose parallel execution of a job step (that is, more than one task in a job step) only when a significant amount of overlap between two or more tasks can be achieved. Both the amount of time the control program takes to establish and control additional tasks and your increased effort to coordinate the tasks and provide for communications between them must be taken into account.

Most of the information concerning subtask creation and control appears in *Supervisor Services and Macro Instructions*. This chapter continues discussion in the following areas:

- Task creation (ATTACH macro instruction)
- Issuing an internal START command (MGCR macro instruction)
- Communicating with a problem program (EXTRACT and QEDIT macro instructions)
- Changing the priority of a task (CHAP macro instruction)

Creating a New Task (ATTACH)

The ATTACH macro instruction causes the control program to create a new task. The complete use of the macro instruction is described in *Supervisor Services and Macro Instructions*.

The macro instruction has parameters that provide the authorized user (protection key 0-7 or supervisor state) flexibility in using the macro instruction's services. If authorized tasks do not specify a particular parameter, the default value for that parameter is assigned. These defaults include:

- JSTCB=NO -- the attached task is a task in the present job step.
- SM=PROB -- the new task is to run in problem program mode.
- SVAREA=YES -- a save area is needed for the new task.

- **KEY=PROP** -- the protection key of the newly created task is to be the same as the task using **ATTACH**.
- **DISP=YES** -- the subtask is to be dispatchable.
- **TID=0** -- the task identifier of the new task is 0.
- **JSCB** -- omission of this parameter specifies that the job step control block of the attaching task is also used for the new task.
- **NSHSPV** and **NSHSPL** -- omission of these parameters specifies that subpools 236 and 237, if they exist, are to be shared with the subtask.
- **RSAPF=NO** -- The APF authorization of the step is to be unchanged.

Changing the Defaults of **ATTACH**

Rather than accepting the default values, (assuming the task is authorized), you can extend the facilities of the **ATTACH** macro instruction by coding the following values:

- **JSTCB=YES** -- the attached task is a new job step task. In this case, the address of the TCB of the newly created task is placed in the **TCBJSTCB** field of the attached TCB. The initiator attaches the first load module of a job as a job step task. For such an attach, the program manager does not search the job library of the attaching task.

Also, only under a job step task can a system program (system key or supervisor state) attach a load module from a non-system library.

In order to attach a job step task, the attaching task (and any of its subtasks) must be job step tasks. If one of these conditions is not met, the new task will not be created.

- **SM=SUPV** -- the system is to run in supervisor mode when executing the attached task.

Supervisor state is a requirement for issuing privileged instructions (for example, **LPSW**). You can specify supervisor mode via this parameter or via the **MODESET** macro instruction.

- **SVAREA=NO** -- the new task does not need a save area.

The save area is obtained from the user's region. Because it might not always be desirable to have a save area (for example, the user's region might not be defined at the time of a system **ATTACH**), this parameter can be used to specify that no save area is to be created.

- **KEY=ZERO** -- the protection key of the newly created task is zero.

Protection key zero allows the new task to reference any defined storage and pass all validity checks.

- **DISP=NO** -- the subtask is to be nondispatchable.

This parameter causes the primary nondispatchability bit **TCBANDSP** to be turned on in the new TCB. As a result, the new TCB will not be dispatched. Thus, specifying **DISP=NO** allows the originating task to alter the new TCB. The new task remains nondispatchable until the originating task issues the **STATUS** macro instruction with the **RESET** option to reset **TCBANDSP**.

Note: **STATUS START TCB** will not make the new TCB dispatchable.

- **TID=task id** -- the task identifier specified is to be placed in the **TCBTID** field of the attached task.

The task identifier can be set to identify critical system tasks. Other uses of this parameter are not recommended.

- **JSCB=job step control block address** -- the address specified for the **JSCB** is to be used for the new task.

This parameter sets the **TCBJSCB** to the address of a job step control block. This action, normally associated with the creation of a job step task, is not required by **ATTACH**.

- **NSHSPV=subpool number and NSHSPL=subpool list address** -- subpools 236 and 237 are not to be shared with the new task.

Subpools 236 and 237 are known as the scheduler work area (SWA). This parameter allows the scheduler to control these subpools.

- **RSAPF=YES** -- reset the step **APF** authorization.

This parameter allows a system program that is not running **APF** authorized to **ATTACH** a subtask and have the **APF** authorization for the step reset according to the attributes of the subtask. The subtask must be attached while in the problem program state and must be in a non-system key. For more information on this parameter see "Authorization Results Under Various Conditions" in the "Protecting the System" section.

Issuing an Internal START or REPLY Command (MGCR)

A program can issue an internal START or REPLY command using the MGCR macro instruction and can pass 31 bits of information, called a token, to the program being started (in the case of the START command). An internal REPLY command is available to reply to a WTOR message. Before issuing the MGCR macro instruction, initialize a buffer for the command and the token, if any, as follows:

1 byte	1 byte	2 bytes	variable length	4 bytes
flags1	length	flags2	text	31 bit token right justified

You must also set register 0 to zero before issuing the MGCR macro instruction.

flags1

If bit 0 of the flags1 byte is one, the flags2 field must contain meaningful information. Bits 1-7 of flags1 must be zero.

length

The length field contains the length of the buffer in bytes, up to but not including the token field.

flags2

If a token is present, flags2 must be set to X'0800', otherwise, it must be set to X'0000'.

text

The text field contains the START or REPLY command followed by operands and, optionally, comments.

token

This field contains any desired information to be communicated to the started program. Token is meaningful only for the START command.

Figure 1 shows how the buffer is set up.

The IEZMGCR mapping macro, in SYS1.MACLIB, is available to map the buffer.

	SR	REGO,REGO	INDICATE SYSTEM ISSUED COMMAND
* MGCRMAC	MGCR	MGRCDATA	
	.	.	
MGRCDATA	EQU	*	
FLG1	DC	X'80'	
LGTH	DC	AL1(TOKEN-MGRCDATA)	
FLG2	DC	X'0800'	
TXT	DC	C'S IMS ***ANY COMMENTS***'	
TOKEN	DC	AL4(ECB)	ECB ADDR

Figure 1. Setting Up the Buffer for MGCR

Communicating with a Problem Program (EXTRACT, QEDIT)

The operator can pass information to the started program by issuing a STOP or a MODIFY command. In order to accept these commands, the program must be set up in the following manner.

The program must issue the EXTRACT macro instruction to obtain a pointer to the communications ECB and a pointer to the first command input buffer (CIB) on the CIB chain for the task. The ECB is posted whenever a STOP or a MODIFY command is issued. The EXTRACT macro instruction is written as follows, and returns what is indicated in Figure 2.

EXTRACT answer area, FIELDS=COMM

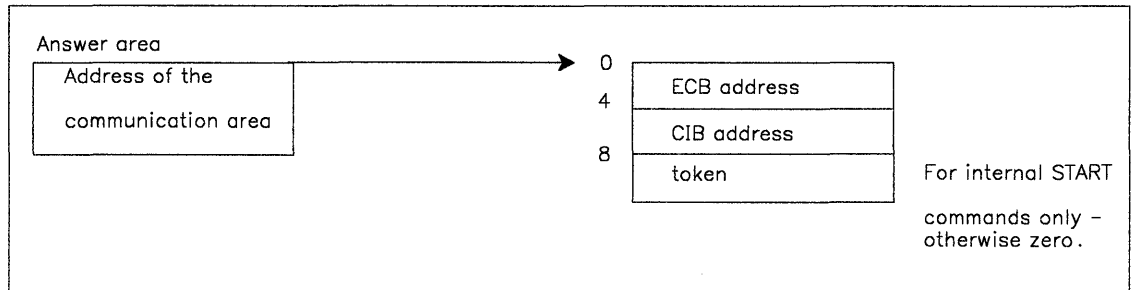


Figure 2. EXTRACT ECB, CIB Pointers, and Token

The CIB contains the information specified on the STOP, START, or MODIFY command, as shown in Figure 3. If the job was started from the console, the EXTRACT macro instruction will point to the START CIB. If the job was not started from the console, the address of the first CIB will be zero.

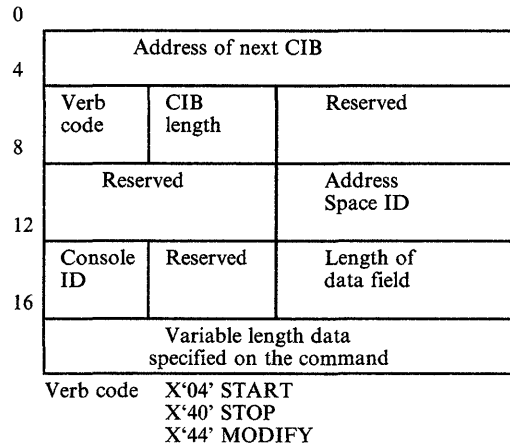


Figure 3. Command Input Buffer Contents

If the address of the START CIB is present, use the QEDIT macro instruction to free this CIB after any parameters passed in the START command have been examined. The QEDIT macro instruction is written as follows:

```
QEDIT ORIGIN=address of pointer to CIB,BLOCK=address of CIB
```

Notes:

1. *The address of the pointer to the CIB is the contents of the answer area plus 4 bytes, as shown in Figure 2.*
2. *The address of the CIB must be the exact address returned by EXTRACT, not an address generated from copying the CIB to another location.*

The CIB counter should then be set to allow CIBs to be chained and MODIFY commands to be accepted for the job. This is also accomplished by using the QEDIT macro instruction:

```
QEDIT ORIGIN=address of pointer to CIB,CIBCTR=n
```

The value of n is any integer value from 0 to 255. If n is set to zero, no MODIFY commands are accepted for the job. However, STOP commands are accepted for the job regardless of the value set for CIBCTR.

Note: When using the address or addresses returned from the EXTRACT macro as input to the QEDIT macro, you must establish addressability via the IEZCOM mapping macro, in SYS1.MACLIB, based on the address returned by the EXTRACT.

For the duration of the job, your program can wait on or check the communications ECB at any time to see if a command has been entered for the program. Check the verb code in the CIB to determine whether a STOP or a MODIFY command has been entered. After processing the data in the CIB, issue a QEDIT macro instruction to free the CIB.

The communications ECB is cleared by QEDIT when no more CIBs remain. Care should be taken if multiple subtasks are examining these fields. Any CIBs not freed by the task are unchained by the system when the task is terminated. The area addressed by the pointer obtained by the EXTRACT macro instruction, the communications ECB, and all CIBs are in protected storage and may not be altered.

The program in Figure 4 follows the procedure outlined in the preceding paragraphs. It shows how you can code the EXTRACT and QEDIT macros to accept MODIFY and STOP commands.

```

QEDITEX  CSECT
          BALR  12,0          PROGRAM...
          USING *,12         ...ADDRESSABILITY
*
* INITIALIZATION PROCESSING - DELETE START CIB
*
* OBTAIN ADDRESS OF CIB
*
          LA      5,ANSRAREA      ADDRESS OF RESPONSE AREA FOR QEDIT
          EXTRACT(5),FIELDS=COMM  OBTAIN ADDRESS OF THE           X
                                     COMMUNICATIONS AREA FOR THE       X
                                     CURRENT TASK
*
          L       5,ANSRAREA      LOAD ADDRESS OF COMMUNICATIONS AREA
          USING COMLIST,5         ESTABLISH ADDRESSABILITY TO IEZCOM
          L       3,COMCIBPT      OBTAIN ADDRESS OF CIB
          USING CIBNEXT,3        ESTABLISH ADDRESSABILITY TO IEZCIB
*
* MOVE DATA FROM CIB TO WORKING STORAGE
*
          LH      4,CIBDATLN      OBTAIN LENGTH OF DATA FIELD
          BCTR   4,0             DECREASE LENGTH BY ONE
          EX     4,DATAMOVE      MOVE DATA TO WORKING STORAGE
*
* FREE THE START CIB, IF PRESENT
*
          CLI    CIBVERB,CIBSTART  FIRST CIB FOR START COMMAND?
          BNE    SETCOUNT        NO, CONTINUE INITIALIZATION
          QEDIT  ORIGIN=COMCIBPT,BLOCK=(3) YES, FREE IT
          LTR   15,15             CHECK RETURN CODE
          BZ     SETCOUNT        IF RETURN CODE IS ZERO, THE CIB           X
                                     WAS FREED, CONTINUE
          WTO   'START CIB NOT FREED'  IF RETURN CODE IS NOT ZERO,           X
                                     NOTIFY THE OPERATOR THAT           X
                                     THE CIB WAS NOT FREED

```

Figure 4 (Part 1 of 2). Example Using the EXTRACT and QEDIT Macros

```

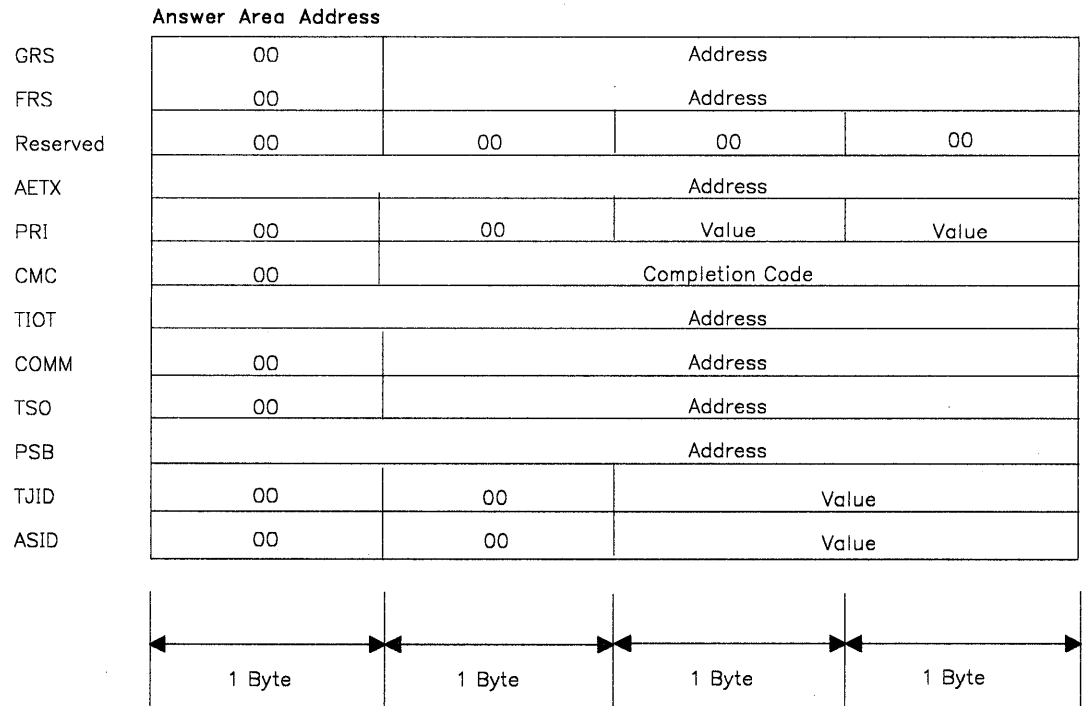
*
* SET THE LIMIT ON MODIFY COMMANDS
*
SETCOUNT EQU *
          QEDIT ORIGIN = COMCIBPT,CIBCTR = 2 SET LIMIT TO 2
*
* COMMAND PROCESSING LOOP
*
* CHECK THE COMMUNICATIONS ECB
*
WAIT     EQU *
          L     4,COMECPBT           OBTAIN ADDRESS OF COMMUNICATIONS ECB
          WAIT  ECB = (4)           WAIT FOR STOP OR MODIFY COMMAND
*                                           NOTE: QEDIT CLEARS THE POSTED ECB
*                                           WHEN THE LAST CIB IS FREED
          L     3,COMCIBPT         OBTAIN ADDRESS OF CIB
*
* TEST FOR MODIFY COMMAND CIB
*
          CLI   CIBVERB,CIBMODFY    IS IT FOR A MODIFY COMMAND?
          BNE   TSTSTOP             NO, GO TEST FOR STOP
*
* PROCESS THE MODIFY COMMAND
*
* FREE THE MODIFY CIB
*
          QEDIT ORIGIN = COMCIBPT,BLOCK = (3) FREE IT
          B     WAIT                WAIT FOR NEXT COMMAND
*
* TEST FOR STOP COMMAND CIB
*
TSTSTOP  EQU *
          CLI   CIBVERB,CIBSTOP     IS IT FOR A STOP COMMAND?
          BNE   ERROR1              NO, GO PROCESS AS ERROR
*
* PROCESS THE STOP COMMAND
*
* FREE THE STOP CIB
*
          QEDIT ORIGIN = COMCIBPT,BLOCK = (3) FREE THE CIB
          B     WAIT                WAIT FOR NEXT COMMAND
*
* ERROR HANDLER FOR UNRECOGNIZED CIB TYPE
*
ERROR1   EQU *
*
* CONSTANTS AND DATA AREAS
*
DATAAREA DS 0F                    WORK AREA FOR CIB DATA
ANSRAREA DS 4F                    ANSWER AREA FOR EXTRACT MACRO
DATAMOVE MVC DATAAREA(0),CIBDATA MOVE DATA FROM CIB TO DATAAREA
          DSECT
          IEZCOM                    MAPPING MACRO FOR COMMUNICATION AREA
          DSECT
          IEZCIB                    MAPPING MACRO FOR CIB
          END

```

Figure 4 (Part 2 of 2). Example Using the EXTRACT and QEDIT Macros

Providing an EXTRACT Answer Area

The EXTRACT macro instruction provides TCB information for either the active task or one of its subtasks. Figure 5 shows the order in which the information from the requested fields is returned. If the information from a field is not requested, the associated fullword is omitted.



Note: See the syntax of the EXTRACT macro instruction in Volume 2 for a description of these fields.

Figure 5. EXTRACT Answer Area Fields

You must provide an answer area consisting of contiguous fullwords, one for each of the codes specified in the FIELDS parameter, with the exception of ALL. If ALL is specified, you must provide a 7-word area to accommodate the GRS, FRS, reserved, AETX, PRI, CMC, and TIOT fields. The ALL code does not include the COMM, TSO, PSB, TJID, and ASID fields.

Most of the addresses are returned in the low-order three bytes of the fullword, and the high-order byte is set to zero; the fields for AETX, TIOT, and PSB could have a nonzero first byte. Fields for which no addresses or values are specified in the task control block are set to zero.

For example, if you code FIELDS=(TIOT,GRS,PRI,TSO,PSB,TJID) you must provide a 6-fullword answer area, and the extracted information appears in the same relative order as shown in Figure 5. (That is, GRS is returned in the first word, PRI in the second word, TIOT in the third word, and so forth.)

If FIELDS=(ALL,TSO,PSB,COMM,ASID) is coded, you need an 11-fullword answer area, and the extracted information appears in the answer area in the relative order shown above.

Changing the Priority of a Task (CHAP)

Programs should not use priority or precedence as a serialization mechanism because they become sensitive to changes in the system's dispatching algorithms. For example, the CHAP macro instruction does not ensure that tasks are dispatched in the expected order, due to dispatching on more than one processor. Also, the PRIORITY and DPRTY JCL parameters cannot be used to accomplish serialization. First, the system resources manager might change the dispatching priority of a task or job, allowing it to execute before a task with a previously higher priority. Second, because tasks can execute on more than one processor, tasks of different priority might be executed on more than one processor simultaneously.

Program Management

You can specify whether you want a program loaded into storage above or below the 16 megabytes line and if you want a program loaded at a specific address. This information along with a description of synchronous exits, the use of checkpoint restart, and the use of re-entrant modules, is described in this chapter.

Load module structures, methods of passing control between programs, and the use of associated macro instructions are described in *Supervisor Services and Macro Instructions*.

Residency and Addressing Mode of Programs

The control program ensures that each load module is loaded above or below 16 megabytes (Mb) as appropriate and that it is invoked in the correct addressing mode (24-bit or 31-bit). The placement of the module above or below 16 megabytes depends on the residency mode (RMODE) that you define for the module. Whether a module executes in 24-bit or 31-bit addressing mode depends on the addressing mode (AMODE) that you define for the module.

When a program is executing in 24-bit addressing mode, the system treats both instruction and data addresses as 24-bit addresses. This allows programs executing in 24-bit addressing mode to address 16 megabytes (16,777,216 bytes) of storage. Similarly, when a program is executing in 31-bit addressing mode, the system treats both instructions and data addresses as 31-bit addresses. This allows a program executing in 31-bit addressing mode to address 2 gigabytes (2,147,483,648 bytes or 128x16 megabytes) of storage. *SPL: 31-Bit Addressing* provides detailed information concerning the AMODE and RMODE of modules.

You can define the residency mode and the addressing mode of a program in the source code. Figure 6 shows an example of the definition of the AMODE and RMODE attributes in the source code. This example defines the addressing mode of the load module as 31 and the residence mode of the load module as 24. Therefore, the program will receive control in 31-bit addressing mode and will reside below 16 megabytes in 24-bit addressable storage.

SAMPLE	CSECT	
SAMPLE	AMODE	31
SAMPLE	RMODE	24

Figure 6. Assembler Definition of AMODE/RMODE

The assembler places the AMODE and RMODE in the output object module for use by the linkage editor. The linkage editor passes this information on to the control program through the directory entry for the partitioned data set that contains the load module. You can also specify the AMODE/RMODE attributes of a load module by using linkage editor control cards. See *Linkage Editor* for information concerning these control cards.

Placement of Modules in Storage

The control program uses the RMODE attribute from the directory entry for the module to load the program above or below 16 megabytes. The RMODE attribute can have one of the following values:

- 24-specifies that the program must reside in 24-bit addressable storage
- ANY-specifies that the program can reside anywhere in virtual storage

Addressing Mode

The AMODE attribute, located in the directory entry for the module, specifies the addressing mode of the module. Bit 32 of the program status word (PSW) indicates the addressing mode of the program that is executing. MVS/XA supports programs that execute in either 24-bit or 31-bit addressing mode.

The AMODE attribute can have one of the following values:

- 24-specifies that the program is to receive control in 24-bit addressing mode
- 31-specifies that the program is to receive control in 31-bit addressing mode
- ANY-specifies that the program is to receive control in either 24-bit or 31-bit addressing mode

Information about the addressing mode as it applies to macro instructions can be found in Volume 2 under the topic "Addressing Mode and the Macro Instructions"

Specifying Where the Module is to be Loaded (LOAD)

There are two mutually exclusive ways of loading a module at a specific location:

- Loading at a specific address - You can request that a module be loaded at a designated address by coding the ADDR parameter on the LOAD macro instruction. You must allocate storage for the module in your key.
- Loading into CSA - Using the GLOBAL parameter on the LOAD macro instruction, you can request that the module be loaded into either fixed or pageable CSA. The module is deleted at task termination time.

Synchronous Exits (SYNCH)

In general, the SYNCH macro instruction is used when a control program in supervisor state gives temporary control to a processing program routine (not necessarily running in supervisor state) where the processing program is expected to return control to the supervisor state control program. This facility should be used only by system programmers or other installation-approved personnel. The program to which control is given must be in virtual storage when the macro instruction is issued. To ensure that a program receives control with a program key mask (PKM) consistent with its key, SYNCH processing forms the PKM using the default key in the TCB along with the key specified by the KEYADDR parameter. If the KEYMASK parameter is coded, the PKM formed thus far is ORed with the specified keymask.

When the processing program returns control, the supervisor state bit, the PSW key bits, the system mask bits, and the program mask bits of the program status word are restored to the settings they had before execution of the SYNCH macro instruction.

The SYNCH macro instruction is similar to the BALR instruction in that you can use register 15 for the entry point address.

SYNCH processing does not save or restore registers when control is returned to the caller unless RESTORE = YES is specified. If you specify RESTORE = NO explicitly or by default, the register contents are unpredictable. When an authorized program uses SYNCH to invoke an exit in an unauthorized program, the general registers returned from the exit might not contain expected data or correct addresses. Therefore, the authorized program must save the registers in a protected save area and then restore them, or validate the contents of the returned registers, or code RESTORE = YES.

Label processing as a result of an OPEN macro instruction is an example of the use of the SYNCH macro instruction. Label processing might proceed to a point at which a user's processing program indicates that it wants or needs private processing. The control program's open routine would then issue a SYNCH macro instruction giving the address of the subroutine required for the user's private label processing.

Using Checkpoint/Restart

When issuing checkpoints and then restarting a task, the restarted task must request control of all resources required to continue processing. Resources are not automatically returned to the task upon restart.

You can use the checkpoint/restart facility with the following restrictions:

- A routine that is restricted from issuing SVCs (for example, a routine running in SRB, disabled, or cross memory mode) is also restricted from establishing checkpoints because programmer-designated checkpoints require the use of the checkpoint SVC.
- An exit routine other than the end-of-volume exit routine cannot request a checkpoint.
- A routine invoked by a program call (PC) cannot request checkpoints because the system environment might be different at the time of the restart from what it was at the time of the checkpoint. This could lead to unpredictable results on the return linkage (PT).
- A routine with a PCLINK STACK request outstanding cannot establish a checkpoint.

- Routines that use both PC/AUTH facilities and checkpoint/restart must reestablish their PC/AUTH environment at restart time. In addition, they must not use any PC/AUTH data (for example, a PC number) that was obtained before the restart.
- Subsystems that use the TCB subsystem affinity service cannot issue checkpoints. This is because the subsystem affinity table (SSAT) index values might change from one system initialization to another.

For additional information concerning the restrictions and use of the checkpoint/restart facility see *MVS/Extended Architecture Checkpoint/Restart User's Guide*, GC26-4012.

Using Re-entrant Modules

When link editing modules as re-entrant, be sure that all the modules and the macro instructions they call are re-entrant. In a multiprocessing system this is important because:

- Two tasks in the same address space making use of the module might cause the module to be executed simultaneously on two different processors.
- Asynchronous appendages can operate on one processor simultaneously with an associated task on another processor.
- Enabled recovery routines can execute on any processor, not necessarily on the one on which the error was detected.

The CSECTs must be unchanged during execution or their critical sections must be explicitly serialized. The general method for ensuring re-entrance of macro instructions is to use the LIST and EXECUTE forms of the macro instructions with a dynamically acquired parameter list.

Serialization

Planning is required when more than one program uses a serially reusable resource. A serially reusable resource is a resource that can be used by another program after the current use has been concluded; that is, a resource that should not be used or modified by more than one program within a given span of processing. Planning is also required when portions of some tasks depend on the completion of events in other tasks.

This chapter discusses some of the services available to control resources, and thus to help you plan ahead for a more efficient installation. The services discussed include:

- Locking (SETLOCK macro instruction)
- Must-complete function (ENQ and DEQ macro instructions)
- Shared DASD (RESERVE and EXTRACT macro instructions)
- Event completion (POST, SPOST, and EVENTS macro instructions)

Global resource serialization (ENQ, DEQ, RESERVE, or GQSCAN macro instructions) is another form of serialization available to an installation. This topic appears in *Supervisor Services and Macro Instructions* and *Planning: Global Resource Serialization*.

When Resource Serialization Is Needed

Resource serialization is used to prevent a program from altering the content or status of a resource while another program is using that resource or is dependent on the content or status of that resource remaining unchanged for a given span of processing. For example, resource serialization prevents a program from issuing an SVC and changing the content of a control block while another SVC is using that control block.

Serialization Requirements

It is necessary to determine and keep track of resources that must be serialized and the routines that access such resources. The only safe method of serialization is one of the following: ENQ/DEQ, WAIT/POST/EVENTS, SUSPEND/RESUME, locking at the TCB level, CS (compare and swap instruction), CDS (compare double and swap instruction), and TS (test and set instruction). Such forms of serialization are required in the following cases:

- Scanning of the command input buffer (CIB) chain. You could use the QEDIT macro instruction to manipulate the CIB chain.
- Using data in subpools shared between tasks.
- Using data referenced by more than one task. (For example, attached tasks can execute at the same time as the attaching task on different processors.)

- Referencing system control block fields that dynamically change after IPL. The serialization technique in this case must match that used by the system. (See the *Debugging Handbook* for information concerning the serialization requirements for a particular system control block.) Also, bits within a byte all require the same serialization technique.
- Accessing of data sets shared between tasks in the same address space, if the tasks update the data and if the access method is not VSAM or BDAM.
- Referencing any common data between an ESTAE exit and asynchronous exits, if ESTAE with ASYNCH= YES is issued.

Locking

A locking mechanism serializes access to resources. This locking technique is only effective, however, if all programs that depend on a resource use the same locking mechanism. Each type of serially reusable resource is assigned a lock. The lock manager controls a hierarchical locking structure with multiple types of locks to synchronize the use of serially reusable resources. The lock manager also handles all functions related to the locks. These functions include obtaining or releasing locks and checking the status of a particular lock on a processor. Use of the lock manager is restricted to key 0 programs running in supervisor state. This prevents unauthorized problem programs from interfering with the system serialization process.

Categories of Locks

There are two categories of locks:

- Global locks -- protect serially reusable resources related to more than one address space. (For example, a unit control block is protected by a global lock because it relates to the entire system. Also, a system-related GETMAIN for a global subpool requires a global lock.)
- Local locks -- protect the resources assigned to a particular address space. When the local lock is held for an address space, the owner of the lock has the right to manipulate the queues and control blocks associated with that address space. (For example, an address space-related GETMAIN for a user subpool requires a local lock.)

All of the locks described in Figure 7, with the exception of the LOCAL and CML locks, are global locks. These global locks provide system-wide services or use control information in the common area and must serialize across address spaces. The local level locks, on the other hand, do not serialize across address spaces, but serialize functions executing within the address space. Figure 7 summarizes the characteristics of MVS/XA locks.

lock	global	local	spin	suspend	single	multiple (class)	shared/exclusive
RSMGL	X		X			X	
VSMFIX	X		X		X		
ASM	X		X			X	
ASMGL	X		X			X	
RSMST	X		X			X	
RSMCM	X		X			X	
RSMXM	X		X			X	
RSMAD	X		X			X	
RSM	X		X		X		X
VSMPAG	X		X		X		
DISP	X		X		X		
SALLO	X		X		X		
IOSYNCH	X		X			X	
IOSUCB	X		X			X	
SRM	X		X		X		
TRACE	X		X		X		X
CPU	X		X			X	
CMS	X			X	X		
CMSEQDQ	X			X	X		
CMSSMF	X			X	X		
CML		X		X		X	
LOCAL		X		X		X	

Note: The CPU lock has no real hierarchy except that once a user obtains it, the user cannot obtain a suspend lock; a user can obtain the CPU lock while holding any spin lock. The CPU lock could be considered a pseudo spin lock. It could also be considered multiple because there is one per processor and any number of requestors can hold it at the same time.

Figure 7. Summary of Locking Characteristics

Types of Locks

The type of lock determines what happens when a function on one processor in an MP system makes an unconditional request for a lock that is held by another unit of work on another processor. There are two major types of locks: spin and suspend. Shared/exclusive locks are a category of spin locks. The CPU lock is in a category by itself but could be considered a pseudo spin lock. Descriptions of these types of locks follow:

- Spin locks -- prevent the requesting function on one processor from doing any work until the lock is freed on another processor. The lock manager enters a loop that keeps testing the lock until it is released on the owning processor. As soon as the lock is free, the lock manager spinning on the requesting processor attempts to obtain the lock for the requesting function. As long as a spin lock (except for shared/exclusive locks and the CPU lock) is held by a function executing on a processor, the ID of that processor is in the lockword. Once the lock is released by the owning function, the lockword is cleared.
- Shared/exclusive locks--serialize the reading or updating of a global resource. More than one processor can own a shared/exclusive lock as shared at one time; only one processor can own a shared/exclusive lock as exclusive at one time.

Code executing under a shared/exclusive lock is physically disabled. Figure 8 summarizes the results of an unconditional request for a shared/exclusive lock that another processor holds. In general, the lock manager gives processors spinning for exclusive ownership of a shared/exclusive lock priority over processors spinning for shared ownership.

Note: The contents of the lockword for a shared/exclusive lock is different from the contents of a spin lockword. In particular, the shared/exclusive lockword does not contain a logical processor ID. For more information about the contents of the lockword for a shared/exclusive lock, see *Diagnostic Techniques*.

Type of Request	How Held by Owning Processor	Results
Shared	Shared	Obtain shared ownership.
Shared	Exclusive	Spin on the lock until the exclusive owner releases it.
Exclusive	Shared	Spin on lock until all shared owners release it. Set the exclusive-pending-request bit in the lockword.
Exclusive	Exclusive	Spin on lock until the exclusive owner releases it. Set the exclusive-pending-request bit in the lockword.

Figure 8. Requests for Shared/Exclusive Locks

- CPU lock--provides system recognized (legal) disablement for units of work (requestors) on a processor level. System recognized (legal) disablement is defined as holding a spin lock or having a super bit set in the PSASUPER field of the PSA. While a requestor holds the CPU lock, the requestor is physically disabled for I/O and external interruptions.

Multiple units of work on the same processor can own the CPU lock. The CPU lockword (in the PSA) contains the cumulative count of requestors who hold the CPU lock. Obtaining the CPU lock increases the ownership count of the CPU lock by 1; releasing the CPU lock decreases the ownership count by 1.

Note: The CPU lockword does not contain a processor ID. See *Diagnostic Techniques* for additional details about the CPU lockword; see Figure 7 for a description of the "hierarchy" of the CPU lock and its other attributes.

- Suspend locks -- prevent the requesting program from doing work until the lock is available, but allow the processor to continue doing other work. The requestor is suspended and other work may be dispatched on that processor. Upon release of the lock, the suspended requestor is given control with the lock or is redispached to retry the lock obtain.

Examples of Lock Types

All of the locks described in Figure 7 with the exception of the CPU, LOCAL, cross memory local (CML), and cross memory services (CMS) locks, are spin locks. The CPU lock can be considered a pseudo spin lock. The LOCAL, CML, CMS, CMSSMF, and CMSEQDQ locks are suspend locks. Their owners receive control enabled and can be interrupted to run higher priority work. If there is another request for the lock while it is held, the requestor is suspended and other work is dispatched. The local lockword contains the ID of the processor on which its owner is dispatched or an indication that the owner is suspended or interrupted. The CMS lockword contains the ASCB address of the locally locked address space that owns the lock. Special IDs are placed in the local lockword whenever the owner of the local lock is

not currently executing on a processor because of an interruption or suspension. See *Diagnostic Techniques* for a description of the contents of a local suspend lockword.

Note: CML (cross memory local) lock means the local lock of an address space other than the home address space. LOCAL lock means the local lock of the home address space. When written in lower case, local lock means any local-level lock, either the LOCAL or a CML lock.

The CMS lock is an enabled global lock for the following reasons:

- Because disabled page faults are not allowed in the system, some global functions need a lock that does not require the functions to fix all their code and control blocks.
- Some functions require significant amounts of time under the lock and could impact the responsiveness of the system. By running these functions enabled under the lock, responsiveness is retained at the expense of some increased contention for the lock.

The other global locks are disabled spin locks because the functions that run under the locks are of short duration and cannot tolerate interruptions. The cost in system overhead to perform the status saving necessary to accept interruptions and allow switching would offset the gain in responsiveness. Also, the more frequently used functions (for example, IOS interruption handler, dispatcher, and storage manager) perform interruption handling and task switching, and have to remain disabled.

If a lock is unconditionally requested, the lock is unconditionally obtained. If the lock is conditionally requested, the requestor is given the lock if it is available; if the lock is unavailable, control is returned to the caller without the lock. (See the COND and UNCOND parameters on the SETLOCK macro instruction.)

Classes of Locks

There are two classes of locks:

- Single locks -- Only one lock exists at a given level of the locking hierarchy. Because there is one lock at a given level, SETLOCK requests for single locks cannot specify the ADDR keyword parameter.
- Multiple locks (commonly referred to as class locks) -- More than one lockword exists at a given level of the locking hierarchy. Because of this, SETLOCK requests for multiple locks must specify the ADDR keyword parameter.

The locks provided in MVS/XA in hierarchical order are:

- RSMGL (real storage management global lock) -- serializes access to all RSM global queues and resources.
- VSMFIX (virtual storage management lock) -- serializes the common area subpools (subpools 226, 227, 228, 231, 239, 241, and 245).
- ASM (auxiliary storage management lock) -- serializes ASM resources on an address space level.
- ASMGL (auxiliary storage management global lock) -- serializes ASM resources on a global level.

- RSMST (real storage management steal lock) -- serializes RSM control blocks on an address space level when it is not known which address space locks are currently held.
- RSMCM (real storage management common lock) -- serializes RSM resources in the common area (such as page table entries, the pageable-frame queue, and the fixed-frame queue).
- RSMXM (real storage management cross memory lock) -- serializes RSM control blocks on an address space level when serialization to a second address space is necessary.
- RSMAD (real storage management address space lock) -- serializes RSM control blocks on an address space level.
- RSM (real storage management lock) -- serializes RSM execution and RSM resources.
- VSMPAG (virtual storage management lock) -- serializes the use of common VSM work area for pageable subpools.
- DISP (global dispatcher lock) -- serializes the use of resources such as address space vector table (ASVT) updating and changes to the address space control block (ASCB) dispatching queue.
- SALLOC (space allocation lock) -- serializes the external receiving routines that enable a processor for either an emergency signal or a malfunction alert.
- IOSYNCH (IOS synchronization lock) -- serializes global IOS functions by means of an IOSYNCH lock table.
- IOSUCB (IOS unit control block lock) -- serializes access and updates to the unit control blocks (UCB)s. There is one lock for each UCB.
- SRM (system resources manager lock) -- serializes use of the SRM control algorithms and associated data.
- TRACE (TRACE lock) -- serializes the system trace buffer structure.
- CPU (processor lock) -- serializes on the processor level, providing system recognized (legal) disablement.
- CMS (general cross memory services lock) -- serializes on more than one address space where this serialization is not provided by one or more of the other global locks.
- CMSEQDQ (ENQ/DEQ cross memory services lock) -- serializes ENQ/DEQ functions and the use of ENQ/DEQ control blocks.
- CMSSMF (SMF cross memory services lock) -- serializes SMF functions and the use of SMF control blocks.
- CML (cross memory local lock) -- serializes resources in an address space other than the home address space.
- Local storage lock (LOCAL) -- serializes functions and storage, used by the local supervisor within an address space. There is one lock for each address space.

You must hold a local lock, either CML or LOCAL, when requesting the CMS, CMSEQDQ, or CMSSMF lock. You cannot release the local lock while holding a cross memory services lock. You need not hold all locks in the hierarchy up to the highest lock needed. Hold only locks that you need.

Locking Hierarchy

The locks are arranged in a hierarchy to prevent a deadlock between functions on the processor(s). An example of a deadlock between functions would be:

- Function A holding the SRM lock and requesting the DISP lock on processor 0
- Function B holding the DISP lock on processor 1 and requesting the SRM lock currently held on processor 0

A function on a processor can request unconditionally only those locks that are higher in the hierarchy than the locks it currently holds, thus preventing deadlocks. The hierarchy is shown in Figure 7, with RSMGL being the highest lock.

The CPU lock has no hierarchical relationship with other spin locks. The CPU lock can be obtained while other spin locks are held; other spin locks can be obtained (in their hierarchical sequence) while the CPU lock is held. The CPU lock is, however, higher in hierarchical order than any of the suspend locks, therefore once you obtain the CPU lock, you cannot obtain any suspend lock. The cross memory services locks (CMS, CMSEQDQ, and CMSSMF) are equal to each other in the hierarchy. The CML and LOCAL locks are also equal to each other in the hierarchy.

With the exception of cross memory services locks, a processor can hold only one lock at the same level of hierarchy. Therefore, if a processor holds an IOSUCB lock, it may not request a different IOSUCB lock at a different address. If a processor holds one cross memory services lock, it can not request another cross memory services lock. However, a processor can hold all cross memory services locks if it unconditionally requests them simultaneously. If the locks are requested at the same time, they must be released at the same time. It is not recommended that all cross memory services locks be held at the same time because it will degrade performance.

Using the Same Lockword for Class Locks at Different Levels

To simplify lockword management, a user can provide the same lockword for certain class locks at different levels of the locking hierarchy (for example, the RSMST, RSMCM, RSMXM, and RSMAD locks). However, the lockword can only represent one lock at any given time.

For example, the RSMXM lock held at location 1000 on processor 0 creates two kinds of locking restrictions:

- No other lock (for example, RSMAD) can be obtained at location 1000 on processor 0 or any other processor, until the RSMXM lock is released (however, another lock, like the RSMAD lock, can be obtained at another location on processor 0).
- An RSMXM lock at another location cannot be obtained on processor 0 until the RSMXM lock at location 1000 is released (however, the RSMXM lock at another location on another processor can be obtained).

The lock manager prevents an interlock by detecting the attempt to simultaneously obtain multiple locks using the same lockword or lock location.

For conditional requests using the same lockword, the lock manager supplies return codes that the user can check. For unconditional requests, if the caller holds the lockword for a different level lock, the lock manager abnormally terminates the caller with an 073 ABEND code. The return codes are described with the syntax of the SETLOCK macro instruction in Volume 2.

There is another situation in which an interlock could occur. This type of interlock is not prevented by the lock manager, but must be solved in the program by using internal hierarchy rules. It involves using the same lockword for a class of locks. For example, if task A and task B are executing on different processors an interlock could occur if:

Task A holds the RSMAD lock located at location 1000 and requests the RSMXM lock located at location 2000 while task B holds the RSMAD lock located at location 2000 and requests the RSMXM lock located at location 1000.

CML Lock Considerations

The cross memory local lock (CML) is provided to allow cross memory services to serialize resources in an address space that might not be the home address space. It has the same attributes as the LOCAL lock. (The LOCAL lock refers only to the home address space pointed to by PSAAOLD.) The owner of a CML lock can be suspended for the same reasons as the owner of the LOCAL lock, such as CMS lock suspension or page fault suspension.

In a multi-tasking environment, it is possible for more than one task or SRB in an address space to obtain a local level lock. For example, task A might own the LOCAL lock of its address space while task B in the same address space owns the CML lock of address space C.

To prevent possible system deadlocks, only one lock at the local level can be held at one time by a unit of work. If a CML lock is requested while owning the LOCAL lock, the requestor will be abended. The same is true if the LOCAL lock is requested while owning a CML lock.

Either a CML lock or the LOCAL lock must be held to request one or all of the cross memory services locks (CMS, CMSEQDQ, or CMSSMF).

The requestor of a CML lock must have authority to access the specified address space prior to the lock request. This is accomplished by setting the primary or secondary address space to that specified on the lock request. The specified address space must be non-swappable prior to the obtain request.

Note: The CML lock of the master scheduler address space cannot be obtained. The master scheduler address space lock can only be obtained as a LOCAL lock.

Obtaining, Releasing, and Testing Locks (SETLOCK)

Use the SETLOCK macro instruction to obtain, release, or test a specified lock or set of locks (using the OBTAIN, RELEASE, and TEST parameters). Users can also obtain the current CPU lock use count for a processor and determine whether a processor holds a spin lock higher in the locking hierarchy than a specified lock. To use SETLOCK, you must be executing in supervisor state with protection key 0. Users of SETLOCK can also be executing in SRB mode, in cross memory mode, as an extension of the interrupt handlers, or as a system service such as the MVS/XA dispatcher.

Disabled/Enabled State for Obtain

When a global spin type lock is successfully obtained, control returns to the caller with the processor disabled for I/O and external interruptions.

When a suspend type lock is successfully obtained via an unconditional request, control returns to the caller with the processor enabled for I/O and external interruptions.

For an unsuccessful conditional request of a spin lock, control returns to the caller disabled only if the caller was disabled on entry. Otherwise, control returns enabled for I/O and external interruptions. If a disabled caller unconditionally requests a suspend type lock that is not immediately available, the caller is abnormally terminated.

Disabled/Enabled State for Release

When a global spin type lock is released, control returns to the caller enabled for I/O and external interruptions unless at least one of the following is true:

- Another global spin lock is held
- A disabled supervisor indicator (PSASUPER) is on
- The DISABLED parameter was specified

If one of the above is true, control returns to the caller disabled for I/O and external interruptions.

When a suspend type lock is released, control returns disabled for I/O and external interruptions if the caller was disabled on entry. Otherwise control returns to the caller enabled for I/O and external interruptions.

For a release request via the SPIN, ALL, or (reg) parameters, the final state is the same as that which would have existed had the locks been released one at a time.

Altering the Dispatching Queue (INTSECT)

The intersect function is the serialization mechanism that the dispatcher and control program functions use to alter the dispatching queues. The LOCAL and dispatcher locks are used in conjunction with the intersect function. Intersect serialization is only between the requestor of the intersect and the dispatcher. The requesting routine must hold the LOCAL or dispatcher lock for serialization with other routines.

A routine can intersect on either the local or global level. The LOCAL lock is required for obtaining the local intersect; it also ensures the proper serialization with other routines requesting the local intersect. The local intersect ensures serialization of an address space with the dispatcher and serialization of routines that modify the TCB dispatching queue or TCB dispatchability. Similarly, the dispatcher lock is required for routines requesting the global intersect. The global intersect ensures serialization of dispatcher functions on a global level.

Using the Must-Complete Function (ENQ/DEQ)

System routines (routines operating under a storage protection key of zero) often update and/or manipulate system resources such as system data sets, control blocks, and queues. These resources contain information critical to continued operation of the system. The task requesting this serialization must successfully complete its processing of the resource. Otherwise, the resource might be left incomplete or might contain erroneous information.

The ENQ service routine ensures that a routine queued on a critical resource(s) can complete processing of the resource(s) without interruptions leading to termination. ENQ places other tasks in a nondispatchable state until the requesting task -- the task issuing an ENQ macro instruction with the set must-complete (SMC) parameter -- has completed its operations on the resource. The requesting task releases the resource and terminates the must-complete condition by issuing a DEQ macro instruction with the reset must-complete (RMC) parameter.

Because the must-complete function serializes operations to some extent, its use should be minimized -- use the function only in a routine that processes system data whose validity must be ensured. Just as the ENQ function serializes use of a resource requested by many different tasks, the must-complete function serializes execution of tasks.

Characteristics of the Must-Complete Function

The must-complete function can be used only at the step level, where only the current problem program task in an address space is allowed to execute. All other problem program tasks, and the initiator task, are made non-dispatchable.

When the must-complete function is requested, the requesting task is marked in "must complete mode" when the resource(s) queued upon are available. All asynchronous exits from the requesting task are deferred. The initiator and all other tasks in the job step are set nondispatchable. Tasks external to the requesting task are prevented from initiating procedures that will cause termination of the requesting task. Other external events, such as a CANCEL command issued by an operator, or a job step time expiration, are also prevented from terminating the requesting task.

The failure of a task that owns a must-complete resource results in the abnormal termination of the entire job step. The programmer and the operator receive a message stating that the failure occurred while the step was in must-complete mode.

Programming Notes

1. All data used by a routine that is to operate in the must-complete mode should be checked for validity to ensure against a program-check interruption.
2. If a routine that is already in the must-complete mode calls another routine, the called routine also operates in the must-complete mode. An internal count is maintained of the number of SMC requests; an equivalent number of RMC requests is required to reset the must-complete function.

3. Interlock conditions can arise with the use of the ENQ function. Additionally, an interlock might occur if a routine issues an ENQ macro instruction while in the must complete mode. Also, a task that is non-dispatchable, because of a must-complete request, might already be queued on the requested resource. In this case, an enabled wait occurs. An enabled wait can be broken by an operator's action (such as the use of the FORCE command).
4. The macro instructions ATTACH, LINK, LOAD, and XCTL should not be used, (unless extreme care is taken) by a routine operating in the must-complete mode. An interlock condition results if a serially reusable routine requested by one of these macro instructions either has been requested by one of the tasks made nondispatchable by the use of the SMC parameter or was requested by another task and has been only partially fetched.
5. The time a routine is in the must-complete mode should be kept as short as possible -- enter at the last moment and leave as soon as possible. One suggested way is to:
 - a. ENQ (on desired resource(s))
 - b. ENQ (on same resource(s)),RET = HAVE,SMC = STEP

Step (a) gets the resource(s) without putting the routine into the must-complete mode. Later, when appropriate, issue the ENQ with the must-complete request (Step b). Issue a DEQ macro instruction to terminate the must complete mode as soon as processing is finished. Tasks set nondispatchable by the corresponding ENQ macro instruction are made dispatchable and asynchronous exits from the requesting task are enabled.

Limiting Global Resource Serialization Requests

Global resource serialization allows an installation to share symbolically-referenced resources between units of work. (*Planning: Global Resource Serialization* explains the function and use of global resource serialization.) A global resource serialization request is an ENQ or RESERVE request that causes an element to be added to any queue in the global resource serialization queue area.

GQSCAN uses the same control blocks as ENQ and RESERVE to obtain the status of resources and requestors of resources during resumption processing. (The GQSCAN macro is described in the section 'Measuring System Characteristics' and in Volume 2.) In order to prevent any one job, started task, or TSO user from generating too many concurrent requests, global resource serialization limits the number of global resource serialization requests in each address space.

Global resource serialization counts the number of ENQ/RESERVE requests and the number of pending GQSCAN requests issued by all TCBs in each address space. Each time a user issues an ENQ/RESERVE, global resource serialization increases the count in that address space by 1 for each resource name and decreases the count by 1 when a user in that address space issues a DEQ. Similarly, when a user issues a GQSCAN request, global resource serialization increases the count in that address space by 1 and decreases the count by 1 when the scan is completed (if resumption is requested).

Global resource serialization compares the computed count of requests to a threshold value (4096) stored in the GVTCREQ field of the global resource serialization vector table (GVT). (See *SPL: System Modifications* for a description of how to change the threshold value.) When

the computed count reaches the threshold value, global resource serialization processes subsequent requests as follows:

- ENQ/RESERVE requests from unauthorized callers are rejected; unconditional requests from these callers are abended and conditional requests receive a return code of X'18'.
- ENQ/RESERVE requests from authorized callers are not rejected until the count exceeds the threshold value by a tolerance value. This higher limit is stored in the GVTCREQA field of the GVT. The tolerance provided by the system is 15, but system programmers can change this value. (See *SPL: System Modifications* for a description of how to change the limit for authorized callers.) This means that an additional 15 concurrent ENQ/RESERVE requests are accepted from authorized callers. This is done to allow for normal termination and to permit error recovery routines to obtain the resources that they need. Once the computed count exceeds the limit for authorized callers, subsequent requests are rejected in the same way as requests from unauthorized callers.
- GQSCAN requests that do not fit into the caller's buffer receive a return code of X'14'; these requests are not queued and a TOKEN is not provided.

ENQ/RESERVE requests from authorized callers use the MASID and MTCB parameters to allow a further conditional control of a resource. One task issues an ENQ or RESERVE for a resource specifying a matching ASID; if the issuing task does not receive control, it is notified whether the matching task has control (which allows the issuing task to use the resource even though it could not acquire the resource itself). This process requires serialization between the issuing and requesting tasks.

Shared Direct Access Storage Devices (Shared DASD)

The shared DASD facility allows systems to share direct access storage devices. Systems can share common data and consolidate data when necessary. No change to existing records, data sets, or volumes is necessary to use the facility. However, reorganization of volumes might be desirable to achieve better performance.

Exercise careful planning in accessing shared data sets or shared data areas. Data integrity can not be assured without proper intersystem communication. This topic, as it relates to macro instructions, is discussed further under "Macro Instructions Used with Shared DASD." Similarly, appropriate security procedures must be performed on each of the multiple systems involved in the sharing of DASD before data can be regarded as secure. Data sets that are intended to be protected via passwords or RACF should be initially protected on each system before sensitive data is placed in them. This topic, as it refers to password protection, is discussed further under "System Configuration."

Devices that Can be Shared

The following control units and devices are supported by the shared DASD option:

- IBM 2835 Storage Control Unit with two-channel switch -- IBM 2305 Fixed Head Storage Facility.
- IBM 3830 Storage Control Unit with two-channel switch -- IBM 3330 Series Disk Storage Drive. The IBM 3330, 3340/3344, 3350 Series devices may also be configured for shared use via the string switch feature.

- IBM 3880 Storage Control Units Models 2 and 3 -- IBM 3380 Direct Access Storage Facility.
- IBM 3880 Storage Control Units Models 1 and 2 -- IBM 3375 Direct Access Storage Facility.

Alternate channels to a device from any one system can only be specified for the IBM 3330, 3340/3344, 3350 Series Storage Unit, the 3375 Direct Access Storage Device, and the 3380 Direct Access Storage Device.

Volume/Device Status

The shared DASD facility requires that certain combinations of volume characteristics and device status be in effect for shared volumes of devices. Figure 9 shows the combinations that must be in effect for a volume or device:

System A	Systems B, C, D
Permanently resident	Permanently resident
Reserved	Reserved
Removable	Offline - Non-JES3 devices
Removable	Removable - JES3 - managed devices
Offline	Removable, reserved, or permanently resident (In JES2, if a device is removable in one system, it must be offline in all others.)

Figure 9. Valid Volume Characteristic and Device Status Combinations

If a volume or device is marked removable on any one system, the device must be either in offline status or removable status on all other systems. The mount characteristic of a volume and/or the status of a device can be changed on one system as long as the resulting combination is valid for other systems sharing the volume or device. No other combinations of volume characteristics and device status are supported.

System Configuration

Operating system configurations do not have to be identical to share a data set. The only additional equipment needed for the Shared DASD option is either a two-channel switch or a 2844 Auxiliary Control unit. The user must also observe certain restrictions about the data sets that are shared. The following data sets cannot be shared:

```

Master catalog
PASSWORD          SYS1.MANx
SYS1.DCMLIB       SYS1.NUCLEUS
SYS1.DUMPxx       SYS1.PAGExx
SYS1.LOGREC       SYS1.STGINDEX
SYS1.LPALIB       SYS1.SVCLIB

```

Because the system does not provide for the sharing of the PASSWORD data set, the PASSWORD data set for each system must contain password records for all protected data sets. Where independent computing systems share common DASD resources, individual installations must ensure that the PASSWORD data set contains records for all protected data sets for each system sharing the DASD. For further details regarding password protection on shared DASD, see *MVS/Extended Architecture System Data Administration*.

Volume Handling

Volume handling with the shared DASD option must be clearly defined because operator actions on the sharing system must be performed in parallel. The following rules should be in effect when using the shared DASD option:

- Operators should initiate all shared volume mounting and demounting operations. The system will dynamically allocate devices unless they are in reserved or permanently resident status, and only the former can be changed by the operator.
- Mounting and demounting operations must be done in parallel on all sharing systems. A VARY OFFLINE must be issued on all systems before a device can be dismounted.
- Valid combinations of volume mount characteristics and device status for all sharing systems must be maintained. To IPL a system, a valid combination must be established before device allocation can proceed. This valid combination is established either by specifying mount characteristics of shared devices in VATLST, or varying all shareable devices offline before issuing START commands and then following parallel mount procedures.

Macro Instructions Used with Shared DASD (RESERVE, EXTRACT)

You can use the RESERVE, ENQ, DEQ, and EXTRACT macro instructions when working with shared DASD. The following paragraphs describe the use of these macro instructions in relation to shared DASD.

- The RESERVE macro instruction reserves a device (identified by its UCB address and symbolic resource name) for use by a particular system. Each task that needs exclusive use of a device must issue the RESERVE macro. When a task issues a RESERVE for a particular device, the system increases the count of outstanding reserve requests (located in the UCBSQC field of the UCB) for that device. If MVS starts I/O for that device while the count is non-zero, it precedes the channel program with a RESERVE channel control word (CCW) that reserves the device to the system that executed the RESERVE CCW.

Notes:

1. *The set-must-complete (SMC) parameter of the ENQ macro instruction can also be used with RESERVE.*
 2. *If a check point restart occurs when a RESERVE is in effect for devices, the system does not restore the RESERVE; the user's program must reissue the RESERVE.*
- The initiator, allocation, and direct access device storage management (DADSM) components of MVS use an ENQ with a major name of SYSDSN to serialize access to datasets, whether these datasets exist on shared DASD or not. The use of global resource serialization in a shared DASD environment allows DADSM to serialize the use of DASD space by all systems sharing a volume so that data integrity is guaranteed. The SYSDSN major name must be a resource known to all systems in the shared DASD environment. Refer to the topic Setting Up Resources Name Lists in the "Global Resource Serialization" section for more information about how to make SYSDSN known to all systems.

- When the task issues a DEQ for the resource named on the RESERVE macro, the system reduces the count in the UCB. When this count reaches zero, the system starts a channel program, consisting of a RELEASE CCW, to free the device.
- If global resource serialization is active, ENQ and DEQ, with SCOPE=SYSTEMS specified, can serialize on a particular shared DASD data set without reserving the entire device. See *Planning: Global Resource Serialization* for details.
- The EXTRACT macro instruction obtains the address of the task input/output table (TIOT) from which the UCB address can be obtained. “Finding the UCB Address for the RESERVE Macro” explains this procedure. EXTRACT provides information, it does not actually serialize a resource.

Releasing Devices

The DEQ macro instruction is used with RESERVE just as it is used with ENQ. It must describe the same resource as the RESERVE and its scope must be stated as SYSTEMS; however, the UCB=pointer address parameter is not required. If the DEQ macro instruction is not issued by a task that has previously reserved a device, the system frees the device when the task is terminated.

Preventing Interlocks

The greater the number of device reservations occurring in each sharing system, the greater the chance of interlocks occurring. Allowing each task to reserve only one device minimizes the exposure to interlock. The system cannot detect interlocks caused by a program's use of the RESERVE macro instruction and therefore, enabled wait states can occur on the system. Global resource serialization can also be used to prevent interlocks by suppressing the hardware RESERVE or simply issuing a global ENQ to serialize the resource. See *Planning: Global Resource Serialization* for additional information on this topic.

Volume Assignment

Because exclusive control is by device, not by data set, consider which data sets reside on the same volume. In this environment it is quite possible for two tasks in two different systems -- processing four different data sets on two shared volumes -- to become interlocked. (If global resource serialization is active and RESERVEs are converted to global ENQs, an interlock does not occur.) For example, as shown in Figure 10, data sets A and B reside on device 124, and data sets D and E reside on device 236. A task in system 1 reserves device 124 in order to use data set A; a task in system 2 reserves device 236 in order to use data set D. Now the task in system 1 tries to reserve device 236 in order to use data set E and the task in system 2 tries to reserve device 124 in order to use data set B. Neither can ever regain control, and neither will complete normally. When the system has job step time limits, the task, or tasks, in the interlock will be abnormally terminated when the time limit expires. Moreover, an interlock could mushroom, encompassing new tasks as these tasks try to reserve the devices involved in the existing interlock.

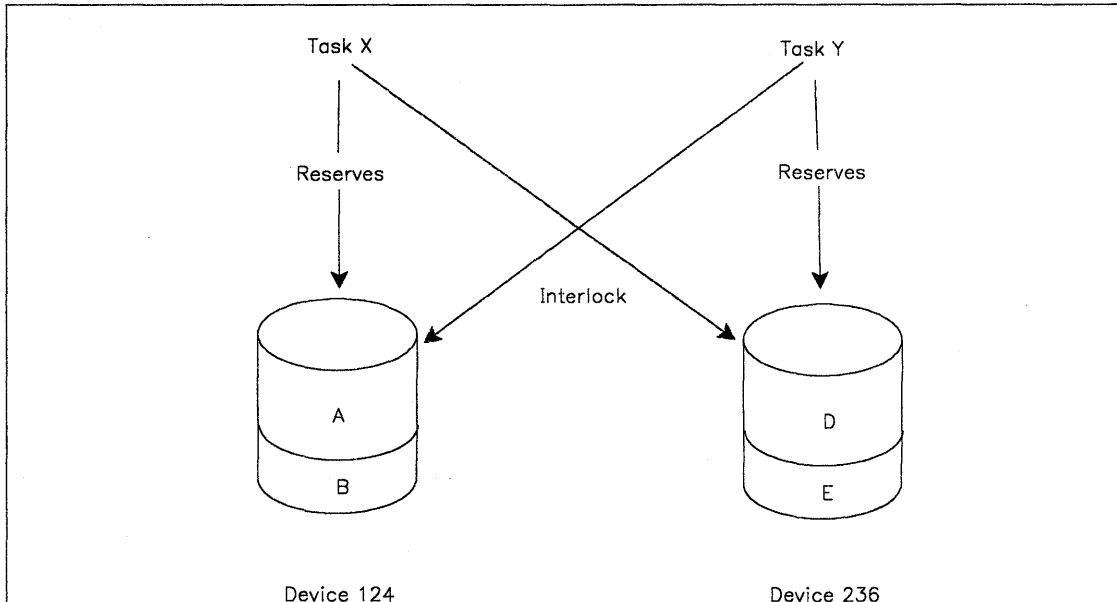


Figure 10. Example of an Interlock Environment

Program Libraries

When assigning program libraries to shared volumes, take care to avoid interlock. For example, SYS1.LINKLIB for system 1 resides on volume X, while SYS1.LINKLIB for system 2 resides on volume Y. A task in system 1 invokes a direct access device space management function for volume Y, causing that device to be reserved. A task in system 2 invokes a similar function for volume X, reserving that device. However, each load module transfers to another load module via XCTL. Since the SYS1.LINKLIB for each system resides on a volume reserved by the other system, the XCTL macro instruction cannot complete the operation. An interlock occurs; because no access to SYS1.LINKLIB is possible, both systems will eventually enter an enabled wait state. (If global resource serialization is active and RESERVEs are converted to global ENQs, an interlock does not occur).

Using Different Serialization Techniques for the Same Volume

A task interlock can occur within a global resource serialization complex when two tasks reserve the same volume and some of the RESERVEs specify resource names that suppress the hardware reserve while other RESERVEs are hardware reserves that lock up the entire volume. The UCB count of outstanding reserves for that volume is manipulated only for the hardware RESERVEs.

If you code a RESERVE macro, the hardware reserve is suppressed when the resource name appears in the reserve conversion resource name list (RNL). See *Planning: Global Resource Serialization* for additional information about RNLs and about preventing interlocks.

Finding the UCB Address for the RESERVE Macro

This topic explains procedures for finding the UCB address for use by the RESERVE macro instruction; it also shows a sample assembler language subroutine that issues the RESERVE and DEQ macro instructions and can be called by routines written in higher level languages.

Job management routines construct the TIOT, which resides in virtual storage during step execution. The TIOT consists of one or more DD entries, each of which represents a data set defined by a DD statement for the jobstep. Each entry includes the DD name. Associated with each DD entry is the UCB address of the associated device. In order to find the UCB address, you must locate the DD entry in the TIOT corresponding to the DD name of the data set for which the RESERVE macro instruction is to be issued.

Providing the Unit Control Block Address to RESERVE: Use the EXTRACT macro instruction to obtain information from the task control block (TCB). The address of the TIOT can be obtained from the TCB in response to an EXTRACT macro instruction. Before issuing an EXTRACT macro instruction, set up an answer area to receive the requested information. One full word is required for each item to be provided by the control program. If you want to obtain the TIOT address, you must specify `FIELDS=TIOT` in the EXTRACT macro instruction.

The control program returns the address of the TIOT, right adjusted, in the full word answer area.

You can also obtain the UCB address via the data extent block (DEB) and the data control block (DCB). The DCB contains data pertinent to the current use of the data set. After the DCB has been opened, offset 44 decimal contains the DEB address. The DEB contains an extension of the information in the DCB. Each DEB is associated with a DCB and the two point to each other.

The DEB contains information about the physical characteristics of the data set and other information that the control program uses. A device-dependent section for each extent is included as part of the DEB. Each such extent entry contains the UCB address of the device to which that portion of the data set has been allocated. In order to find the UCB address, you must locate the extent entry in the DEB for which you intend to issue the RESERVE macro instruction. (In disk addresses of the form MBBCCCHHR, the M indicates the extent number starting with 0).

Procedures for Finding the UCB Address of a Device:

- For data sets using the queued access methods in the update mode or for unopened data sets:
 1. Extract the TIOT from the TCB.
 2. Search the TIOT for the DD name associated with the shared data set.
 3. Add 16 to the address of the DD entry found in step 2. This results in a pointer to the UCB address in the TIOT.
 4. Issue the RESERVE macro specifying the address obtained in step 3 as the parameter of the UCB keyword.

Note: This procedure can be used for non-concatenated DD statements and for data sets that reside on a single volume.

- For opened data sets:
 1. Load the DEB address from the DCB field labeled DCBDEBAD.
 2. Load the address of the field labeled DEBDVMOD in the DEB obtained in step 1. The result is a pointer to the UCB address in the DEB.
 3. Issue the RESERVE macro specifying the address obtained in step 2 as the parameter of the UCB keyword.

- For BDAM data sets, you can reserve the device at any point in the processing in the following manner:
 1. Open the data set.
 2. Convert the block address used in the READ/WRITE macro to an actual device address of the form MBBCCHHR.
 3. Load the DEB address from the DCB field labeled DCBDEBAD.
 4. Load the address of the field labeled DEBDVMOD in the DEB.
 5. Multiply the “M” of the actual device address (step 2) by 16.
 6. The sum of steps 4 and 5 is the address of the correct extent entry in the DEB for the next READ/WRITE operation. The sum is also a pointer to the UCB address for this extent.
 7. Issue the RESERVE macro specifying the address obtained in step 6 as the parameter of the UCB keyword.

- If the data set is an ISAM data set, QISAM in the load mode should be used only at system update time. Further, if it is a multivolume ISAM data set, it must be assumed that all jobs will access the data set through the highest level index. The indexes should never reside in virtual storage when the data set is being shared. In this case, by issuing a RESERVE macro for the volume on which the highest level index resides, the user effectively reserves the volumes on which the prime data and independent overflow areas reside. The following procedures can be used to achieve this:
 1. Open the data set.
 2. Locate the actual device address (MBBCCHHR) of the highest level index. This address can be obtained from the DCB.
 3. Load the DEB address from the DCB field labeled DCBDEBAD.
 4. Load the address of the field labeled DEBDVMOD in the DEB.
 5. Multiply the “M” of the actual device address located in step 2 by 16.
 6. The sum of steps 4 and 5 is the address of the correct extent entry in the DEB for the next READ/WRITE operation. The sum is also a pointer to the UCB address for this extent.

7. Issue the RESERVE macro specifying the address obtained in step 6 as the parameter of the UCB keyword.
- For information concerning how to find the UCB address when using the VSAM access method, see *Virtual Storage Access Method (VSAM) Programmer's Guide*.

RESDEQ Subroutine: The assembler language subroutine in Figure 11 can be used by assembler language programs to issue the RESERVE and DEQ macro instructions. Parameters that must be passed to the RESDEQ routine, if the RESERVE macro instruction is to be issued, are:

DDNAME - the eight character name of the DD statement for the device to be reserved.

QNAME - an eight character name.

RNAME LENGTH - one byte (a binary integer) that contains the RNAME length value.

RNAME - a name from 1 to 255 characters in length.

The DEQ macro instruction does not require the UCB=*ucb addr* as a parameter. If the DEQ macro is to be issued, a fullword of binary zeroes must be placed in the leftmost four bytes of the DDNAME field before control is passed.

```

RESDEQ  CSECT
RESDEQ  AMODE 24
RESDEQ  RMODE 24
        SAVE  (14,12),T          SAVE REGISTERS
        BALR  12,0              SET UP ADDRESSABILITY
        USING *,12
        ST   13,SAVE+4
        LA   11,SAVE            ADDRESS OF MY SAVE AREA IS
                                STORED IN THIRD WORD OF CALLER'S
                                SAVE AREA
        ST   11,8(13)
        LR   13,11            ADDRESS OF MY SAVE AREA
        LR   9,1              ADDRESS OF PARAMETER LIST
        L    3,0(9)          DDNAME PARAMETER OR WORD OF
                                ZEROS
        CLC  0(4,3),=F'0'    WORD OF ZEROS IF DEQ IS
                                REQUESTED
        BE   WANTDEQ
*PROCESS FOR DETERMINING THE UCB ADDRESS USING THE TIOT
        XR   11,11            REGISTER USED FOR DD ENTRY
        EXTRACT ADDR TIOT, FIELDS=TIOT
        L    7,ADDR TIOT      ADDRESS OF TASK I/O TABLE
        LA   7,24(7)          ADDRESS OF FIRST DD ENTRY
NEXTDD  CLC  0(8,3),4(7)      COMPARE DD NAMES
        BE   FINDUCB
        IC   11,0(7)          LENGTH OF DD ENTRY
        LA   7,0(7,11)        ADDRESS OF NEXT DD ENTRY
        CLC  0(4,7),=F'0'    CHECK FOR END OF TIOT
        BNE  NEXTDD
        ABEND 200,DUMP        DDNAME IS NOT IN TIOT, ERROR
FINDUCB LA   8,16(7)          ADDRESS OF WORD IN TIOT THAT
*                                CONTAINS ADDRESS OF UCB
*PROCESS FOR DETERMINING THE QNAME REQUESTED
WANTDEQ L    7,4(9)          ADDRESS OF QNAME
        MVC  QNAME(8),0(7)    MOVE IN QNAME
*PROCESS FOR DETERMINING THE RNAME AND THE LENGTH OF RNAME
        L    7,8(9)          ADDRESS OF RNAME LENGTH
        MVC  RNLEN+3(1),0(7)  MOVE BYTE CONTAINING LENGTH
        L    7,RNLEN
        STC  7,RNAME          STORE LENGTH OF RNAME IN THE
                                FIRST BYTE OF RNAME PARAMETER
*                                FOR RESERVE/DEQ MACROS
*
        L    6,12(9)          ADDRESS OF RNAME REQUESTED
        BCTR 7,0              SUBTRACT ONE FROM RNAME LENGTH
        EX   7,MOVERNAM      MOVE IN RNAME
        CLC  0(4,3),=F'0'
        BE   ISSUEDEQ
RESERVE (QNAME,RNAME,E,0,SYSTEMS),UCB=(8)
        B    RETURN
ISSUEDEQ DEQ (QNAME,RNAME,0,SYSTEMS)
RETURN  L    13,SAVE+4        RESTORE REGISTERS AND RETURN
        RETURN (14,12),T
MOVERNAM MVC  RNAME+1(0),0(6)
ADDR TIOT DC  F'0'
SAVE      DS  18F
QNAME     DS  2F
RNAME     DS  CL256
RNLEN     DC  F'0'
END

```

Figure 11. Example of Subroutine Issuing RESERVE and DEQ

Note: This example assumes that non-concatenated DD statements and single volume data sets are used.

Indicating Event Completion (POST)

The POST macro instruction signifies the completion of an event by one routine to another. Usually the system posts the completion of the event in the user's address space. The user can, however, cause the system to post completion of the event in another address space.

Cross Memory POST

The authorized user (executing in supervisor state, under protection key 0-7, or APF-authorized) of the POST macro instruction can use the ASCB and ERRET parameters to schedule an SRB to be dispatched to perform a POST in an address space other than his own. If the caller is authorized to specify the ASCB and ERRET parameters, no check is made to determine if the requested address space is the issuing address space. This use of the POST macro instruction is sometimes known as "cross memory post."

The ERRET routine is given control in the issuer's address space when an error condition is detected. It receives control enabled, unlocked, in SRB mode, and with the following register contents:

Register	Contents
0	ECB address
1	ASCB address
2	original completion code
3	completion code from failing address space
4-13	unpredictable
14	return address
15	ERRET address

The ERRET routine will receive control in the addressing mode of the caller of the cross memory POST. The ERRET routine must return control to the address in register 14, unlocked and enabled.

If cross-memory post is being used, a synchronization problem arises when it becomes necessary to eliminate an ECB that is a potential target for a cross memory post request. To ensure that all outstanding cross memory post requests for the ECB have completed, the user must invoke the SPOST macro instruction. The ECB might or might not be posted, depending on existing conditions. Because SPOST invokes the PURGEDQ SVC, see the description of PURGEDQ for the restrictions on its use.

The serialization method used to control modifications to an ECB depends on whether or not the ECB is waiting. If the ECB is not waiting (the high order bit of the ECB is off), it may be 'quick posted' via the compare-and-swap instruction using the technique described in "Bypassing the POST Routine" If the ECB is waiting (the high order bit of the ECB is on), the LOCAL lock serializes updates to the ECB.

Bypassing the POST Routine

The problem programmer can bypass the POST routine whenever the corresponding WAIT has not yet been issued if the wait bit is not on. In this case, a compare-and-swap (CS) instruction can be used to quick post the ECB. The compare operand should reflect the ECB content with the wait and post bits off, and the swap operand should have the post bit on and contain the desired post code. If the wait bit is on in the ECB, the CS will fail (giving a non-zero condition code), and the normal POST routine must be executed. If the wait bit is not on, the CS will, in effect, post the completion of the event. Note that holding the LOCAL lock does not eliminate

the requirement to use the CS instruction. Figure 12 demonstrates an example of how to 'Quick Post' an ECB.

```
L  RX,ECB           Get contents of ECB.
N  RX,=X'3FFFFFFF' Turn off wait and post bits
L  RY,=X'40000000' Post bit and post code
CS RX,RY,ECB       Compare and swap to post ECB
BZ POSTDONE        Branch if CS is successful
LTR RX,RX           Wait bit on?
BM DOPOST           If yes, then execute POST
N  RX,=X'40000000' Is ECB posted?
BNZ POSTDONE        If yes, do not execute POST
DOPOST  POST ECB
POSTDONE EQU  *
```

Figure 12. Bypassing the POST Routine

Waiting for Event Completion (WAIT)

The EVENTS macro instruction allows a user to wait for the completion of one of a series of events and be directly informed by the system which of the events have completed. Branch entry to this function, significantly more efficient than SVC entry, is available to users executing in key 0, supervisor state, and holding only the LOCAL lock.

Branch entry specifies `BRANCH=YES` on the EVENTS macro instruction. If this parameter is used, the branch entry routine performs all normal WAIT processing and ECB initialization. You can specify `BRANCH=YES` in conjunction with either `WAIT=YES`, `WAIT=NO`, or `ECB=`.

- If you specify `WAIT=YES`, control will later be returned to the dispatcher, even though there might be ECBs posted to the EVENTS table. EVENTS frees the LOCAL lock. Before issuing the EVENTS macro instruction with the `WAIT=YES` option, you must establish the return environment (the PSW and registers in the RB and TCB). EVENTS stores a pointer to the first completed EVENTS entry into the TCB register 1 save location. (This service is not available to Type 1 SVCs or SRBs.)
- If you do not specify `WAIT=YES`, control returns to you. EVENTS does not free the LOCAL lock.

Writing POST Exit Routines

The POST exit function provides authorized system routines with a service that allows them to receive control immediately upon each completion of an outstanding event. Thus, the user can write a routine that receives control between the time the ECB is marked completed and the return by POST to the caller.

This function defines a special type of ECB known as an extended ECB. When initialized, these extended ECBs identify potential work requests rather than waiting tasks. A purpose of an extended ECB is to notify a process (for example, a subsystem) of an additional work

request. Thus when an extended ECB is posted, a subroutine of the process receives control and updates a queue to identify the current work request.

When using the POST exit function, your routine must follow this sequence:

- Identify POST exit routines.
- Initialize extended ECBs and ECB extensions.
- Wait for work requests.
- Delete POST exit routines before terminating.

Identifying and Deleting Exit Routines

IEA0PT0E is the entry point to POST. It performs exit identification and deletion through a function code that indicates whether the input exit address should be added to or deleted from the POST exit address queue for the current address space. A function code of 4 indicates an exit creation request, while 8 indicates an exit deletion request. Details of this interface are in “Branch Entry to the POST Macro Instruction.”

You cannot provide the same exit routine as input to IEA0PT0E on separate invocations in different addressing modes. A 24-bit caller of the POST-exit-delete function can only delete an exit below 16 megabytes; a 31-bit caller must pass a valid 31-bit address and can delete an exit above or below 16 megabytes.

The process that establishes a POST exit is responsible for deleting that exit before its normal or abnormal termination.

Initializing Extended ECBs and ECB Extensions

The user must obtain and initialize the extended ECBs and ECB extensions. A system service is not available to perform these functions.

The ECB extension must be obtained and initialized before the initialization of the extended ECB. This sequence avoids the possibility of an initialized extended ECB being posted before the initialization of the ECB extension.

The ECB extension is two words long, begins on a word boundary, and can be from any subpool. However, the POST routine must be able to read from the ECB extension in the PSW key of the issuer of the POST macro instruction. The ECB extension must also be accessible in the addressing mode of the POST’s caller. More than one extended ECB can point to it. The mapping for the ECB extension is available via the EXT = YES parameter on the IHAECB mapping macro. It has the format shown in Figure 13.

VALUE (1 byte)	MODE (1 byte)	RESERVED (2 bytes)
POST DATA (4 bytes)		

Figure 13. ECB Extension (ECBE)

The register contents upon return to POST from the user's routine must be:

Register 11, 14-Unchanged
Register 0-10, 12, 13, 15-Irrelevant

The user's exit routine must return control in supervisor state, PSW key zero and with the LOCAL lock still held.

For performance reasons, the user's routine should not cause page faults (that is, the routine's code and the data that it references should be available when the routine receives control).

Re-entry to POST from a POST Exit

A POST exit routine can issue POST only via the POST entry point, IEA0PT03. Details of the interface are in "Branch Entry to the POST Service Routine."

Because of the save area recursion within POST, a POST exit routine cannot post another extended ECB unless it does so by specifying a cross memory post. Any attempt to activate another POST exit before the completion of the current exit causes a 702 abend. If you must post another extended ECB from a POST exit routine, you should either have your routine issue a cross memory post or schedule your own SRB so that your routine enters POST by branching to it.

Example of Using a POST Exit Function

A subsystem allocates and initializes extended ECBs, ECBEs, and EQTs. These data areas appear in Figure 15. Once initialized, the subsystem dispatcher waits on a list of ECBs. Each list entry identifies an ECB in an EQT.

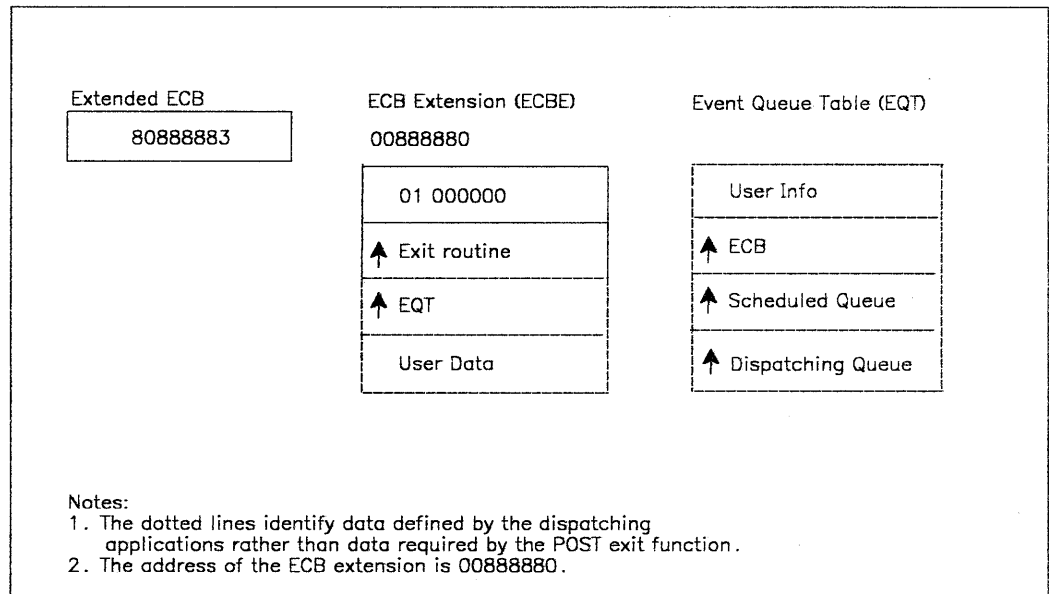


Figure 15. Data Areas Post Exit Example

As soon as any system routine posts an extended ECB, the subsystem exit routine identified in the ECBE receives control. The exit routine receives control in the addressing mode specified by the first bit in byte 1 of the ECBE. If this bit is on, the exit routine receives control in 31-bit addressing mode; if this bit is off, the exit routine receives control in 24-bit addressing mode. After checking the validity of the work request, the exit routine places a work element on the EQT schedule queue identified by the ECBE. The exit routine then posts the ECB associated with that EQT, which completes the queuing of the work and the activation of the dispatching task.

The subsystem dispatcher then scans the ECB list to locate posted ECBs (that is, an EQT with work queued). The subsystem dispatcher then uses compare and swap to switch the schedule queue to the dispatching queue values. Then the subsystem dispatcher dequeues work from the dispatching queue until the queue is depleted. Then the subsystem dispatcher clears the post bit in that EQT ECB and again uses compare and swap to move the schedule queue to the dispatching queue. If the dispatching queue is still empty, the subsystem dispatcher checks the next ECB in the ECB list. After having processed the entire ECB list, the subsystem dispatcher again awaits requests for work.

The subsystem dispatcher can use the USERINFO field in the EQT for serialization where multiple system dispatcher tasks can wait on the same EQT.

Branch Entry to the POST Service Routine

Branch entry to the POST service routine provides all the normal ECB and RB POST processing. The BRANCH parameter on the POST macro instruction uses entry point IEA0PT01 only. To use the other entry points, shown in Figure 16, you must write your own code. In general, the caller must hold the LOCAL lock and be in supervisor state, PSW key zero. Upon completion of the POST process, control returns to the caller in supervisor state, PSW key zero with the LOCAL lock.

Note: CML (cross memory local) lock means the local lock of an address space other than the home address space. LOCAL lock means the local lock of the home address space. When written in lower case, local lock means any local-level lock, either the LOCAL or a CML lock.

You can use branch entry to the POST service routine in cross memory mode for cross memory POST. If you hold the LOCAL lock of the home address space and if bit 0 of register 12 is 0, then the current address space must be the home address space and registers 0-9 and 14 are preserved. If you do not hold home's local lock or if bit 0 of register 12 is 1, then the current address space can be any address space and only registers 9 and 14 are preserved.

Note: If the high-order bit of register 12 is 0 and an error routine is invoked, the error routine is dispatched in the home address space.

Figure 16 shows the POST function and the branch entry points through which those functions can be performed. Figure 17 shows the input parameters to POST. Figure 18 shows the output parameters from POST.

Functions	Entry Points			
	IEA0PT01 (CVT0PT01)	IEA0PT02 (CVT0PT02)	IEA0PT03* (CVT0PT03)	IEA0PT0E (CVT0PT0E)
Local ECB POST	X	X	X	
Local POST without ECB	X		X	
Cross address space POST	X**		X	
Post exit creation/deletion				X
* This entry point performs processing identical to entry point IEA0PT01. It is designed for use only by POST exit routines (that is, routines that receive control from POST as the result of having established that exit via entry point IEA0PT0E).				
** The local lock does not need to be held for a cross address space POST at this entry point.				

Figure 16. POST Function and Branch Entry Points

Registers	IEA0PT01 (CVT0PT01)	IEA0PT02 (CVT0PT02)	IEA0PT03 (CVT0PT03)	IEA0PT0E (CVT0PT0E)
0	ECB storage protect key ¹		ECB storage protect key ¹	Func. Code
1				Exit Routine Address
10	Completion Code ²	Completion Code	Completion Code ²	
11	ECB Address ³	ECB Address	ECB Address ³	
12	Error Routine Address ⁴		Error Routine Address ⁴	
13	ASCB Address ⁴		ASCB Address ⁴	
14	Return Address	Return Address	Return Address	Return Address
15	Entry Point Address	Entry Point Address	Entry Point Address	Entry Point Address
¹ If cross address space post, optionally contains the storage protection key of the ECB in bits 24-27. ² If POST-without-ECB, contains RB address; if cross address space post and the storage protection key of the ECB is supplied in register 0, then the high order bit must be set to one. ³ If POST-without-ECB, set to zero; if local address space POST, ensure high-order bit of register is zero; if cross address space POST, set high-order bit of register to 1. ⁴ Only necessary when performing cross address space POST. If performing a cross address space POST and the high order bit in register 12 is on, only registers 9 and 14 are retained, and the error routine executes in the master scheduler's address space.				

Figure 17. POST Branch Entry Input

Entry Points	Registers Saved and Restored
IEA0PT01 ¹	0-9, 12 ² , 13 ² , 14
IEA0PT02	0-9, 12-14
IEA0PT03	0-14
IEA0PT0E	2-14
¹ The contents of only registers 9 and 14 are retained during a cross address space POST when either the LOCAL lock is not held or the high order bit in register 12 is on; all other register contents are unpredictable. ² The contents of these registers will not be saved and restored during a cross address space POST; their contents are therefore unpredictable in these circumstances.	

Figure 18. POST Branch Entry Output

Branch Entry to the WAIT Service Routine

Branch entry to the WAIT service routine provides all the normal ECB and RB WAIT processing. This function is not available, however, to Type 1 SVCs or SRBs. The caller must hold home's LOCAL lock and be in key zero, supervisor state with current addressability to the home address space. While holding home's LOCAL lock and before branching to WAIT, the caller must establish the PSW and register return environment in its RB and TCB. When WAIT is invoked, the caller should hold only the LOCAL lock. WAIT performs the following functions:

- Stores the ECB/ECBLIST address into the register 1 location of the TCB register save area, (user data cannot be passed through this field or register).
- Releases home's LOCAL lock.
- Returns control to the dispatcher (control does not return to the caller even though all previously pending events have already occurred). The dispatcher ensures that all FRRs have been deleted.

Branch entry to WAIT can occur without identification of any ECBs. This process sets the wait count in the current RB to the specified value. The corresponding POSTs-without-ECB then activate the RB. If you use this process, make sure that the WAIT-without-ECB precedes the POST-without-ECB in order to avoid causing the RB to wait indefinitely.

The following registers contain parameters for branch entry to WAIT:

Register	Contents
0	The wait count in the low order byte. When the high order bit is one, it indicates long-wait (The LONG = YES specification).
1	The ECB pointer value. If only one ECB is being waited on, place that ECB address in register 1. If a list of ECBs is being waited on, place the complemented ECBLIST address in register 1. If the WAIT-without-ECB function is being requested, set register 1 to a value of zero.
15	The branch entry address to WAIT (IEAVWAIT), which in turn is obtained from the CVT (CVTVWAIT).

You can use branch entry to the WAIT service routine in cross memory mode if you hold the LOCAL lock of the home address space and if the current address space is the home address space.

Suspension and Resumption of Request Blocks

An alternate method of waiting for an event and indicating its completion is available on a restricted basis for systems programming. This method gives faster performance than the normal method of using the WAIT and POST macro instructions. The summary below outlines the functions that provide this alternative:

Macro	Description
SUSPEND	Wait for an event to complete.
RESUME	Indicate the completion of the event.
TCTL	Give control directly to a ready task.
CALLDISP	Give up control so that an event can complete.

Waiting for an Event to Complete (SUSPEND)

The SUSPEND macro instruction provides an efficient means of waiting for an event to complete. It is analogous to the WAIT macro instruction, and is used in a SUSPEND-RESUME sequence, which is analogous to the WAIT-POST sequence. The SUSPEND macro instruction causes the wait for event completion through the wait count field (RBWCF) in the request block (RB). This field is the same one the WAIT macro instruction uses. When used with the SUSPEND macro instruction, however, the wait count field is known as the suspend count field, even though the function it performs for both macro instructions is the same.

The SUSPEND macro instruction does not have an immediate effect on the issuer as the WAIT macro instruction does. Instead, the effect is delayed, depending on the type of suspension the macro instruction user requests. If the previous RB is suspended, the effect takes place when the current RB exits. If the current RB is suspended, the suspended state occurs when the RB passes control to the dispatcher.

RBs that issue the SUSPEND macro instruction with the RB=CURRENT option should hold the suspended state time to a minimum. As soon as possible after SUSPEND completion, the RB that issues a SUSPEND RB=CURRENT should exit to the dispatcher (for example, issue a CALLDISP macro instruction with the BRANCH=YES option). Using the SUSPEND macro instruction this way minimizes potential performance problems because the RB in this case must either be disabled or must hold the LOCAL lock or a CML lock. Minimizing suspension time also minimizes other potential problems the program might experience by limiting the time in which the RB is unable to cause any synchronous interrupts (such as SVCs and page faults) or provide interfaces to the WAIT, POST, or EVENTS macro instructions.

RBs that issue SUSPEND RB=PREVIOUS, on the other hand, do not require the same synchronization because they are operating on behalf of another RB. The suspension of the previous RB does not require disabled execution or the holding of the LOCAL lock or a CML lock.

The following scenarios show typical SUSPEND macro instruction sequences:

Scenario 1:

SUSPEND RB = PREVIOUS

1. Type 2 SVC routine receives control.
2. The SVC suspends the macro issuer's RB.
3. The process that will eventually issue the RESUME is started.
4. The SVC completes processing and exits.
5. Event completion occurs; process started in step 3 resumes issuer of the macro instruction.
6. The macro issuer's task resumes (at return from the SVC routine).

Scenario 2:

SUSPEND RB = CURRENT

1. User acquires the LOCAL lock or a CML lock.
2. The macro suspends processing of the current RB.
3. The process that will eventually issue the RESUME is started.
4. Macro issuer issues CALLDISP BRANCH = YES, which releases the LOCAL lock or CML lock before going to the dispatcher.
5. Event completion occurs; process started in step 3 resumes issuer of the macro instruction.
6. Normal processing resumes.

Consider the following when using the SUSPEND macro instruction:

- The SUSPEND macro instruction can be issued in cross memory mode.
- Only a routine executing under protection key 0 can issue SUSPEND.
- The SUSPEND macro instruction requires that the CVT mapping macro be included.
- When the issuer requests (explicitly or by default) the SUSPEND RB = PREVIOUS option, there must be a previous RB on the chain to prevent a task abend.
- Only task-related users can issue SUSPEND, and then only for the current task. SUSPEND cannot be issued for another TCB or by an SRB.
- SUSPEND RB = PREVIOUS is intended for use by Type 2, 3, and 4 SVCs to place the issuer of the SVC in a suspended state.
- The SUSPEND function user must ensure that the SUSPEND and RESUME sequence takes place in proper order. The user must issue SUSPEND, then event completion must occur, and then the RESUME function must take place. One way to ensure proper sequencing is to issue SUSPEND before scheduling the asynchronous process on which the RB must wait.

- When using the `SUSPEND RB=CURRENT` option, the issuer must either execute disabled or hold the LOCAL lock or a CML lock. The issuer must remain in this state until the program initiates the stimulus for event completion in order not to lose control, which could result in never being redispached. Because the issuer must also coordinate the `SUSPEND` and `RESUME` sequence, the event completion must not occur until after the `SUSPEND RB=CURRENT` macro takes effect. The caller that is in the key 0 supervisor state and EUT (enabled unlocked task) mode and that uses a local lock to serialize the `SUSPEND` and `RESUME` processing sequence can use issue `CALLDISP FRRSTK=SAVE` to enter the dispatcher. The `CALLDISP` routine releases the local lock, which serialized the `SUSPEND/RESUME` processing of the caller. Because an EUT FRR exists, the current FRR stack is saved.

When a Type 1 or Type 6 SVC issues a `SUSPEND RB=CURRENT`, the top RB (the caller of the SVC) is suspended. Whenever the SVC exits (via `EXIT PROLOGUE` or `T6EXIT`), the caller is suspended until `RESUME` occurs. A TYPE 1 SVC must not issue the `CALLDISP` macro instruction or release home's LOCAL lock, and it must exit via its exit mechanism. A Type 6 SVC must not issue the `CALLDISP` macro instruction or become enabled, and it also must exit via its exit mechanism.

When a TYPE 2, 3, or 4 SVC issues a `SUSPEND RB=CURRENT`, the top RB (the SVC itself) is suspended. The `SUSPEND` routine returns control to the SVC. The SVC can continue to execute as long as it remains locally locked or disabled. Once the SVC releases the LOCAL lock or enables, an interrupt or an entry to the dispatcher (via `CALLDISP`) suspends the SVC until it is resumed. While the SVC is enabled and before it is resumed, it cannot incur a page fault, issue an SVC, or branch enter any supervisor service that makes local work ready or places the caller in a wait state (for example, `WAIT`, `POST`, `EVENTS`, or `STATUS`).

- The `SUSPEND` and `RESUME` sequence must not be intermixed with the `WAIT` and `POST` sequence on a single RB because both sequences use the same count field for control of the functions. Because the `SUSPEND-RESUME` sequence is a restricted-use function, it does only minimal validity checking. For example, if an RB were already waiting on 255 events and someone issued a `SUSPEND` against it, the count would be reset to one.
- An RB can have only one `SUSPEND` outstanding at a time. There can be no subsequent `SUSPEND` macros issued until a `RESUME` occurs for the outstanding `SUSPEND` macro.
- A program that has invoked the `SUSPEND RB=CURRENT` option must not be suspended again (for example, through a page fault or a lock suspension) after releasing home's LOCAL lock, a CML lock, or enabling until a `RESUME` is issued counteracting the outstanding `SUSPEND` macro instruction.

Resuming Execution of a Suspended Request Block (RESUME)

The `RESUME` macro instruction, which is supported in cross memory mode, provides an efficient means for indicating the completion of an event. The `RESUME` macro instruction specifies the TCB and RB that were previously suspended by the `SUSPEND` macro instruction. The specified TCB and RB must be addressable in the currently addressable address space. Only routines executing in supervisor state and PSW key zero can issue the `RESUME` macro instruction.

The `RESUME` macro instruction and the service routine it calls must serialize the use of the task that is being resumed. This serialization might require the local lock of the task's address space, called the target address space. Because disabled or locked callers of `RESUME` are not

allowed to obtain a local lock, the RESUME macro instruction has the MODE and ASYNC options to handle these types of situations.

Note: The ASYNC parameter for the RESUME macro instruction is spelled differently from similar parameters on other macro instructions.

The MODE option specifies whether or not the RESUME operation must complete (MODE = UNCOND) or not (MODE = COND). MODE = UNCOND requires that certain system locks can be obtained.

The ASYNC option specifies whether or not RESUME can schedule an SRB to perform the resume if necessary. These RESUME options can be combined in four ways:

- MODE = UNCOND and ASYNC = N
 - RESUME attempts to obtain the necessary task serialization to complete the function synchronously. If it can obtain serialization, RESUME completes its function and returns to its caller. If it cannot obtain serialization, RESUME requests the local lock of the target address space to serialize the operation. The caller of RESUME must be able to obtain the target address space's local lock or already hold it when RESUME is issued. This means that, with one exception, the caller of RESUME must either be running enabled and unlocked or disabled and holding the target address space's local lock when the RESUME is issued. The exception is the disabled caller that resumes the TCB under which it is running, that is, the currently executing TCB. This situation could occur if, for example, a routine became disabled, executed a SUSPEND RB = CURRENT macro instruction, and then determined that there was more work to be done. The disabled, unlocked routine could issue a RESUME macro instruction for the TCB and RB to counteract the SUSPEND.

If the local lock is required but not available, the caller will be suspended waiting for the local lock. Control returns to the caller after the RESUME has occurred.

- Disabled interrupt exits cannot issue the RESUME macro instruction with the MODE = UNCOND and ASYNC = N options.
- The RETURN = N option on the RESUME macro instruction is allowed only with this combination of options. The RETURN = N option cannot be used with the ASCB parameter. To use RETURN = N, the caller must be running in SRB mode, must be in home addressing mode, and must not hold any locks. If these three conditions are met, the TCTL service is entered to transfer control to the task that was just resumed. If these three conditions are not met, that caller is abended with an X'070' abend code.
- MODE = UNCOND and ASYNC = Y
 - RESUME attempts to obtain the necessary task serialization and complete the function synchronously. If RESUME cannot obtain serialization, RESUME does not obtain the local lock. RESUME unconditionally obtains an SRB from the supervisor SRB pool and schedules it to complete the RESUME asynchronously.
 - The caller can be enabled or disabled, however, the VSMFIX lock must be available. The caller must not hold any locks higher in the lock hierarchy than the VSMFIX lock or the caller must hold the VSMFIX lock when the RESUME macro instruction is issued.

- **MODE = COND and ASYNC = N**
 - RESUME attempts to obtain the necessary task serialization to complete the function synchronously. If serialization is available, the task is resumed and control returns to the caller. If serialization is not available, RESUME returns to the caller without completing the RESUME operation.
 - The caller can either be enabled or disabled and can hold any combination of locks. RESUME does not attempt to obtain any locks. The caller must be prepared to handle the situation when the RESUME operation can not be performed because the necessary serialization is not available.
- **MODE = COND and ASYNC = Y**
 - RESUME attempts to obtain the necessary serialization to complete the function synchronously. If serialization is available, the task is resumed and control returns to the caller. If serialization is not available, RESUME conditionally obtains an SRB from the supervisor SRB pool and schedules it to perform the RESUME asynchronously. If the supervisor SRB pool is empty, RESUME returns to the caller without completing the RESUME operation.
 - The caller can be either enabled or disabled, and can hold any combination of locks. RESUME does not attempt to obtain any locks. The caller must be prepared to handle the situation when the RESUME operation cannot be performed because the necessary serialization is not available.

RESUME provides return codes in register 15 to indicate the result of the RESUME attempt. See the RESUME macro instruction in Volume 2 for details on the return codes.

The RESUME macro instruction requires the IHAPSA mapping macro. If the ASCB option is not specified, then the MODE = UNCOND and ASYNC = Y combination requires the CVT mapping macro. All other combinations require the IHASVT mapping macro.

Transferring Control for SRB Processing (TCTL)

The TCTL (transfer control) macro instruction allows an SRB routine to exit from its processing and to pass control to a task with minimal dispatcher overhead. When an SRB specifies RESUME RETURN = N, control transfers to the resumed TCB. Control then passes to the top RB on the TCB/RB chain, but only if it passes all the dispatchability tests the dispatcher normally makes

Some other considerations for using the TCTL macro instruction are:

- The TCTL macro may be used only by SRB programs, but they may be in any key. If a non-SRB routine issues either the TCTL macro or a RESUME RETURN = N, the routine will abnormally terminate with a X'070' system completion code.
- The TCTL constitutes an exit from the issuing routine, which therefore causes cleanup of the SRB.
- The TCTL service requires inclusion of the CVT mapping macro.

The TCTL service requires that the SRB requesting the TCTL must not hold any locks and must be in home addressing mode.

Using the **BRANCH = YES** Option of **CALLDISP** (**CALLDISP**)

The **CALLDISP** macro instruction with the **BRANCH = YES** option is supported in cross memory mode. The **BRANCH = YES** option allows an issuer of the **SUSPEND** macro with its **RB = CURRENT** option to exit while leaving the current **RB** in the wait state. This option causes the supervisor to save status and control to pass to the dispatcher.

Some considerations for using the **BRANCH = YES** option on the **CALLDISP** macro instruction are:

- The issuer of **CALLDISP** must be executing in supervisor state with **PSW** key zero.
- The issuer must be in task mode rather than in **SRB** mode.
- The **BRANCH = YES** option requires inclusion of these mapping macros:

```
IHASVT  
IHAPSA
```

- The **FIXED = YES** or **FIXED = NO** option can be specified with **BRANCH = YES**.
- When **FRRSTK = SAVE** is specified:

The caller must not hold any locks or an abend results.

Note: For **MVS/System Product Version 2 Release 1.3 Vector Facility Enhancement** or **MVS/System Product Version 2 Release 1.3 Availability Enhancement** and later releases:

- If any **EUT** (enabled unlocked task) **FRRs** exist, the current **FRR** stack is saved and the caller may hold either the **LOCAL** or **CML** lock. **CALLDISP** releases the lock before going to the dispatcher.
- If no **EUT FRR** exists, the caller cannot hold any locks. Otherwise, an abend occurs.
- When **FRRSTK = NOSAVE** is specified:
 - The current **FRR** stack is purged.
 - The caller may hold either the **LOCAL** or **CML** lock. **CALLDISP** releases the lock before going to the dispatcher.

Notes:

1. A type 1 and type 6 **SVC** must not issue the **CALLDISP** macro instruction.
2. The **LOCAL** or **CML** lock can be used to serialize the **SUSPEND** processing and establish the **RESUME** processing. If a local lock is not used to serialize the **FRR** stack, the caller can use the **CALLDISP FRRSTK = SAVE** option to serialize the stack. For more information, see "Suspension and Resumption of Request Blocks" in this volume.

Reporting System Characteristics

This chapter describes three ways to report system characteristics:

- Using GQSCAN to obtain resource usage reports
- Using SRM to obtain subsystem measurement reports
- Using SYMREC to obtain software error reports

Collecting Information About Resources and Their Requestors (GQSCAN)

Global resource serialization enables an installation to share symbolically-named resources between units of work. Programs issue ENQ, DEQ, and RESERVE macro instructions to request access to resources. Global resource serialization runs in its own address space and maintains the resource queues in this address space, which cannot be swapped, cancelled, or forced. The only way you can extract information from the resource queues is by using the GQSCAN macro instruction. You can extract this information whether or not global resource serialization is active. See *Planning: Global Resource Serialization* for information about how global resource serialization functions.

Using GQSCAN, you can inquire about a particular scope of resources (such as STEP, SYSTEM, or SYSTEMS), a specific resource by name, a specific system's resources, a specific address space's resources, or resources requested by the RESERVE macro instruction. The GQSCAN service routine collects the information you request from the resource queues and consolidates that information before returning it.

In order to specify a scope of LOCAL or GLOBAL, you must be a supervisor state, key zero user. SCOPE=LOCAL and SCOPE=GLOBAL are restricted because when you specify these parameters the GQSCAN service routine serializes the use of the GRS control blocks to stop any other ENQ, DEQ, or RESERVE instructions from changing them.

The information you request is returned in an area whose location and size you specify using the AREA parameter on the GQSCAN macro instruction. You can also specify a fullword location in which the GQSCAN service routine can return a token by specifying the TOKEN parameter. If the amount of information exceeds the size of your area, the GQSCAN service routine gives you as much as the area will hold and returns a token in the specified location. On subsequent invocations of the GQSCAN macro instruction, if you provide the token, you can obtain the remaining information. You can invoke the GQSCAN macro instruction again and again using the same token, until all of the information has been returned. You can also invoke the GQSCAN macro instruction using the same token and QUIT=YES to terminate the scan.

The information is returned in the form of resource information blocks and resource information block extensions, as shown below.

RIB	A	Resource information block (RIB) describes a resource
RIBE	A1	RIB extension (RIBE) describes resource requestor
RIBE	A2	
RIBE	A3	
RIB	B	
RIBE	B1	
RIBE	B2	

The RIB and RIBE are described in the *Debugging Handbook*.

The amount of information you get about a particular resource depends on the scope you specify on the GQSCAN macro instruction, the size of the area you provide, and whether or not you specify a token.

Whether you specify a scope of STEP, SYSTEM, SYSTEMS, or ALL with or without a token, the information returned the first time you issue the GQSCAN macro instruction is the same. You get the first RIB and as many of its associated RIBEs as will fit in your area. The RIB has an entry that tells how many RIBEs it has and, of those, how many appear in your area. Any RIBEs that do not fit are lost. If there is room for another RIB and all of its RIBEs, it also is returned and so on until there is not enough room for the next RIB and all of its RIBEs.

The second return of information occurs the next time you issue the GQSCAN macro instruction. The contents of this return differ depending on whether or not you specify a token. If you specified a token on the first GQSCAN and supply that token on the next invocation, the information returned continues with the next RIB, and as many of its associated RIBEs as will fit in your area. Any RIBEs that do not fit are lost. If there is room for another RIB and all of its RIBEs, it too is returned, and so on, until there is not enough room for the next RIB and all of its RIBEs. If you do not specify a token, the information starts over with the first RIB.

The example in Figure 19 shows three returns. The scope is either STEP, SYSTEM, SYSTEMS, or ALL and a token is specified each time.

For the example, assume that there are four resources, A, B, C, and D. Resource A has three requestors, resource B has six requestors, resource C has two requestors, and resource D has one requestor.




Figure 19. GQSCAN Results with STEP, SYSTEM, SYSTEMS, or ALL

Whether you specify a scope of LOCAL or GLOBAL with or without a token, the information returned the first time you issue the GQSCAN macro instruction is the same. You get the first RIB and as many of its associated RIBEs as will fit in your area. The RIB has an entry that tells the number of RIBEs associated with it, and, of those, how many appear in your area. Any RIBEs that do not fit will appear in your area the next time you issue the GQSCAN macro instruction if a token is provided. If there is room for another RIB and at least one RIBE, it is also returned, and so on, until there is not enough room for a RIB and RIBE combination. If you do not specify a token, the information starts over with the first RIB.

The example in Figure 20 shows two returns. The scope is either LOCAL or GLOBAL and a token is specified each time. Note that the GQSCAN service routine does not truncate any of the RIBEs. On every return after the first, in which you supply the token that was returned, the previous RIB is repeated to put the remaining RIBEs in context and more resource information, if any, continues.

First return

RIB _A 3 RIBEs total 3 here
RIB _{A1}
RIB _{A2}
RIB _{A3}
RIB _B 6 RIBEs total 4 here
RIB _{B1}
RIB _{B2}
RIB _{B3}
RIB _{B4}

Second return

RIB _B 6 RIBEs total 2 here
RIB _{B5}
RIB _{B6}
RIB _C 2 RIBEs total 2 here
RIB _{C1}
RIB _{C2}

Figure 20. GQSCAN Results with LOCAL or GLOBAL

In scanning the information returned, be sure to use the size of the fixed portion of the RIB and the RIBE that is returned in register 0. The size of the fixed portion of the RIB is in the high-order half of register 0 and the size of the RIBE is in the low-order half.

The first RIB starts at the beginning of the area you specify. The first RIBE is pointed to by the current RIB pointer plus the size of the fixed portion of RIB plus the size of the variable portion of RIB (RIBVLEN). To find the second RIBE, add the size of the RIBE.

To find the second RIB, use the current RIB pointer plus the size of the fixed portion of the RIB plus RIBVLEN and then add the number of RIBEs times the RIBE size.

Using the SRM Reporting Interface to Measure Subsystem Activity

The reporting interface allows an IBM or user-written interactive subsystem to pass transaction performance data to the system resources manager (SRM). The data collected by the SRM can be reported through the RMF workload activity report or the transaction activity report. The Resource Measurement Facility (RMF) program product must be installed to obtain these reports. The data is reported according to the subsystem identifiers (subsystem name, transaction name, transaction class, or userid) specified in the installation control specification (IEAICSxx parmlib member). For more information on the installation control specification, see *Initialization and Tuning*.

The reporting interface is necessary because, except for TSO, the SRM does not normally recognize the individual transactions of an interactive subsystem. For example, the SRM considers a subsystem that consists of an address space created by a START command to be a single long transaction, and the RMF workload activity report indicates the total service for the address space but does not indicate the average transaction response time. However, when a subsystem uses the interface and the subsystem is specified in the installation control specification, the RMF workload activity report provides the average transaction response time.

The reporting interface consists of a SYSEVENT macro instruction, which the subsystem must issue at the completion of each transaction. Issuing the macro instruction allows the subsystem to pass the transaction start time or elapsed time and, optionally, its resource utilization. The SRM does not use data collected through the reporting interface to dynamically adjust resource distribution to subsystems. However, the installation can review the RMF reports to determine which, if any, SRM parameters need to be changed.

Reporting Software Error Symptoms (SYMREC)

SYMREC is a facility that can store and format information about non-abend errors that occur in an authorized user application. Authorized applications may be written to detect their own errors and to invoke the SYMREC macro, which stores information about each error on SYS1.LOGREC, the on-line repository where error information is collected. The unit of information that the SYMREC macro stores in the repository is called a **symptom record**. The data in the symptom record is a description of some programming failure combined with a description of the environment where the failure occurred.

The SYMREC facility can interpret the symptom records stored on the repository, and formatting them into various kinds of reports. SYMREC consists of the following elements:

1. Two macros, ADSR and SYMREC.
2. Formatting capabilities that are invoked through EREP and IPCS.

Writing Applications That Use SYMREC

The use of SYMREC implies that an application is written in a programming style in which the program monitors execution-time conditions that denote errors. To write an application that uses SYMREC, you use the ADSR macro to obtain a DSECT for an area in your virtual storage that is called a **symptom record**. This area has six sections. When the application starts to execute, it first stores zeroes into the entire area, then initializes certain fields in the first two sections. These fields identify the application's environment and supply essential addressing information for the application as it executes. After identifying the symptom record, the application continues with its main processing.

If the application detects an error during its processing, it collects information about the error, then stores the error information in the symptom record by referencing the data fields in the ADSR DSECT. When the application stores the error information, the information is in a special format called the **SDB** (structured data base) format.

After storing data into the symptom record, the application issues the SYMREC macro, which writes data from the symptom record to an on line repository. Before SYMREC processing writes the data, it updates the symptom record in the application's virtual storage (even if the storage is store-protected) with various data about the execution environment. After SYMREC writes the data, it returns control to the application with return and reason codes respectively in registers 15 and 0. The application can continue processing if the error is not fatal. The return and reason codes, in hexadecimal, are:

Return Code	Reason Code	Explanation
0000		Symptom record component completed successfully and the symptom record was recorded.
	0000	Successful completion of the SYMREC macro service routine.
0004		Error(s) detected on the SYMREC macro statement. The entire input record was recorded. Following are specific reasons why the symptom record component processed unsuccessfully:
	0164	The input symptom record was successfully copied. However, an attempt to write section 1 information from the completed symptom record failed. The area was found non-accessible to a write request.

Return Code	Reason Code	Explanation
0008		Error(s) detected on the SYMREC macro statement. A partial symptom record was recorded. Following are specific reasons why the symptom record component processed unsuccessfully:
	0158	Total length of the input symptom record exceeds the maximum.
	015C	Optional segments of the input symptom record were found non-accessible. The record includes the accessible entries of the input symptom record.
000C		Serious error on the SYMREC macro statement. No symptom record was recorded. Following are specific reasons why the symptom record component processed unsuccessfully:
	0104	The first 2 bytes of the input symptom record does not contain the SR operand.
	0108	The input symptom record does not contain the required entries for section 2.
	010C	The input symptom record does not contain the required entries for section 2.1.
	0114	The input symptom record does not contain the required entries for section 3.
	0128	Portions of the input symptom record were found non-accessible to a write request.
	012C	Required portions of the input symptom record were found non-accessible to a write request.
	0134	Input symptom record address is in non-accessible storage.
	0144	Program attributes of the job issuing the SYMREC macro are not written in accordance with the symptom record component standards.
0010		Serious error in the symptom record component. Error is not related to SYMREC macro statement. No symptom record was recorded. Following are specific reasons why the symptom record component processed unsuccessfully:
	0F04	Insufficient space in the LOGREC buffer to accommodate the symptom record.
	0F08	SYMREC macro service routine could not acquire storage for its workarea and a copy of the symptom record.
	0F0C	Failure occurred while moving the symptom record to the LOGREC buffer.
	0F10	SYMREC macro service routine has a logic error.
0014		Symptom record component is not operable.

To invoke the SYMREC macro, an application must be authorized, enabled for interrupts, and running in primary addressing mode (bit 16 of the PSW is equal to zero). The AMODE attribute of the program can be set for either 24-bit or 31-bit addressing, but any address passed to SYMREC must be a 31-bit address. A program that invokes SYMREC can be running in cross memory mode, and it can be holding a suspend lock.

The Format of the Symptom Record

The symptom record consists of six sections that are structured according to the format of the ADSR DSECT. These sections are numbered 1 through 5, including an additional section that is numbered 2.1. Because sections 2.1, 3, 4, and 5 of the symptom record are non-fixed, they do not need to be sequentially ordered within the record. In section 2, the application supplies the offset and the length of the non-fixed sections. The ADSR format is described in the *MVS/XA Debugging Handbook*, and the purpose of each section is as follows:

Section 1 (Environmental data): Section 1 contains the record header with basic environmental data. The SYMREC caller initializes this area to zero and stores the characters, 'SR', into the record header. The environmental data of section 1 is filled in automatically whenever the SYMREC macro is invoked. The environmental data that SYMREC stores in this section provides a system context within which the software errors can be viewed. Section 1 includes items such as:

- CPU model and serial numbers
- Date and time, with a time zone conversion factor
- Name of the customer installation
- Product ID of the control program
- Special features installed, if any

Section 2 (Control data): Section 2 contains control information with the lengths and offsets of the sections that follow. The application must initialize the control information before invoking SYMREC for the first time. Section 2 immediately follows section 1 in the symptom record structure.

Section 2.1 (Component data): Section 2.1 contains the name of the component in which the error occurred, as well as other specific component-related data. The application must also initialize section 2.1 before invoking SYMREC.

Section 3 (Primary SDB symptoms): Section 3 contains the primary string of problem symptoms, which may be used to perform tasks such as duplicate problem recognition. Whenever an application detects a given error, the string that it stores in section 3 becomes uniquely associated with that error incident as its **primary symptom string**. Note that an application does not store any primary symptom string or invoke SYMREC unless it detects an error in its own processing.

Section 4 (Secondary SDB symptoms): Section 4 contains an optional **secondary symptom string**. The purpose of the secondary string is to save any critical diagnostic values that existed at the time of the incident.

Section 5 (Free-format data): Section 5 contains logical segments of optional problem related data to aid in problem diagnosis. However, the data in section 5 is not in the SDB format, which is only found in sections 3 and 4. Each logical segment in section 5 is structured in a **key-length-data** format. This format consists of a two-byte key field, which is immediately followed by a two-byte length field, which is immediately followed by a variable length data field. The length that is indicated in the two-byte length field should be equal to the length of the data field.

Symptom Strings — SDB Format

The symptom strings placed in sections 3 and 4 of the symptom record must be in the SDB (structured data base) format. In this format, the individual symptoms in sections 3 and 4 of the symptom record are separated by syntactical dividers called **prefixes**. For more information on the prefixes that SYMREC recognizes, see *MVS/XA Debugging Handbook*. Also see *A Structured Approach to Describing and Searching Problems* for more general information on symptom strings and prefixes. Examples of typical prefixes are:

PIDS/	a component name follows the slash
RIDS/	a routine name follows the slash
AB/	an abend code follows the slash
PRCS/	a return code follows the slash
REGS/	a register name follows the slash
LVLS/	a release level follows the slash
FLDS/	a data field name follows the slash
VALU/	an error-related value follows the slash

An SDB symptom string is a continuous character string. Each prefix must end in a slash, and a blank is required between successive symptoms in the string. For example:

```
LVLS/B10 PIDS/5752SASR RIDS/ASRSERVX PRCS/00000001 REGS/GR14  
VALU/H81ABCDEF FLDS/XSTATUS VALU/CUSED
```

Symptom strings in SDB format can be interpreted as meaningful remarks about program incidents. The symptom string shown above is interpreted as follows:

LVLS/B10	In release level B10
PIDS/5752SASR	of the component, 5752SASR,
RIDS/ASRSERVX	routine ASRSERVX was executing,
PRCS/00000001	and it received return code 00000001.
REGS/GR14	At that time, general register 14
VALU/H81ABCDEF	contained hex 81ABCDEF, and
FLDS/XSTATUS	the field, XSTATUS,
VALU/CUSED	contained characters USED.

Using EREP and IPCS to Format Symptom Record Reports

To format and print various reports from the data on the repository, you use EREP and IPCS programs, which invoke the SYMREC post-processing services. The four basic reports generated by SYMREC post processing services are:

The System Summary Report: The system summary EREP report is an overview of processor, channel, subchannel, operating system, and I/O subsystem errors. The report contains a count per CPU of the symptom records that were stored.

The Event History Report: The event history report consists of one-line abstracts of selected information from each record, listed in chronological order. The one-line abstract is the primary symptom string. The event history also has a count of the symptom records classified by processor.

The Detail Edit Report: The detail edit report shows the entire contents of an error record, except for section 2. Optional fields that have not been completed, which contain hexadecimal zeroes, are not included in the report.

The Detail Summary Report: The detail summary report shows each unique primary symptom string in the repository; it does not duplicate symptom strings that recur. For each unique string, the report indicates the number of occurrences and the date and time of the first and last occurrence. SYMREC uses the first 80 bytes of the symptom string when comparing for duplicate strings.

Note: For detailed information on these reports, see *MVS/XA Debugging Handbook*.

Through EREP, you can generate all four reports, which are obtained when you specify the TYPE=S parameter. IPCS can generate only one kind of report, the detail edit report. To cause IPCS to generate this report, you specify the LOGDATA verb. LOGDATA causes one report to be generated for each symptom record in its recording buffer.

Programming Notes for SYMREC Applications

This section contains programming notes on how the various fields of the ADSR data area (symptom record) are set. Some fields of the ADSR data area (symptom record) must be set by the caller of the SYMREC macro, and other fields are set by the system when the application invokes the SYMREC service. The fields that the SYMREC caller must always set are indicated by an RC code in the following sections. The fields that are set by SYMREC are indicated by an RS code.

The RA code designates certain flag fields that need to be set only when certain kinds of alterations and substitutions are made in the symptom record after the incident occurs. These alterations and substitutions must be obvious to the user who interprets the data. Setting these flag fields is the responsibility of the program that alters or substitutes the data. For example, if a program changes a symptom record that is already written on the repository, it should set the appropriate RA-designated flag-bit fields as an indication of how the record has been altered.

The remaining fields, those not marked by RC, RS, or RA, are optionally set by the caller of the SYMREC macro. When SYMREC is invoked, it checks that all the required input fields in the symptom record are set by the caller. If the required input fields are not set, SYMREC issues appropriate return and reason codes.

Programming Notes for Section 1

Notes in this section pertain to the following fields, which are in section 1 of the ADSR data area.

ADSRID	Record header	(RC)
ADSRGMT	Local Time Conversion Factor	
ADSRTIME	Time stamp	(RS)
ADSR TOD	Time stamp (HHMMSSSTH)	
ADSRDATE	Date (YYMMDD)	
ADSRSID	Customer Assigned System/Node Name	(RS)
ADSRSYS	Product ID of Base System (BCP)	(RS)
ADSR CML	Feature and level of Symrec Service	(RS)
ADSRTRNC	Truncated flag	(RS)
ADSRP MOD	Section 3 modified flag	(RA)
ADSRSGEN	Surrogate record flag	(RA)
ADSRSMOD	Section 4 modified flag	
ADSRNOTD	ADSR TOD & ADSRDATE not computed flag	(RS)
ADSRASYN	Asynchronous event flag	(RA)
ADSRDTP	Name of dump	

Notes:

1. *SYMREC stores the TOD clock value into ADSRTIME when the incident occurs. However, it does not compute ADSRTOD and ADSRDATE when the incident occurs, but afterwards, when it formats the output. When the incident occurs, SYMREC also sets ADSRNOTD to 1 as an indication that ADSRTOD and ADSRDATE have not been computed.*
2. *SYMREC stores the customer-assigned system node name into ADSRSID.*
3. *SYMREC stores the first four digits of the base control program component id into ADSRSYS. The digits are 5752, 5759 and 5745 respectively for MVS, VM, and DOS/VSE.*
4. *The ADSRDTP field is not currently used by the system.*
5. *If section 3 of a symptom record is changed or extended after SYMREC is invoked, the program that makes the change should set the ADSRPMOD flag to 1. If section 4 of a symptom record is changed or extended after SYMREC is invoked, the program that makes the change should set the ADSRSMOD flag to 1. The purpose of these flags is to indicate whether the symptom strings are original, intact, and exactly the same as when SYMREC was invoked.*
6. *If some application creates the record asynchronously, that application should set ADSRSYN to 1. 1 means that the data is derived from sources outside the normal execution environment, such as human analysis or some type of machine post-processing.*
7. *If SYMREC truncates the symptom record, it sets ADSRTRNC to 1. This can happen when the size of the symptom record provided by the invoking application exceeds SYMREC's limit.*
8. *ADSRSGEN indicates that the symptom record was not provided as 'first time data capture' by the invoking application. Another program created the symptom record. For instance, the system might have abended the program, and created a symptom record for it because the failing program never regained control. Setting the field to 1 means that another program surrogate created the record. The identification of the surrogate might be included with other optional information, for example, in section 5.*
9. *The application invoking SYMREC must provide the space for the entire symptom record, and initialize that space to hex zeroes. The application must also store the value 'SR' into ADSRID.*
10. *The fields ADSRCPM through ADSRFL2, which appear in the record that is written on the repository, are also written back into the input symptom record as part of the execution of SYMREC.*

Programming Notes for Section 2

Notes in this section pertain to the following fields, which are in section 2 of the ADSR data area.

ADSRARID	Architectural level designation	(RS)
ADSRRL	Length of section 2	(RC)
ADSRCSL	Length of section 2.1	(RC)
ADSRCSO	Offset of section 2.1	(RC)
ADSRDBL	Length of section 3	(RC)
ADSRDBO	Offset of section 3	(RC)
ADSRROSL	Length of section 4	
ADSRROSA	Offset of section 4	
ADSRRONL	Length of section 5	
ADSRRONA	Offset of section 5	
ADSRRISL	Length of section 6	
ADSRRISA	Offset of section 6	
ADSRRES	Reserved for system use	

Notes:

1. *The invoking application must ensure that the actual lengths of supplied data agree with the lengths indicated in the ADSR data area. The application should not assume that the SYMREC service validates these lengths and offsets.*
2. *The lengths and offsets in section 2 are intended to make the indicated portions of the record indirectly addressable. Invoking applications should not program to a computed absolute offset, which may be observed from the byte assignments in the data area.*
3. *The value of the ADSRARID field is the architectural level at which the SYMREC service is operating. This field is supplied by the SYMREC service.*
4. *Section 2 has a fixed length of 48 bytes. Optional fields (not marked with RC, RS, or RA) will contain zeroes when the invoking application provides no values for them.*

Programming Notes for Section 2.1

Notes in this section pertain to the following fields, which are in section 2.1 of the ADSR data area.

ADSRC	C'SR21' Section 2.1 Identifier	(RC)
ADSRCRL	Architectural Level of Record	(RC)
ADSRCID	Component identifier	
ADSRFLC	Component Status Flags	
ADSRFLC1	Non-IBM program flag	(RC)
ADSRVLV	Component Release Level	(RC)
ADSRPTF	Service Level	
ADSRPID	PID number	(RC)
ADSRPIDL	PID release level	(RC)
ADSRCDSC	Text description	
ADSRRET	Return Code	(RS)
ADSRREA	Reason Code	(RS)
ADSRPRID	Problem Identifier	
ADSRID	Subsystem identifier	

Notes:

1. This section has a fixed length of 100 bytes, and cannot be truncated. Optional fields (not marked with RC, RS, or RA) will appear as zero if no values are provided.
2. ADSRCID is the component ID of the application that incurred the error, without the imbedded punctuation that normally appears when the component id is seen in print.

Under some circumstances, there can be more than one component ID involved. For ADSRCID, select the component ID that is most indicative of the source of the error. The default is the component ID of the detecting program. In no case should the component ID represent a program that only administratively handles the symptom record. Additional and clarifying data (such as, other component ID involved) is optional, and may be placed in optional entries such as ADSRCDS of section 2.1, section 4, or section 5.

For example: if component A receives a program check; control is taken by component B, which is the program check handler. Component C provides a service to various programs by issuing SYMREC for them. In this case, the component ID of A should be given. Component B is an error handler that is unrelated to the source of the error. Component C is only an administrator. Note that, in this example, it was possible for B to know A was the program in control and the source of the program check. This precise definition is not always possible. B is the detector, and the true source of the symptom record. If the identity of A was unknown, then B would supply, as a default, its own component ID.

ADSRPID is not a required field in this section, although it is required in section 3 after the PIDS/ prefix of the symptom string. Repeating this value in section 2.1 is desirable but not required. Where the component ID is not given in section 2.1, this field must contain zeroes.

ADSRPID is the Program Information Department (PID) number assigned to the program that incurred the error. It appears as a seven-character value with no punctuation and one byte of padding. ADSRPID must be provided only by IBM programs that do not have an assigned component ID. Therefore, ADSRCID contains hex zeroes if ADSRPID is provided.

3. ADSRLVL is the release level of the assigned component ID. It is required even if the assigned component ID value is not given in ADSRCID for IBM products. All release level values are numeric values. Therefore, this field normally has a blank (X'40') as the rightmost pad character.
4. ADSRPIDL is the release level of the program designated by ADSRPID, and it should be formatted using the characters, V, R, and M as separators, where V, R, and M represent the version, release, and modification level respectively. For example, V1R21bbbb is Version 1 Release 2.1 without any modification. No punctuation can appear in this field, and ADSRPIDL must be provided only when ADSRRPID is provided.
5. ADSRPTF is the service level. It may differ from ADSRLVL because the program may be at a higher level than the release. ADSRPTF can contain any number indicative of the service level. For example, a PTF, FMID, APAR number, or user modification number. This field is not required, but it should be provided if possible.
6. ADSRCDS is a 32-byte area that contains text, and it is only provided at the discretion of the reporting component. It provides clarifying information. For example, 'IOS IOSB ANALYSIS ROUTINE'.

7. *ADSRREA* is the reason code, and *ADSRRET* is the return code from the execution of *SYMREC*. *SYMREC* places these values in registers 0 and 15 and in these two fields as part of its execution. The fields are right justified, and identical to the contents of registers 0 and 15.
8. *ADSRCRL* is the architectural level of the record. Note that *ADSRARID* (section 2) is the architectural level of the *SYMREC* service.
9. *ADSRPRID* is a value that can be used to associate the symptom record with other symptom records. This value must be in *EBCDIC*, but it is not otherwise restricted.
10. *ADSRNIBM* is a flag indicating that a non-IBM program originated the symptom record. When this flag is 1, *ADSRPID*, *ADSRPIDL*, and *ADSRPTF* are interpreted respectively as program name, program major level, and program fix level. For IBM programs originating a symptom record, this flag must be '0'.
11. *ADSRSSID* is the name of a subsystem. The primary purpose of this field is to allow IBM subsystems to intercept the symptom record from programs that run on the subsystem. They may then add their own identification in this field as well as additional data in sections 4 and 5. The subsystem can then pass the symptom record to the system via *SYMREC*. A zero value is interpreted as no name.
12. The *ADSRVSE6* flag is set to 1 when the user has added another section to the symptom record after section 5 (that is, the contents of *ADSRRISL* are non-zero).

Programming Notes for Section 3

Section 3 of the symptom record contains the primary symptoms associated with the error incident, and it is provided by the application that incurred the error, or some program that acts in its behalf. The internal format of the data in section 3 is the *SDB* format, with a blank separating each entry. Once this data has been passed to *SYMREC* by the invoker, it may not be added to or modified without setting *ADSRPMOD* to '1'. The data in this section is *EBCDIC*, and no hex zeros may appear. The symptoms are in the form *K/D* where *K* is a keyword of 1 to 8 characters and *D* is at least 1 character. *D* can only be an alphanumeric or @, \$, and #.

Notes:

1. The symptom *K/D* can have no imbedded blanks, but the '#' can be used to substitute for desired blanks. Each symptom (*K/D*) must be separated by at least one blank. The first symptom may start at *ADSRRSCS* with no blank, but the final symptom must have at least one trailing blank. The total length of each symptom (*K/D* combination) can not exceed 15 characters.
2. This section is provided by the component that reports the failure to the system. Once a *SYMREC* macro is issued, the reported information will not be added to or modified, even if the information is wrong. It is the basis for automated searches, and even poorly chosen information will compare correctly in such processing because the component consistently produces the same symptoms regardless of what information was chosen.

3. *If section 3 is modified, then a flag in section 1 (ADSRPMOD) should be set to warn users (such as detecting programs) that this data has been modified (perhaps to correct it), and may be useless in SDB automated searches.*
4. *The PIDS/ entry is required, with the component ID following the slash. It is required from all programs that originate a symptom record and have component a ID assigned. Further, it must be identical to the value in ADSRCID (section 2.1) if that is provided. (ADSRCID is not a required field).*

Programming Notes for Section 4

Section 4 of the symptom record contains the secondary symptoms associated with the error incident, and it is provided by the application that incurred the error, or some program that acts in its behalf. The internal format of the data in section 4 is the SDB format, with a single blank separating each entry. Once this data has been passed to SYMREC by the invoker, it may not be added to or modified without setting ADSRSMOD to 1. The data in this section is EBCDIC, and no hex zeroes may appear. The symptoms are in the form, K/D, where K is a keyword of one to eight characters and D is at least one character. D must be alphanumeric or @, \$, or #.

Notes:

1. *The secondary symptom string is in the same SDB format as the primary symptom string.*
2. *If any changes are made in this section, ADSRSMOD must be set to 1.*

Programming Notes for Section 5

Section 5 of the symptom record contains logical segments of data that are provided by the component or some program that acts in its behalf. The component may store data in section 5 when SYMREC is invoked, or the system may add notes in this section at the time of SYMREC execution. Further, section 5 may be added to by a general edit of the record or by other programs operating on the entry any time after SYMREC is invoked.

Notes:

- 1. The first segment must be stored at symbolic location ADSR5ST. In each segment, the first two characters are a hex key value, and the second two characters are the length of the data string, which must immediately follow the two-byte length field. Adjacent segments must be packed together. The length of section 5 is in the ADSRRONL field in section 2, and this field should be correctly updated as a result of all additions or deletions to section 5.*
- 2. There are 64K key values grouped in thirteen ranges representing thirteen potential SYMREC user categories. The data type (that is, hexadecimal, EBCDIC, etc.) of section 5 is indicated by the category of the key value. Thus, the key value indicates both the user category and the data type that are associated with the information in section 5. Because the component ID is a higher order qualifier of the key, it is only necessary to control the assignment of keys within each component ID or, if a component ID not assigned, within each PID number.*

Key Value	User Category and Data Type
0001-00FF	Reserved
0100-0FFF	MVS system programs
1000-18FF	VM system programs
1900-1FFF	DOS/VSE system programs
2000-BFFF	Reserved
C000-CFFF	Product programs and non-printable hex data
D000-DFFF	Product programs and printable EBCDIC data
E000-EFFF	Reserved
F000	Any program and printable EBCDIC data
F001-F0FF	Not assignable
F100-FEFF	Reserved
FF00	Any program and non-printable hex data
FF01-FFFF	Not assignable

Communication

The following types of communication are included in this chapter:

- Interprocessor communication
- Writing operator messages
- Inter-address space communication

Interprocessor Communication

Interprocessor communication (IPC) is a function that provides communication between processors sharing the same control program. Those executing functions that require a processor or program action on one or more processors use the IPC interface to invoke the desired action. The IPC function uses the signal processor (SIGP) instruction to provide the necessary hardware interface between the processors.

Based on the condition code of the SIGP instruction, the IPC function may invoke the excessive spin routine. The excessive spin routine may cause message IEE331A to be issued. This message either requires the operator to initiate alternate CPU recovery (ACR) or continue with processing. For more information concerning the SIGP instruction, see *Principles of Operation*.

Service Classes

IPC services divide the SIGP order codes into two classes, direct and remote. The SIGP instruction and the valid order codes are documented in *Principles of Operation*.

Direct service class is defined for those control program functions that require the modification or sensing of the physical state of one of the configured processors. The following SIGP order codes can be invoked via the DSGNL macro instruction.

Sense: The specified processor presents its status to the issuing processor. No other action is caused at the specified processor.

Start: The specified processor is placed in the operating state. The processor does not necessarily enter the operating state during the execution of the SIGP instruction. No action is caused at the specified processor if that processor is in the operating state when the order code is accepted.

Stop: The specified processor stops. The processor does not necessarily enter the stopped state during the execution of the SIGP instruction. No action is caused at the specified processor if that processor is in the stopped state when the order code is accepted.

Restart: The specified processor restarts. The processor does not necessarily perform the function during the execution of the SIGP instruction.

Stop and Store Status: The specified processor stops and stores status. The processor does not necessarily complete the operation, or even enter the stopped state, during the execution of the SIGP instruction.

Store Status at Address: The specified processor stores status starting at a specified location. If the specified processor is not stopped, it does not accept the order. The processor does not necessarily complete the operation during the execution of the SIGP instruction.

Initial CPU Reset: The specified processor performs initial processor reset. The execution of the reset does not affect other processors and does not cause any channels, including those reconfigured to the specified processor, to be reset. The reset operation is not necessarily completed during the execution of the SIGP instruction.

CPU Reset: The specified processor performs processor reset. The execution of the reset does not affect other processors and does not cause any channels, including those configured to the specified processor, to be reset. The reset operation is not necessarily completed during the execution of the SIGP instruction.

Set Prefix: The specified processor's prefix register is set to the value passed to it by the control program. If the specified processor is not stopped, it does not accept the order. This function is not necessarily completed during the execution of the SIGP instruction.

Remote class services are defined for those control program functions that require the execution of a software function on one of the configured processors. The remaining SIGP functions are defined as remote services. External call is a remote pendable service that can be invoked via the RPSGNL macro and emergency signal is a remote immediate service that can be invoked via the RISGNL macro. A description of these services follows:

External Call: An "external call" external-interruption condition is generated to the specified processor. The interruption condition becomes pending during the execution of the SIGP instruction. The associated interruption occurs when the processor is interruptible for that condition. Only one external-call condition can be kept pending in a processor at a time. Issue the RPSGNL macro instruction to invoke the external-call function.

Emergency Signal: An "emergency-signal" external-interruption condition is generated at the specified processor. The interruption condition becomes pending during the execution of the SIGP instruction. The associated interruption occurs when the processor is interruptible for that condition. At any one time the receiving processor can keep pending one emergency-signal condition for each processor of the multiprocessing system, including the receiving processor itself. Issue the RISGNL macro instruction to invoke the emergency signal function.

Status Indicators

If the user receives a return code of 8 from the DSGNL macro, register 0 contains a status indicator describing the state of the specified processor. The status indicators describe the following conditions:

Equipment Check: This condition exists when the processor executing an instruction detects equipment malfunctioning that has affected only the execution of the instruction and the associated hardware function. The order code may or may not have been transmitted, and may or may not have been accepted, and the status bits provided by the specified processor may be in error.

Incorrect State: This condition exists when an order has been directed to a processor that is not stopped. The condition, when present, is indicated only in response to status or prefix.

Invalid Parameter: This condition exists when an address exception occurs. This happens when the storage referenced by the status or prefix function is not installed or not configured on the system. The condition, when present, is indicated only in response to status or prefix.

External Call Pending: This condition exists when an external-call interruption condition is pending in the specified processor because of a previously issued SIGP instruction. The condition exists from the time an external-call function is accepted until the resulting external interruption is accepted. The condition may exist on the issuing processor or another processor. The condition, when present, is indicated only in response to sense and to external call.

Stopped: This condition exists when the specified processor is in the stopped state. The condition, when present, is indicated only in response to sense.

Operator Intervening: This condition exists when the specified processor is executing certain operations initiated from the console or the remote operator control panel. The particular manually initiated operations that cause this condition to be present depend on the model and on the specified functions. This condition, when present, can be indicated in response to all functions. Operator intervening is indicated in response to sense if the condition is present and precludes the acceptance of any of the installed orders (or SIGP hardware functions). The condition might also be indicated in response to unassigned or uninstalled orders.

Check Stop: This condition exists when the specified processor is in the check-stop state. The condition, when present, is indicated only in response to sense, external call, emergency signal, start, stop, restart, and stop and store status. The condition may also be indicated in response to unassigned or uninstalled functions.

Invalid Function: This condition exists during the communications associated with the execution of SIGP when the specified processor decodes an unassigned or uninstalled function code.

MSSF Failure: This condition exists when the MSSF (3082) is currently inoperative. The MSSF performs the SIGP between two processors.

Receiver Check: This condition exists when the specified processor detects malfunctioning of equipment during the communications associated with the execution of SIGP. When this condition is indicated, the function has not been initiated and, because the malfunction may have affected the generation of the remaining receiver status bits, these bits are not necessarily valid. A machine-check condition may or may not have been generated at the specified processor.

Writing and Deleting Messages (WTO, WTOR, DOM, and WTL)

The WTO and WTOR macro instructions allow you to write a message to a display device or a printer at the operator console. Besides writing a message, WTOR allows you to request a reply from the operator who receives the message. The DOM macro instruction allows you to delete a message that is already written to the operator. The standard printable EBCDIC characters that constitute messages are shown in Figure 21. All other characters, which are not printable, are replaced by blanks. If the terminal does not have dual-case capability, it prints lowercase EBCDIC characters as uppercase EBCDIC characters.

Hex Code	EBCDIC	Hex Code	EBCDIC	Hex Code	EBCDIC	Hex Code	EBCDIC
40	(space)	7B	#	99	r	D5	N
4A	¢	7C	@	A2	s	D6	O
4B	.	7D	'	A3	t	D7	P
4C	<	7E	=	A4	u	D8	Q
4D	(7F	"	A5	v	D9	R
4E	+	81	a	A6	w	E2	S
4F		82	b	A7	x	E3	T
50	&	83	c	A8	y	E4	U
5A	!	84	d	A9	z	E5	V
5B	\$	85	e	C1	A	E6	W
5C	*	86	f	C2	B	E7	X
5D)	87	g	C3	C	E8	Y
5E	;	88	h	C4	D	E9	Z
5F	⌋	89	i	C5	E	F0	0
60	-	91	j	C6	F	F1	1
61	/	92	k	C7	G	F2	2
6B	,	93	l	C8	H	F3	3
6C	%	94	m	C9	I	F4	4
6D	-	95	n	D1	J	F5	5
6E	>	96	o	D2	K	F6	6
6F	?	97	p	D3	L	F7	7
7A	:	98	q	D4	M	F8	8
						F9	9

Figure 21. EBCDIC Characters Printed or Displayed on an MCS Console

Notes:

1. If the display service or printer is defined to JES3, the following characters are translated to blanks:

| ! ; ⌋ : "

2. The system recognizes the following hexadecimal representations of the U.S. national characters: @ as X'7C'; \$ as X'5B'; and # as X'7B'. In countries other than the U.S., the U.S. national characters represented on terminal keyboards might generate a different hexadecimal representation and cause an error. For example, in some countries the \$ character generates a X'4A'.

Routing the Message

The ROUTCDE parameter allows you to specify the routing code or codes for a WTO and WTOR message. The routing codes determine which MCS console or consoles receive the message. Each code represents a predetermined subset of the consoles that are attached to the system, and that are capable of displaying the message. It is up to the user to define the consoles that belong to each routing code. WTO and WTOR allow routing codes from 1 to 128. Routing codes 29 through 41 are reserved, and are ignored if specified. Routing codes 42 through 128 are available to authorized programs only, although the ROUTCDE parameter itself is available to non-authorized as well as authorized users.

The parameters, MSGTYP and MCSFLAG, which are associated with message routing, should only be used by programmers familiar with multiple console support (MCS). The MSGTYP parameter is typically used for messages related to the monitor command. The MCSFLAG parameter is used to specify various attributes of the message, such as:

- Whether the message is for a particular console
- Whether the message is for all active consoles
- Whether the message is a command response
- Whether the message is for the hardcopy log

The MCSFLAG = BUSYEXIT parameter determines what happens if no message buffers are available. If BUSYEXIT is specified and there are no console buffers for either MCS or JES3 are available, or if BUSYEXIT is specified and there is a JES3 WTO staging area excess, the WTO is terminated. If BUSYEXIT is not specified, the WTO invocation will be placed in a wait state until WTO buffers are again available. BUSYEXIT is available to authorized programs only.

When MSGTYP = Y is specified, the MCSFLAG field indicates that the MSGTYP field is to be used for the message routing criteria. In this case, the issuer of the WTO(R) must set a message identifier bit in the MSGTYP field of the macro expansion. The macro expansion is mapped by IEZWPL. When a message identifier bit is recognized by the system, the message is routed to all consoles and TSO terminals (in operator modes) that have requested the particular type of information represented by the identifier bit. If there are no consoles or terminals requesting that kind of information, a WTO message is not sent anywhere; however, a WTOR message is sent to the master console. The routing codes and REG0 MCSFLAG field, if present, are ignored when MSGTYP = Y is specified.

Another alternative for routing a message is to use the CONSID parameter. This parameter lets you specify a field or register that contains the four-byte id of the console that is to receive the message. This is a handy alternative to the MCSFLAG option of placing the console id in register zero. When you issue a WTO or WTOR macro that uses both the CONSID and the ROUTCDE parameters, the message(s) will go to all of the consoles specified by both parameters.

Notes:

1. *By using the various parameters of WTO(R), messages can be routed by route code, console id, descriptor code, and message type. Messages can be sent on multiple paths. See the description of the WTO(R) macro instruction in Volume 2 for additional information.*
2. *For the convenience of the operator, messages can be associated with individual keynames. A keyname consists of 1 to 8 alphanumeric characters, and it appears with the message on the console. The keyname can be used as an operand in the D R console command, which operators can issue at the console. Use the KEY parameter on the WTO or WTOR macro for this purpose.*

Writing a Multiple-Line Message

Messages consisting of multiple lines should be issued using the WTO multiple-line capability to assure that all lines of a multiple-line message appear together and are not broken up by other single-line messages.

Using the WTO macro instruction, a program can write a multiple-line message to one or more operator consoles. System programs (supervisor state, PSW key 0-7, or APF-authorized) can create a message that consists of up to 255 lines with one WTO request. If more than 255 lines are needed, the authorized user can issue more than one WTO.

When issuing more than one request, the first WTO supplies the first lines of the message, up to a limit of 255 lines. Subsequent WTO requests can then add lines to the message. The additional lines appear at the end of the message and continue until an "END" line is specified. For the first request, you must ensure that the left most three bytes of register zero are zero. If the bytes are not zero, WTO assumes that the multiple-line request is adding lines to an existing message, and no new message is created.

After processing the first request, the system places a message identifier in register 1. For each additional request, you must pass this identifier to the subsequent lines via the CONNECT parameter of WTO.

Embedding Label Lines in a Multiline Message

Label lines provide column headings in tabular displays. You can change the column headings used to describe different sections of a tabular display by embedding label lines in the existing multiline WTO message for a tabular display.

System programs (supervisor state, PSW key 0-7, or APF-authorized) that are authorized to add lines to an existing multiline WTO message are also permitted to embed label lines within that existing multiline WTO message. The label line does not have to appear immediately following the control line and before the data lines. At most two label lines can appear consecutively without an intervening data line.

Note: You cannot use the WTO macro instruction to embed label lines. The WTO macro instruction handles label lines at the beginning of the message only.

Using the Authorized Parameters of WTO and WTOR

CONNECT, JOBID, JOBNAME, SUBSMOD, SYSNAME, PRTY, and SYSNAME are authorized parameters used with WTO or WTOR.

The CONNECT parameter is used to connect a subsequent message to a previous message. For example, if your program develops a large, multi-line message of unknown length, it can issue different WTOs for the different parts of the message at different times. The CONNECT parameter can force all these WTOs to use the same message id, and cause all the different parts of the message to be physically reunited at the display console as a single message. To use CONNECT, you save the message id that is returned in general register 1 after you issue the first part of the multi-line message. On subsequent invocations of WTO for the remaining message parts, you supply the returned message id as an input parameter by using the CONNECT parameter. The end-of-message character in the text allows the system to recognize the last WTO in the sequence, and to start issuing the message. CONNECT is mutually exclusive with CONSID and SYSNAME, and it is not available with WTOR.

The SYSNAME parameter, which is mutually exclusive with CONNECT, is used to provide an eight-byte system name that appears on the console with the message. SYSNAME is available with WTO and WTOR.

The WQEBLK parameter, which is similar in principle to the DOMCBLK parameter of the DOM macro, is used when the control information for WTO(R) is in a table instead of in the input parameters. WQEBLK is mutually exclusive with all other parameters. As an example of how WQEBLK can be used, an application might capture the internally-generated control table resulting from the invocation of a previous WTO(R). Then it might supply this table as an input parameter in the subsequent WTO(R), by using the WQEBLK parameter.

The JOBNAME and JOBID parameters are used by JES2 and JES3 programmers to correlate WTO or WTOR macros and their resulting messages with the jobs that are passing through the system.

The SUBSMOD parameter is used to indicate whether the message can be modified by a subsystem, and the PRTY parameter is used to give the message a priority that is visible when the message appears on the console. These two parameters are only available with WTOR.

Deleting Messages Already Written

The DOM macro instruction allows you to delete an operator message or group of messages from a display that is to appear on the operator console. The messages that DOM deletes are the same messages that were written earlier via the WTO or the WTOR macro.

When a message is displayed on the operator screen, it has a certain format, which indicates to the operator whether or not the operator has responded yet to the message. When the operator responds to a message, the message format changes to remind the operator that a response was already given. The actual message, however, remains displayed until it is rolled off the screen. When DOM deletes a message, it does not actually erase the message. It only changes its format, making it appear as if the operator has already responded to it.

Depending on the timing of a DOM macro instruction relative to the WTO or WTOR, the message may not yet be on the operator screen. If the message is already on the screen, it is deleted. If the message is not yet on the screen, DOM deletes the message before it appears.

Identifying Messages to be Deleted

To identify the message(s) that you want to delete, you normally use the MSG, MSGLIST, or TOKEN parameters. When you issue a WTO or WTOR macro instruction to write a given message to the operator, the system generates a message id, which it returns in general register 1. To delete the message, you can issue the DOM macro instruction with a MSG or MSGLIST parameter specifying the same system-generated message id that WTO or WTOR returned in general register 1. If you specify MSGLIST (message list), then several message ids can be associated with the delete request. The number of message ids in the message list is defined by the COUNT parameter or it is defined by an 1 in the high order bit position of the last message id in the list. The count parameter cannot exceed 60.

On the other hand, the TOKEN parameter allows the message id to be generated by the user rather than the system. When you issue WTO or WTOR with a TOKEN parameter, the system associates your TOKEN parameter with all the message(s) that are written by this particular WTO or WTOR. Then you can issue DOM with the same TOKEN parameter to delete all the message(s) associated with the token.

Limiting the Extent of Message Deletion

DOM allows you to limit the extent of message deletion. One way to limit the extent of message deletion is to use the `SYSID` parameter, which deletes only messages from a particular system whose id you specify. Another way to limit the extent of deletion is to use the `SCOPE` parameter. If you specify `SCOPE=SYSTEMS`, the delete request is sent to all other processors. To delete messages only within the host system, you can specify `SCOPE=SYSTEM`. `SYSID` and `SCOPE` can only be coded by authorized users.

Custom-Built Delete Functions

When DOM executes, it builds a control block from the specified parameters to control the execution of the delete. All the input parameters specified on the DOM macro are represented by fields in this control block. However, the control block can be built directly by an authorized user. To build the DOM control block without specifying parameters, simply specify the address of the control block by using the `DOMCBK` parameter. When this parameter is specified, the system does not build any delete control block, and it substitutes the user-provided control block instead.

Note: Specifying the `REPLY=` parameter of the DOM macro causes an MNOTE warning message to be issued at assembly time. The MNOTE warns you that you are coding the `REPLY=` parameter, which is a function no longer supported in the system. If you code the `REPLY=` parameter and receive the MNOTE warning, remove the `REPLY=` parameter from your program and reassemble it. Programs containing the `REPLY=` parameter that are already assembled do not need to be reassembled.

Writing to the System Log

The WTL macro instruction causes a message to be written to the system log. The message can include any character that can be used in a C-type (character) DC statement, and is assembled as a variable-length record. An optional prefix specifies whether the WTL text contains an identifier for the SYSLOG record. If this prefix option is omitted, a 2-character prefix is automatically added.

Inter-Address Space Communication

There are many advantages to the use of multiple virtual address spaces in MVS/XA. Virtual addressing permits an addressing range that is greater than the real storage capabilities of the system. The use of multiple virtual address spaces provides this virtual addressing capability to each job in the system by assigning each job its own separate virtual address space. The potentially large number of address spaces provides the system with a large virtual addressing capacity.

With multiple virtual address spaces, errors are confined to one address space, except for errors in commonly addressable storage, thus improving system integrity and making error recovery easier. Programs in separate address spaces are protected from each other. Isolating data in its own address space also protects the data. In addition, having a separate address space for data increases the amount of data that can be addressed.

However, multiple virtual address spaces require efficient ways to communicate between address spaces. There are two basic methods of inter-address space communication: asynchronous and synchronous.

Asynchronous Inter-Address Space Communication

MVS/XA can use a dispatchable unit of work, the service request block (SRB), for asynchronous communication between address spaces. SRBs are most efficient when there is no need to synchronize the processes in the two address spaces.

Scheduling an SRB (SCHEDULE)

An SRB represents a request to execute a routine that performs a particular function or service in a specified address space. Supervisor state, PSW key 0 functions inform the system of such a request by using the SCHEDULE macro instruction to schedule an SRB. SCHEDULE places the SRB on a dispatcher queue to be dispatched when it becomes the highest priority work in the system. The caller must specify the priority of the request relative to other requests in the system. This can be either local or global.

The SRB is a control block that must be allocated from commonly addressable fixed storage either above or below the 16 megabytes line. The SCHEDULE macro does not obtain storage for the SRB. The fields in the SRB identify the routine to be executed, the address space in which it is to execute, the priority level of the SRB, and additional information for recovery and control. SRB routines should be functions that require only a few instructions to complete because an SRB routine is not preempted by I/O interruptions once the SRB is dispatched. After an I/O interruption is processed, the I/O interruption handler returns to the interrupted SRB routine instead of passing control to the dispatcher to select the highest priority ready work.

Although SRBs can also be used for synchronous communication between address spaces, their major advantages are that they are basically asynchronous in nature and execute independently of the routine that schedules them. Once the SRB has been scheduled, it is usually available for immediate dispatch. These attributes make SRBs very useful in the following situations where the routine scheduling the SRB does not require synchronizing with the completion of the SRB routine:

- **Parallelism:** In a multi-processor environment, the SRB routine, after being scheduled, can be dispatched immediately on another processor and can execute concurrently with the routine that scheduled it. The scheduling routine can continue to do other processing in parallel with the SRB routine.
- **Serialization Avoidance:** Because the SRB represents a separate unit of work, the unit of work that schedules the SRB routine is not serialized, delayed, or encumbered by any delay the SRB routine might encounter in completing its function. The following types of delays can usually be avoided:

- Page fault resolution
- Address space swap-ins
- Lock suspensions - wait time

- **Resource Accounting:** Because the SRB represents a separate unit of work, the processor time spent accomplishing that work can be charged to the address space in which it is executing.
- **State Changes:** In some instances, a routine might be executing in some state (disabled, for example) that prevents certain functions from being performed (requesting a suspend-type lock or referencing pageable storage, for example). These restrictions can be avoided by scheduling an SRB to complete the function.

- **Priority Association:** Because the SRB represents a separate unit of work, the SRB has its own dispatching priority. The SRB can execute at a priority higher than that of any address space or at the priority of the address space in which it is scheduled.

The mechanisms required to use SRBs follow.

Service Management Facilities

Service management facilities provide a basic set of system services that allow internal system components to structure themselves to run enabled, nonserialized, and in parallel on a multiprocessing system with considerably less overhead than ordinary task management services require. Service management facilities:

- Introduce a service request to execute a service routine into the queue of work within the system.
- Perform priority dispatching of the requested service routine.
- Support the needs of recovery/termination for cleaning up the asynchronous processes.

The SCHEDULE macro instruction performs the first two services; the PURGEDQ macro instruction performs the third.

Service management operates as follows:

- A control block, called a service request block (SRB), represents a service request. This block is similar to a TCB but requires less information. It identifies a unit of work to the dispatcher.
- The SCHEDULE macro instruction enters SRBs into the dispatchable work queue with a minimum of overhead.
- The dispatcher operates under the service request control structure in addition to the task structure. In addition, SRBs can be scheduled to different address spaces either at a priority equal to the specified address space or at an independent priority higher than that of any address space.
- The service management facilities provide an increase in parallel processing on multiple processors. The programs benefiting most from these facilities are those that have independently dispatchable units of work, but are forced to run as a single task. In addition, service management facilities make the system aware of the smaller units of work and allow the service requests to be dispatched in parallel on multiple processors at a higher priority than the task work.

The service management facilities also provide a mechanism that is used for asynchronous communications between address spaces. It is used to run some parts of interruption handlers as service requests, allowing more enablement and parallelism for these services. For example, when an I/O interruption occurs, the interruption handler collects the necessary information about the I/O interruption and schedules a service request block. The interruption handler can then start I/O requests that were waiting for the I/O path and accept any additional pending interruptions. Delaying complete processing of the interruption allows faster reuse of channels and lower disabled interruption time.

The scheduling of the SRB provides the ability to:

- Complete the interruption process on any processor, not just on the one that took the interruption.
- Process the interruption enabled except where specific serialization through locks is used.
- Switch from the random address space where the interruption was taken to the address space of the user who originally requested the I/O. This latter capability provides the interruption handler routine with addressability to the user's control blocks that is necessary to complete the interruption processing.

Service Request Blocks (SRBs)

Service requests are represented by service request blocks (SRBs). SRBs must reside in fixed, commonly-addressable storage. These control blocks are supplied by the function requesting a service. The basic information contained in an SRB is documented in the *Debugging Handbook*. The information required in the SRB at the time the SCHEDULE macro instruction is issued as follows:

SRBASCB	contains the address of the ASCB indicating the address space in which the asynchronous routine will be dispatched.
SRBPKF	indicates, in the high-order 4 bits, the PSW protect key that the routine assumes. The low-order 4 bits must be zero.
SRBEP	specifies the address of the entry point of the asynchronous routine. This field is pointer defined. See the note below for the definition of a pointer defined field.
SRBSAVE	contains the address of a status save area. This field must contain all zeroes.
SRBPARAM	contains a fullword to be loaded into register 1 when the asynchronous routine is dispatched.
SRBCPAFF	defines the processor affinity. If all zeroes or all ones, no affinity is implied. Otherwise, this field contains a bit mask in which the bits that are set "on" to indicate on which processors the SRB can be dispatched. (The nth bit set "on" indicates that the SRB can be dispatched on the processor with physical address n.)
SRBRMTR	contains the address of a resource manager termination routine. This routine is responsible for cleaning up an SRB that has been scheduled but not yet dispatched. The address also serves as a function identifier to PURGEDQ when purging SRBs from the queues. Because it is possible for the SRB to be purged, either directly by a system or user resource manager, or indirectly by task or memory termination, this field must contain a valid nonzero address of a routine to clean up this SRB. The routine must be commonly-addressable by all address spaces. This field is pointer defined. See the note below for the definition of a pointer defined field.
SRBPTCB	contains the address of a TCB associated with the asynchronous routine, which serves two purposes: <ul style="list-style-type: none">• If the asynchronous routine encounters a double failure (error in routine and in recovery routine), the task is scheduled for abnormal termination.• Identifier to PURGEDQ during the purge process. If this SRB is to be purged during the termination of that task, the address of the associated TCB must be specified in this field. If this SRB is not related to any task, or purging is not necessary, specify a zero value.
SRBPASID	contains the ASID of an address space associated with the asynchronous routine. If a nonzero value was specified in the SRBPTCB field, this value must be specified and must contain the ASID of the address space containing that TCB. Otherwise, this field can be zeroes.

Note: The SRBEP and SRBRMTR fields are pointer defined. That is, the first bit indicates whether the routine is to execute in 31-bit or in 24-bit addressing mode. If the bit is set to 1, the routine executes in 31-bit addressing mode; and if the bit is set to 0, the routine executes in 24-bit addressing mode.

Priorities

Services can be scheduled for execution at either global or local priorities. Service requests queued at the global level are given a priority that is above that of any address space, regardless of the actual address space in which they are dispatched. Service requests at the local level are given a priority equal to that of the address space in which they are dispatched, but higher than that of any task within that address space.

Within the global and local priorities, there are two additional priority levels. One of the levels is for general system use (system level); the other (nonquiesceable level) is for specialized functions necessary to perform a quiesce of SRBs.

Nonquiesceable level service requests continue to be dispatched while the address space is being quiesced. System level service requests are not dispatched while the quiesced status is in effect. Use of the nonquiesceable level is restricted in the following manner: At times, it is necessary to stop the dispatching of SRBs in an address space -- that is, to prevent new SRBs from being dispatched and allow all SRBs that have already been dispatched to complete their processing. However, some of the dispatched SRBs might have been suspended because of lock requests or page faults. Because page fault processing and rescheduling of the suspended SRBs themselves make use of SRBs, it is necessary to have a nonquiesceable level at which these SRBs can be scheduled while the other SRBs have been stopped.

Characteristics of Service Requests

Service routines have the following characteristics:

- Routines are entered in supervisor state, enabled and unlocked. The routines cannot enter problem program mode and must establish their own recovery environment. However, by specifying the LOCAL lock required option and/or the FRR required option on the SCHEDULE macro instruction, a service routine is dispatched with the LOCAL lock held and its recovery environment already established. Upon entry to the service routines, the registers contain the following:

Register	Contents
0	address of the SRB
1	user field
14	return address
15	entry point address

- The routines cannot issue SVCs except ABEND.
- The routines are non-preemptable -- that is, although the routines run enabled and can be interrupted by an asynchronous interruption, they do not lose control to higher priority tasks or SRBs until they give up control voluntarily. However, service routines might lose control because of synchronous events that cause suspension of the program in control, such as page faults and unconditional requests for suspend-type locks. In this case, full status of the process is saved and other work is dispatched; the service request is redispached when the problem is resolved.

- The routines can take page faults. Page faults encountered in the unlocked state are handled per service request. Page faults encountered in the locked state prevent other processing that requires the lock from proceeding until the page fault is resolved.
- The routines can request any lock through the locking interface (SETLOCK).
- The routines must return control to the address supplied in register 14, in supervisor state with no locks held. All cleanup must be performed before exiting.

Purging Service Requests (PURGEDQ)

When a task or address space terminates, all outstanding requests for the task or address space must also be terminated. The PURGEDQ macro instruction establishes a standard mechanism for purging these requests. To use this macro instruction, you must provide the following information:

- The address space identifier of the address space in which the SRB is scheduled to be dispatched. If none is specified, the current address space is assumed.
- The address space of the task control block of the task associated with the SRB for which the purge is to be performed. If none is specified, the current TCB is assumed in the current address space.
- The address of the resource manager termination routine. If no address is supplied, the purge is performed for all resource managers.

If the PURGEDQ was for the current address space, the PURGEDQ routine waits for completion of any active SRBs and then dequeues all nondispatched SRBs. After all of the SRBs have been dequeued or completed, the resource manager termination routine specified in the SRB is given control to perform the required cleanup for each dequeued SRB. No locks should be held when PURGEDQ is invoked.

PURGEDQ Parameters

The ASID, ASIDTCB, and RMTR parameters specify input to the PURGEDQ macro instruction.

The ASID parameter specifies the address of a half word containing an address space identifier. PURGEDQ searches for SRBs scheduled to be dispatched into the address space specified by this parameter. If an address space other than the current address space is indicated, only SRBs that have not yet been dispatched are affected because PURGEDQ does not wait for SRBs already dispatched but not completed. If this parameter is omitted, the current address space is assumed.

The PURGEDQ caller can purge the SRBs associated with a specific task by coding the ASIDTCB parameter. The ASIDTCB parameter specifies the address of a doubleword that describes the TCB for which SRBs are to be purged. If the parameter is omitted, the purge

occurs for SRBs associated with the current task in the current address space. The following table describes three acceptable values for the ASIDTCB parameter, and the meaning of the values:

First:			
Bytes 0 - 7	zero		All SRBs that meet the criteria specified in the ASID and RMTR parameters are to be purged regardless of their task association.
Second:			
Bytes 0 - 1	reserved		All SRBs that meet the criteria specified in the ASID and RMTR parameters and are associated with the specified address space (SRBPASID field) are to be purged.
Bytes 2 - 3	nonzero ASID		
Bytes 4 - 7	zero		
Third:			
Bytes 0 - 1	reserved		All SRBs that meet the criteria specified in the ASID and RMTR parameters and are associated with the specified address space (SRBPASID field) and the specified task (SRBPTCB field) are to be purged.
Bytes 2 - 3	nonzero ASID		
Bytes 4 - 7	nonzero TCB address		

All other values are unacceptable and produce unpredictable results.

The RMTR parameter specifies the address of the resource manager termination routine (RMTR). Only those SRBs whose SRBRMTR field contains the address of the RMTR will be purged.

The interface to the RMTR is as follows:

Register	Contents
0	Contents of register 0 of the caller of PURGEDQ at the time the PURGEDQ SVC was issued. This provides the ability to pass information from the caller of PURGEDQ to the RMTR routine.
1	Address of the dequeued SRB.
2	Contents of SRBPARM of the dequeued SRB.
14	Return address of PURGEDQ.
15	Entry point of RMTR.

The PSW is enabled, in supervisor state with PSW key 0 and no locks held.

The RMTR must not leave supervisor state. The RMTR must return control using a BR 14, enabled, in supervisor state with PSW key zero and no locks held. It may, however, acquire locks, issue SVCs and destroy input registers.

This routine must be commonly-addressable from all address spaces. It must be unique to the resource manager (that is, it should not be the address of some common system service) if the PURGEDQ issuer wants the routine to receive control only on a unique PURGEDQ. These restrictions are required to ensure the uniqueness of the identifier and to allow the RMTR to be invoked from any address space.

Synchronous Inter-Address Space Communication

MVS/XA provides a synchronous method of communication between address spaces that is called cross memory. Using cross memory, programs can pass control directly to programs in other address spaces and can move data directly from one address space to another. At any time, a program has associated with it two addressable address spaces, the primary address space and the secondary address space. It is between these two address spaces, which may be the same, that synchronous communication occurs.

Three of the ways cross memory can be used are program sharing, data movement, and data access.

Program Sharing: A program residing in the private area of a particular address space can be directly called by programs residing in a number of different address spaces. These address spaces must have such calling ability defined for them. MVS/XA provides a set of macro instructions to establish the required access structures. Thus, a service needed by a number of address spaces no longer has to reside in commonly addressable storage, or have multiple images in many address spaces.

Data Movement: A program can move data directly between the primary and secondary address spaces, within an address space, and between storage areas of differing storage protection keys. Thus, programs can pass data directly between address spaces.

Data Access: A program residing in commonly addressable storage can choose to access data from either the primary or secondary address space. Being able to access data in two address spaces increases the amount of data that the program can handle.

Program sharing, data movement, and data access enable cross memory to provide:

- Storage isolation and protection of code and data structures. By moving their programs and data structures from commonly addressable storage to their own address spaces, MVS/XA components and subsystems avoid the accidental destruction of their information by unrelated processes.
- Migration of code and data from commonly addressable storage to private storage. With code and data migrating to private storage, the common storage requirements of the system decrease. This decrease effectively expands the private area addressing range and provides each user with more virtual storage.
- Creation of data address spaces. Partitioning data into address spaces isolates sensitive data and provides restricted access to it. This also enables a program to address greater amounts of data.

Cross Memory

Cross memory is a very complex concept, and there are a number of warnings and restrictions associated with its use. Before listing the restrictions, however, some definitions are needed.

Cross Memory Terminology: The following terms are associated with cross memory.

- Cross memory environment: The environment in which synchronous inter-address space communication can take place.
- Home address space: The home address space, whose address space identifier (ASID) is called the HASID¹, is the address space defined by PSAAOLD. The home address space contains the address space local control blocks that describe a unit of work to the control program. On initial dispatch of a unit of work, the home address space and the primary address space are the same.

¹ *Principles of Operation* uses the terms HASN, PASN, SASN, and primary-space mode. MVS/XA publications use the equivalent terms HASID, PASID, SASID, and primary mode.

- **Primary address space:** The primary address space, whose ASID is called the PASID², is the address space whose segment table is used to access data and instructions in primary mode.
- **Secondary address space:** The secondary address space, whose ASID is called the SASID², is the address space whose segment table is used to access data in secondary mode. In secondary mode, instructions must be in common storage because they might be fetched from either the primary or secondary address space.
- **Current address space:** The current address space is the primary address space when in primary mode and the secondary address space when in secondary mode.
- **Primary mode:** In primary mode, instructions and data are fetched from the primary address space.
- **Secondary mode:** In secondary mode, data is fetched from the secondary address space. Instructions might be fetched from either the primary or secondary address space.
- **Home mode:** A unit of work is in home mode if it is in primary mode and HASID and PASID are the same.
- **CML lock:** The cross memory local (CML) lock is the local level lock of an address space other than the home address space.
- **Cross memory mode:** Cross memory mode exists when at least one of the following is true:
 - The home address space is not the primary address space.
 - The home address space is not the secondary address space.
 - Secondary mode is active.
 - A CML lock is held.
- **Active addressing bind to an address space:** An executing unit of work has an active addressing bind to an address space if that address space is the current PASID or SASID.
- **Active bind to an address space:** An executing unit of work has an active bind to an address space if the unit of work holds the CML lock of that address space or if the address space is associated with the current HASID, PASID, or SASID.

A series of macro instructions create a cross memory environment. The macro instructions establish the necessary linkage and authorization information for synchronous inter-address space communication. The following System/370-XA instructions actually accomplish the communication:

- **PC - program call** - causes another program to get control. The program can be in another address space.
- **PT - program transfer** - returns control from the program called by the PC instruction to the calling program.

² *Principles of Operation* uses the terms HASN, PASN, SASN, and primary-space mode. MVS/XA publications use the equivalent terms HASID, PASID, SASID, and primary mode.

- SSAR - set secondary ASN³ - sets the secondary address space to any desired address space.
- MVCP - move to primary - moves data from the secondary address space to the primary address space.
- MVCS - move to secondary - moves data from the primary address space to the secondary address space.
- MVCK - move with key - moves data between storage areas that have different protection keys.
- SAC - set address space control - explicitly sets either the primary or secondary mode.
- IAC - insert address space control - indicates in a general purpose register whether primary or secondary mode is in effect.
- EPAR - extract primary ASN³ - places the primary ASID into a general purpose register.
- ESAR - extract secondary ASN³ - places the secondary ASID into a general purpose register.

Warnings and Restrictions: The design and implementation of programs using synchronous cross memory communication is extremely complex. System services use cross memory on a user's behalf; the user can obtain the benefits of cross memory without having to know the details. Using cross memory services improperly could cause severe system problems. Therefore, it is very important to consider all the implications of using cross memory. Some general considerations that apply to users of cross memory are:

- Real storage requirements might increase.
- Resource management is different.
- Accounting methods might be affected.

Cross memory has the following specific restrictions:

- Services are not available in cross memory mode unless their description specifically states that they are available.
- Code running in cross memory mode cannot issue any SVCs except ABEND. That is, any system service that depends on SVCs is not available in cross memory mode. TSO test, for example, provides only limited testing for programs that execute in cross memory mode because TSO test uses the TEST SVC.
- Only one step of a job can establish ownership of space switch entry tables. Subsequent job steps cannot issue the LXRES, AXRES, or ETCRE macro instructions.

³ *Principles of Operation* uses the term ASN. MVS/XA publications use the equivalent term ASID.

- MVS/XA does not support cross memory accesses to a swapped-out address space; such accesses cause an ASID translation exception-program interruption that is treated as an error. Thus, in order to be accessed, the address space must be one of the following:

- The home address space
- A non-swappable address space
- An address space whose local lock is held

This restriction must be a major consideration when using cross memory because it might increase the storage requirements of the system.

- Some MVS/XA services require an active addressing bind to the address space in which processing is to occur. Such an address space must be one of the following:

- The home address space
- A non-swappable address space
- An address space whose local lock is held

If none of these three requirements are met, the address space might be swapped out and a unit of work that referenced the address space would be abnormally terminated.

- Storage acquired in a cross memory environment is attributed to the job step task of the address space in which it was obtained if the subpool it comes from is task related. A program that acquires such a resource should provide a task termination/address space termination resource manager to clean up any resources obtained on behalf of the terminating task or address space but attributed to another address space's job step task. For more considerations on resource management see "Designing a PC Routine" later in this section.
- Execution time is attributed to the home address space, not necessarily the address space in which the program executes.
- Routines that get control as the result of a PC instruction must not use the checkpoint/restart facility.

Summary of MVS/XA Facilities Available in Cross Memory Mode

The MVS/XA facilities available in cross memory mode can be divided into two categories: those services that are available to cross memory mode programs without restriction and those services that have special cross memory options or restrictions associated with their use. A list of the macro instructions available without restrictions to cross memory mode callers and a list of the macro instructions that have special options or restrictions for cross memory callers are provided in Volume 2 under the topic "Cross Memory Restrictions for Macro Instructions."

In addition to the services provided by the macro instructions, the following functions are available to cross memory mode programs without restriction:

Segment and page faults - The system function of resolving segment and page faults is supported for a unit of work executing in cross memory mode.

Dispatcher, interrupt handling - The MVS/XA dispatcher and interrupt handlers, except the SVC interrupt handler, support programs executing in cross memory mode, and these functions save and restore the additional status required by cross memory mode programs.

System tracing also traces cross memory information.

The macro instruction descriptions in Volume 2 give details about cross memory support.

Cross Memory Structures

Cross memory uses a set of programming and data structures that can be divided into three functional areas: cross memory authorization, cross memory linkage, and linkage conventions.

“Cross Memory Authorization” describes how address spaces and programs are authorized to use PT, SSAR, and PC instructions and how the user can request that the system provide this authorization.

“Cross Memory Linkage” describes the structures and tables used by the PC instruction and how a user can request that the system create and connect these structures to particular address spaces.

“Linkage Conventions” describes a set of programming conventions that must be used to preserve register information and maintain system serviceability when using cross memory.

Cross Memory Authorization: Cross memory uses a more flexible authorization mechanism for inter-address space communication than the PSW key zero, supervisor state requirement for scheduling SRBs. There are multi-level authorization facilities that permit both supervisor and problem state programs in an address space to access programs and data in a selected set of address spaces, and also to restrict an individual problem state program’s access to only a selected set of programs in other address spaces.

Programs have a selected set of PSW keys to which they are authorized, and, in problem state, this set controls the program’s authority to access data in the secondary address space. Users request authorization by invoking a series of macro instructions.

Address Space Authorization: An address space’s authorization to access other address spaces is based upon the authorization index (AX). Each address space has an AX. A program runs with the AX of the primary address space. The AX indicates the authority of the program to set another address space as its primary address space using the PT instruction and to set another address space as its secondary address space using the SSAR instruction.

The PT instruction is the mechanism used to return control from a PC routine (a routine that gets control as the result of a PC instruction). A program should use the PT instruction only to return to a program that called it using the PC instruction because instruction processing continues at the specified virtual address in the new primary address space, and system integrity and serviceability might be exposed by using the PT instruction in any other way.

Once a program has established another address space as its secondary address space, the program is authorized to move data between the secondary address space and the primary address space if the storage protection key of the data permits. The program can also directly reference data in the secondary address space by switching to secondary mode if the storage protection key of the data permits.

Each address space has associated with it an authorization table (AT). The AT contains one entry for every AX in use and indicates the authority of programs running with that AX to issue PT and SSAR instructions to the address space. The AX is used to index into the AT of the target address space on a PT or SSAR instruction to check if the issuing program has the authority to set the target address space as its primary or secondary address space. Authorization checking is required for both supervisor state and problem programs.

A particular address space, then, can selectively obtain PT and SSAR authority to a specific set of address spaces based on the ATs of those address spaces.

MVS/XA provides macro instructions to supervisor state or PKM 0-7 (the PKM is described below under "Program Authorization") programs to:

- Reserve an AX value for an address space (AXRES macro instruction)
- Free an AX value (AXFRE macro instruction)
- Set an address space's AX to a specified value (AXSET macro instruction)
- Set an address space's AT to indicate authorization levels for a specified AX value (ATSET macro instruction)
- Determine the AX value of an address space (AXEXT macro instruction)

All address spaces start with an AX of 0. An AX of 0 is an unauthorized AX value that prevents the address space from using PT and SSAR instructions. An AX of 1 is a fully authorized AX value that permits the address space to issue PT and SSAR instructions to any active address space. Certain system services that functionally serve all address spaces have an AX of 1. To have any other AX value, the user must explicitly reserve and set the AX with the AXRES and AXSET macro instructions, respectively. An address space to which this AX value is to be authorized (the address space to be accessed using PT and SSAR) must have its AT set using the ATSET macro instruction.

Program Authorization - PKM (PSW Key Mask): Each program has associated with it a PSW key mask (PKM) value. The PKM value can authorize individual programs to use cross memory. The PKM is a 16-bit string value that represents storage protection keys that are valid for the program to use, where bit n equal to 1 indicates that the program is authorized to use key n. The PKM is used only to perform authorization checking for problem state programs; supervisor state programs do not require PKM authority.

The PKM value is checked to see whether a problem state program can use the secondary access key specified on the MVCP and MVCS instructions to access storage in the secondary address space. It is also checked to see whether a problem program can use the secondary access key on the MVCK instruction.

The PKM value is also checked to see whether a problem program can issue a PC instruction. The PC instruction looks up an entry table entry (described later under "Cross Memory Linkage") that contains information for the PC instruction. Part of this information consists of an authorization key mask (AKM) value. The AKM is a 16-bit string value that indicates authorized keys in which a problem program can use a particular PC instruction. If the program's PKM indicates that it is authorized to use any of the keys indicated by the AKM, then the program can use the PC instruction. The PKM value also indicates whether a problem program can set a particular PSW key using the SPKA instruction.

All programs are initially dispatched with a PKM value equal to the bit mask representation of the field TCBPKF or SRBPKF. For example, X'0080' represents key 8 and X'8000' represents key 0. The PKM value can be changed using the PC and PT instructions and the MODESET SVC instruction.

The MODESET SVC sets the PKM value to the bit mask representation of the PSW key value when control returns to the program in problem state.

The entry table that contains information for the PC instruction also contains an execution key mask (EKM). The EKM is a 16-bit string value like the PKM and could contain additional keys to which the PC service is to be authorized. The EKM is ORed into the PKM when the PC routine receives control.

A program that issues a PT instruction specifies a 16-bit string value that indicates the PSW keys the program is authorized to use when the PT instruction is completed. This 16-bit string is ANDed with the original PSW key mask, and the result is placed in the PSW key mask in control register 3.

Cross Memory Linkage: Synchronous cross memory transfer of control is done with the program call (PC) instruction and the program transfer (PT) instruction. The PC instruction uses a PC number as input. The PC number is composed of two concatenated indexes, the linkage index (LX) and the entry index (EX). The PC instruction uses these indexes to perform a two level table lookup that causes a specific program to get control in the address space and mode specified in the table.

The first level table is the linkage table; the second level table is the entry table. Figure 22 shows how the PC number is used to access a particular entry table entry. The first portion of the PC number is the linkage index (LX), which selects a specific entry in the linkage table. The low order byte of the PC number is the entry index (EX), which is an index into the entry table pointed to by the linkage table entry. The entry table entry contains information that describes the program to receive control.

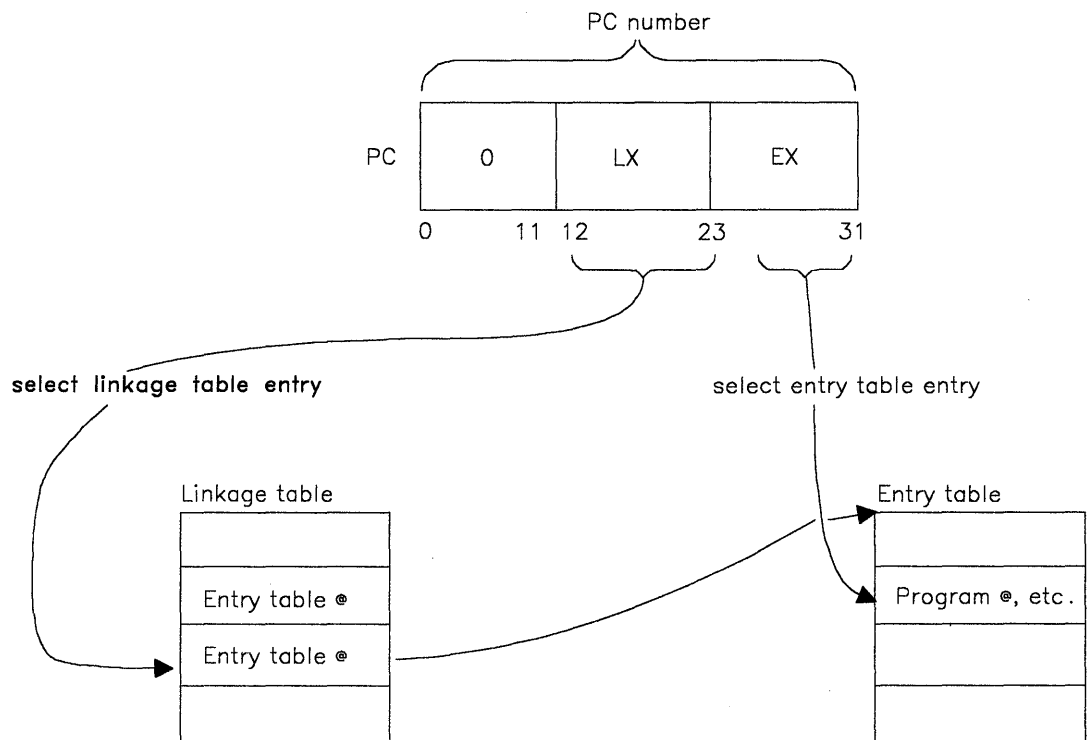


Figure 22. PC Number Indexing Linkage and Entry Tables

Linkage Tables: There is a linkage table associated with each address space in the system. An address space can have its own unique linkage table that gives it a set of cross memory services that is different from the set of services for any other address space. The linkage table can contain up to 1024 entries. Each linkage table entry can point to an entry table that describes a subset of the services available to the address space.

When a program wants to provide services via the PC instruction, it reserves an index into every linkage table in the system. The program connects an entry table to the reserved index for every linkage table whose address space is to have access to the services. To reserve an index, the program invokes the reserve linkage index service by issuing the LXRES macro instruction, which returns the reserved linkage index (LX). For further details, see "How to Establish a Cross Memory Environment" later in this section.

Entry Tables: Each program that provides services accessed via a PC instruction owns one or more entry tables. These entry tables are connected to the linkage tables of those address spaces that require access to the programs. Each entry in the entry table contains the following information about the program to be given control:

- Instruction address - specifies the addressing mode and virtual address in which the service is to receive control. (The addressing mode bit specifies the addressing mode of the called program as 24-bit or 31-bit.) For those entry table entries that do not describe user defined programs, this entry points to a special abend routine.
- ASID - specifies the ASID of the address space in which the called program will execute. If the value is zero, the program executes in the caller's primary address space.
- Problem state bit - specifies whether the called program will operate in problem state or supervisor state.
- Authorization key mask (AKM) - the AKM and the EKM are described earlier under "Cross Memory Authorization."
- Execution key mask (EKM) - the EKM and AKM are described earlier under "Cross Memory Authorization."
- Latent parameter address - specifies the address of a double word to be passed to the called program. The entry table creator supplies the first word. The second word is used by the PCLINK macro instruction. (PCLINK is described later in this section under "Linkage Conventions.")

An entry index (EX) is associated with each entry created in the entry table; the first entry has an EX of X'00' and subsequent entries have EXs of X'01' through X'FF'.

A program creates an entry table by issuing the ETCRE (create entry table) macro instruction, supplying all necessary details about the programs to receive control. These details go into the entry table. The ETCON and ETDIS macro instructions, respectively, connect and disconnect entry tables from linkage tables. The ETDES macro instruction destroys an entry table by removing it from the system. The uses of these macro instructions are described in greater detail in "How to Establish a Cross Memory Environment" later in this section. Figure 23 summarizes the macro instructions used to establish authorization and linkage.

Macro Instructions For Authorization	Function
AXRES (Reserve AX)	Reserve authorization index
AXFRE (Free AX)	Return an AX for reuse
AXEXT (Extract AX)	Determine the AX of an address space
AXSET (Set AX)	Set the AX for an address space
ATSET (Set AT)	Set PT and SSAR authority in an authorization table entry
Macro Instructions For Linkage	
LXRES (Reserve LX)	Reserve a linkage index
LXFRE (Free LX)	Return an LX for reuse
ETCRE (Create ET)	Create an entry table
ETDES (Destroy ET)	Destroy an entry table
ETCON (Connect ET)	Connect an entry table to a linkage table at the specified LX
ETDIS (Disconnect ET)	Disconnect an entry table from a linkage table

Figure 23. Authorization and Linkage Macro Instructions

PC Numbers: PC numbers are not permanently associated with a particular service the way SVC numbers are. The LX portion of the PC number is assigned by the control program and is not known before IPL. The EX portion is assigned by the component that owns a particular entry table. (While the component could make the EX portion of the PC number known by convention to the callers of its services, this is neither necessary nor desirable.)

Because the PC numbers themselves are not known before program execution, macro instructions and control program services cannot use PC numbers directly. Instead, PC numbers are determined indirectly by a table lookup process. For example, the PC numbers corresponding to many system functions are contained in a system function table (SFT) pointed to by the CVT. A macro instruction that invokes one of these PC services uses a permanently assigned index into the SFT to obtain the PC number for the service. A program that provides PC services must use a similar indirect method to give its callers the PC numbers they need to invoke its services. The caller of a service is not dependent on the actual PC number that is issued to obtain the service, on which module performs the service, or on where that module is located.

Linkage Conventions

In a cross memory environment linkage conventions are more important than in other environments because the “how did I get here” information is essential. Therefore, users must save and restore status and diagnostic information in a consistent way for every program call/program transfer sequence.

When a program gets control as the result of a PC instruction, and uses PT to return, there are several things to be aware of:

- The called program must preserve registers 3 and 14 in order to return control with PT.
- The called program must preserve the PSW key and program mask across the PC/PT interface.

- If there is a dump when the called program is executing, the following information might be needed for the dump:

Who called the currently executing program?
 What were the original contents of the caller's registers?
 Where is the caller's save area chain?

In order to preserve the above information, a program that is about to issue a PC instruction does the following:

- Saves registers 2 through 12 in words 8 through 18 of a standard save area pointed to by register 13. You must save registers before issuing a PC because the PC instruction updates registers 3, 4, and 14, and the address space where the save area resides might no longer be the currently addressable address space.
- Saves the current SASID in bits 16-31 of save area word 5.
- Optionally loads registers 0, 1, and 15 as parameter registers.
- Loads register 2 with a PC number.
- Issues a PC specifying register 2.

The program that receives control as a result of the PC issues the PCLINK macro instruction with the STACK option to save linkage information. The PCLINK macro instruction can only be issued in supervisor state. The PCLINK macro instruction creates an area called a stack element (STKE), which contains the following information:

- Caller's save area address from caller's register 13
- AMODE, return address, and PSW problem state bit from caller's register 14
- Parameter registers 0, 1, and 15
- Caller's PSW key and other information from caller's register 2 as follows:

In bits 0-23, bits 8-31 of caller's register 2
 In bits 24-27, PSW key
 In bits 28-31, zeroes

- Caller's PSW key mask and PASID from caller's register 3
- Latent parameter list address for this entry from caller's register 4
- Return address from the PCLINK service routine to the program that issued PCLINK STACK. This point is just after the PC routine entry point.
- Program mask from current PSW

After issuing PCLINK STACK, the program begins processing. It can, if it needs to, get information from the stack element using the PCLINK macro instruction with the EXTRACT option.

When the program is about to return control to its caller, it loads any data to be passed back to the caller into registers 0, 1, and 15 and then issues PCLINK with the UNSTACK,THRU option. This option restores registers 3, 13, 14, the program mask and, optionally, the original PSW protection key. The program then issues a PT instruction to return control. The caller restores its own registers and its SASID.

The PCLINK stack element is described in the *Debugging Handbook. Diagnostic Techniques* describes how the PCLINK stack elements and register save areas are chained together.

Figure 24 summarizes the PC/PT linkage conventions. In the figure, a program in ASID 8 issues a PC that invokes a program in ASID 7.




Figure 24. PC/PT Linkage Conventions

How to Establish a Cross Memory Environment

This section contains three examples that show three ways to establish a set of services for access via a PC instruction. The term “subsystem” is used in this section but note that the functions providing cross memory services are not limited to those functions that use the subsystem interface. The required operations are grouped into five categories:

1. **SETTING UP** initializes the structure that cross memory needs so the transfers of control can take place.
2. **ESTABLISHING ACCESS** sets up the linkage necessary for an address space to use cross memory services.
3. **PROVIDING SERVICE** consists of designing a service for cross memory use. Refer to “Designing a PC Routine” later in this section.
4. **REMOVING ACCESS** disconnects the linkage that enabled an address space to use cross memory services.
5. **CLEANING UP** removes the structures established in the initialization step.

The first example shows how to make a cross memory subsystem’s services available to a select group of users. (The code actually shows only one user, but the extra steps for adding users are pointed out.) The second example shows how to make a subsystem’s services available system-wide to all address spaces. The third example shows how a subsystem can provide a series of non-space switch services that operate on data in the user’s address space. A non-space switch service is one that does not cause an address space switch. See “Designing a PC Routine” later in this section.

Assume for all the examples that the subsystem has obtained common storage that can be accessed via the CVT. In this area it would store the PC numbers corresponding to its services. It could also store some of the lists that are needed to invoke PC/AUTH services, and that must be available to different address spaces. Assume also that SSBLOCK, shown in Figure 25, is in common storage accessible via the hypothetical CVT field, CVTXXXX. All examples use the declared storage areas shown in Figure 25.

```

SSBLOCK      DS   0D   SUBSYSTEM'S BLOCK
LXL          DS   0F   LX LIST
LXCOUNT      DS   F    NUMBER OF LXs REQUESTED
LXVALUE      DS   F    LX RETURNED BY LXRES
AXL          DS   0F   AX LIST
AXCOUNT     DS   H    NUMBER OF AXs REQUESTED
AXVALUE      DS   H    AX RETURNED BY AXRES
TKL          DS   0F   TOKEN LIST
TKCOUNT     DS   F    NUMBER OF ETS CREATED
TKVALUE      DS   F    TOKEN RETURNED BY ETCRE
PCTAB        DS   0F   TABLE OF PC NUMBERS
SERV1PC      DS   F    PC NUMBER FOR SERVICE 1
SERV2PC      DS   F    PC NUMBER FOR SERVICE 2

```

Figure 25. Declared Storage For Cross Memory Examples

Example 1 - Making Services Available to Selected Address Spaces

Setting Up: To make its services available to other address spaces via a PC instruction, the subsystem sets up the linkage and entry tables and the authorization structures.

To request that the control program reserve an LX for later use, use the LXRES macro instruction to reserve a 4-byte LX across the entire system. When LXRES is issued, the home address space becomes the owner of the LX.

```

.
.
LA          2,1
ST          2,LXCOUNT          REQUEST 1 LX
GETLX      LXRES  LXLIST=LXL,RELATED=(FREE LX, CONET)
.
.

```

To set up the entry table describing the services and their entry points, use the ETCRE macro instruction. An entry table describes all the services the subsystem makes available to users through a PC instruction. The home address space, at the time the ETCRE macro instruction is issued, becomes the owner of the entry table.

First construct a list of entry table descriptors. Each descriptor, mapped by the IHAETD mapping macro instruction, describes a program that gets control when a PC is issued. Figure 26 shows an entry table descriptor list with two entries.

```

.
.
CET1       ETCRE  ENTRIES=ETDESC,RELATED=(CONET,DISET1,DESET2)
           ST     0,TKVALUE          SAVE RETURNED TOKEN
.
.

```

ETDESC	DS	0D	ENTRY TABLE DESCRIPTION LIST (MAPPED BY IHAETD)
*	----- ETD HEADER		
	DC	X'00'	ETDFMT - MUST BE ZERO
	DC	X'00'	ETDRSV1 - RESERVED, MUST BE ZERO
	DC	H'2'	ETDNUM - NUMBER OF ENTRY DESCRIPTIONS THAT FOLLOW
*	----- ENTRY 1		
	DC	X'00'	ETDEX - ENTRY INDEX (EX)
	DC	X'C0'	ETDFLG - PROGRAM WILL EXECUTE SUPERVISOR STATE (ETDSUP ON) AND ENTRY WILL CAUSE SPACE SWITCH (ETDXM ON)
*			ETDRSV3 - RESERVED, MUST BE ZERO
	DC	H'0'	ETDPR01
	DC	F'0'	ETDPR02 - VIRTUAL ADDRESS TO BE GIVEN CONTROL
	DC	A(SERVICE1)	ETDAKM - CALLER CAN BE IN ANY KEY
*	DC	X'FFFF'	ETDEKM - SERVICE1 CAN ACCESS ONLY
	DC	X'8000'	KEY 0 PLUS KEYS AUTHORIZED IN CALLER'S PKM
*	DC	F'0'	ETDPR - LATENT PARAMETER PASSED TO CALLED PROGRAM
*	----- ENTRY 2		
	DC	X'01'	ETDEX
	DC	X'80'	ETDFLG - PROGRAM WILL EXECUTE IN SUPERVISOR STATE (ETDSUP ON) AND ENTRY IN NON-SPACE SWITCH (ETDXM OFF)
*			ETDRSV3
	DC	H'0'	ETDPRO (PROGRAM NAME) PROGRAM MUST BE IN LPA
*	DC	CL8'SERVICE2'	ETDAKM - CALLER MUST BE KEY 8-15
	DC	X'00FF'	ETDEKM - SERVICE CAN ACCESS ONLY
*	DC	X'00FF'	KEYS 8-15 PLUS KEYS AUTHORIZED IN CALLER'S PKM
*			ETDPR - LATENT PARAMETER PASSED TO CALLED PROGRAM
*	DC	F'0'	

Note: Upon entry, the PC routine receives a pointer, in general purpose register 4, to the latent parameter list. The first word of the latent parameter list is the value from the ETDPR field.

Figure 26. Entry Table Descriptions for Examples

To request that the control program reserve an authorization index (AX) for the service, use the AXRES macro instruction. The AX is reserved across the entire system. The home address space at the time the AXRES macro instruction is issued becomes the owner of the AX.

```


LA      2,1
STH    2,AXCOUNT      REQUEST 1 AX
GETAX  AXRES  AXLIST=AXL,RELATED=(AXSET,FREEAX)

```

To set the AX of the subsystem's address space to the AX value the control program reserved, use the AXSET macro instruction.

```
SETAX  AXSET  AX=AXVALUE,RELATED=(GETAX,SETAX)
```

At this point, you can construct the PC numbers that will be used to invoke the services. A PC number is a fullword value formed from an LX and an EX.



The linkage index returned by LXRES is in the following format so that you can OR it with an EX to form a PC number:

0	LX	0
---	----	---

L	1,LXVALUE	LX=PC# WITH EX OF 0
LA	2,0(,1)	CONSTRUCT EX=0 PC#
ST	2,SERV1PC	SAVE PC# FOR FIRST SERVICE
LA	2,1(,1)	CONSTRUCT EX=1 PC#
ST	2,SERV2PC	SAVE PC# FOR SECOND SERVICE

Establishing Access: The next two steps make the subsystem's services available to a user. The instructions used for these two steps must be issued from the user's address space but they must be invoked by a supervisor state or PKM 0-7 routine. If the user is a problem state program, the subsystem must arrange for the instructions to be executed on its behalf with the user's address space as the home address space. These two steps, of course, must be repeated for each user.

1. Set the PT and SSAR authority in the user's authorization table entry that corresponds to the subsystem's AX value so that the subsystem can issue a PT or SSAR instruction to the user's address space. This action allows the subsystem to access user data and return control to the user.

```
SETAT  ATSET  AX=AXVALUE,PT=YES,SSAR=YES,RELATED=(GETAX,
          SETAX,RESETAT)
```

2. Connect the subsystem's entry table to the user's linkage table at the entry that corresponds to the subsystem's LX. The linkage table now points to the subsystem's entry table.

```
LA      2,1
ST      2,TKCOUNT  SET COUNT OF ETS TO BE CONNECTED
CONET  ETCON  TKLIST=TKL,LXLIST=LXL,RELATED=(GETLX,CET1)
```

Now all the user needs to get to the subsystem is the correct PC number. The subsystem devises a method of making its PC numbers available and makes the method known. The subsystem could use an executable macro instruction that expands into code that locates the PC number and then executes the PC instruction to invoke the desired service. The subsystem could keep the PC numbers in a table that each address space can locate in commonly addressable storage.

At this point in the example, the subsystem has set up two services that the user can access using PC instructions. The subsystem has also established authority to issue PT and SSAR instructions to the user. The user's linkage table is connected to the subsystem's entry table as shown in Figure 27.

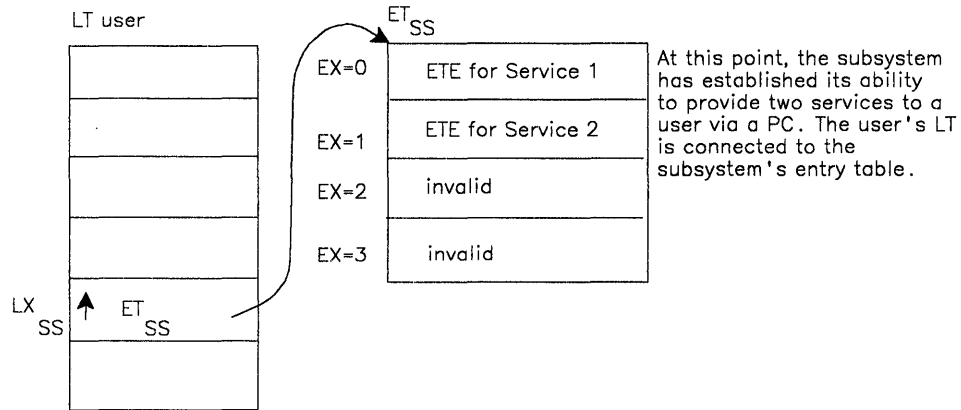


Figure 27. Linkage Table and Entry Table Connection

Providing Service: The PC instruction gives control to a PC routine that might run in cross memory mode. The PC routine must have been designed following the guidelines described in “Designing a PC Routine” later in this section.

To invoke the PC routine, the subsystem might provide a macro instruction. In the example, a macro instruction to invoke the first service would generate the following code:

```

STM 14,12,12(13)           SAVE REGISTERS
ESAR 2                     SAVE CALLER'S SASID IN THE
ST 2,16(,13)              REG 15 SLOT OF SAVEAREA
L 15,FLCCVT              ACCESS CVT
L 15,CVTXXXX(,15)       ACCESS SUBSYSTEM BLOCK
L 2,SERV1PC(,15)        OBTAIN SERVICE1 PC NUMBER
PC 0(2)                  ISSUE THE PC
L 14,12(,13)            RESTORE REG 14
L 2,16(,13)             LOAD SAVED SASID
SSAR 2                   RESTORE CALLER'S SASID
LM 2,12,28(13)          RESTORE REGS 2-12

```

Removing Access: The next two steps remove access to subsystem services. The steps are performed with the user's address space as the home address space. These steps are essentially the opposite of the steps used to establish access. First, remove the subsystem's PT and SSAR authority to the user's address space.

```
RESETAT    ATSET AX=AXVALUE,PT=NO,SSAR=NO,RELATED=(SETAT)
```

Second, disconnect the subsystem's entry table from the user's linkage table.

```
DISET1     ETDIS TKLIST=TKL,RELATED=CET1
```

Cleaning Up: When the subsystem is about to shut down, it must remove all cross memory connections and release any cross memory resources it owns. Destroy the subsystem's entry table, first making sure that all connections to it have been disconnected.

```
DESET1      ETDES  TOKEN=TKVALUE,RELATED=CET1
```

Free the subsystem's linkage index so that another subsystem can reuse it.

```
FREELX      LXFRE  LXLIST=LXL,RELATED=GETLX
```

Reset the subsystem's AX to zero.

```
RESETAX     SR      2,2          ZERO VALUE  
            AXSET  AX=(2)       RESET AX TO ZERO
```

Free the AX value so the system can reuse it. This action removes PT and SSAR authority corresponding to the subsystem's AX in all authorization tables in the system.

```
FREEAX      AXFRE  AXLIST=AXL,RELATED=GETAX
```

Example 2 - Making Service Available to All Address Spaces

This example shows how a subsystem makes global services available system-wide to all users. The example uses the same storage areas as example 1, however, it does not need the AX list. Figure 25 and Figure 26 earlier, show the areas. The main differences between example 1 and example 2 are example 2's use of a system linkage index and a system AX value. There are only ten slots available in the system linkage table for the user to use. A system linkage index allows the subsystem to globally connect an entry table to all address spaces, and a system AX value gives the subsystem PT and SSAR authority to all address spaces.

Setting Up: The first step in the set up operation is obtaining a "system" linkage index. The control program sets aside part of the available linkage indexes for use as system LXs. When an entry table is connected to a system LX, the entry table is automatically connected to all present and future address spaces.

Unlike ordinary LXs, system LXs cannot be freed for reuse. When an address space that owns a system LX terminates, the LX becomes "dormant." The system allows a dormant system LX to be reconnected to an address space different from the original owning address space. This is an important consideration for a subsystem that can be terminated and then restarted. The subsystem must have a way to "remember" the system LX it owned so that it can connect the LX to an entry table when it is restarted.

There are two ways subsystems or components become owners of a system LX. Many IBM-supplied global services use a pre-assigned system LX and the PC numbers that correspond to their services are in the system function table (SFT).

The second way that a subsystem can get a system LX is by issuing the LXRES macro instruction with the SYSTEM=YES option.

The code shown in the following three steps runs with the subsystem's address space as the home address space. The first step obtains a system LX. If the subsystem is coming up for the first time since IPL and does not have a system LX preassigned in the SFT, issue the LXRES macro instruction with the SYSTEM=YES option. Save the LX somewhere, probably in common storage, so that it is accessible if the subsystem is restarted. On a subsystem restart, this step is not necessary.

```

LA      2,1
ST      2,LXCOUNT      REQUEST 1 SYSTEM LX
GETSLX LXRES  LXLIST=LXL,SYSTEM=YES

```

Next, set the subsystem AX to 1. This allows the subsystem to issue a PT or SSAR instruction to all other address spaces because an AX of 1 is authorized to all address spaces. The subsystem that is providing a global service does not need to obtain a unique AX.

```

LA      2,1
AXSET   AX=(2)

```

Define the subsystem's entry table as follows:

```

ETCRE   ENTRIES=ETDESC
ST      0,TKVALUE      SAVE THE ET TOKEN

```

Next construct the PC numbers in the same way as in example 1.

Establishing Access: The following ETCON macro instruction, issued once from any address space, connects the subsystem's entry table to all address spaces in the system, current and future.

```

LA      2,1
ST      2,TKCOUNT      SET COUNT OF ETS TO BE CONNECTED
ETCON   LXLIST=LXL,TKLIST=TKL

```

All address spaces in the system now have access to the subsystem's services. All linkage tables are connected to the subsystem's entry table and, because the subsystem's AX is 1, it can issue PT and SSAR to any address space. Figure 28 shows how the linkage and entry tables appear at this point.




Figure 28. Linkage and Entry Tables for a Global Service

Providing Service: The subsystem provides service in the same way as in Example 1. The users of the services need to determine the PC number associated with each service. For many IBM-supplied services, the PC numbers are in the SFT. Other services must use a similar method.

Removing Access: To remove access, disconnect all users and destroy the entry table by issuing the ETDES macro instruction with the PURGE= YES option. This disconnects the entry table from all linkage tables in the system and then destroys it. (You cannot issue an ETDIS macro instruction for an entry table connected to a system LX.)

```
ETDES  TOKEN=TKVALUE ,PURGE=YES
```

Cleaning Up: To clean up, reset the subsystem's AX to 0 as follows:

```
SR      2,2  
AXSET  AX=(2)
```

Example 3 - Providing Non-Space Switch Services

In this example, a subsystem is providing a series of non-space switch services. Non-space switch services are described in "Designing a PC Routine" later in this section. This example is like example 1 except that the macro instructions for AX authorization are not used because an address space switch does not occur.

Designing a PC Routine

PC routines get control as the result of a PC instruction. A PC routine serves the same general purpose as an SVC routine in that it is a means of providing a function at an increased level of authority. While the SVC always increases the authority of the SVC routine to key 0, supervisor state, the PC instruction allows greater flexibility in the authority that a PC routine can have. A PC instruction can switch to supervisor state, increase PKM authority, and switch to a different primary address space that has greater AX authority. The PC instruction can also prevent problem state programs from calling PC services based on PKM authority. For more information on the PKM and AX, see "Cross Memory Authorization" earlier in this section.

The SVC instruction requires that the issuer be in enabled task mode and hold no locks, but the PC instruction does not have these restrictions.

The PC instruction also allows the PC routine and the data it manipulates to reside in its own address space and be isolated from its callers.

The PC instruction must be issued in primary mode. The PC instruction can cause an address space switch. When you set up the entry table descriptor as input to the ETCRE macro instruction, one of the attributes specified is whether or not the PC routine is space switch.

When a non-space switch PC routine gets control, it executes in primary mode, and both the primary and secondary addresses are set to the address space in which the PC was issued.

When a space switch PC routine gets control, it executes in primary mode in the address space in which the entry table by which it is accessed was created, not the address space in which the PC instruction was issued. Secondary mode is set to the PC issuer's address space. Thus, the space switch PC routine gets control in cross memory mode.

All PC routines, both space switch and non-space switch, must:

- Be loaded under the job step task of the address space that created the entry table(s) that describe the routine or else be permanently resident (in PLPA or the nucleus).

- Preserve and restore PC linkage information.
- Use a PT instruction to return to their caller.
- Not use the checkpoint/restart facility.

PC routines that cause an address space switch have the following additional requirements:

- The address space in which the PC routine runs must be non-swappable.
- The PC routine must perform its functions using only the MVS/XA services supported in cross memory mode. Refer to “Warnings and Restrictions” earlier.

A space switch PC routine can access data in another address space by using the MVCP and MVCS instructions, or by using secondary mode. When executing in secondary mode, remember that all data is accessed from the secondary address space and that you cannot predict which address space (primary or secondary) instructions will be fetched from. Thus, all space switch PC routines that run in secondary mode must reside in common storage.

In deciding whether to make the PC routine a space switch or non-space switch routine, consider the nature of the routine itself and the data it manipulates. If either the PC routine or the data it manipulates needs to be isolated from the routine’s callers, then the program or data, whichever requires isolation, should be located in the private area of an address space and the PC routine should be a space switch PC. If the program and data reside in commonly addressable storage or in the caller’s address space, then the PC routine can be a non-space switch PC. A non-space switch PC can increase PKM authority and switch to supervisor state, but it cannot increase AX authority because no address space switch occurs.

Figure 29 shows the possible locations of the non-space switch PC routine and the data it manipulates and also lists the types of users who can invoke a non-space switch PC routine.

Location of PC routine:	Location of data to be manipulated:	PC routine can be connected to address space:
Common	Common or caller’s private area	All address spaces, via connection to a system LX or any address space via connection to an LX.
An entry table owner’s private area	Common or caller’s private area	Only entry table owner’s address space

Figure 29. Characteristics of a Non-Space Switch PC Routine

Figure 30 shows the possible locations of the space switch PC routine and the data it manipulates, what types of users can invoke the PC routine, and whether or not the routine can run in secondary mode.

Location of PC routine	Location of data	PC routine can be connected to:	PC routine can run in secondary
Common	<ul style="list-style-type: none"> „ Common „ Particular address 	<ul style="list-style-type: none"> „ All address spaces with a system LX 	Yes
Particular address space that owns an ET	<ul style="list-style-type: none"> space that owns ET „ Associated data address space „ Each calling address space 	<ul style="list-style-type: none"> „ Any address space (to connect to a specific address space, the address space that creates the ET must be able to issue PT and SSAR instructions to the connected address space) 	No

Figure 30. Characteristics of a Space Switch PC Routine

Recovery Considerations

There are special recovery considerations when you write a space switch PC routine. A PC routine executing in cross memory mode has active binds to address spaces other than home. If one or more of these address spaces terminates, then the PC routine will incur a program check and its recovery routine might get control. The SETFRR macro instruction provides options that specify the cross memory mode in which the recovery routine must get control. There are also options that enable a recovery routine to get control as a resource manager when the requested cross memory mode cannot be established in order to recover resources serialized by local (CML) or global locks. Refer to “Providing Recovery Routines” for details on recovery in cross memory mode. See “Locking” in the Serialization section for more information about the CML lock.

Linkage Conventions: The linkage conventions for PC and PT transfers of control have already been described under “Linkage Conventions” earlier. You will recall that the PCLINK macro instruction provides a standard method for saving status. The stack entries created by PCLINK are formatted like standard save areas so that you can trace the flow of control across address spaces in the event of a dump.

Resource Management: PC routines should be loaded under the job step task of the address space that created the associated entry tables. If the task under which the PC routine was loaded fails and it is not the job step task of the address space that created the entry tables, the PC routine is freed even though users are still connected to the inoperative PC routine, and results are unpredictable. If the PC routine is loaded under the job step task, any failure of the task causes a program check when any program issues a PC to the inoperative PC routine.

When a job step task that owns entry tables providing space switch PC services terminates, whether normally or abnormally, the space switch event mask for the address space is turned on. If this indicator is on, no unit of work can execute in cross memory mode in the address space. A unit of work currently using the space switch PC services or a unit of work attempting to issue a PC to the address space causes a space switch event program check. Subsequent job steps execute normally except that they cannot reestablish space switch PC services. If a unit of work in a subsequent job step attempts to reestablish space switch PC services (that is, issues an LXRES, AXRES, or ETCRE macro instruction), it causes a X'052' abend.

When a job step task that owns space switch entry tables terminates (normally or abnormally), the address space of the task is not terminated. The ASID representing the address space of the terminated task is retained and evaluated for possible reuse before the next IPL takes place. The reuse of address space spares an installation the burden of scheduling IPLs at frequent intervals to recover the lost space. The automatic recovery of used address spaces is an important consideration in the installation's choice of a MAXUSERS parameter. It is also a factor in how an installation controls the creation and termination of cross memory environments.

The system maintains a history of cross memory binds and address spaces. When all cross memory binds have terminated, address spaces that created space switch entry tables are generally reused, although there are a few special cases in which system integrity cannot be guaranteed if the spaces are reused. The system recognizes those cases (there are two) and prevents the reuse of the corresponding address space for the duration of the current IPL.

- The first case involves circular PC chains. For example, consider programs p1, p2, and p3 running in respective address spaces s1, s2, and s3. If PC instructions are issued by p1 to p2, and by p2 to p3, and then by p3 back to p1, a **circular PC chain** exists. Upon termination, address spaces such as s1-s3 that are used in a circular PC chain are considered non-reusable for the duration of the IPL.
- The second case involves any address space that has a cross memory connection to a **system linkage index (LX)**. When this kind of address space terminates, it is considered non-reusable for the duration of the IPL. In addition, if this space is connected to any other address spaces, upon their termination they would also become non-reusable for the duration of the IPL.

A TCB for any job step task that owns a cross memory resource imposes a restriction on other TCBs that are higher up. The higher TCBs (that is, TCB for the initiator, RCT, DUMP, or STC) are restricted; they can only use **system PC** services. When the TCB that represents the task terminates, any connections between the higher TCBs and non-system entry tables are severed. Subsequent PCs that depend on those connections will not be successful.

While a PC routine is running, execution time is attributed to the home address space whether or not it is the same address space in which the PC routine executes.

Virtual Storage Management

Virtual storage management (VSM) allocates and releases blocks of virtual storage on request, ensures that real frames exist for SQA, LSQA, and V=R pages, and protects storage with fetch and storage protection keys. In addition, VSM provides the following services through the use of the macro instructions specified:

- List the starting address and the size of the private area regions associated with a given TCB -- VSMREGN
- Verify that a given area has been allocated via a GETMAIN macro instruction -- VSMLOC
- List the ranges of virtual storage allocated in a specified area -- VSMLIST

These VSM services are especially useful when determining available storage, coding recovery procedures, or specifying areas to be included in a dump. VSMREGN enables you to determine the amount of storage that you have for potential use. If you need to check whether a GETMAIN was issued to allocate a given block of storage, you can use the VSMLOC macro instruction to perform this check. If the given block is located in private area storage, you can also request the address of the TCB that issued the GETMAIN macro. VSMLOC enables you to verify control blocks or storage locations when coding recovery procedures. You can use VSMLOC to check whether a control block has been allocated and to verify that the control block is located in the correct subpool. VSMLIST enables you to obtain detailed information about virtual storage that could be useful in determining the areas that you might need in a dump and thereby limit the size of the dump. Limiting the size of a dump is especially critical when executing in 31-bit addressing mode because of the amount of storage involved. The use of VSMLIST is described later in this topic under the heading "Obtaining Information about the Allocation of Virtual Storage."

Allocating and Freeing Virtual Storage (GETMAIN, FREEMAIN)

The GETMAIN and FREEMAIN macro instructions respectively allocate and free one or more areas of virtual storage. The KEY parameter allows a user executing in PSW key zero to specify the storage key for storage requests involving subpools 227, 228, 229, 230, 231, and 241.

You can use the GETMAIN and FREEMAIN macro instructions when your program is executing in either 24-bit or 31-bit addressing mode. If you specify the options R, LC, LU, VC, VU, V, EC, EU, or E (provided by SVC 4, 5, and 10), storage addresses and lengths are treated as 24-bit addresses and lengths. If you want to specify 31-bit address and lengths, you must use the options RU, RC, VRC, or VRU. You can use the keyword LOC with these options to indicate the location of both virtual and real storage. See Figure 31 for a list of valid subpools and the location of these subpools when backed in real storage.

Most of the functions of GETMAIN and FREEMAIN (including the options mentioned above) are available to all users. However, some of the GETMAIN and FREEMAIN functions are available only to programs executing in supervisor state and PSW key zero. The restricted functions are provided by the parameters BRANCH and KEY.

The BRANCH Parameter

In addition to the normal SVC entries to the GETMAIN and FREEMAIN macros, there are also branch entries, which are available through the BRANCH parameter. Although the branch entries require the user to do more work, they are more efficient than the SVC entries.

Branch entry to the GETMAIN or FREEMAIN macro instructions is accomplished by specifying BRANCH=YES on the macro instructions. If the BRANCH=YES parameter is used, the caller must preload register 4 with the TCB address, preload register 7 with the ASCB address, and hold the LOCAL lock. The contents of register 3 are destroyed if RC, RU, VRC, or VRU are specified with this parameter.

Callers in cross memory mode can use the BRANCH=YES parameter of the GETMAIN and FREEMAIN macro instructions. If the caller is in cross memory mode, the storage that is allocated or freed is located in the currently addressable address space. The caller must hold the CML lock for the currently addressable address space; load register 7 with the address of the ASCB of the currently addressable address space; and load register 4 with zero or the address of a TCB in the currently addressable address space. If register 4 contains a zero, the storage is associated with the current job step task that owns the cross memory resources in the currently addressable address space (that is, the TCB anchored in ASCBXTCB).

An additional branch entry point is provided to obtain global storage without the need for holding the LOCAL lock. This entry point is available to programs that contain no references to particular address spaces (for example, timer routines). The caller must be in key zero, supervisor state, and be disabled. In addition, the caller must hold no locks higher in the locking hierarchy than the VSMFIX lock for global subpools or the VSMPAG lock for subpools 231 and 241. Although the TCB address and ASCB address are not required for this entry, the macro expansion loads register 4 with the address of the global save area pointed to by the CVT.

Global branch entry can be obtained by coding `BRANCH=(YES,GLOBAL)` on the `GETMAIN` or `FREEMAIN` macro instruction that includes the positional parameter `RC`, `RU`, `VRC`, or `VRU`. The subpools that are supported by this entry are limited to the global subpools: common service area (CSA) subpools 227, 228, 231, and 241, and system queue area (SQA) subpools 226, 239, and 245. Any other subpool is considered an error.

The KEY Parameter

The `KEY` parameter allows a user executing in PSW key zero to specify the storage key for storage he requests. Because branch entry users must be executing in PSW key zero at entry time, the `KEY` parameter satisfies the need to specify the actual key in which the requested storage is to be obtained.

The `KEY` parameter applies only to six subpools: 227, 228, 229, 230, 231, and 241. These subpools allow the requestor to obtain both global and local storage in key 0. (The `KEY` parameter allows an override of the PSW key.) Subpools 227 (fetch protected) and 228 (not fetch protected) are fixed global storage in the common service area, and must be freed explicitly. Subpools 229 (fetch protected) and 230 (not fetch protected) are local storage allocated from the top of the private area downward and intermixed with LSQA and SWA, and are freed automatically when the task terminates. Subpools 231 (fetch protected) and 241 (not fetch protected) are global storage in the common service area, and must be freed explicitly.

Using Cell Pool Services (CPOOL)

The cell pool macro instruction provides users with another way of obtaining virtual storage. This macro instruction provides centralized, high performance cell management services.

Cell pool services obtain a block of virtual storage (called a cell pool) from a specific subpool at the user's request. The user can then request smaller blocks of storage (called cells) from this cell pool as needed. If the storage for the requested cells exceeds the storage available in the cell pool, the user can also request that the cell pool be increased in size (extended) to fill all requests.

The `CPOOL` macro instruction makes the following cell pool services available:

- Create a cell pool (`BUILD`)
- Obtain a cell from a cell pool if storage is available (`GET,COND`)
- Obtain a cell from a cell pool and extend the cell pool if storage is not available (`GET,UNCOND`)
- Return a cell to the cell pool (`FREE`)
- Free all storage for a cell pool (`DELETE`)

The CPOOL macro instruction, with the exception of the TCB, KEY, and LINKAGE=BRANCH parameters, is available to all users. Note, however, that in order to provide high performance, cell pool services do not attempt to detect most user errors. For example, the following user errors are not detected by cell pool services:

- The user is executing in a non-zero key that does not match the key of the pool being manipulated.
- The user attempts to free a cell from a pool that has already been deleted.
- When trying to free a cell, the user passes cell pool services a bad cell address. (This might damage the cell pool, preventing subsequent requests from being properly handled.)
- A disabled user requests that a cell pool be built in a pageable subpool.

Because user errors often produce unpredictable results, the user should try to avoid them.

Using Storage Subpools

Both the GETMAIN and the CPOOL instructions allow users to allocate storage from specified storage subpools.

The chart in Figure 31 lists the valid MVS/XA subpools and the characteristics of the subpools. It indicates the type of storage, whether the storage is fixed or fetch protected, where the storage is backed when fixed, and the storage key associated with the storage.

The storage map in Figure 32 shows the location of the storage areas listed in Figure 31. Virtual storage management allocates low private area storage beginning at the start of the private area or the start of the extended private area and it allocates high private area storage beginning at the upper end of the private area or the upper end of the extended private area.

The storage keys listed are:

Key	Meaning
0	MVS/XA system control program
1	Job scheduler and job entry subsystem (JES2 or JES3)
USER	The storage key is taken from the PSW at the time of the GETMAIN or can be specified on the GETMAIN/FREEMAIN macro instructions.
JOB	The storage key is from the TCB associated with the request at the time of the first GETMAIN request. All subsequent GETMAIN requests use this key regardless of the key currently in the TCB.

Subpool	Type of Storage	Fixed	Fetch Protected	Where Backed	Storage Key
0-127	Low private	No	Yes	Below 16 Mb	Job
226	Common-SQA	Yes	No	Below 16 Mb	0
227	Common-CSA/ECSA	Yes	Yes	Below 16 Mb	User
228	Common-CSA/ECSA	Yes	No	Below 16 Mb	User
229	High Private	No	Yes	Below 16 Mb	User
230	High private	No	No	Below 16 Mb	User
231	Common-CSA/ECSA	No	Yes	Below 16 Mb	User
233	Private-LSQA/ELSQA	Yes	No	Anywhere	0
234	Private-LSQA/ELSQA	Yes	No	Anywhere	0
235	Private-LSQA/ELSQA	Yes	No	Anywhere	0
236	High private	No	No	Anywhere	1
237	High private	No	No	Anywhere	1
239	Common-SQA/ESQA	Yes	Yes	Anywhere	0
240	Low private	No	Yes	Below 16 Mb	Job
241	Common-CSA/ECSA	No	No	Below 16 Mb	User
245	Common-SQA/ESQA	Yes	No	Anywhere	0
250	Low private	No	Yes	Below 16 Mb	Job
251	Low private	No	Yes	Below 16 Mb	Job
252	Low private	No	No	Below 16 Mb	0
253	Private-LSQA/ELSQA	Yes	No	Anywhere	0
254	Private-LSQA/ELSQA	Yes	No	Anywhere	0
255	Private-LSQA/ELSQA	Yes	No	Anywhere	0

Figure 31. Characteristics of the Valid Storage Subpools

Notes:

1. *All private area subpools are swappable. Common area subpools are not swappable.*
2. *All subpools allocated virtually in the extended area can be backed anywhere.*
3. *Subpools 0-127, 229, 230, 231, 240, 241, 250, 251, and 252 can be backed anywhere. However, if a page fix is requested for allocation in the nonextended areas, these subpools are backed below 16 megabytes real unless LOC is specified with ANY for real allocation.*
4. *Subpool 226 is valid only for allocating virtual storage below 16 megabytes.*
5. *Subpools 227 and 228 are backed anywhere for virtual addresses above 16 megabytes. For virtual addresses below 16 megabytes, they are backed below 16 megabytes unless the user of the GETMAIN instruction specifies the LOC parameter with ANY for real allocation.*
6. *Callers executing in key 0 and supervisor state, who request storage from subpool 0, via the GETMAIN macro instruction, obtain that storage from subpool 252. Therefore, if they want to dump the storage using the SDUMP macro instruction, they must specify subpool 252 rather than subpool 0.*

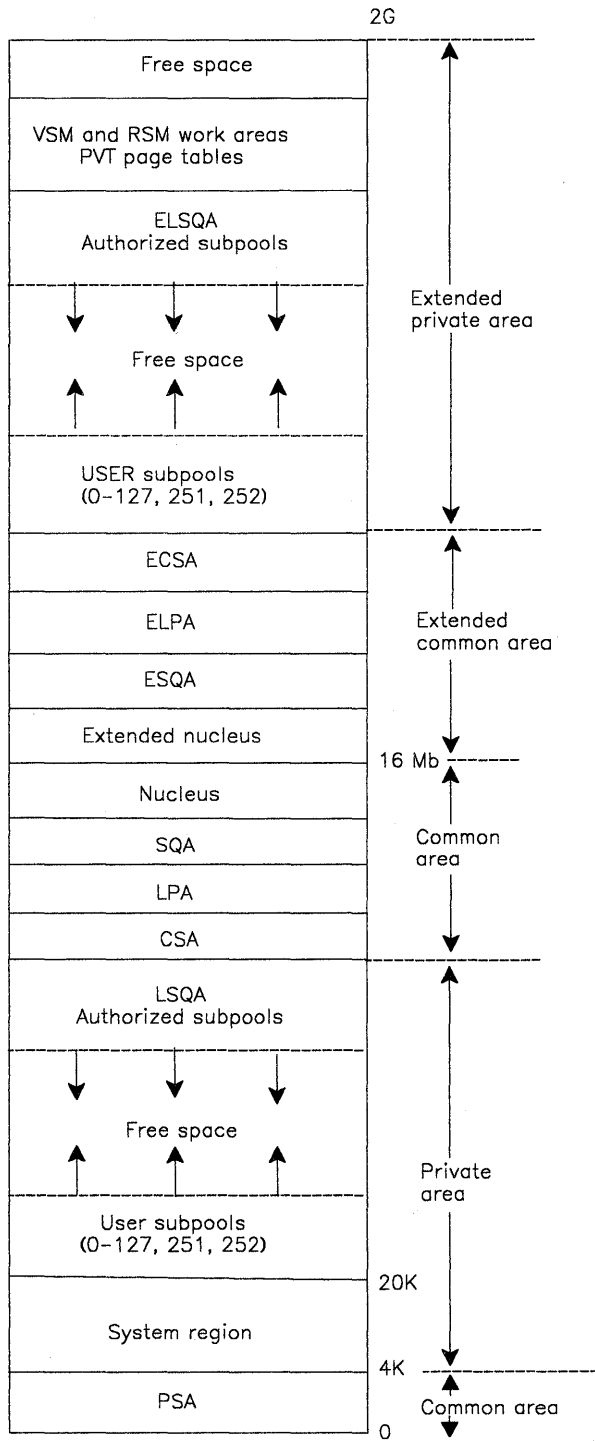


Figure 32. MVS/XA Virtual Storage Map

Obtaining Information about the Allocation of Virtual Storage

The VSMLIST macro instruction provides information about the allocation of virtual storage. The VSMLIST service routine returns the information in a user-supplied work area specified as a parameter of the VSMLIST macro instruction. The length of the work area varies but it must be a minimum length of 4K bytes. Figure 33 shows the format of the VSMLIST work area.




Figure 33. Format of the VSMLIST Work Area

Using the VSMLIST Work Area

Prior to the first invocation of the VSMLIST macro instruction for a single request, you must set the first four bytes of the work area to zero. This field will contain the return code of the VSMLIST macro instruction after control returns to the issuer.

The VSMLIST service routine updates the work area and places the requested information in the data area located at the end of the work area. If the macro instruction was executed successfully and all of the requested information fit into the data area, the VSMLIST service routine returns to the caller with a return code of 0 in the first four bytes of the work area. If the macro instruction was executed successfully, but all of the requested information could not fit into the data area, the service routine returns to the caller with a return code of 4 in the first four bytes of the work area. In this case, the caller can reissue the macro instruction as many times as necessary to obtain all of the information.

For multiple invocations of VSMLIST, the service routine continues supplying the information, starting where it left off on the previous invocation, provided the work area is not changed. However, multiple invocations do not provide cumulative results. For each invocation of a set of multiple invocations for a specific request:

- The count fields are relative to the current invocation of the macro instruction (for example, the number-of-subpools field contains the number of subpool descriptors in the current invocation only).
- The output in the data area describes the current invocation only.

You can avoid multiple invocations by enlarging the work area to hold all of the information. If you do enlarge the work area, be sure to set the first four bytes of the work area (the return code area) to zero before reissuing the macro instruction.

Bytes	Field name	Description
0-3	Return code	This field contains the return code from the previous invocation of the VSMLIST macro instruction. You must set this field to zero before the first invocation of the VSMLIST macro instruction for a single request.
4-7	Address of data area	The data area is located at the end of the work area and contains the information that you requested.
8-11	Length of the data area	The data area varies in length and is limited in size by the length of the work area that you specified as a parameter of the VSMLIST macro instruction.
12-15	Parameter list	This section of the work area is constructed by the VSMLIST service routine according to the parameters that you specified when you issued the VSMLIST macro instruction.
		Bytes Contents
		12-15 Length of work area
		16 SP operand represented as follows: X'00' -- SQA X'01' -- CSA X'02' -- LSQA X'03' -- PVT X'FF' -- Subpool list provided
		17 SPACE operand represented as follows: X'00' -- ALLOC X'01' -- FREE X'02' -- UNALLOC
		18 Information about the TCB, LOC, and REAL operands represented as follows: X'80' -- ALL specified for the TCB operand X'40' -- ANY specified for the LOC operand X'20' -- REAL operand specified
		19 Set to zero
		20-23 TCB address or zero
		24-27 Subpool list address or zero
		28-31 Set to zero
32-x	Control information	The control information is used by the VSMLIST service on multiple invocations for a single request. This area varies in size.
x-y	Data area	This area contains the actual output of the VSMLIST macro instruction. The area varies in size and is limited by the length of the work area specified as a parameter of the macro instruction.

Figure 34. Description of VSMLIST Work Area

The information returned in the data area depends on the parameters specified on the macro invocation. You can use the VSMLIST macro instruction to obtain information about the following types of storage:

- Allocated
- Free
- Unallocated

Except for subpool 245, an allocated block of storage is a multiple of 4K, some of which has been allocated via a GETMAIN macro instruction. Free space within that block is the area that has not been allocated via a GETMAIN macro instruction. An unallocated block of storage is some multiple of 4K none of which has been allocated via a GETMAIN macro instruction.

VSMLIST reports all SQA pages not allocated to subpools 226 and 239 as allocated to subpool 245. These pages of subpool 245 may not have been allocated via a GETMAIN macro.

The format of the information returned in the data area for each of these three types of requests follows.

Allocated Storage Information

You can request allocated storage information by coding the `SPACE=ALLOC` parameter of the `VSMLIST` macro instruction. The format of the output varies according to what you specify for the `SP` parameter.

If you specify `SP=SQA`, `SP=CSA`, or `SP=LSQA`, the output consists of the allocated storage information for the subpools in the specified area. The subpools listed in each of these areas are:

SQA: 226, 239, 245
CSA: 227, 228, 231, 241
LSQA: 255

Figure 35 shows the format of the output for a request for information about the allocated storage in a specified area.

If you specify `SP=PVT`, the output consists of the allocated storage information for subpools in the private area according to the owning TCB. These subpools are 0-127, 229, 230, 236, 237, 251, and 252. Figure 38 shows the format of the allocated storage information for the private area.

If you specify a subpool list, the output consists of the allocated storage information for each of the subpools in the list. Figure 39 shows the format of the allocated storage information for a subpool list request.

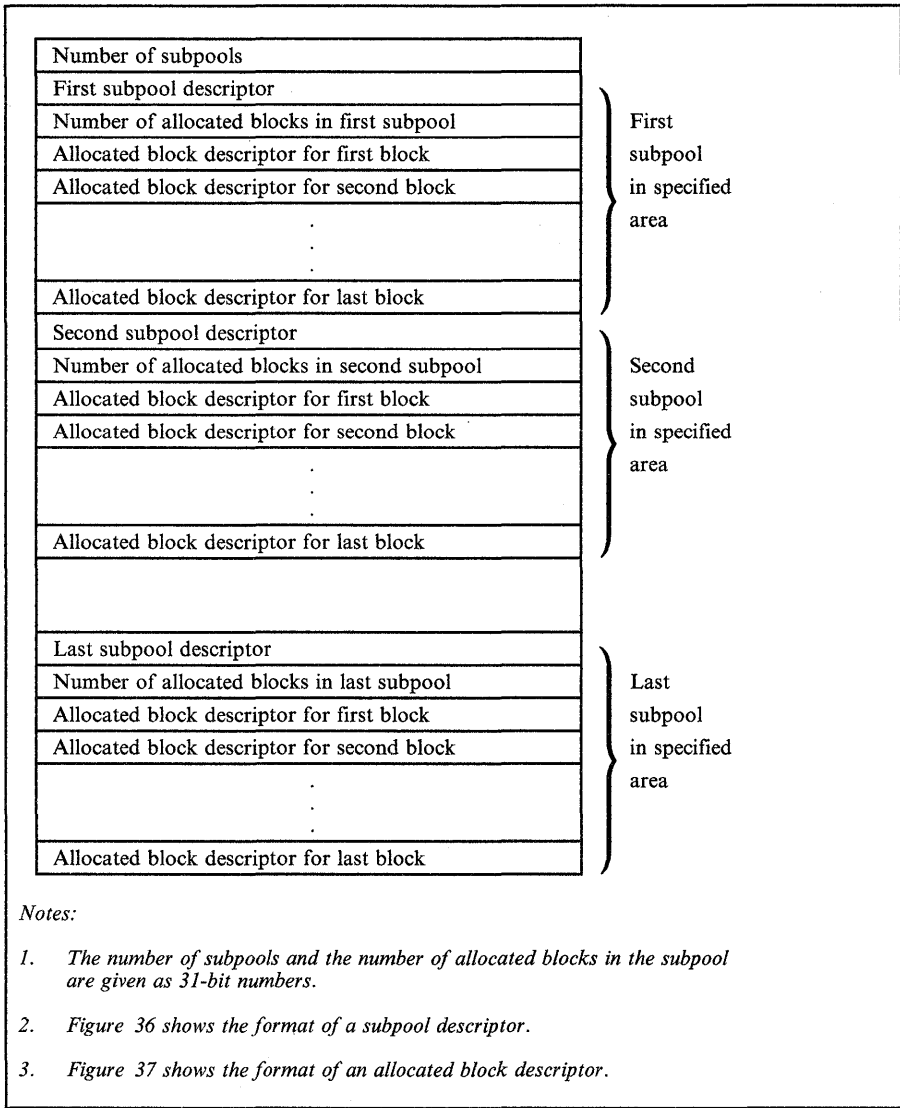


Figure 35. Allocated Storage Information for Subpools in a Specified Area

Byte	Content										
0	X'00' to identify a subpool descriptor										
1	Length of subpool descriptor										
2	Subpool ID										
3	Miscellaneous flags and storage key as follows:										
	<table border="1"> <thead> <tr> <th>Bit</th> <th>Meaning When Set</th> </tr> </thead> <tbody> <tr> <td>0-3</td> <td>Storage key</td> </tr> <tr> <td>4</td> <td>The TCB with which this descriptor is associated owns the storage described by this descriptor This is meaningful for private area storage only.</td> </tr> <tr> <td>5</td> <td>The storage described by this descriptor is shared. This is meaningful for private area storage only.</td> </tr> <tr> <td>6-7</td> <td>Reserved</td> </tr> </tbody> </table>	Bit	Meaning When Set	0-3	Storage key	4	The TCB with which this descriptor is associated owns the storage described by this descriptor This is meaningful for private area storage only.	5	The storage described by this descriptor is shared. This is meaningful for private area storage only.	6-7	Reserved
Bit	Meaning When Set										
0-3	Storage key										
4	The TCB with which this descriptor is associated owns the storage described by this descriptor This is meaningful for private area storage only.										
5	The storage described by this descriptor is shared. This is meaningful for private area storage only.										
6-7	Reserved										
4-7	Owning TCB address (if PVT subpool), otherwise zero.										

Figure 36. Format of Subpool Descriptor

Byte	Content				
0-3	The virtual address of the allocated block				
	<table border="1"> <thead> <tr> <th>Bit</th> <th>Meaning When Set</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>The caller specified the REAL option and this allocated block could be backed in real storage above (bit 0=1) or below (bit 0=0) 16 megabytes.</td> </tr> </tbody> </table>	Bit	Meaning When Set	0	The caller specified the REAL option and this allocated block could be backed in real storage above (bit 0=1) or below (bit 0=0) 16 megabytes.
Bit	Meaning When Set				
0	The caller specified the REAL option and this allocated block could be backed in real storage above (bit 0=1) or below (bit 0=0) 16 megabytes.				
4-7	The length of the allocated block				

Figure 37. Format of Allocated Block Descriptor

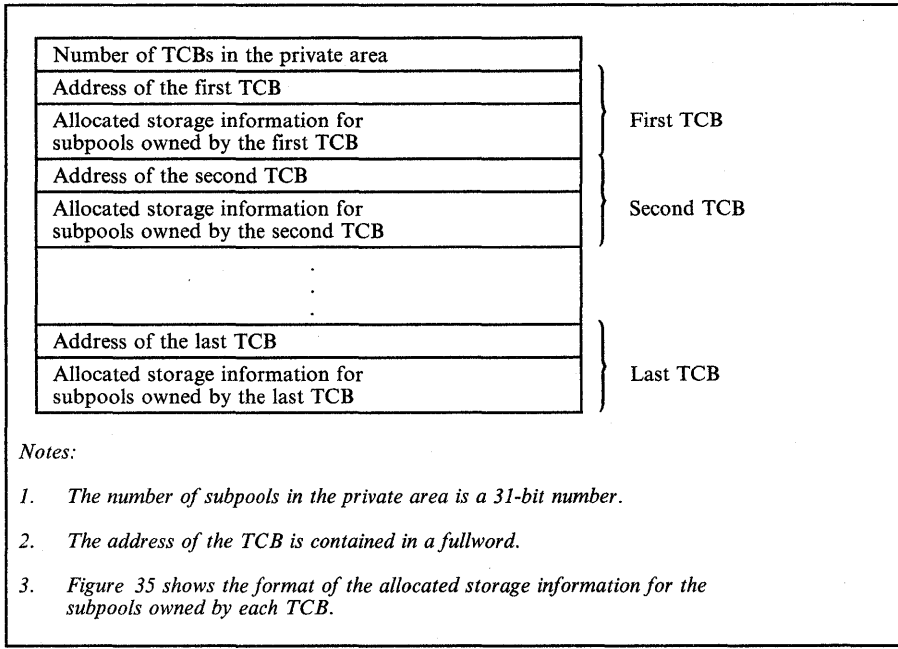


Figure 38. Allocated Storage Information for the Private Area

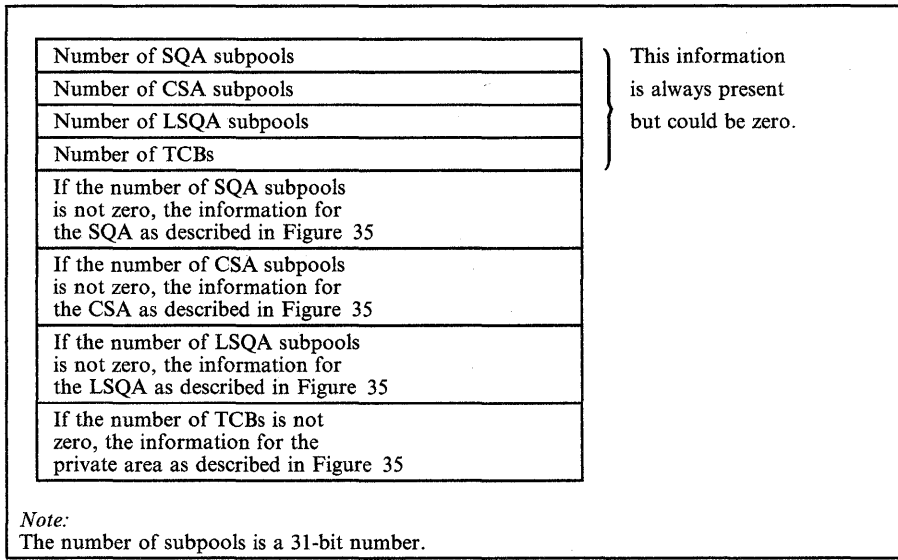


Figure 39. Allocated Storage Information for a Subpool List

Free Storage Information

A request for free storage information is specified by the `SPACE=FREE` parameter of the `VSMLIST` macro instruction. The `VSMLIST` service routine returns information about both allocated and free virtual storage. The information is returned in the same manner as allocated storage information except that each allocated block descriptor is followed by the number of pieces of contiguous free storage contained within the allocated block and the free space descriptors for each of these areas. Figure 40 shows the format of a free space descriptor.

Byte	Content
0-3	The virtual address of the start of the free space
4-7	The length of the free space

Figure 40. Format of Free Space Descriptor

Unallocated Storage Information

You can request information about unallocated storage by specifying the `SPACE=UNALLOC` parameter of the `VSMLIST` macro instruction. You can obtain this information for CSA and private area subpools only, by specifying `SP=CSA` or `SP=PVT`. Figure 41 shows the format of the output for a `SPACE=UNALLOC` request for CSA or PVT subpools.

Number of region descriptors	} First region in specified area
First region descriptor	
Number of unallocated blocks in the first region	
Unallocated block descriptor for the first unallocated block	
Unallocated block descriptor for the second unallocated block	
.	
Unallocated block descriptor for the last unallocated block	} Second region in specified area
Second region descriptor	
Number of unallocated blocks in the second region	
Unallocated block descriptor for the first unallocated block	
Unallocated block descriptor for the second unallocated block	
.	
Unallocated block descriptor for the last unallocated block	} Last region in specified area
.	
.	
Last region descriptor	
Number of unallocated blocks in the last region	
Unallocated block descriptor for the first unallocated block	
Unallocated block descriptor for the second unallocated block	
.	
Unallocated block descriptor for the last unallocated block	

Notes:

1. The number of region descriptors and the number of unallocated blocks in each region are given as 31-bit numbers.
2. Figure 42 shows the format of a region descriptor.
3. Figure 43 shows the format of an unallocated block descriptor.

Figure 41. Unallocated Storage Information for CSA and PVT Subpools

Byte	Content
0-3	The virtual address of the region (CSA, ECSA, RCT area, V=V area, extended V=V area, or V=R area)
4-7	The length of the region

Figure 42. Format of Region Descriptor

Byte	Content
0-3	The virtual address of the unallocated block
4-7	The length of the unallocated block

Figure 43. Format of Unallocated Block Descriptor

Managing SWA Control Blocks

The control blocks that reside in the scheduler work area (SWA) contain information from specifications on JCL DD statements or dynamic allocation requests. The control blocks reside in virtual storage below 16 megabytes in the user's high-end private area.

However, an installation has the option of assigning specific, pre-determined SWA control blocks to virtual storage above 16 megabytes and using macro instructions to access those SWA control blocks.

By assigning SWA control blocks to virtual storage above 16 megabytes, installations that run applications containing a large number of JCL DD statements may reduce the occurrences of out-of-space abends, such as 80A.

Assigning SWA Control Blocks

An installation can control which jobs may assign SWA control blocks above 16 megabytes by

- Using the JES2 JOBCLASS initialization statement to assign job classes.
- Using the JES2 TSUCLASS initialization statement to control SWA for TSO users.
- Using the JES2 STCCLASS initialization statement to control SWA for started jobs.
- Using the JES3 CIPARM initialization statement.
- Using the SMF exit IEFUJV for job-by-job selectivity.

JES2 supports the assignment of the selected SWA control blocks above 16 megabytes except for those blocks associated with SYSIN/SYSOUT data sets.

Checkpoint/Restart Considerations

Jobs with SWA control blocks that reside above 16 megabytes cannot be restarted, either by an automatic or a deferred restart. However, when the job is resubmitted, assigning the SWA control blocks below 16 megabytes does allow a deferred restart from a checkpoint.

SWA Control Blocks

All SWA control blocks have the same general format, as shown in Figure 44. Figure 45 shows the format of a specific SWA control block, the JFCB.




Figure 44. Format of a SWA Control Block




Figure 45. Format of the JFCB

As shown in Figure 44, all SWA control blocks are preceded by a 16-byte prefix. Not all SWA control blocks, however, can reside above 16 megabytes. Figure 46 lists those blocks that can reside above 16 megabytes, and Figure 47 lists those blocks that must reside below 16 megabytes. In both cases, the prefix has the same format.

<u>Control Block</u>	<u>Acronym</u>	<u>Block ID</u>
Data Set Enqueue Table	DSEN	0F
Data Set Name Table	DSNT	07
DD Name Table	DDNT	27
GDG Name Table	GDGN	23
Interpreter Work Area Block	IWAB	25
Job File Control Block	JFCB	1C
Job File Control Block Extension	JFCX	1D
Job File Control Block Extension	JFCE	29
Passed Data Set Information Block	PDIB	21
Procedure Override Table	POT	0A
Scheduler Work Block	SWB	32
Step Input Output Table	SIOT	03
Volume Unload Table	VUT	26

Figure 46. SWA Control Blocks That Can Reside Above 16 Megabytes

<u>Control Block</u>	<u>Acronym</u>	<u>Block ID</u>
Account Control Table	ACT	01
AMP Keyword Extension	AMPX	28
Job Control Table	JCT	00
Job Control Table Extension	JCTX	30
Job Management Record	JMR	1B
Step Control Table	SCT	02
Step Control Table Extension	SCT	0C
Subsystem Scheduler Work Area	SSWA	31

Figure 47. SWA Control Blocks That Must Reside Below 16 Megabytes

SWA Manager

To access control blocks in SWA requires the services of the SWA manager. A routine may call the SWA manager in either move mode or locate mode. Move mode requests cause actual movement of data to or from control blocks residing in SWA. Locate mode requests return to the calling routine either an SVA or a pointer to a SWA control block; no actual movement of data occurs.

The following definitions may be helpful prior to using either move mode or locate mode.

- EPA — external parameter area. Used to communicate information to and from the SWA manager. The EPA is mapped by macro IEFZB506 for move mode and by IEFZB505 for locate mode.
- QMPA — queue manager parameter area. A control block passed by the caller that contains input and output parameters for the SWA manager. The QMPA is mapped by macro IEFQMNGR.
- QMAT — SWA manager address table. Used to maintain the correspondence between the SVA of a SWA control block and the actual 31-bit virtual address of a block that resides above 16 megabytes. The QMAT is mapped by macro IEFQMAT.
- QMST — SWA manager storage table. Used as a translate table to indicate where a block is assigned, above or below 16 megabytes. The QMST is mapped by macro IEFQMST.

To allow a caller to keep parameters in fetch-protected storage, the SWA manager reads the QMPA and the EPA in the key of the caller.

If your installation allows SWA control blocks to reside in virtual storage above 16 megabytes, any program that accesses these blocks must:

- For a move mode request, use the IEFQMREQ macro, placing the SVA of the block in the EPA.
- For a locate mode request, run in AMODE 31 and use the SWAREQ macro, placing the SVA of the block in the EPA.
- If the required block is the JFCB, use the RDJFCB macro to access the block.

The SWA manager provides a unique reason code for each error situation that occurs. If there is relevant data associated with the error, the SWA manager stores the data in an area pointed to by register 0.

SWA Manager Move Mode

The IEFQMREQ macro is used to make move mode requests to the SWA manager. In addition, the calling routine must pass a QMPA for all requests and, except for a delete or an assign/start request, one or more EPAs. The QMPA and the EPA contain the relevant information for the requested function(s). To request a particular function, the caller provides a function code in the QMPA. Figure 48 lists the valid function codes for a move mode request.

<u>Function Code</u>	<u>Function</u>	<u>Explanation</u>
'00'X	Assign/Start	Requests Assign processing for a job that is just starting
'01'X	Assign	Requests the initialization of a SWA control block in a SWA subpool.
'02'X	Write/Assign	Requests the movement of the contents, except for the prefix, of a calling routine's buffer into a SWA control block and the initialization of another control block in SWA.
'03'X	Write	Requests the movement of the contents, except the prefix, of a calling routine's buffer into a SWA control block.
'04'X	Read	Requests the movement of a SWA control block, except for the prefix, into a calling routine's buffer.
'05'X	Readall/Move	Requests the movement of a SWA control block, including the prefix, into a calling routine's buffer.
'06'X	Writeall/Move	Requests the movement of the contents, including the prefix, of a calling routine's buffer into a SWA control block.
'08'X	Delete	Requests a FREEMAIN for a SWA subpool

Figure 48. Valid Function Codes for a Move Mode Request

To indicate to the SWA manager whether an EPA resides below or above 16 megabytes, a caller must set particular fields in the QMPA:

- If an EPA is to reside below 16 megabytes, store the 3-byte address in QMPACL.
- If an EPA is to reside above 16 megabytes, store the 4-byte address in QMPACLX and set the associated indicator in QMACLEX.

Several actions require the caller to pass an extended EPA (a 16-byte EPA) and set the associated indicator in QMEPAX in the QMPA. The actions that require an extended EPA are:

- Assigning a control block above 16 megabytes
- Assigning a control block with a length of other than 192 bytes
- Doing a write/assign in which the block ID of the control block being written is different from the block ID of the control block being assigned

If a caller does not pass a 16-byte EPA, the SWA manager assigns the control block below 16 megabytes and defaults the block length to 192 bytes.

Input to the SWA Manager for a Move Mode Request

The input to the SWA manager is as follows:

- Register 0 - not applicable
- Register 1 - points to a QMPA that points to an EPA
- Registers 2-12 - not applicable
- Registers 2-12 - not applicable
- Register 13 - save area address
- Register 14 - return address
- Register 15 - entry point address

Output from the SWA Manager for a Move Mode Request - Successful Completion

The SWA manager assigns, reads, writes, or deletes the specified control blocks and makes the appropriate updates to the block prefixes, the EPAs, and the QMAT. Upon successful completion of a move mode request, the SWA manager

- Restores registers 0 - 14
- Stores a return code of '0'X in register 15

For the actions of the SWA manager on unsuccessful completion, see "SWA Manager Error Processing" on page 1-125.

Size and Contents of an EPA for a Move Mode Request

The information that a caller passes in an EPA and the information returned by the SWA manager, together with the size of the EPA, for each move mode request is shown in Figure 49.

Function	EPA Size (in bytes)	Inputs in the EPA	Output in the EPA
Assign or Assign/Start	4	none	SVA
Assign or Assign/Start	16	block ID or 0 block length	SVA
Read	8	SVA buffer address	block ID
Write	8	SVA buffer address block ID or 0	none
Write/Assign	8	SVA for Write buffer address block ID or 0	SVA for Assign
Write/Assign	16	SVA for Write buffer address block ID for Write or 0 block length for Assign block ID for Assign or 0	SVA for Assign
Readall/Move	8	SVA buffer address	block ID
Writeall/Move	8	SVA buffer address block ID or 0	none
Delete	0	not applicable	not applicable

Figure 49. EPA Size and Contents for Move Mode Requests

Format of an EPA for a Move Mode Request

An EPA can have three formats for a Move Mode request:

1. When the caller requests an assign and does not pass a block ID or a block length, the input EPA is a 4-byte field containing zero. In this case, the SWA manager outputs a 4-byte EPA containing the SVA (three bytes) followed by a zero (one byte).
2. When the caller requests a read, write, readall/move, writeall/move, or write/assign, the input EPA is 8 bytes long. For example, a write request has the following input EPA format:
 - The buffer address (four bytes)
 - The SVA (three bytes)
 - The block ID (one byte)

In this case, the output EPA has the same contents as the input EPA.

3. When the caller requests an assign and passes a block ID or block length, the input EPA is 16 bytes long. It is also 16 bytes long when the caller requests a write/assign, and is passing a block ID different from the one being written.

For example, a write/assign request that also passes a block ID has the following input EPA format:

- The buffer address (four bytes)
- The SVA (three bytes)
- The write ID (one byte)
- The block length (four bytes)
- The assign ID (one byte)
- The version number (two bytes)
- Reserved field (one byte)

In this case, the SWA manager outputs sixteen bytes containing:

- The SVA (three bytes)
- Zero (one byte)
- Unused data (twelve bytes)

Move Mode Summary

To issue a move mode request, take the following steps:

1. Build a QMPA (mapped by macro IEFQMNGR), which includes specifying the function code and setting a pointer to the EPA.
2. Build an EPA (mapped by macro IEFZB506).
3. Set register 1 to point to the QMPA.
4. Issue the IEFQMREQ macro.

SWA Manager Locate Mode

The SWAREQ macro is used to make locate mode requests to the SWA manager. Using the macro, a caller passes a function code and an EPA for all requests. In the EPA, the caller may pass a pointer to a QMPA for a Locate/All request. If the caller does not pass a pointer to a QMPA, the SWA manager uses the default QMPA, the one whose address is contained in the active JSCB.

For a list of valid function codes, see Figure 50.

<u>Function Code</u>	<u>Function</u>	<u>Explanation</u>
AC	Assign/Conditional	Requests the initialization of a SWA control block in a SWA subpool (conditional GETMAIN).
AL	Assign/Locate	Requests the initialization of a SWA control block in a SWA subpool (unconditional GETMAIN).
DB	Delete Block	Requests a FREEMAIN for a SWA control block.
LA	Locate/All	Requests the return in the caller's EPA of the addresses of the prefix and data portion of the control block.
RL	Read/Locate	Requests the return in the caller's EPA of the address of the data portion of the control block, the block length, and the block ID.
WL	Write/Locate	Requests the updating of the prefix of the SWA control block.

Figure 50. Valid Function Codes for a Locate Mode Request

Input to the SWA Manager for a Locate Mode Request

The input to the SWA manager is as follows:

- Register 0 - not applicable
- Register 1 - points to a 2-word parameter list that contains the addresses of:
 - a pointer to the EPA chain
 - the function code
- Registers 2-12 - not applicable
- Register 13 - save area address
- Register 14 - return address
- Register 15 - entry point address

Output from the SWA Manager for a Locate Mode Request - Successful Completion

The SWA manager locates, reads, writes, or deletes the specified control blocks; makes the appropriate updates to the block prefixes, the EPAs, and the QMAT; and returns any requested information to the caller. Upon successful completion of a locate mode request, the SWA manager

- Restores registers 0 - 14
- Stores a return code of '0'X in register 15

For the actions of the SWA manager on unsuccessful completion, see "SWA Manager Error Processing" on page 1-125.

Size and Contents of an EPA for a Locate Mode Request

The information that a caller passes in an EPA and the information returned by the SWA manager, together with the size of the EPA, for each locate mode request is shown in Figure 51.

Function	EPA Size (in bytes)	Inputs in the EPA	Outputs in the EPA
Assign	16	block length block ID or 0	SVA pointer to block none
Delete Block	16	SVA	pointer to block none
Locate/All	28	SVA pointer to QMPA or 0	pointer to prefix block length prefix length block ID
Read/Locate	16	SVA	pointer to block block length block ID
Write/Locate	16	SVA pointer to block block ID or 0	none

Figure 51. EPA Size and Contents for Locate Mode Requests

Locate Mode Summary

To issue a locate mode request, take the following steps:

1. If necessary, build a QMPA (mapped by macro IEFQMNGR). If you do not provide a QMPA and the request is locate/all, the system uses, as a default, the QMPA pointed to by the current JSCB.
2. Build an EPA (mapped by macro IEFZB505).
3. Issue the SWAREQ macro instruction, specifying the address of the EPA pointer and the required function code.

SWA Manager Error Processing

If an abend occurs during move or locate mode processing, the SWA manager:

- Sets register 0 to point to a 4-word parameter area that contains the addresses of the:
 - QMPA
 - QMST
 - QMAT
 - EPA (the address of the particular EPA that is invalid, not the address of the complete EPA chain)
- Stores abend code 0B0 in register 1

- Stores reason codes in register 15:
 - ‘4’X - invalid function requested
 - ‘8’X - invalid self-pointer in the block prefix
 - ‘C’X - an attempt was made to read a block not yet written
 - ‘10’X - invalid block length
 - ‘14’X - invalid count field
 - ‘1C’X - invalid block ID
 - ‘24’X - error occurred while translating the address of a block residing above 16 megabytes
 - ‘38’X - conditional GETMAIN failed

Real Storage Management

The real storage manager (RSM) administers the use of real storage and directs the movement of virtual storage pages between auxiliary storage slots and real storage frames in blocks of 4096 bytes. It makes all addressable virtual storage in each address space appear as real storage. Only the virtual pages necessary for program execution are kept in real storage. The remainder reside on auxiliary storage. RSM employs the auxiliary storage manager (ASM) to perform the actual paging I/O necessary to transfer pages in and out of real storage. ASM also provides DASD allocation and management for paging space on auxiliary storage. RSM relies on the system resource manager (SRM) for guidance in the performance of some of its operations.

RSM assigns real storage frames upon request from pools of available frames, thereby associating virtual addresses with real storage addresses. Frames are repossessed when freed by a user, when a user is swapped-out, or when needed to replenish the available pool. While a virtual page occupies a real storage frame, the page is considered pageable unless it is fixed by the FIX option of the PGSER macro instruction, a PGFIX or PGFIXA macro instruction, or obtained from a fixed subpool. RSM also allocates virtual equals real (V = R) regions upon request by those programs that cannot tolerate dynamic relocation. Such a region is allocated contiguously from a predefined area of real storage and is non-pageable.

The PGSER macro instruction in MVS/XA provides all the paging services through the use of parameters rather than separate macro instructions as in MVS/370. PGSER handles virtual addresses above or below 16 megabytes. The macro instructions, PGFIX, PGFIXA, PGFREE, PGFREEA, PGLOAD, PGANY, PGOUT, and PGRLSE are supported by MVS/XA to maintain compatibility with MVS/370, but it is recommended that you use the PGSER macro instruction.

Users should note that MVS/XA paging services function differently from MVS/370 paging services in the following ways:

- The end address (EA) specified on the PGSER macro instruction is the address of the last byte on which the page service is to be performed (not the last byte + 1).
- In the register format SVC entry for the PGSER macro instruction, register 14 is used in addition to registers 0, 1, and 15.
- If an ECB is supplied, with a page-fix or page-load request and the caller invokes PGSER, then the return code must be checked because the ECB is not posted for a return code of 0.
- If an ECB is not supplied the return code need not be checked. Control will not be returned until the request is successfully completed. If the request fails, the caller will be abnormally terminated.
- Users of the PGSER macro instruction do not need to hold the local lock.

- Users of **BRANCH=Y** or **BRANCH=SPECIAL** options of the **PGSER** macro instruction must provide an 18-word savearea; this savearea must be in non-pageable storage if **BRANCH=SPECIAL** is specified.

The paging services provided include the following:

- Fix virtual storage contents -- **PGFIX**, **PGFIXA**, or the **FIX** option of **PGSER**
- Fast path to fix virtual storage contents -- the **FIX** and **BRANCH=SPECIAL** options of **PGSER**
- Free real storage -- **PGFREE**, **PGFREEA**, or the **FREE** option of **PGSER**
- Fast path to free real storage -- the **FREE** and **BRANCH=SPECIAL** options of **PGSER**
- Load virtual storage areas into real storage -- **PGLOAD** or the **LOAD** option of **PGSER**
- Page out virtual storage areas from real storage -- **PGOUT** or the **OUT** option of **PGSER**
- Release virtual storage contents -- **PGRLSE** or the **RELEASE** option of **PGSER**
- Page anywhere (above or below the 16 megabyte (megabytes) line of real storage) -- **PGANY** or the **ANYWHERE** option of **PGSER**

The **PGFIX**, **PGFIXA**, **PGFREE**, and **PGFREEA** functions as well as the **FIX** and **FREE** options of **PGSER** are available only to authorized system functions and users and are described in the following topics. **PGANY**, **PGLOAD**, **PGOUT**, and **PGRLSE** as well as the **ANYWHERE**, **LOAD**, **OUT**, and **RELEASE** options of **PGSER** are not restricted and are available to all users. **PGSER** and **PGANY** are described in this publication. **PGLOAD**, **PGOUT**, and **PGRLSE** are described in *Supervisor Services and Macro Instructions*.

Fixing/Freeing Virtual Storage Contents

Fixing virtual storage and freeing real storage are complementary functions. The **PGFIX** and **PGFIXA** macro instructions and the **FIX** option of **PGSER** make specified storage areas resident in real storage and ineligible for page-out as long as the requesting address space remains in real storage. Note that page fixing ties up valuable real storage and is usually detrimental to system performance unless the use of the fixed pages is extremely high.

The **PGFREE** and **PGFREEA** macro instructions and the **FREE** option of **PGSER** make specified storage areas, which were previously fixed via the **PGFIX** macro instruction or the **FIX** option of **PGSER**, eligible for page-out. Pages fixed by **PGFIX**, **PGFIXA**, or the **FIX** option of **PGSER** are not considered pageable until the same number of page free and page-fix requests have been issued for any virtual area. The fix and free requests for a page must be issued by the same task (unless **TCB=0** is specified), otherwise the page will not be freed.

When using the fix function, you have the option of specifying the relative real time duration anticipated for the fix. Specify **LONG=Y**, if you expect that the duration of the fix will be relatively long. (As a rule of thumb, the duration of a fix is considered long if the interval can be measured on an ordinary timepiece—that is, in seconds.) Additional processing might be required to avoid an assignment of a frame to the **V=R** area or an area that might be varied

offline. Specify `LONG=N`, if you expect the time duration of the fix to be relatively short. A long-term fix is assumed if you do not specify this option.

In both the fix and free functions, you have the option of specifying that the contents of the virtual area are to remain intact or be released. If the contents are to be released, specify `RELEASE=Y`; otherwise, specify `RELEASE=N`. If you specify `PGFIX` or the `FIX` option of `PGSER` with `RELEASE=Y`, the release function is performed before the fix function. If you specify `PGFREE` or the `FREE` option of `PGSER` with `RELEASE=Y`, the free function is performed and those pages of the virtual subarea with zero fix counts are released; that is, the contents of virtual areas spanning entire virtual pages that were fixed are expendable and no page-outs for these pages are necessary.

The `BRANCH=SPECIAL` and the `FIX` or `FREE` options of `PGSER` provide the fast path version of `PGSER`. The fast path version of `PGSER` with the `FIX` option ensure that specific storage areas are resident in real storage and ineligible for page-out. These functions execute only short-term, synchronous page fixes.

Notes:

1. *`PGFIX` and the `FIX` option of `PGSER` do not prevent pages from being paged out when an entire virtual address space is swapped out of real storage. Consequently, the user of `PGFIX` and the `FIX` option of `PGSER` cannot assume a constant real address mapping for fixed virtual areas in most cases.*
2. *When using the `PGFIXA` macro instruction or the fast path version of `PGSER` with the `FIX` option, or a branch entry to `PGSER` with the options `FIX` and `TCB=0`, fixed areas will not automatically be freed at the end of a job; to free them, issue a `PGFREEA` macro instruction or the `PGSER` macro instruction with the `FREE` and `BRANCH=SPECIAL` options.*

PGFIX/PGFREE Completion Considerations

Under normal circumstances, you can reverse the effect of a `PGFIX` by using a `PGFREE` when the need for a page fix ceases. You can also reverse the effect of the `FIX` option of `PGSER` by using the `FREE` option of `PGSER` when the need for a page fix ceases. However, a page-fix request sometimes completes asynchronously if, for example, it requires a page-in operation. In such cases, you might need to explicitly purge page-fix operations.

For this reason, the page-fix function provides a mechanism for signalling event completion. The mechanism is the standard ECB together with `WAIT/POST` logic. The requestor supplies an ECB address and waits on the ECB after a request if the return code indicates that all of the pages were not immediately fixed. The ECB is posted when all requested pages are fixed in real storage.

Note: Callers who supply an ECB and use `PGSER` must check the return code before waiting since the ECB is not posted for a return code of 0.

There are two ways to explicitly purge a page fix:

- If the page fix is known to be complete, the page fix is reversed through the page-free function.

- If there is any possibility that the page fix has not been posted as complete, issue PGFREE or PGSER with FREE and supply an ECB address. This ECB parameter identifies the event control block that was supplied as an input parameter with the page fix being purged. Note that for the purpose of canceling a page-fix request that has not yet completed, the ECB must uniquely identify the page-fix request. Consequently, to provide for explicit purging, you must ensure that the ECB for any incomplete page fix can be located in a purge situation, and that the ECB has not been reused at the time the page fix is to be canceled.

The page-free function always completes immediately and requires no ECB address except for purging considerations.

The issuer of the following instructions is responsible for freeing the fixed frames:

- PGFIXA
- PGSER, with the FIX, BRANCH, and TCB=0 options
- PGSER, with the FIX and BRANCH=SPECIAL options

This can be accomplished by using PGFREEA; PGSER with FREE, BRANCH, and TCB=0; or PGSER with FREE and BRANCH=SPECIAL.

An FRR (functional recovery routine) or ESTAE recovery routine should be established during the period these fixes are outstanding. The recovery routine should free the frames in case there is an unexpected error.

Input to Page Services

There are two formats for providing input to page services. These are the register (R) and list (L) formats. If you specify R, page services uses the input information supplied in registers to perform the requested function; if you specify L, page services uses the input information provided in a parameter list to perform the requested function. The information that you must provide in the parameter list includes the starting and the ending addresses for which you want the page service to be performed and an indication of the end of the list.

The list used depends on which page services macro instruction you code. Descriptions of the parameter lists and the macros that use them follow.

Virtual Subarea List (VSL)

The virtual subarea list provides the basic input to the page service functions: PGFIX, PGFIXA, PGFREE, PGFREEA, PGLOAD, PGRLSE, and PGOUT.

The list contains one or more doubleword entries; each entry describes an area in virtual storage. The list must be non-pageable and located in the address space to be processed. The VSL is not required to be on a word boundary.

See *Debugging Handbook* for an exact description of the VSL.

Page Service List (PSL)

The page service list provides the basic input to the page service functions of the PGSER macro with the exception of the `BRANCH = SPECIAL` option. Each entry in the list specifies a range of addresses to be processed, or specifies the address of the next list entry to be processed, or is null. The first entry also indicates the paging service that is to be performed on all the ranges specified in the list.

The PSL has the following characteristics:

- The list must be in non-pageable storage.
- The PSL is not required to be on a word boundary.
- All addresses specified are 31-bit addresses.

See *Debugging Handbook* for an exact description of the PSL.

Short Page Service List (SSL)

The short page service list provides the basic input to the PGSER macro instruction with the `BRANCH = SPECIAL` option. The list contains entries for the 31-bit starting and 31-bit ending addresses of the virtual area to be fixed or freed.

The SSL has the following characteristics:

- The list must be in non-pageable storage.
- The SSL is not required to be on a word boundary.
- All addresses specified are 31-bit addresses.

See *Debugging Handbook* for an exact description of the SSL.

Branch Entry to the PGSER Routine

Branch entry to the PGSER macro instruction is available in both cross memory mode and non-cross memory mode for the FIX, FREE, OUT, LOAD, ANYWHERE, and RELEASE options. The caller must be enabled, in supervisor state and key 0, and must set up the PSL as shown in *Debugging Handbook*. The caller does not need to hold the local lock, but must ensure that register 13 contains the address of an 18-word savearea when the PGSER macro instruction is issued.

The macro uses the registers as follows:

Register(s)	Bit(s)	Contents
0		ECB address or 0 if no ECB
1	0 1-31	0 for register format 1 for list format Start of virtual area for register format Pointer to the first PSL for list format
2		31-bit address of the last byte of the virtual area for register format Irrelevant for list format
3	0-15 16-23 24-31	Reserved for register format Same as FUNC in PSL for register format Same as FLAG2 in PSL for register format Irrelevant for list format
4		TCB address or 0 for register format Irrelevant for list format
5-12		Not used
13		Address of standard 72-byte save area, required for branch entry only. For BRANCH = SPECIAL, the save area must be non-pageable.
14		Pointer defined return address (The first bit indicates the AMODE. If this bit is 1, the AMODE is 31-bit; if this bit is 0, the AMODE is 24-bit.
15		Entry point address

On return from the PGSER macro instruction, the registers are set as follows:

Register	Contents
0-4	The contents are destroyed and unpredictable.
5-13	The contents are unchanged.
14	The contents are destroyed and unpredictable.
15	Return code

Branch Entry to MVS/370 Page Services

Branch entry is available for all MVS/370 page services (page-fix, page-free, page-load, page-release, page-any, and page-out) in non-cross memory mode; and for all but the page-out service in cross memory mode. The caller must be in key 0, supervisor state, and must hold the local lock of the currently addressable address space.

Note: LOCAL lock means the local lock of the home address space. When written in lower case, the local lock refers to either the LOCAL or CML (cross memory local) lock.

Cross Memory Mode

The pages that are candidates for page services must be addressable in the current address space. The caller must set up registers as follows:

Register(s)	Bit(s)	Contents
0		0
1	0 1-7 8-31	0 for register format 1 for list format Same as bits 1-7 of VLSFLAG1 field of VSL for register format; irrelevant for list format 24-bit starting address on which the service is to be performed for register format; 24-bit address of user's first VSL for list format
2	0-7 8-31	Same as VSLFLAG2 field in VSL for register format; irrelevant for list format 24-bit ending address + 1 for which the service is to be performed for register format; irrelevant for list format
3		Irrelevant
4		0
5-6		Irrelevant
7		ASCB address of current address space
8-13		Irrelevant
14		Return address
15		Entry point to page services (contents of CVTPSXM)

On return, the page service sets the registers as follows:

Register	Contents
0-14	Unchanged
15	Return code

The only return code possible is 0. This indicates that the requested function was processed successfully.

Note: PGFIXA and PGFREEA can be invoked in cross memory mode.

Non-Cross Memory Mode

The caller must set up registers as follows:

Register(s)	Bit(s)	Contents
0		ECB address or 0 if no ECB is specified
1	0 1-7 8-31	0 for register format 1 for list format Same as bits 1-7 of VSLFLAG1 field of VSL for register format; irrelevant for list format 24-bit starting address on which the service is to be performed for register format; 24-bit address of user's first VSL for list format
2	0-7 8-31	Same as VSLFLAG2 field in VSL for register format; irrelevant for list format 24-bit ending address + 1 for which the service is to be performed for register format; irrelevant for list format
3		Irrelevant
4		TCB address or 0
5-13		Irrelevant
14		Return address
15		Entry point to page-anywhere service (contents of CTVPSIB or PVTPSIB)

On return, the page service sets the registers as follows:

Register	Contents
0-14	Unchanged
15	Return code

The return codes are as follows:

Code	Meaning
0	The requested function was processed successfully. If the function was page-fix or page-load, and an ECB was supplied, it will be posted.
8	The requested function was page-fix or page-load with an ECB. The function will be processed asynchronously and the ECB will be posted upon completion.
12	The requested function was page-out and the function was unsuccessful for at least one of the specified pages.

The Nucleus

The nucleus contains routines that execute with dynamic address translation (DAT) turned off and routines that execute with DAT on. These routines are located in two separate load modules. Load module IEAVNUC0n (n identifies the particular load module) contains the DAT-ON nucleus and load module IEAVEDAT contains the DAT-OFF nucleus. See *System Initialization Logic* for information concerning the manner in which the nucleus is loaded into storage.

There are two macro instructions that provide services for the nucleus. These macro instructions are:

Macro	Function
DATOFF	Provides a means of linking to routines in the DAT-OFF nucleus
NUCLKUP	Provides a means of obtaining information about CSECTs in the DAT-ON nucleus

Linking to Routines in the DAT-OFF Nucleus (DATOFF)

The DAT-OFF nucleus is not mapped in virtual storage. IPL processing loads the DAT-OFF nucleus into consecutive real storage located at the highest available real address. Because the DAT-OFF nucleus is not mapped in virtual storage, a special method is used to link to routines in this area. The DATOFF macro instruction provides the means of linking to routines in the DAT-OFF nucleus.

When using the DATOFF macro instruction, the caller specifies an index that identifies the routine that is to receive control in the DAT-OFF nucleus. The index, entry point, and purpose of the routines available to users in the DAT-OFF nucleus are shown in Figure 52.

Index	Entry Point	Purpose
INDCDS	IEAVCDS	Compare Double and Swap
INDMVCL0	IEAVMVC0	General DAT-OFF MVCL function
INDMVCLK	IEAVMVKY	General DAT-OFF MVCL function in user key
INDXC0	IEAVXC0	General DAT-OFF XC function
INDUSR1	IEAVEUR1	User defined function
INDUSR2	IEAVEUR2	User defined function
INDUSR3	IEAVEUR3	User defined function
INDUSR4	IEAVEUR4	User defined function

Figure 52. DAT-OFF Routines Available to Users

All routines that execute with DAT turned off must be located in the DAT-OFF nucleus. These routines receive control and execute in 31-bit addressing mode and must be capable of residing either above or below the 16 megabytes line. Therefore routines that execute in the DAT-OFF nucleus must have the attributes `AMODE=31`, `RMODE=ANY`. For information concerning 24-bit/31-bit compatibility, see *SPL: 31-Bit Addressing*.

Using System Provided DAT-OFF Routines (DATOFF)

The system defined index values, INDMVCL0, INDMVCLK, INDXC0, and INDCDS are available to users. INDMVCL0 initiates the move character long (MVCL) function, INDMVCLK initiates the MVCL function in user key, INDCDS initiates the compare double and swap function, and INDXC0 initiates the exclusive OR (XC) function. The register usage and linkage for these functions follows.

In all cases, the DATOFF macro instruction destroys the contents of general registers 0, 14, and 15.

INDMVCL0- Move Character Long

All register values must be 31-bit addresses. Before issuing the macro instruction, the user must load the registers as follows:

Register	Use
0	Used by macro
2	Real location into which the characters are to be moved
3	Length of the area into which the characters are to be moved
4	Real location of the area from which the characters are to be moved
5	Length of the area from which the characters are to be moved
14	Used by macro
15	Used by macro
1,6-13	Unused

The user invokes the MVCL function by coding the following macro instruction:

```
DATOFF INDMVCL0
```

INDMVCLK- Move Character Long in User Key

All register values must be 31-bit addresses. Before issuing the macro instruction, the user must load the registers as follows:

Register	Use
0	Used by macro
2	Real location into which the characters are to be moved
3	Length of the area into which the characters are to be moved
4	Real location of the area from which the characters are to be moved
5	Length of the area from which the characters are to be moved
6	Bits 24-27 contain the PSW key in which the MVCL function is to be performed.
14	Used by macro
15	Used by macro
1,7-13	Unused

The user invokes the MVCL in user key function by coding the following macro instruction:

```
DATOFF INDMVCLK
```

INDXC0 - Exclusive OR

All register values must be 31-bit addresses. Before issuing the macro instruction, the user must load the registers as follows:

Register	Use
0	Used by macro
2	Real location of first operand and location for results of exclusive OR character operation
3	Length, in bytes, of operand pointed to by register 2. The length must be in bits 24-31 of register 3. Allows a maximum length of 256 bytes
4	Real location of the operand to be exclusive orred with the operand pointed to by register 2.
14	Used by macro
15	Used by macro
1,5-13	Unused

The user invokes the XC function by coding the following macro instruction:

```
DATOFF INDXC0
```

IEAVMVC0- Compare Double and Swap

All register values must be 31-bit addresses. Before issuing the macro instruction, the user must load the registers as follows:

Register	use
0	Used by macro
1	Unchanged
2,3	First 64 bit operand in even-odd pair of registers (target data)
4,5	Third 64 bit operand in even-odd pair of registers (source data)
6	Real address of second operand, a doubleword in storage (target address)
7-13	Unchanged
14	Used by macro
15	Used by macro

The user invokes the CDS function by coding the following macro instruction:

```
DATOFF IEAVCDS
```

Writing User DAT-OFF Routines

As shown in Figure 52, there are four DAT-OFF indexes that users can define. These indexes are INDUSR1, INDUSR2, INDUSR3, and INDUSR4. The entry points corresponding to these indexes are IEAVEUR1, IEAVEUR2, IEAVEUR3, and IEAVEUR4, respectively.

User written DAT-OFF routines are restricted as follows:

- The user of the DATOFF macro instruction must be in key 0, supervisor state, and executing with DAT turned on.
- The DAT-OFF routine must have the attributes AMODE=31 and RMODE=ANY.
- The DAT-OFF routine must preserve register 0 because register 0 contains the return address of the module that issued the DATOFF macro.
- The DAT-OFF routine must use branch instructions to link to other DAT-OFF routines.
- The DAT-OFF routine must use BSM 0,14 to return.

See *SPL: System Modifications* for information about how user-written DAT-OFF routines are placed in the DAT-OFF nucleus.

Obtaining Information about CSECTs in the DAT-ON Nucleus (NUCLKUP)

IPL processing places the CSECTs located in the DAT-ON nucleus in virtual storage and creates a map of them. The real addresses do not equal the virtual addresses and the real addresses are not necessarily contiguous. IPL processing loads the CSECTs into storage according to residency mode and according to whether they are read only or read/write. If the CSECT is assembled with `RMODE=ANY`, it is placed in the extended nucleus. Figure 53 shows the virtual storage map of the DAT-ON nucleus.

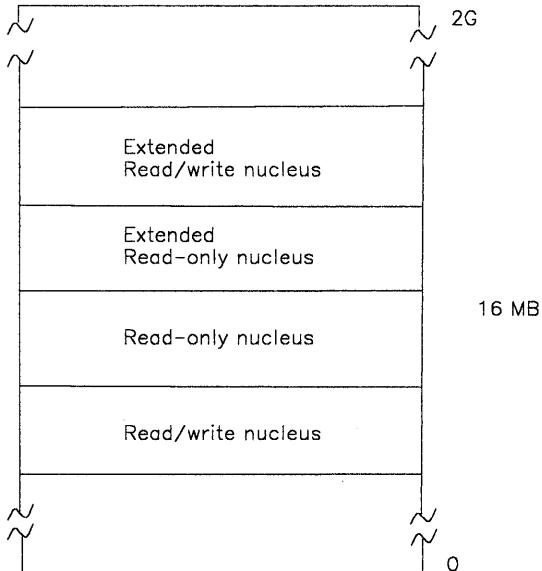


Figure 53. Virtual Storage Map of DAT-ON Nucleus

The nucleus map look up service provides users with information about these CSECTs. Through the use of the NUCLKUP macro instruction, users can perform two functions:

- Retrieve the address and addressing mode of a nucleus CSECT, given the name of the CSECT
- Retrieve the name and entry point address of a nucleus CSECT, given an address within the CSECT.

Normal and Abnormal Program Termination

The supervisor offers many services that help to detect and process abnormal conditions during system execution. The hardware detects certain types of abnormal conditions (such as an attempt to execute an instruction with an invalid operation code) and causes program interruptions to occur. The software detects other abnormal conditions (such as an attempt to open a data set that is not defined to the system) and causes abnormal terminations.

The supervisor enables you to write recovery routines to handle interruptions and abnormal conditions. The supervisor initiates the recovery termination process of your program either when you request it (for example, by issuing an ABEND macro instruction) or when MVS/XA detects a condition that will degrade the system or destroy data.

The services described in this section include:

- Invoking recovery termination (CALLRTM and ABEND macro instructions)
- Processing program interruptions (SPIE and ESPIE macro instructions)
- Intercepting system errors (SLIP command)
- Using system trace facilities (PTRACE)
- Dumping virtual storage (SDUMP macro instruction and CHNGDUMP command)
- Providing recovery routines (ESTAE, ATTACH with the ESTAI parameter, FESTAE, and SETFRR macro instructions)
- Uses of resource managers

Recovery Termination Manager

The recovery termination manager (RTM) controls the flow of software recovery processing by handling all normal and abnormal terminations of tasks and address spaces. RTM gets control in response to events such as the following:

- Unanticipated program checks (except those protected by SPIE routines)
- Machine checks
- Invalid use of an SVC (issuing an SVC while locked, disabled, in SRB mode, or in cross memory mode)
- I/O error on page-in request

- ABEND or CALLRTM macro instruction requesting termination of a task or address space (see “Invoking the Recovery Termination Manager” later in this topic)

When one of these events occurs, RTM initiates recovery processing before proceeding with abnormal termination.

Your installation-written functions can use RTM’s recovery processing by providing recovery routines for the functions. A recovery routine is a routine that you establish to get control if your main function terminates abnormally. The recovery routine can perform such processing as:

- Documenting the error
- Providing a dump of the storage needed to diagnose the error
- Freeing resources acquired by the main function
- Requesting a retry -- returning control to an appropriate point in the main function
- Requesting that RTM continue with the abnormal termination

To provide recovery for tasks and SRBs, RTM recognizes two types of recovery routines: functional recovery routines (FRRs) and ESTAE-type recovery routines. See “Uses of Resource Managers” later in this chapter for a full description.

When a function terminates abnormally, RTM gets control and generally invokes the most recently-established recovery routine to recover for the process that was in control. If this recovery routine cannot recover from the error (it fails or requests that termination continue), RTM invokes the next most recently established-recovery routine. The passing of control from one recovery routine to another is called *percolation*.

Note: MVS/XA functions provide their own recovery routines; thus, percolation can pass control to both installation-written and system-provided recovery routines. If all recovery routines percolate -- that is, no recovery routine can recover from the error -- then the process in control (an SRB or a task) is terminated.

RTM invokes recovery routines only during abnormal termination of tasks or SRBs. RTM also invokes resource manager routines during both normal and abnormal termination of a task or an address space. The major purpose of a resource manager is to release any resources held by the task or address space and make these resources available to other users. See “Uses of Resource Managers” later in this chapter for a description of the processing such routines can perform.

Invoking the Recovery Termination Manager

RTM can be called to perform its recovery and termination services on behalf of the caller or on behalf of another routine. Two macro instructions -- CALLRTM and ABEND -- invoke RTM.

CALLRTM

A routine issues the CALLRTM macro instruction to direct the recovery termination services to a task or address space other than itself or its callers. Only key 0 supervisor state routines can issue CALLRTM. Control returns to the issuer of the macro instruction if TYPE = ABTERM or TYPE = MEMTERM is specified.

TYPE = ABTERM: If the CALLRTM macro instruction specifies TYPE = ABTERM, RTM processing is directed toward the specified task, and you should consider locking and save area requirements:

- If the TCB parameter is specified as 0 (or defaulted to 0) and the ASID parameter is omitted, the current task in the current address space is abnormally terminated. In this situation, the caller must be disabled (for example, hold any of the spin locks) and need not provide a work area via register 13. If dump options are supplied, they must be contained in fixed pages. The routine must exit to the dispatcher without changing the TCB or RB and without enabling.
- If the TCB parameter is specified as an address and the ASID parameter is omitted, the task associated with the specified TCB in the current address space is abnormally terminated. In this situation, the caller must own the LOCAL lock, and need not provide a work area. If the caller specifies a TCB equal to the current TCB address, the caller must also be disabled.
- If the ASID parameter is specified, the ABTERM function is scheduled as a service request block (SRB) to terminate the task in the specified address space. The caller, who specifies ASID, must pass the address of an 18-word save area in register 13.

TYPE = MEMTERM: If the CALLRTM macro instruction specifies TYPE = MEMTERM, RTM processing is directed toward an address space and you should consider the following locking and work area information:

- If the ASID parameter is nonzero, the specified address space is abnormally terminated. The caller need not be disabled or own any locks. The caller must pass the address of an 18-word work area in register 13.
- If the ASID parameter is specified as 0 or is omitted, the current address space is abnormally terminated. The caller need not be disabled or own any locks. The caller must pass the address of an 18-word work area in register 13.

Note: The required work area is not the standard 18-word save area; therefore, standard IBM linkage conventions do not apply to it. One aspect of this difference is that CALLRTM does not save registers in this work area in the same order as it would in a standard save area. All 18 words are used.

Because TYPE = MEMTERM processing circumvents all task recovery and task resource manager processing, its use is restricted to a select group of routines that can determine that task recovery and task resource manager clean-up are either not warranted or will not successfully operate in the address space being terminated. These routines include:

- Paging supervisor, when it determines that it cannot swap in the LSQA for an address space
- Memory create, when it determines that an address space cannot be initialized
- RTM or supervisor control functional recovery routine (FRR), when it determines that uncorrectable translation errors
- RTM, when it determines that are occurring in the address space task recovery and termination cannot take place in the current address space

- Region control task, when it has determined that the address space might become permanently deadlocked -- that is, unusable -- or that the status of the address space is unpredictable because of an error during swap-out processing
- RTM, when all tasks in the address space have terminated
- Auxiliary storage management (ASM) recovery, when it has an indeterminate error from which it cannot recover while handling a request for either swap-in or swap-out
- SVC 34, in response to a FORCE command

In addition, the terminal control address space (TCAS) specifies TYPE=MEMTERM when the system operator replies "FSTOP" (forced stop) to certain messages that can occur when TSO/VTAM time sharing starts or stops. The messages are IKT001D (replying "FSTOP" cancels terminal users already active when TSO/VTAM is starting) and IKT010D (replying "FSTOP" cancels terminal users still active when TSO/VTAM is being stopped). In both cases, the system operator should reply "FSTOP" to cancel users only if "SIC" (system-initiated cancellation) is ineffective. Replying "SIC" does not cause task resource manager processing to be bypassed.

ABEND

Any routine, including supervisor state, locked, disabled, or SRB routines, can issue the ABEND macro instruction to direct the recovery termination services to itself (cause entry into its recovery routine) or to its callers. The issuer of ABEND should remove its own recovery routine if it wishes its caller to be abended or to enter recovery. Control never returns to the issuer of the macro (except as a result of a retry). See *Supervisor Services and Macro Instructions* for a description of the ABEND macro instruction.

Processing Program Interruptions (SPIE, ESPIE)

The SPIE macro instruction enables a problem program executing in 24-bit addressing mode to specify an error exit routine to get control in response to one or more program error interruptions. The ESPIE macro instruction extends the function of SPIE to callers in 31-bit addressing mode. Callers in both 24-bit and 31-bit addressing mode can use the ESPIE macro instruction.

Each succeeding SPIE/ESPIE macro instruction completely overrides any previous SPIE/ESPIE macro instruction specifications for the task. The specified exit routine gets control in the key of the TCB (TCBPKF) when one of the specified program interruptions occurs in any problem program of the task. When a SPIE macro instruction is issued from a SPIE exit routine, the program interruption element (PIE) is reset (zeroed). Thus, a SPIE exit routine should save any required PIE data before issuing a SPIE.

If a caller issues an ESPIE macro instruction from within a SPIE exit routine, it has no effect on the contents of the PIE. However, if an ESPIE macro instruction deletes the last SPIE/ESPIE environment, the PIE is freed, and the SPIE exit cannot retry.

If the current SPIE environment is cancelled during SPIE exit routine processing, the control program will not return to the interrupted program when the SPIE program terminates. Therefore, if the SPIE exit routine wishes to retry within the interrupted program, a SPIE cancel should not be issued within the SPIE exit routine.

The SPIE macro instruction can be issued by any problem program being executed in the performance of the task. The control program automatically deletes the SPIE exit routine when the request block (RB) that created the SPIE macro instruction terminates.

The SPIE and ESPIE macro instructions and their related services are discussed in detail in *Supervisor Services and Macro Instructions*. The syntax of both the SPIE and the ESPIE macro instructions appears in Volume 2.

Interruption Types

The problem programmer can specify interruptions 1-15 using either the SPIE or the ESPIE macro instruction. The installation-authorized system programmer can also specify interruption 17. Interruption 17 designates page faults and can be specified so that a user-written SPIE/ESPIE exit routine gets control before a supervisor routine when a problem state page fault occurs. The user-provided SPIE/ESPIE exit routine gets control in problem program state and in the key of the TCB (TCBPKF) when a page fault occurs for the program that issued the SPIE/ESPIE macro instruction. The exit routine gets control in the addressing mode that was in effect when the SPIE or ESPIE macro instruction was issued. (If a SPIE macro instruction was issued this is 24-bit addressing mode.) The SPIE/ESPIE exit routine for interruption type 17 handles page faults at the task level. This includes all RBs executing under the task for which the SPIE/ESPIE was issued. The exit routine resolves page faults by invoking the paging supervisor.

A caller in supervisor state, who issues the SPIE macro instruction is abnormally terminated with a 30E abend completion code. A caller in supervisor state, who issues the ESPIE macro instruction is abnormally terminated with a 46D-18 abend completion-reason code. If the caller takes a page fault while in supervisor state, the exit routine does not get control even if a SPIE/ESPIE macro instruction specifying interruption type 17 is in effect. Supervisor routines resolve the page fault and continue program processing without abending the caller.

If a program fault occurs while a SPIE/ESPIE specifying interruption type 17 is in effect, the program check first level interrupt handler (FLIH) passes control to a SPIE/ESPIE service routine, which then passes control to the SPIE/ESPIE exit routine via an LPSW. The SPIE/ESPIE service routine sets up functional recovery routines (FRRs) to handle possible page faults caused by PIE/PICA references (for a SPIE) or EPIE references (for an ESPIE) during this set-up processing. If such a page fault occurs, the SPIE/ESPIE service routine returns to the FLIH, which invokes the paging supervisor to handle the original page fault. In this case the SPIE/ESPIE exit routine does not handle the original page fault because the SPIE/ESPIE service routine cannot provide the information that the exit routine needs. After the page fault is resolved, processing continues in the problem program.

There is another situation in which the exit routine might not get control when a page fault occurs. To use interruption type 17, you must place the PIE/PICA (for a SPIE macro instruction), the EPIE (for an ESPIE macro instruction), the SPIE/ESPIE program, and data areas in fixed storage. Page faults can occur after issuing the SPIE/ESPIE macro instruction and before placing this information in fixed storage. If a page fault occurs at this point, the SPIE/ESPIE service routine performs set-up processing and, if it can reference the PIE/PICA (for SPIE) and the EPIE (for ESPIE), passes control to the exit routine. If the exit routine encounters a page fault, the paging supervisor resolves the page fault unless the routine is running disabled. A disabled page fault causes an 0C4 abend. Once the page fault is resolved, normal processing continues in the exit routine.

Note: In MVS/370 the SPIE environment existed for the life of the task. In MVS/XA the SPIE environment is deleted when the request block containing the macro is deleted. That is,

when a program running under MVS/XA completes, any SPIE environments created by the program are deleted. This might create an incompatibility with MVS/370 for programs that depend on the SPIE environment remaining in effect for the life of the task rather than the request block.

Intercepting System Errors

Intercepting system events provides a way to get information about software conditions, in addition to the information normally supplied by dumping services during abnormal termination. The intercepting process, known as serviceability level indication processing (SLIP) is a major debugging tool.

SLIP definitions, called traps, specify the system conditions at the time of interception and the action that is to be taken after the interception. There are two types of SLIP traps: non-PER traps and PER traps. A non-PER trap obtains information about an error condition normally handled by the recovery/termination manager (RTM), such as a program check, abend, or restart interrupt. A PER trap uses program event recording (PER) hardware to obtain information about any of the following PER events:

- Successful branch: Successful execution of a branch taken within a user defined range of virtual addresses.
- Storage alteration: Alteration of the contents of a virtual storage location within a user-defined range of virtual addresses.
- Instruction fetch: Fetching and execution of an instruction within a user-defined range of virtual addresses.

Once the trap is defined and enabled, SLIP processing checks current system conditions for each SLIP event. PER traps are checked each time a PER event occurs; non-PER traps are checked when RTM processes an error. When current system conditions match the conditions specified in the trap, the action specified in the trap occurs. Normally, the action specified, such as scheduling an SVC dump or writing a GTF trace record, is designed to collect diagnostic data.

SLIP users can minimize the effect of PER traps on system performance by limiting the scope of a SLIP trap to, for example, a particular address space or job. A SLIP trap can also include controls that automatically disable the trap when it is causing excessive overhead or has collected the required amount of data. The action to be taken when the trap matches can also be tailored. For example, SLIP users can tailor the SVC dump to include only diagnostic material essential for debugging. In addition, SLIP users can specify the type and contents of the GTF trace records. For more information on the SLIP command operands, see *Operations: System Commands* and *SPL: TSO*.

Using the SLIP Command

Use of the SLIP command should be restricted to system programmers.

The command uses three operands to control SLIP traps:

- SET -- establish SLIP traps
- MOD -- modify SLIP traps
- DEL -- delete SLIP traps

It is also possible to display information about SLIP traps by using the DISPLAY command at the operator's console or from a TSO terminal. For specific information about how to enter both the SLIP and DISPLAY commands, refer to *Operations: System Commands* or *SPL: TSO*. The descriptions in those publications also explain the operands. For more information on designing effective SLIP traps, see *Diagnostic Techniques*.

Obtaining an SVC Dump During Slip Processing

SLIP processing invokes SDUMP (if A = SVCD or A = TRDUMP is specified) to obtain an SVC dump at the time the SLIP trap is entered. It uses SDURGPSA (a system parameter of the SDUMP macro instruction, not available for general use) to specify the following information:

- The PSW at the time of the error
- The contents of control register 3 and 4
- The contents of general registers 0-15

SVC dump puts this information into the dump header record. Summary dump dumps 4K of storage (for one address space only) around the PSW at the time of the error. The address space used depends on the setting of the S-bit. (If the S-bit is on, it dumps this area in the secondary address space; if off, it dumps this area in the primary address space. In addition, it dumps 4K of storage around the addresses in all general registers for both address spaces.

SDUMP copies the storage specified by the SDURGPSA parameter into SDUMP storage before returning to SLIP. This allows SLIP to continue without waiting for SDUMP to complete.

Bypassing Dump Suppression

Dump analysis and elimination (DAE) suppresses duplicate SDUMPs and SYSMDUMPs under certain conditions. However, you can use the ACTION keyword of the SLIP command to override DAE. When all of the match conditions are met for a non-PER SLIP trap, the ACTION=NOSUP keyword prevents suppression of a recovery routine's dump request.

The SLIP actions of SVCDUMP and TRDUMP (trace dump) also override DAE. These actions trap critical problems that always require a dump.

"Suppressing SDUMPs and SYSMDUMPs" later in this section provides more information about DAE and contains a list of other publications that describe DAE.

System Trace Facilities

MVS/XA system trace provides a record in storage of significant software events. System tracing is in effect, by default, in MVS/XA. The system trace facility consists of branch tracing, address space tracing, and explicit tracing. Each of these types of tracing causes entries to be made in the trace table. The trace table consists of one queue of trace buffers for each processor. These buffers are located in the private area of the trace address space. *Diagnostic Techniques* provides diagnostic information about the system trace table.

Performing Branch Tracing

Branch tracing causes entries to be made in the trace table when the following branch instructions execute successfully:

- Branch and link register (BALR)
- Branch and save register (BASR)
- Branch and save and set mode (BASSM)

Performing Address Space Tracing

Address space tracing causes entries to be made in the trace table when the following instructions execute successfully:

- Program call (PC)
- Program transfer (PT)
- Set secondary ASID (SSAR)

Performing Explicit Tracing (PTRACE)

System services use the PTRACE macro instruction to explicitly trace normal system events. Callers in key 0 and supervisor state can also use the PTRACE macro instruction to make explicit entries in the trace table. These entries consist of an event identifier, the contents of a designated range of general registers or storage locations and system supplied status information. See the *Debugging Handbook, Volume 1* for examples of trace records.

The TRACE operator command dynamically controls tracing. It allows the operator to start or stop tracing, set the trace table size, and select the type of tracing. The TRACE command is described in *Operations: System Commands*.

Dumping Virtual Storage

There are a variety of ways in which users can obtain dumps of virtual storage. The macro instructions that produce dumps are ABEND, SNAP, and SDUMP. Any user can issue the ABEND or SNAP macro instruction to request a dump of virtual storage. These unprivileged macro instructions are described in *Supervisor Services and Macro Instructions*. A system routine that encounters an error can use the SDUMP macro instruction to obtain a dump. The syntax of the SDUMP macro instruction is given in Volume 2 and a description of its use follows.

There are several operator commands related to controlling, taking, displaying, and suppressing dumps. These commands are: CHNGDUMP, DISPLAY DUMP, DUMP, DUMPDS, and SET DAE. The syntax and description of these operator commands are given in *Operations: System Commands*. A brief description of the way in which a system programmer can use these commands is given later under the heading "Using the Dumping Services Commands."

Additional information on dumps, of interest to system programmers, can be found in *SPL: System Modifications, Diagnostic Techniques, Debugging Handbook, IPCS User's Guide and Reference, and Service Aids*. *SPL: System Modifications* describes pre-dump suppression, post-dump installation exits, and dump data sets. *Diagnostic Techniques* provides information on the analysis and use of dumps. *Debugging Handbook* describes the types of dumps and also gives samples of the output from dumps. *Service Aids* describes Print Dump (PRDMP), including the improved and reformatted summary dump output that it produces in MVS/XA. *IPCS User's Guide and Reference* describes the various subcommands that allow you to analyze a dump online.

This topic contains information that a system programmer needs to control and use MVS/XA dumping facilities. It includes the following subtopics:

- Using the IPCS Macro Instructions
- Using the SDUMP macro instruction
- Obtaining an SVC dump
- Obtaining a summary dump
- Using dump data sets
- Using the dumping services commands
- Cancelling and restarting the DUMPSRV address space
- Getting more than one SYSMDUMP
- Suppressing SDUMPs and SYSMDUMPs

Using the IPCS Macro Instructions

Interactive Problem Control System (IPCS) provides several macro instructions that PRDMP, SNAP, IPCS and user-written exit routines can use to tailor dump output.

The BLSQMDEF and BLSQMFLD macro instructions are used together to create a control block model. The format model processor service and the control block formatter service under IPCS and PRDMP use this model to format a control block in a dump. The model is a read-only structure that resides in a load library or in a CSECT within an exit program, and describes the control block to the formatter. The model consists of a header, defined by the BLSQMDEF macro instruction, and an array of entries, defined by the BLSQMFLD macro instruction, that describe individual fields.

The BLSRESSY macro instruction maps the IPCS symbol table record for use in the get symbol and equate symbol services. With the BLSRESSY macro instruction, users of the get symbol and equate symbol services can retrieve definitions described in the table and can create definitions for later use by the IPCS user or by other routines.

The BLSABDPL macro instruction maps the exit parameter list (BLSABDPL), a data area that enables IPCS, PRDMP, SNAP, and user-written exit routines to tailor dumps. With the BLSABDPL macro instruction, users can access different parameter lists within the BLSABDPL parameter list and then invoke the corresponding exit service routine.

For the syntax of the macro instructions, see *SPL: System Macros and Facilities*, Volume 2. For information about the exit service routines, see *IPCS Planning and Customization*.

Using the SDUMP Macro Instruction

System routines use the SDUMP macro instruction to obtain fast unformatted dumps of virtual storage. The SDUMP macro instruction invokes SVC dump to provide these services. Only one SVC dump can be taken in the system at one time.

Users can issue the SDUMP macro instruction while executing in either 24-bit or 31-bit addressing mode. The routine that issues the macro instruction regains control in the same addressing mode it was in when it requested the dump. SVC dump can handle both 24-bit and 31-bit addresses. It uses the value of the PSW, at the time of the error, to determine the addressing mode of the storage to be dumped. Users who switch addressing mode should note that SVC dump interprets the contents of registers and the addresses of parameter lists according to the caller's addressing mode at the time of the error.

Using SDUMP in a Reentrant Program

Callers who want to generate reentrant code must code the list and execute forms of the SDUMP macro instruction. In order to use the list and execute form of the SDUMP macro instruction, the caller needs to know the length of the parameter list that is generated by the list form of the SDUMP macro instruction. The length of the parameter list varies depending on the options that you specify.

The SDUMP parameter list is a minimum of 40 bytes in length. If any of the following options are specified, the parameter list is at least 48 bytes long:

```
TYPE = XMEM or XMEME
SDATA = GRSQ
LISTA
SUMLSTA
SUSPEND = YES or NO
```

If any of the following options are specified, the parameter list is at least 68 bytes long:

```
TYPE = NOLOCAL
SUBPLST
KEYLIST
SDATA = ALLNUC
```

If the HDR or STORAGE option is specified, the length of these areas must be added to the value 40, 48, or 68, as indicated above, to obtain the total length of the parameter list.

A caller can determine the amount of storage needed for the parameter list of the SDUMP macro instruction in the following ways:

- Include the IHASDUMP mapping macro in the program. This provides the mapping of all of the parameters (whether specified or not) and is the maximum length, excluding the length of the HDR and STORAGE options.
- Dynamically compute the amount of storage needed and place the result in SDUMPLEN:

```
SDUMPBEG SDUMP SDATA=(SUM) ,SUBLIST=SLIST ,MF=L
SDUMPEND EQU   *
SDUMPLEN DC    A(SDUMPEND-SDUMPBEG)
```

SDUMP Options

The parameters of the SDUMP macro instruction provide several important options including: the type of entry to the dumping routine, the address spaces to be dumped, the use of the SQA buffer, and a reason code option for dumps that fail.

BRANCH Options: Issuers of the SDUMP macro instruction can specify branch entry (BRANCH=YES) or SVC entry (BRANCH=NO) to the dumping routine. Callers who specify an SVC entry must satisfy one of the following conditions:

- Be APF authorized
- Be in supervisor state
- Be executing in a system key

Callers who cannot issue an SVC must specify branch entry on the SDUMP macro instruction to obtain an SVC dump. The branch entry caller must be in key 0, supervisor state, and must also satisfy one of the following conditions:

- Be in SRB mode
- Hold any lock
- Be disabled (with a supervisor bit on in the PSASUPER field of the prefixed save area)
- Have an enabled-unlocked-task FRR on the FRR stack

The branch entry interface uses standard linkage conventions. On entry, register 13 must point to a 72-byte save area. Branch entry callers must include the CVT mapping macro instruction with the PREFIX=YES parameter.

Address Space Options: There are several options that allow a user to dump one or more address spaces up to a limit of 15. These options are obtained by specifying the following parameters of the SDUMP macro instruction:

```
ASID
ASIDLST
LISTA
TYPE=XMEM
TYPE=XMEME
SUBPLST
```

Except for SUBPLST, these parameters automatically cause an SRB to be scheduled to produce the dump. This type of dump is called a scheduled dump. If any ASID in the subpool list specified by SUBPLST is different from the current ASID, SUBPLST also produces a scheduled dump; if all of the ASIDs in the subpool list are not different from the current

ASID, SUBPLST produces a synchronous dump. A synchronous dump is a dump that is finished when control returns to the caller rather than a dump that is still in progress when control is returned to the caller.

SQA Buffer Option: Callers who specify the BUFFER= YES parameter of the SDUMP macro instruction will obtain a dump of a 4K buffer reserved in the system queue area (SQA) for the callers of SVC dump. A user can reserve the buffer (by setting the high order bit of the CVTSDBF field of the communication vector table (CVT)) and fill it with information before invoking SVC dump. The buffer should be used by routines that are involved with volatile data that would be changed or must be changed before SVC dump can dump it.

The CVTSDBF field of the CVT points to the buffer. Before using the buffer, callers of SVC dump must check that the high order bit of CVTSDBF is off, using compare and swap logic. If the bit is set, assume that a dump is in progress and continue processing as if a dump could not be taken. If the bit is not set, set the bit before filling the buffer and calling SVC dump.

Failing Dump option: Callers who specify the TYPE= FAILRC parameter on the SDUMP macro receive special information from SVC DUMP whenever the dump fails. When control returns to the caller after a dump failure and TYPE= FAILRC was specified, the reason code is combined with the return code and passed to the caller in both register 15 and the ECB. When the return code is in the ECB, the POST flag is set on. The values and meanings of these codes follow:

Value	Meaning
00000208	Another dump was in progress at the time of request.
00000308	DUMP=NO was specified at IPL time or CHNGDUMP SET,SDUMP,NODUMP was specified.
00000408	Dump was suppressed by the SLIP NODUMP facility.
00000508	No SYS1.DUMP data set was available.
00000608	I/O error occurred writing the first record.
00000808	No SRBs could be scheduled to start the dump.
00000908	A termination error occurred in SVC DUMP before the first record could be written.
0000A08	No dump could be taken because of a STATUS STOP SRBs condition, which prevents SVC DUMP from scheduling any SRBs or doing I/O.
0000B08	The dump was suppressed because the dump analysis and elimination facility determined that the problem was a duplicate of a known problem.

Determining the Initial Status of an SVC Dump Request

Users can determine the initial status of an SVC dump request by examining the SDWASDRC byte in the system diagnostic work area (SDWA). The possible values of this indicator and their meanings follow:

Value	Meaning
0	No SVC dump was requested.
1	An SVC dump was successfully started.
2	An SVC dump was suppressed because another SVC dump was in progress.
3	An SVC dump was suppressed by a request by the installation (for example: DUMP=NO at IPL or CHNGDUMP SET,SDUMP,NODUMP).
4	An SVC dump was suppressed by a SLIP NODUMP command.
5	An SVC dump was suppressed because a SYS1.DUMP data set was not available. (Only for synchronous dumps)
6	An SVC dump was suppressed because an I/O error occurred during the initialization of the SYS1.DUMP data set. (Only for synchronous dumps)
8	An SVC dump was suppressed because an SRB could not be scheduled to activate the dump tasks in the requested address spaces.
9	An SVC dump was suppressed because a terminating error occurred in SDUMP before the first dump record was written.
10	An SVC dump was suppressed because a status stop SRB condition was detected. (This prevents dump I/O from completing.)
11	An SVC dump was suppressed by DAE.
12-254	Reserved.
255	An SVC dump was suppressed for some other unspecified reason.

Obtaining an SVC Dump

The two types of SVC dumps are a scheduled dump and a synchronous dump. The type of dump produced depends on the dump options in the SDUMP parameter list and the mode of entry (SVC entry or branch entry).

Scheduled Dump

A scheduled dump is a dump that is scheduled by an SRB. It is produced in the following cases:

- **BRANCH= YES** is specified on the **SDUMP** macro instruction.
- **BRANCH= NO** is specified (or used as a default) on the **SDUMP** macro instruction along with one or more of the following parameters:
 - **ASID** or **ASIDLST**
 - **LISTA**
 - **TYPE= XMEM** or **TYPE= XMEME**
 - **SUBPLST** (with any **ASID** different from the **ASID** of the current address space)

Callers interested only in the current address space should not use the **ASID** option. If the current **ASID** is specified as a parameter on the **SDUMP** macro instruction, the caller will obtain a scheduled dump and not a synchronous dump.

If a caller obtaining a scheduled dump does not specify the **ECB** option and wait for the dump to complete, the dump is produced asynchronously with the caller's recovery. For this reason, a scheduled dump is sometimes called an asynchronous dump. When **BRANCH= YES** is specified, the dump is also called a branch-entry **SDUMP**. The dump taken by **SLIP** is a branch-entry **SDUMP**.

A scheduled dump can contain one of three types of summary dumps. See the section "Obtaining a Summary Dump" for a description of the types of summary dumps.

Synchronous SDUMP

A synchronous **SDUMP** is a non-scheduled dump that is completed before the caller regains control from **SVC 51**. This type of dump is always entered by an **SVC (BRANCH= NO)** and is sometimes called an **SVC-entered SDUMP**. A **SYSMDUMP** is an example of a synchronous dump. It is requested by an **SVC 51** during **ABDUMP** processing.

In the case of a synchronous dump, the caller should not specify the **ECB** option, because there is nothing to wait for.

A synchronous dump can contain only an enabled summary dump. See the topic "Obtaining a Summary Dump" for a description of this type of dump.

Obtaining a Summary Dump

A summary dump will be obtained by default unless the **CHNGDUMP** command specifies **NOSUM** or **SDATA= NOSUM** is specified as a parameter on the **SDUMP** macro instruction. In order to avoid duplicate data in a summary dump, the dumping routine uses a range table to accumulate the addresses of the storage areas to be dumped. Duplicate storage is eliminated only around the **PSW** and registers. If the **SDWA** address is in a register, it is dumped once as the **SDWA** in formatted output, and once as unformatted storage.

The type of summary dump depends on two SDUMP parameters, BRANCH and SUSPEND. (The default for both of these parameters is NO.) The summary dumps and the parameters that produce them are:

Summary Dump	Parameters
DISABLED	BRANCH = YES, SUSPEND = NO
SUSPEND	BRANCH = YES, SUSPEND = YES
ENABLED	BRANCH = NO (SUSPEND cannot be specified)

A description of each of the three types of summary dumps follows.

Disabled Summary Dump: The purpose of the disabled summary dump is to save volatile system information before returning control to the user. The caller can specify the SUMLIST and SUMLSTA parameters to save specific information in the summary dump. However, because the system is disabled, the dump includes only data that is paged in.

A disabled summary dump contains the following information:

- The cross memory status record giving the home, primary, and secondary ASIDs and the CML ASID (if the caller holds a CML lock)
- The storage areas specified by the parameters SUMLIST and SUMLSTA
- The PSA, PCCA, and LCCA for each processor
- The current PCLINK stack (pointed to by PSASEL)
- The IHSA and its associated XSB and PCLINK stack (The PSW and registers from the IHSA are added to the range table causing 4K of storage around each address to be dumped.)
- If it exists, the caller's SDWA (The PSW and register addresses from the SDWA are added to the range table causing 4K of storage around each address to be dumped.)
- The SUPER FRR stacks
- The global, local, and CPU work save area (WSA) vector tables and the save areas pointed to by entries in each WSA vector table
- 4K of storage on either side of the address portion of the I/O old PSW, the program check old PSW, the external old PSW, and the restart old PSW

Suspend Summary Dump: The purpose of the suspend summary dump is to save volatile system information before returning to the caller. The caller can specify the SUMLIST and SUMLSTA parameters to save specific information in the summary dump. The difference between the suspend summary dump and the disabled summary dump is that the suspend summary dump can save pageable data.

A suspend summary dump contains the following information:

- The ASID record, PSA, PCCA, LCCA, IHSA, XSB, and PCLINK stack
- The storage area specified by the parameters SUMLIST and SUMLSTA
- The caller's ASCB

- The caller's unit of work (a TCB, RB, and XSB or SSRB, XSB, and PCLINK stack)
- For TCB mode callers, the caller's SDWA, RTM2 work areas, the associated SDWAs, the PSW, and registers
- For SRB mode callers, the SDWA (The PSW and register addresses from the SDWA)
- The caller's register save area is added to the range table causing 4K of storage around each register address to be dumped.

Enabled Summary Dump: The purpose of the enabled summary dump is to group information for debugging dumps by specifying a particular option on the SDUMP macro instruction. This is the only type of summary dump that can be produced by an SVC entry to SDUMP. If the dump is a scheduled dump, the summary information is saved for each address space specified.

An enabled summary dump contains the following information for each address space:

- The ASID record (contains ASID, jobname, and stepname)
- The storage areas specified by the parameters SUMLIST and SUMLSTA
- The RTM2 work areas pointed to by all TCBs in the address space
- 4K of storage before and 4K after the PSW in each RTM2WA at the time of the error
- 4K of storage before and 4K after each register in each RTM2WA at the time of the error

Suppressing SDUMPs and SYSMDUMPs

You can use dump analysis and elimination (DAE) to suppress SDUMPs and SYSMDUMPs when the symptoms of the dump match existing symptoms recorded in the SYS1.DAE dataset and the VRADAE indicator in the SDWAVRA is set. To set the VRADAE indicator issue the following macro instruction in your recovery routine:

```
VRADATA KEY (VRADAE)
```

Specification of this key indicates to DAE that the SDWA contains sufficient data to uniquely identify the problem.

Use of DAE reduces the number of duplicate dumps and avoids tracking and screening of unnecessary dumps. System performance is improved because duplicate problems are recognized before the dump is taken. Another benefit is that the symptom data, stored in the SYS1.DAE dataset, provides a consistent set of data for identifying a failure.

The following publications contain additional information about DAE:

- *Operations: System Commands* contains the syntax and use of the SET DAE command.
- *SPL: System Modifications* provides information about how an installation can modify DAE to fit its needs.
- *Debugging Handbook* contains sample symptom output and DAE control block information.
- *System Logic Library* provides a description of the logic and an explanation of symptom strings.
- *SPL: Initializing and Tuning* contains information about the ADYSETxx parmlib members that are used by DAE.

Using Dump Data Sets

The SDUMP macro instruction, the DUMP command, and the ACTION = SVCD or TRDUMP parameter of the SLIP command produce SVC dumps that are stored in dump data sets. A dump data set can be placed on a direct access storage device (allocated with RECFM = F, LRECL = 4104) or on a tape. The system obtains these data sets by using the DUMP parameter of the system parameter list (at IPL) or by using the DUMPDS operator command (after IPL). (For details on the DUMP parameter of the system parameter list, see *Initialization and Tuning*. For a description of the DUMPDS operator command, see *Operations: System Commands*.)

Dump data sets are dynamically allocated to the DUMPSRV address space with DISP = SHR. Users must do one of the following:

- If you are running another job using these data sets (for example, AMDPRDMP service aid), specify DISP = SHR when allocating the dump data set (This will share the dump data set with the DUMPSRV address space.)
- Issue the DUMPDS operator command to remove the data set from use by the system

Using the Dumping Services Commands

There are several operator commands related to controlling, taking, displaying, and suppressing dumps. For complete details and the syntax of each of these commands, see *Operations: System Commands*.

A brief description of the use of these commands follows:

Command	Use
CHNGDUMP	Modifies or overrides the default dump options used on SYSABEND, SYSUDUMP, SYSMDUMP, and SDUMP.
DISPLAY DUMP	Displays the current dump options set by the CHNGDUMP command and some specific information from the dump header record. It also indicates whether the dump data sets are full or empty and includes a timestamp.
DUMP	Obtains a scheduled dump for a job or address space. It can also obtain a dump of the system for performance analysis.
DUMPDS	Modifies the SDUMP data set queue and thereby adds or deletes data sets from use by SDUMP. It also clears full dump data sets.
SET DAE = xx	Depending on the keywords specified (by the SYS1.PARMLIB member ADYSETxx), starts or stops DAE processing or requests that SDUMPs or SYSMDUMPs be suppressed, matched, and/or updated.

Canceling and Restarting the DUMPSRV Address Space

If an installation is using post dump exits, cancelling the DUMPSRV address space can be useful. The cancellation, followed by the resource manager's restart causes new versions of the post dump exits to be loaded and used. See *SPL: System Modifications* for additional information concerning an installation's use of post dump exits.

Getting More Than One SYSMDUMP

MVS/XA can handle a SYSMDUMP in two ways. One makes only the most recent SYSMDUMP available; each dump after the first overlays the preceding one. The other makes only the first SYSMDUMP available; any subsequent dumps are lost. If you choose this second method, however, you can take additional steps to avoid losing subsequent SYSMDUMPs.

To make only the first SYSMDUMP available, you must specify `DSNAME=SYS1.SYSMDPxx,DISP=SHR` on the SYSMDUMP DD statement. (See *Job Control Language* for more information about the SYSMDUMP DD statement.) Any other value for DSNAME or DISP causes the system to make only the most recent SYSMDUMP available.

When the data set is created, it must be cleared and an EOF must be written as the first record on the SYSMDPxx data set. EOF means that the data set is empty; anything other than EOF means that the data set is full and subsequent dumps are lost. To obtain these subsequent SYSMDUMPs, you must intercept the message

```
IEA993I SYSMDUMP TAKEN TO data set name
```

and pass control to an installation-written routine that will copy the dump onto another data set, clear the data set, and write another EOF as the first record on the SYSMDPxx data set. One way to intercept message IEA993I is to use the WTO exit routine IEECVXIT. (See *SPL: User Exits* for information on message routing exit routines. *SPL: System Modifications* contains information on post dump exit processing, which can be used to offload SYSMDUMP data sets.)

Providing Recovery Routines

When a unit of work (either a task or an SRB) terminates abnormally, RTM invokes a recovery routine, if one exists. The recovery routine might determine where the error occurred and decide whether to continue with termination or to continue processing (retry) at some appropriate point in the main routine. Other functions that your recovery routine might perform are to document the error, take a dump of the storage necessary to determine the cause of the error, and, when required, free resources that are no longer needed. See "Providing Information for DAE" later in this section for information about documenting the error.

The decision about whether or not your particular function requires recovery depends, naturally enough, on the nature of the function. One major factor in this decision is the set of resources that the function acquires. If a function acquires resources that might be requested by another function or that are not known to be related to the task or SRB then a recovery routine should be established to free the resources. Examples of resources that must be freed are queues of control blocks serialized for exclusive use, private locks, cells obtained from cell pools, and storage that must be explicitly freed.

Once you have determined that your particular function requires recovery, you must select the type of recovery routine that your function needs -- either a functional recovery routine (FRR) or an ESTAE-type recovery routine.

Because of the complexity of the MVS/XA environment and the various uses your installation might make of that environment, it is difficult to provide many hard and fast rules about which type of recovery routine is best for a particular situation. If the recovery routine requires that a

lock held by the main routine not be freed, or if cross memory mode is a factor, the recovery routine must be an FRR. In other situations generally grouped as task recovery, an ESTAE-type recovery routine can provide the function you need.

In general, establish an FRR to provide recovery for locked, disabled, or SRB mode routines or routines executing in cross memory mode. Identify an FRR to RTM by coding the SETFRR macro instruction. This macro instruction creates an entry in a system recovery area called an FRR stack. When RTM invokes an FRR, the FRR runs with the locks that were held at the time of the error, or as modified by previous FRRs, and the enablement implied by those locks.

In general, establish an ESTAE-type recovery routine to provide recovery for unlocked tasks that do not execute in cross memory mode. Identify an ESTAE-type recovery routine to RTM by coding the ESTAE or FESTAE macro instruction or the ESTAI parameter on the ATTACH macro instruction.

Note: The STAE macro instruction and the STAI parameter of the ATTACH macro instruction are available for compatibility with Release 1 of VS2 and with MFT and MVT. “STAE/STAI Exit Routines” later in this chapter provides information on STAE and STAI; however, it is recommended that you use ESTAE, FESTAE, or ESTAI for all new development.

“Selecting a Recovery Routine” later in this chapter describes the factors, particularly the system environment, that you must consider in choosing the type of recovery routine for a particular situation. Which type of recovery routine you choose also affects the order in which your recovery routine can get control. When percolation occurs, recovery routines get control in LIFO (last-in, first-out) order, starting with FRRs (if any). If all FRRs percolate, then all ESTAE-type recovery routines get control in the reverse of the order in which they were established. See “Percolation, Retry, and Resume” later in this chapter for detailed information on the process of percolation, the factors that go into making a request for a retry to the main routine, and a definition of the special situation when an FRR can request a resume.

In some situations, the function that your recovery routine must perform is itself one that requires recovery. See “Recovery for Recovery” later in this chapter for a description of such situations and some of the considerations involved in providing recovery for a recovery routine.

“Recovery Routine Guidelines” pulls together various pieces of information for you to use as a basis for making decisions about recovery routines and designing recovery routines.

Providing Information for Dump Analysis and Elimination

Dump analysis and elimination (DAE) depends on information that users provide in ESTAE and FRR recovery routines to construct unique symptom strings needed to describe software failures. DAE uses these symptom strings to analyze dumps and suppress duplicates as requested. The symptom string contains symptoms (specific pieces of information) that DAE obtains from the system diagnostic work area (SDWA), SDWA extensions, ABDUMP symptom area, and the SDWA variable recording area (SDWAVRA).

If you are placing information into the SDWAVRA for use by DAE, you must provide the information in key/length/data format. You can use the VRADATA macro to create the entry in this format.

Users must select symptoms carefully. If the data they supply is too precise, no other failure will have the same symptoms; if the data is too general, many failures will have the same symptoms. See “Suppressing SDUMPs and SYSMDUMPs” earlier in this section for additional information about DAE and for a list of other publications that document DAE.

Selecting a Recovery Routine

Several basic restrictions and requirements govern your choice between an FRR and an ESTAE-type recovery routine. These are:

- An FRR can be established only when you make the request in key 0 and supervisor state.
- If an enabled unlocked task establishes an FRR using EUT = YES, it cannot issue any SVCs except SVC 13. Also, the system does not dispatch any new asynchronous exits on that task until all FRRs for the task have been deleted.
- From the time an FRR is established until the time it is deleted, at least one of the following must be true.
 - Some lock is held.
 - The unit of work is executing either disabled or as an SRB.
 - An FRR with EUT = YES exists.
- The size of each FRR stack satisfies the recovery needs of the control program. If additional FRRs placed on the stack cause the size to be exceeded, the routine issuing the SETFRR macro instruction terminates abnormally. Any user-written routines outside of the control program may add one, and only one, FRR to the stack; if they add more than one, abnormal termination occurs if the size of the stack is exceeded. This applies to all of the recovery stacks, including the normal stack. The normal FRR stack is used by control program routines that are invoked on behalf of the user.
- You cannot use the ESTAE macro instruction or the ESTAI parameter of the ATTACH macro instruction to establish a recovery routine when executing in cross memory mode or as an SRB.
- You can use the FESTAE macro instruction in cross memory mode as long as data addressability at the time the macro is issued is to the dispatched address space (home).

Beyond these basic restrictions and requirements, your choice of an FRR or an ESTAE-type recovery routine depends on the environment for which recovery is to be established and on the environment in which the recovery routine is to get control. Other major differences are the register interface and the availability of a system diagnostic work area (SDWA) for use by the recovery routine.

The following information under “System Environment,” “Register Interface,” and “System Diagnostic Work Area (SDWA)” describes the differences between FRRs and ESTAE-type recovery routines in more detail. And, if you choose an ESTAE-type recovery routine, you must decide whether to establish it with the ESTAE macro instruction, the FESTAE macro instruction, or the ESTAI parameter of the ATTACH Macro instruction. See “ESTAE-type Recovery Routines” for the basic differences between the three methods of establishing an ESTAE-type recovery routine.

System Environment

The environment -- the state of the system -- consists of such factors as locks held, enablement or disablement, supervisor or problem program state, PSW key, PSW key mask (PKM), authorization index, and cross memory mode. To enable you to determine whether an FRR or an ESTAE-type recovery routine best meets your needs, identify the system environment your recovery routine requires and use the following information to select the type of recovery routine.

LOCKING

ESTAE-Type Recovery Routines: All locks are freed before the first ESTAE-type recovery routine gets control.

FRRs: The locks held when the first FRR gets control are the same as they were at the time of the error.

If the recovery routine is to free locks held by the main routine, the FRR should issue the SETRP macro instruction to use the lock-freeing functions of RTM.

The lock-freeing functions of RTM are effective only for percolation, not for retry. In any event, the FRR must not free the last global lock and it must not free any local lock because RTM depends on these locks to serialize the system diagnostic work area (SDWA) passed to the FRR. Freeing these locks can cause different units of work to use the same SDWA simultaneously.

For more information on how RTM manipulates locks, see "Decisions Made in a Recovery Routine" later in this chapter.

DISABLEMENT

ESTAE-Type Recovery Routines: All ESTAE-type recovery routines are entered enabled.

FRRs: An FRR gets control and is entered disabled if, at the time of the error, the main routine is disabled and any of the following states exist:

1. Any global spin lock is held.
2. A super FRR stack was active or the control program set a bit in the PSASUPER area of the PSA. (Generally, this condition occurs when an interrupt handler is running.)

ADDRESSING MODE

The addressing mode for all recovery routines (both ESTAE and FRR) is the addressing mode of the caller at the time the routine was established.

SUPERVISOR/PROBLEM PROGRAM STATE

ESTAE-Type Recovery Routines: An ESTAE-type recovery routine is entered in the state -- either supervisor or problem program -- that existed at the time it was established.

FRRs: All FRRs are entered in supervisor state.

AUTHORIZATION INDEX (AX)

The authorization index for all recovery routines -- both ESTAE-type recovery routines and FRRs -- is the effective AX for the address space in which the recovery routine is to get control. Refer to "Cross Memory Authorization" in the Inter-Address Space Communication section for more information about the AX.

PSW KEY

ESTAE-Type Recovery Routines: An ESTAE-type recovery routine is entered with key 0 whenever it was established with a key less than 8. Otherwise, it is entered with the key that existed at the time it was established.

FRRs: All FRRs are entered in key 0.

PSW KEY MASK (PKM)

ESTAE-Type Recovery Routines: All ESTAE-type recovery routines are entered with a PSW key mask (PKM) that is the ORing of the following:

- The PSW key under which the exit is to get control.
- The TCBPKF.
- The PKM that existed when the ESTAE-type recovery routine was established.

However, if the recovery routine was established with a FESTAE macro instruction, the system uses the PKM that existed at the time of the error.

FRRs: See the following information on "Cross Memory State" for information about the PKM that exists when an FRR is entered.

CROSS MEMORY STATE

ESTAE-Type Recovery Routines: All ESTAE-type recovery routines are entered in primary mode with the primary address space (PASID) and the secondary address space (SASID) the same as the dispatched address space (HASID).

FRRs: The cross memory environment for an FRR depends on the values coded for the MODE parameter in the SETFRR macro instruction that established the FRR.

NORMAL Addressing Environments

Specifying HOME, PRIMARY, or FULLXM for the MODE parameter of the SETFRR macro instruction indicates to RTM the normal or expected addressing environment of the FRR.

MODE = HOME

If you specify MODE=HOME or omit the MODE parameter, the FRR gets control in home mode; that is, the FRR is entered with PASID=SASID=HASID, in primary mode. The PSW key mask (PKM) for an FRR that covers SRB code is the same as the PKM at the time of the error. The PKM for an FRR that covers a task is the TCBPKF transformed into a PKM.

MODE=PRIMARY

If you code **MODE=PRIMARY**, the FRR gets control in primary mode with the primary and secondary address space the same as the primary address space that existed when the **SETFRR** macro instruction was issued. The PKM is the PKM that existed when the **SETFRR** macro instruction was issued.

MODE=FULLXM

If you specify **MODE=FULLXM**, the FRR gets control in the cross memory environment that existed when the **SETFRR** macro instruction was issued. That is, the primary address space (PASID), secondary address space (SASID), the mode, and PKM are the same as those that existed when the **SETFRR** macro instruction was issued.

If RTM cannot enter an FRR with its normal addressing environment established as defined by the **MODE** parameter on the **SETFRR** macro instruction, RTM bypasses the FRR and percolates to the next FRR on the FRR stack, unless the FRR was established to execute in a restricted addressing environment (also specified by the **MODE** parameter).

RESTRICTED ADDRESSING ENVIRONMENTS

Specifying either **LOCAL** or **GLOBAL** (or both) for the **MODE** parameter of **SETFRR** indicates to RTM that the FRR can run in a restricted addressing environment. When an FRR is entered in a restricted addressing environment, it is often called a resource manager because its only purpose to recover resources. These resources can be critical system resources (**GLOBAL**) or critical address space resources (**LOCAL**). RTM does not allow an FRR to retry.

When RTM enters an FRR to recover critical resources, it sets bits in the SDWA to indicate to the FRR in which of the restricted addressing environments it is being entered.

If you specify both **LOCAL** and **GLOBAL**, RTM first tries to enter the FRR in the **LOCAL** restricted addressing environment. If it cannot, RTM then tries to enter the FRR in the **GLOBAL** restricted addressing environment. These environments, and the bits that RTM sets to indicate the environment to the routine, are:

MODE=GLOBAL: **MODE=GLOBAL** should be used by services that need to clean up global resources if the main routine terminates abnormally. When **GLOBAL** is specified and RTM cannot enter the FRR in its normal mode, it enters the FRR in the restricted **GLOBAL** mode if one of the following system conditions exists:

- The main routine holds a global spin lock.
- A “super FRR stack” was active or the control program has set a bit in the PSASUPER word of the PSA. (Generally this condition occurs when an interrupt handler is running.)

A global FRR must reside in commonly addressable storage.

When RTM enters an FRR in restricted **GLOBAL** mode, the entry environment is:

- PASID=SASID, in primary mode. The PASID can be that of **any** address space; thus, when entered in this mode, the routine must not reference private storage.
- RTM sets the SDWAGLBL bit in the SDWA to one to indicate that RTM is entering the routine in restricted **GLOBAL** mode (otherwise, this bit is set to zero).

- RTM sets the SDWACLUP bit in the SDWA to one to indicate that the routine is not allowed to retry, although, if system conditions permit, subsequent FRRs are permitted to retry.

MODE=LOCAL: MODE=LOCAL should be used by services that need to clean up critical address space related resources serialized by means of a local lock (including the CML lock). When LOCAL is specified and RTM cannot enter the FRR in its normal mode, it enters the FRR in the restricted LOCAL mode as long as a local lock is held and the address space whose lock is held has not terminated or suffered a DAT error.

If it is possible for the FRR to get control in one address space in normal mode and in another address space in restricted mode, the FRR must reside in commonly addressable storage.

RTM enters an FRR established with the LOCAL parameter in LOCAL restricted mode for two different reasons:

1. RTM tried to establish the environment required to enter the FRR in normal mode but could not; this problem can occur, for example, when the SASID is no longer valid.
2. An address space is terminating and at least one unit of work in that address space is holding the local lock for another address space (CML lock).

In both cases, the entry environment is:

- Primary mode, the home address space can be any address space, and the PASID and SASID are the same as the locked address space.
- RTM sets the SDWALCL bit in the SDWA to one to indicate that RTM is entering the routine in LOCAL restricted mode (otherwise, this bit is set to zero).
- RTM sets the SDWALCL and SDWAGLBL bits if the SETFRR macro instruction included both the LOCAL and GLOBAL parameters, and RTM is entering the FRR in LOCAL restricted mode while the LOCAL lock is held and one of the following is true:
 - A spin lock is held
 - A super bit is set
 - The FRR is on a super FRR stack
- RTM sets the SDWACLUP bit in the SDWA to one to indicate that the routine is not allowed to retry, although, if system conditions permit, subsequent FRRs are permitted to retry.

The following considerations apply when RTM enters the FRR in LOCAL restricted mode as a result of the second reason stated earlier.

1. If the FRR issues a SETRP macro instruction to request that RTM free the CML lock, subsequent FRRs are not entered because the resources in the address space are no longer serialized and therefore no further LOCAL resource clean up can be done.
2. The FRR must not depend on executing in task mode because, even though the SETFRR macro instruction was issued in task mode, an FRR entered in LOCAL restricted mode executes in SRB mode. In this case, the information in the SDWA reflects the interrupted process (from IHSA or SSRB) that originally held the CML lock, and the FRR is not permitted to retry.

3. The FRR created with the LOCAL parameter must be prepared to be suspended during its normal mode recovery processing and then be entered a second time in LOCAL restricted mode to recover critical address space resources. If the executing FRR has established another FRR, specifying MODE=LOCAL, the newer FRR gets control in LOCAL restricted mode followed by the FRR that was in control at the time of suspension.

REGISTER INTERFACE

ESTAE-Type Recovery Routines: Before entering an ESTAE-type recovery routine, RTM attempts to obtain and initialize a system diagnostic work area (SDWA). The SDWA contains information about the error; see "System Diagnostic Work Area (SDWA)" later in this section. The first word of the SDWA contains the address of the parameter list specified on the ESTAE-type recovery request. The register interface to the recovery routine varies depending on whether or not RTM can obtain an SDWA.

If RTM can obtain an SDWA, the register contents on entry to the recovery routine are:

Register	0	a code indicating the type of I/O processing performed:
	0	active I/O has been quiesced and is restorable.
	4	active I/O has been halted and is not restorable
	8	no active I/O at ABEND time
	16	no I/O processing was performed
Register	1	address of the SDWA
Register	13	save area address (72 bytes)
Register	14	return address
Register	15	entry point address of the ESTAE recovery routine

The contents of all other registers are unpredictable.

If RTM cannot obtain an SDWA, the register contents on entry to the recovery routine are:

Register	0	a decimal 12 to indicate that an SDWA was not obtained.
Register	1	ABEND completion code
Register	2	address of user-supplied parameter list
Register	13	unpredictable
Register	14	return address
Register	15	entry point address of the ESTAE recovery routine

The contents of all other registers are unpredictable.

When a recovery routine entered without an SDWA completes, it must set a return code for RTM in register 15. This return code must be one of the following:

Hexadecimal Code	Meaning
0	RTM is to continue with termination (percolate).
4	RTM is to schedule a retry. The recovery routine has placed the address of the retry routine in register 0.
16	Valid only for an ESTAI recovery routine. RTM is to continue with termination, and no further ESTAI processing is to be performed.

FRRs: All FRRs are entered with an SDWA. The register contents on entry to the FRR are:

Register	0	address of a 200-byte work area for the FRR
Register	1	address of the SDWA
Register	14	return address
Register	15	address of the FRR

The contents of all other registers are unpredictable.

The FRR can use any register without saving its contents; however, the routine must maintain the return address supplied in register 14.

SYSTEM DIAGNOSTIC WORK AREA (SDWA)

The system diagnostic work area (SDWA) is a communication area between RTM and the recovery routine. RTM always supplies an SDWA for an FRR. If for any reason RTM cannot supply an SDWA, it bypasses the FRR and percolates to the next FRR on the FRR stack. Because RTM cannot always supply an SDWA for an ESTAE-type recovery routine, an ESTAE-type recovery routine must always check the contents of register 0 to determine if there is an SDWA available.

The mapping macro IHASDWA provides the field names and describes their content and use. Detailed information on the name, offset, and meaning of each field appears in the *Debugging Handbook*. To enable you to determine what kinds of information RTM supplies to the recovery routine in the SDWA and what kinds of information the recovery routine must return to RTM, Figure 54 lists the names and meanings of some of the key fields.

It is recommended that you use the SETRP macro instruction, which manipulates fields in the SDWA, to change fields in the SDWA before returning control from your recovery routine to RTM.

Field Name	Use
SDWAPARM	This 4-byte field, located at offset 0, contains the pointer to the user parameter list supplied by the user for an ESTAE-type recovery routine. For an FRR, this field contains the address of the six-word parameter area returned by a SETFRR macro instruction with the PARMAD parameter.
SDWACMPC	This field contains the ABEND completion code that existed when RTM entered the recovery routine. The recovery routine can change the ABEND code by altering this field. The system code appears in the first twelve bits, and the user code appears in the second twelve bits.
SDWACRC	This field contains the error code associated with the ABEND completion code in SDWACMPC. The recovery routine can change the reason code by altering this field by means of the SETRP macro instruction and its REASON keyword. This field is not defined if the SDWARCF flag field is zero.
SDWAGRSV	This field shows the contents of the general purpose registers (0-15) as they were at the time of the error.
SDWAEC1	This field contains the extended control (EC) PSW that existed at the time of the error.
SDWAEC2	The contents of this field vary according to the type of recovery routine: <ul style="list-style-type: none">• For an ESTAE/FESTAE routine, the field contains the PSW of the RB that created the recovery routine at the time the RB last incurred an interruption. However, if this RB was disabled or was holding a lock when the interrupt occurred, the field does not contain the last interrupt address, but the address of the second to last interrupt.• For an ESTAI routine, this field contains zeros.• For an FRR, the field contains the extended control PSW used to give control to the FRR.

Figure 54 (Part 1 of 2). Key Fields in the SDWA

Field Name	Use
SDWASRSV	<p>The contents of this field vary according to the type of recovery routine:</p> <ul style="list-style-type: none"> ● For an ESTAE/FESTAE routine, this field contains the general purpose registers (0-15) of the RB that established the recovery routine as they were at the time the RB last incurred an interrupt. ● For an ESTAI routine, this field contains zeros. ● For an FRR, this field has the same contents as SDWAGRSV. <p>If the recovery routine requests a retry, RTM uses the contents of this field to load the registers for the retry routine. To update the contents of the registers for the retry routine, you must make the required changes to SDWASRSV and request a register update using the RETREGS= YES parameter on the SETRP macro instruction. You can update the registers directly or with the RUB parameter on SETRP.</p>
SDWASPID	This field contains the subpool ID of the storage used to obtain the SDWA.
SDWALNTH	This field contains the length, in bytes, of this SDWA, the SDWA extensions, and the variable recording area. (This allows the user to free the extensions along with the SDWA.)
SDWACOMU	The recovery routines can use this eight-byte field to communicate with each other when percolation occurs. RTM copies this field from one SDWA to the next on all percolations. When the field contains all zeros, either no information is passed or RTM has not been able to pass the information.
SDWAFAIN	This 12-byte field contains the six bytes of the instruction stream that both precede and follow the failing instruction pointed to by the PSW. The SDWAFAIN field contains zeroes if RTM cannot access the failing instruction stream pointed to by the time-of-error PSW. For example, if the time-of-error PSW is not valid, the SDWAFAIN field contains zeroes.
SDWADAET	This eight-byte field contains DAE status and error flags for this dump.
SDWAOCUR	This two-byte field contains the current count of the number of previous occurrences of these symptoms in other SDWAs.
SDWAVRAL	This field contains the length of the variable recording area (VRA) for this SDWA.
SDWAHEX	This bit is set by the recovery routine to indicate that EREP is to print the data in the VRA in hexadecimal form.
SDWAEBC	This bit is set by the recovery routine to indicate that EREP is to print the data in the VRA in EBCDIC.
SDWAURAL	This one-byte field is used by the recovery routine to indicate the length of the VRA used. The field is always zeroes initially, and it must be set whenever a recovery routine uses any part of the VRA.

Figure 54 (Part 2 of 2). Key Fields in the SDWA

ESTAE-Type Recovery Routines

You establish an ESTAE-type recovery routine (also known as an ESTAE exit or an ESTAE environment) by means of the ESTAE macro instruction, the FESTAE macro instruction, or the ESTAI parameter of the ATTACH macro instruction. Any ESTAE-type recovery routine, regardless of how it is established, executes under its own program request block created by the SYNCH service routine. The recovery routine executes in the same addressing mode as the issuer of the ESTAE macro instruction.

Before an ESTAE-type recovery routine gets control, the control program performs any purge and asynchronous processing that was specified as a macro option when the routine was established. The control program performs the requested I/O processing only for the first recovery routine. Subsequent routines receive an indication of the I/O processing previously performed, but no additional I/O processing is done. However, the control program performs asynchronous processing for each routine.

An ESTAE macro instruction creates a task-related recovery routine. A fast ESTAE (FESTAE) macro instruction performs the same function with minimal processor overhead. A recovery routine established by the FESTAE macro instruction gets control in the same sequence and under the same conditions as a recovery routine established by the ESTAE macro instruction. However, only an SVC routine executing under an SVRB (type 2,3 or 4 SVC) can issue FESTAE; other factors to be considered for using FESTAE are described later under "Using the FESTAE Macro Instruction." The ESTAI parameter on the ATTACH macro instruction establishes a recovery routine that gets control if the attached subtask encounters an unrecovered abnormal termination; that is, a recovery routine established by the ESTAI parameter gets control under the subtask. Also, any recovery routine established by the ESTAI parameter is propagated to any subsequent subtasks. Figure 55 shows the queuing structure of the ESTAE routines and the propagation of routines created by ESTAI to a subtask.

Because ESTAE-type recovery routines are associated with RBs, they are removed when their RBs terminate. The fact that they are removed is important because a program expects one of its own ESTAE-type recovery routines to get control rather than one left behind by a subprogram. A program might, however, invoke a service routine that does not create an RB. If that routine then issues an ESTAE macro instruction and fails to delete the resulting ESTAE-type recovery routine, a problem could develop if the original program is scheduled for abnormal termination. The ESTAE-type recovery routine left behind by the service routine would receive control rather than the ESTAE-type recovery routine associated with the program, because the recovery routine specified by the most recently-issued ESTAE macro instruction gets control.

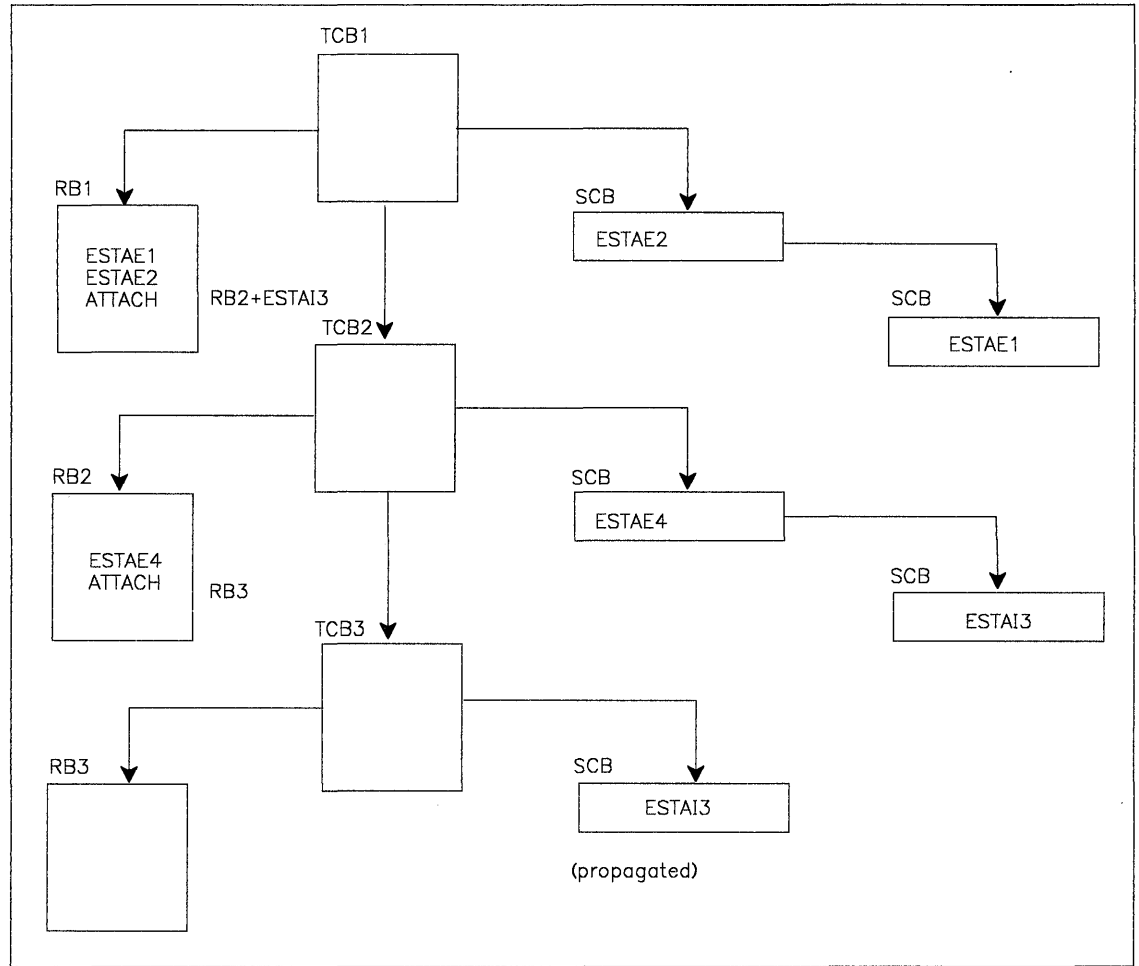


Figure 55. ESTAE Environment

This potential problem can be avoided by using the `TOKEN` parameter on the `ESTAE` macro instruction to associate a token with that `ESTAE` routine. In order to delete or overlay an `ESTAE` routine that was created with `TOKEN`, the same token must be specified in the `ESTAE` cancel or overlay macro instruction. All more recent `ESTAE` exits are deleted.

If a program issues an `ESTAE` macro instruction that specifies both the `TOKEN` parameter and `XCTL = YES` and then issues `XCTL`, the token must be passed as part of the parameters to the called routine so that the routine can delete the `ESTAE` routine.

Using the FESTAЕ Macro Instruction

The FESTAЕ macro instruction enables a type 2, 3, or 4 SVC (an SVC for which the SVC FLIH creates an SVRB) to establish an ESTAE-type recovery routine with minimal processor overhead. However, the following restrictions apply:

- Only a type 2, 3, or 4 SVC in key 0 can use FESTAЕ.
- The SVC can use FESTAЕ only once to create a recovery routine. Therefore, any SVC needing to change its exit address must use branch entry ESTAE services, and any SVC needing more than a single recovery routine must use SVC 60 or branch entry to get the additional recovery routines.
- FESTAЕ can be issued in cross memory mode if the home address space is addressable at the time the FESTAЕ recovery routine is entered. In addition to the parameter area that you can supply by coding the PARAM parameter on the FESTAЕ macro instruction, a 24-byte parameter area is also available as an option. The name of the optional parameter area is RBFEPARM, and it is in the SVRB. The recovery routine receives this parameter area when an error occurs. Hence, the main routine can clear (set to zero) and initialize the parameter area with appropriate information (such as tracking data) that might be useful to the recovery routine. You must clear the parameter area before using it to ensure that no spurious data remains in it from previous processing.

FESTAЕ users must also include the following DSECTs for the FESTAЕ macro expansion:

IHARB
IKJTСB
IHAPSA
IHASCB

Special Considerations

When writing an ESTAE-type recovery routine, consider the following:

- When an ESTAE-type recovery routine receives control, it should first examine the code in register 0 to see if an SDWA was provided. If an SDWA was not provided (that is, register 0 contains a decimal 12), register 13 does not point to a save area, and your routine must not save the registers.
- An ESTAE-type recovery routine can request, via a SETRP macro instruction parameter, that the control program free the SDWA instead of freeing it in a retry routine. When the retry routine is to free the SDWA, note that an ESTAE-type recovery routine created under any control program protection key (key 0-7) receives an SDWA in key 0 storage. Therefore, if the retry routine is executing under a key other than key 0, it must issue the MODESET macro instruction to become key 0 before issuing the FREEMAIN macro instruction to free the SDWA.

- If an ESTAE-type recovery routine itself requests termination or fails, RTM percolates and does the following:
 - Accumulates dump options
 - Resets the asynchronous exit indicator according to the request of the next recovery routine
 - Ignores the I/O options for the next recovery routine
 - Initializes a new SDWA
 - Gives control to the next recovery routine

If all recovery routines fail or indicate termination, the task is terminated.

- If a non-jobstep task issues an ABEND macro instruction with the STEP parameter, RTM enters recovery for the non-jobstep task. If the recovery routines do not request a retry, the jobstep is terminated with the specified ABEND code. RTM enters subsequent recovery routines for the jobstep task only when the macro instruction that established the recovery routine specified the TERM = YES parameter.
- For some situations, RTM enters ESTAE-type recovery routines only when the TERM = YES parameter was specified when the ESTAE macro instruction was issued. The situations are:
 - System initiated logoff
 - Job step timer expiration
 - Wait time limit for job step exceeded
 - ABEND occurred because a DETACH macro instruction was issued for an incomplete subtask
 - Operator cancel
 - Error occurred on a task higher in the tree

When RTM enters the recovery routines established with the TERM = YES parameter as a result of the above errors, RTM takes the following actions:

- Gives control to all such routines in LIFO order
- Does not enter any ESTAI routine previously suppressed by a return code of 16 or any previously-entered recovery routine that specified a return code of 0
- Ignores any request for retry
- Ignores the TERM = YES parameter if it is specified on a nested ESTAE macro instruction

Decisions Made in a Recovery Routine

When a recovery routine gets control, it determines why it has been entered and decides either to percolate (continue with termination), to retry, or in a special case, to resume. To convey its decision to RTM, the recovery routine issues the SETRP macro instruction, which manipulates appropriate fields in the SDWA. When the recovery routine returns to RTM, RTM honors the request, if appropriate.

The MVS/XA recovery scheme provides a parameter area for communication between the main routine and its recovery routine. The parameter area varies according to the type of recovery routine:

1. For an FRR, RTM supplies a six-word parameter area.
2. For a recovery routine established by an ESTAE or FESTAE macro instruction the user can supply a parameter area by coding the PARAM parameter on the macro instruction.

When a recovery routine is established, RTM saves a pointer to the parameter area and makes the pointer available to the recovery routine when it is entered. Usually, the main routine uses the parameter area to leave a footprint, that is, to set indicators that let the recovery routine know where in the main process the failure occurred. The recovery routine can examine the footprint to determine what action to take.

Note: For a recovery routine established by a FESTAE macro instruction, there is also a six-word parameter area available in the SVRB. (RTM does not preserve the pointer to this area).

RESUME

When an operator presses the RESTART key on the processor to break a spin loop, RTM gives control to the FRR established by the routine executing on that processor. In this situation, the first -- and only the first -- FRR to get control can request RESUME. As a result of a RESUME request, RTM terminates the unit of work executing on the processor identified by the recovery routine. Execution resumes with the next sequential instruction on the interrupted processor.

RESUME should be specified only when the main routine is deliberately spinning waiting for a resource to be freed by the another processor.

RETRY

A retry request from a recovery routine asks RTM to continue execution of the code that established the recovery routine at some appropriate point. The retry routine executes in the same addressing mode as the issuer of the ESTAE macro instruction. The recovery routine cannot change the addressing mode of the retry routine. Note that retry is not always permitted. Whenever the system cannot permit a retry, RTM sets the SDWACLUP bit in the SDWA to one. If a recovery routine requests retry when it is not allowed, RTM ignores the request and continues with termination (percolates).

Any recovery routine that requests a retry must include logic designed to avoid recursion, to prevent the creation of a tight loop between the recovery routine and its retry routine. For example, if the recovery routine supplies a bad retry address to RTM, and the execution of the first instruction at the given address causes a program check, the first recovery routine to get

control is the one that just requested the retry. If the recovery routine requests another retry at the same address, the loop is created.

The environment in which the retry routine gets control from an FRR differs from the environment for a retry routine that gets control from an ESTAE-type recovery routine.

RETRY FROM AN FRR

An FRR can request a valid retry whenever the SDWACLUP bit in the SDWA is set to zero; to request a retry, the FRR must supply a retry address, the entry point of the retry routine. The retry address is the point in the main routine that is to get control in order to continue its processing. In response to a valid retry request, RTM gives control to the retry address that the recovery routine supplies. The retry routine executes as a continuation of the unit of work that encountered the error. Note that RTM does not delete the FRR that requests a retry; the FRR remains valid and can be entered again.

The environment that exists when the retry routine gets control is described in the following topics.

Registers: Upon entry to the retry routine, the contents of the registers are the same as the contents of the SDWASRSV field in the SDWA. The FRR that requests the retry manipulates SDWASRSV to set the contents of the registers for the retry routine. Register 15 always contains the retry address.

Locks: The status of locks held is the same on entry to the retry routine as it was when the FRR requesting retry completed its processing.

Disablement: The retry routine is entered disabled, if, and only if, the FRR requesting the retry returns to RTM disabled. Otherwise, the retry routine is entered enabled.

Supervisor/Problem Program State: The retry routine from an FRR is always entered in supervisor state.

PSW Key: The retry routine from an FRR is always entered in key zero.

Cross Memory Environment: The PKM, PASID, SASID, and mode for the retry routine can be either those that existed at the time of the error or those that existed at the time of the entry to the FRR that is requesting the retry. The FRR makes the choice. Use caution in choosing to establish the environment that existed at the time of the error because a problem with the environment might very well be the cause of the error.

The authorization index (AX) is the current one for the PASID of the retry routine.

SDWA: When the retry request comes from an FRR, the SDWA is not available to the retry routine.

Retry From an ESTAE-Type Recovery Routine

An ESTAE-type recovery routine can request a valid retry whenever the SDWACLUP bit in the SDWA is set to zero; to request a retry, the recovery routine must supply a retry address, the entry point of the retry routine. The retry address is the point in the code that established the recovery routine that is to get control in order to continue its processing. In response to a valid retry request, RTM gives control to the retry address that the recovery routine supplies. The retry routine executes as a continuation of the code that established the recovery routine. That is, the retry routine executes under the same RB that established the ESTAE-type

recovery, and RTM purges all RBs more current than the retry RB before giving control to the retry routine. Note that ESTAI is an exception; a retry request from a recovery routine established by the ESTAI parameter of the ATTACH macro instruction must execute under a PRB. If there is a PRB whose RBLINK field points to an RTM2 SVRB, that PRB is used. If there is no previous RTM2 SVRB on the queue, then the retry RB is the newest PRB that is older than the oldest non-PRB. If there is no PRB, retry is suppressed.

RTM purges the RB queue to attempt to cancel the effects of partially-executed programs that are at a lower level in the program hierarchy than the program under which the retry occurs. However, the RB purge does not cancel certain effects on the system such as:

- Subtasks created by an RB to be purged
- Resources allocated by the ENQ macro instruction
- DCBs that exist in dynamically-acquired virtual storage

If there are quiesced restorable I/O operations, the retry routine can restore them. RTM supplies a pointer to the purged I/O request list (PIRL). The retry routine can use SVC RESTORE to have the system restore all I/O requests on the PIRL.

Note that RTM does not cancel the ESTAE-type recovery routine that requests a retry; the recovery routine remains valid and can be entered again.

The environment that exists when the retry routine gets control is described in the following topics.

Registers: The contents of the general purpose registers upon entry to the retry routine depend on whether or not RTM was able to obtain an SDWA for the ESTAE-type recovery routine.

If RTM could not obtain an SDWA for the recovery routine, the register contents on entry to the retry routine are:

Register	Contents
0	A decimal 12.
1	Address of the user parameter list established using ESTAE or ATTACH with ESTAI.
2	A pointer to the PIRL, if I/O was quiesced and is restorable; otherwise zero.
14	Address of supervisor-assisted exit linkage (SVC 3)
15	Entry point address of the retry routine.

The contents of all other registers are unpredictable.

If RTM could obtain an SDWA for the recovery routine, the register contents depend on whether or not the recovery routine requested a register update and/or requested that the SDWA be freed. (See the SETRP macro instruction in Volume 2 of this publication for information on the RETREGS parameter used to request a register update and the

FRESDDWA parameter used to request that RTM free the SDWA.) The register contents are one of the following:

1. If the recovery routine did not request register update and did not request that the SDWA be freed, the register contents on entry to the retry routine are:

Register	Contents
0	Zero
1	Address of the SDWA
14	Address of supervisor-assigned exit linkage (SVC 3)
15	Entry point address of the retry routine

The contents of all other registers are unpredictable.

2. If the recovery routine did not request update but did request that the SDWA be freed, the register contents on entry to the retry routine are:

Register	Contents
0	A decimal 20
1	Address of the user parameter list established using ESTAE or ATTACH with ESTAI
2	A pointer to the PIRL, if I/O was quiesced and is restorable; otherwise zero
14	Address of supervisor-assisted linkage (SVC 3)
15	Entry point address of the retry routine

The contents of all other registers are unpredictable.

3. If the recovery routine requests register update, the contents of the 16 general purpose registers on entry to the retry routine are the same as the contents of the 16 words in SDWASRSV. In this case, the recovery routine provides the contents of the registers for the retry routine by updating any or all of the register slots in the SDWASRSV before returning control to RTM with the retry request. If the recovery routine does not also request that the SDWA be freed, it must keep a pointer to the SDWA; this pointer enables the retry routine to reference and subsequently free the SDWA. Note that register 15 does not contain the entry point address of the retry routine unless the recovery routine sets it up that way.

Locks

The retry routine is entered with no locks held. If the ESTAE-type recovery routine obtains any locks, it must free those locks before returning to RTM. Otherwise, an SVC error occurs and the retry routine does not get control.

Disablement: The retry routine is always entered enabled.

Supervisor/Problem Program State: If the recovery routine was established by an ESTAE macro instruction, the retry routine is entered in the state that existed when the macro instruction was issued. If the recovery routine was established by a FESTAE macro instruction, it is entered in supervisor state. If the recovery routine was established by the ESTAI parameter of the ATTACH macro instruction, the retry routine is entered in supervisor state if, and only if, the RBOPSW of the retry RB is in supervisor state and the main routine was authorized at the time of the error. Otherwise, the retry routine is entered in problem program state.

The main routine is considered to be authorized at the time of the error when at least one of the following is true:

1. The program is APF-authorized.
2. The TCBPKF of the task in error is less than 8.
3. The bit TCBFSM of the task in error (in TCBFLGS3 of the TCB) is on, indicating that all RBs for that TCB execute in supervisor state.

PSW Key

If the recovery routine was established by an ESTAE or FESTAE macro instruction, the retry routine is entered with the same PSW key that existed when the macro instruction was issued.

If the recovery routine was established by the ESTAI parameter of the ATTACH macro instruction, the retry routine is entered with the same PSW key as the one in RBOPSW of the retry RB when one of the following is true:

1. The main routine was authorized (as defined earlier under "Supervisor/Problem Program State") at the time of the error.
2. The RBOPSW of the retry RB has a key greater than or equal to 8 and is in problem program state, and the PKM of that RB (XSBKM field in the XSB control block for that RB) does not have authority to keys less than 8.

Otherwise, the PSW key of the retry routine is that of the TCBPKF.

PSW Key Mask (PKM): If the recovery routine was established by the ESTAE macro instruction, the retry routine is entered with the PKM that existed when the macro instruction was issued.

If the recovery routine was established by the FESTAE macro instruction, the retry routine is entered with the PKM that existed at the time of the error.

If the recovery routine was established by the ESTAI parameter of the ATTACH macro instruction, the retry routine is entered with the PKM from the XSBKM field in the XSB control block of the retry RB only if one of the following is true:

1. The main routine was authorized (as defined earlier under "Supervisor/Problem Program State") at the time of the error.
2. The RBOPSW of the retry RB has a key greater than or equal to 8 and is in problem program state, and the PKM of that RB does not have authority to keys less than 8.

Otherwise, the PKM of the retry routine only has authority that is equivalent to that of the TCBPKF.

Authorization Index (AX): The retry routine is entered with the authorization index that is in effect for the dispatched address space (HOME or PSAAOLD).

Cross Memory Mode: The retry routine is entered in primary mode with PASID = SASID = HASID.

Addressing Mode: The Retry routine is entered in the same addressing mode that existed when the recovery routine was entered.

SDWA: When RTM obtains an SDWA for an ESTAE-type recovery routine, the recovery routine can make the SDWA available to the retry routine. In this case, the retry routine must free the SDWA using the pointer to it and its length. The length appears in the SDWALNTH field, and the subpool in which the SDWA resides appears in the SDWASPID field.

Percolation

When a recovery routine gets control and cannot recover from the error (that is, it does not retry), it must free the resources held by the main routine and request that RTM continue with termination (percolate). Note that a recovery routine entered with the SDWACLUP bit set to one in the SDWA, indicating that retry is not permitted, has no choice but to percolate. When the recovery routine requests percolation, the previously-established recovery routine gets control. When a retry is not requested and RTM has entered all possible recovery routines, the unit of work (either an SRB or a task) terminates abnormally. Figure 56 shows the decisions RTM makes to determine which recovery routine is to get control in a particular situation.

When a recovery routine requests percolation, it is cancelled; that is, RTM effectively removes that recovery routine from the environment. A cancelled recovery routine is not entered again unless that recovery routine is established again after a retry. There are two types of percolation: percolation for the same unit of work and SRB-to-task percolation.

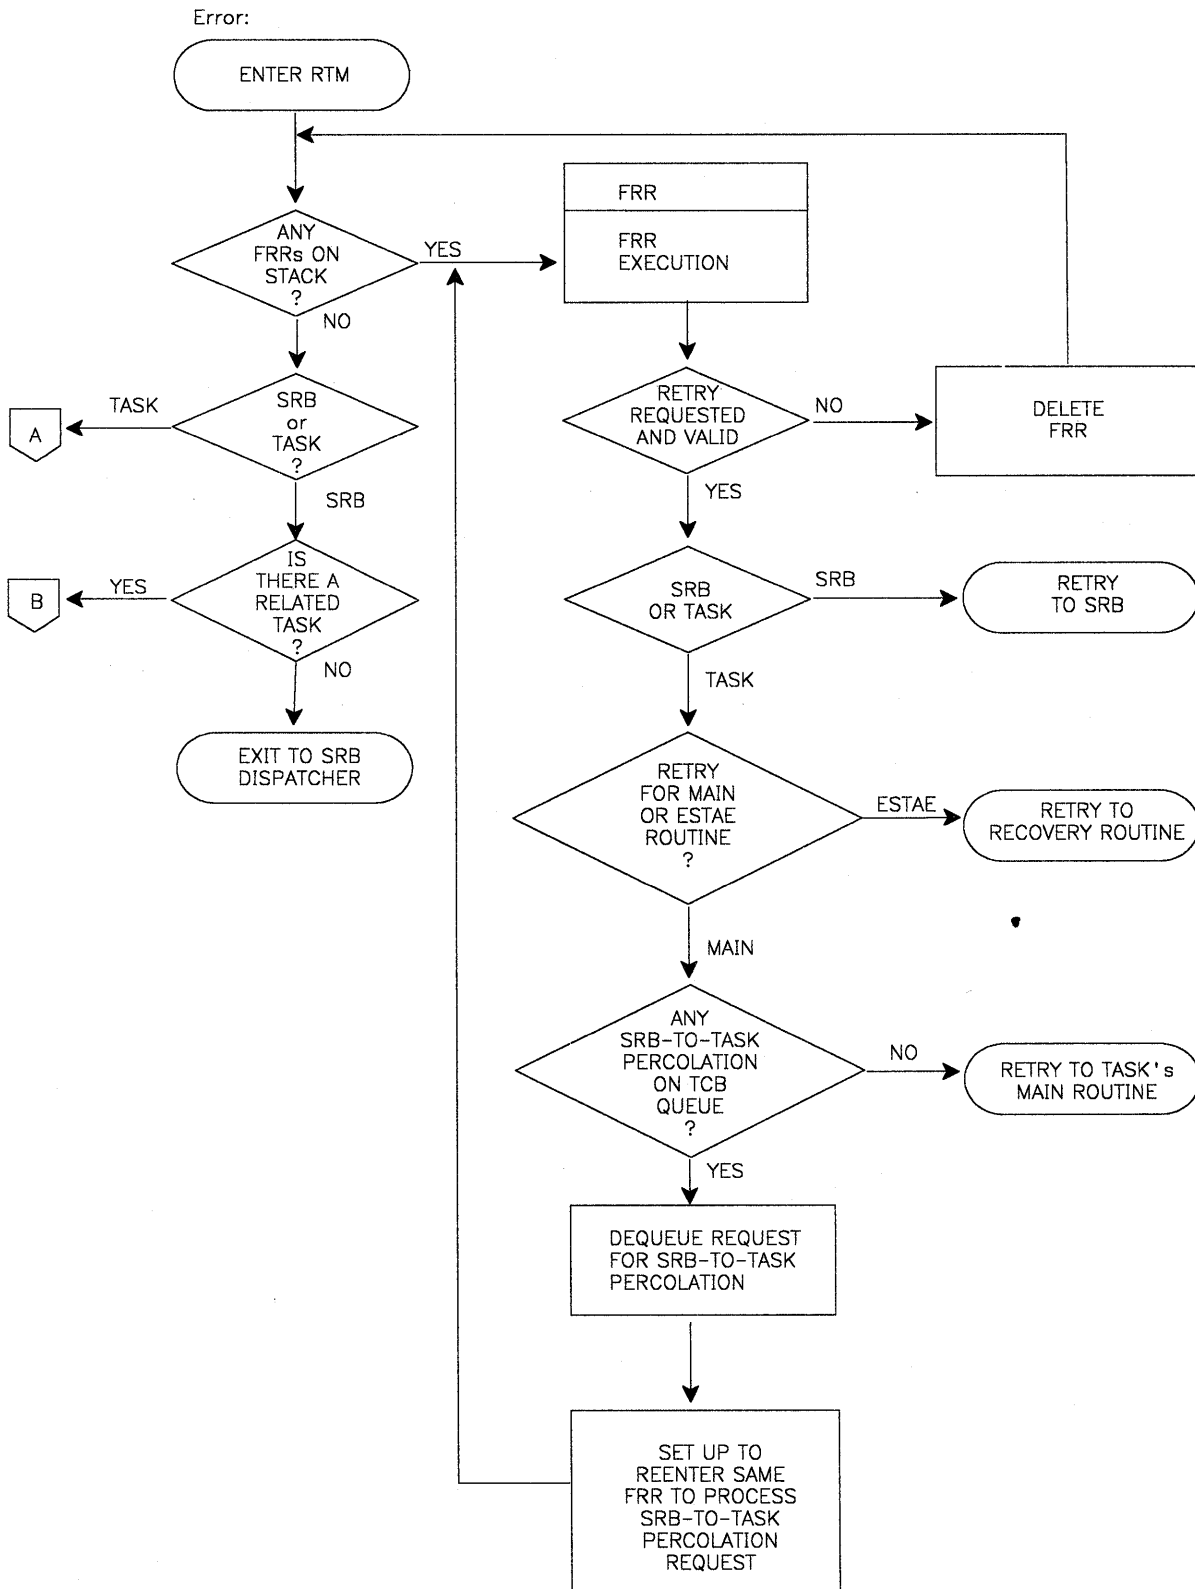


Figure 56 (Part 1 of 3). Routing Control to Recovery Routines

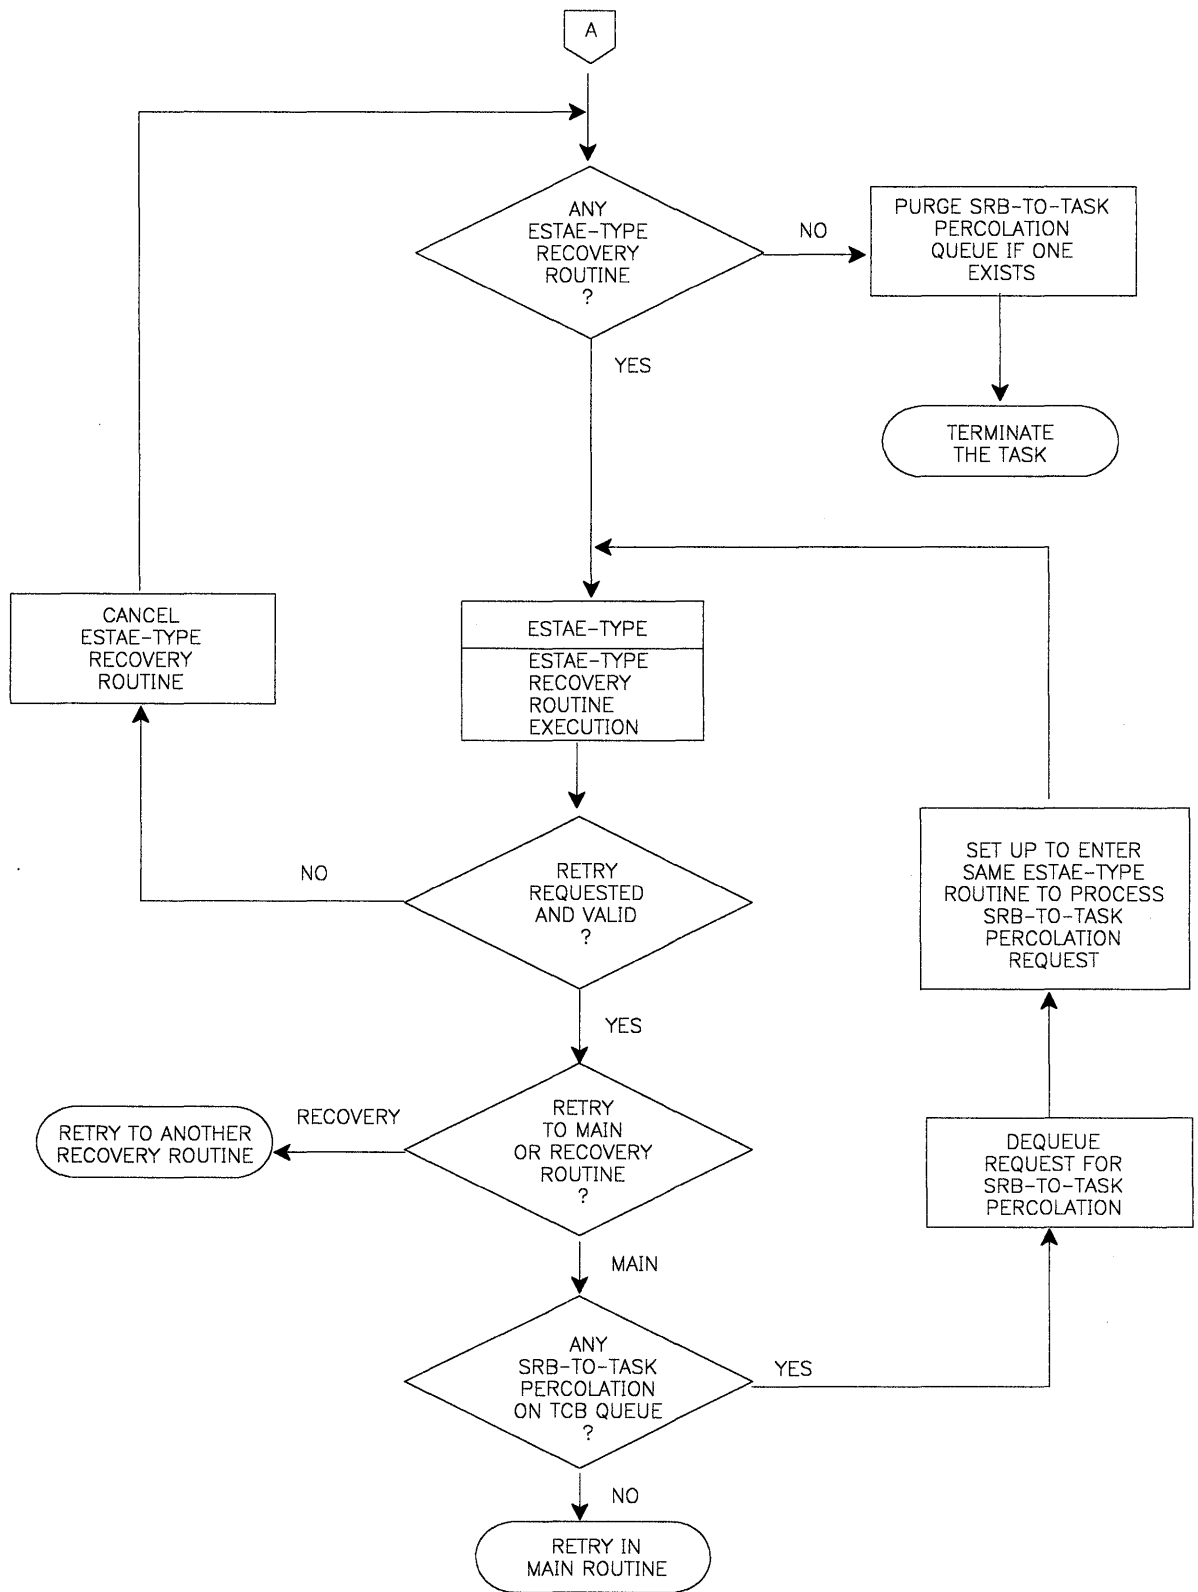


Figure 56 (Part 2 of 3). Routing Control to Recovery Routines

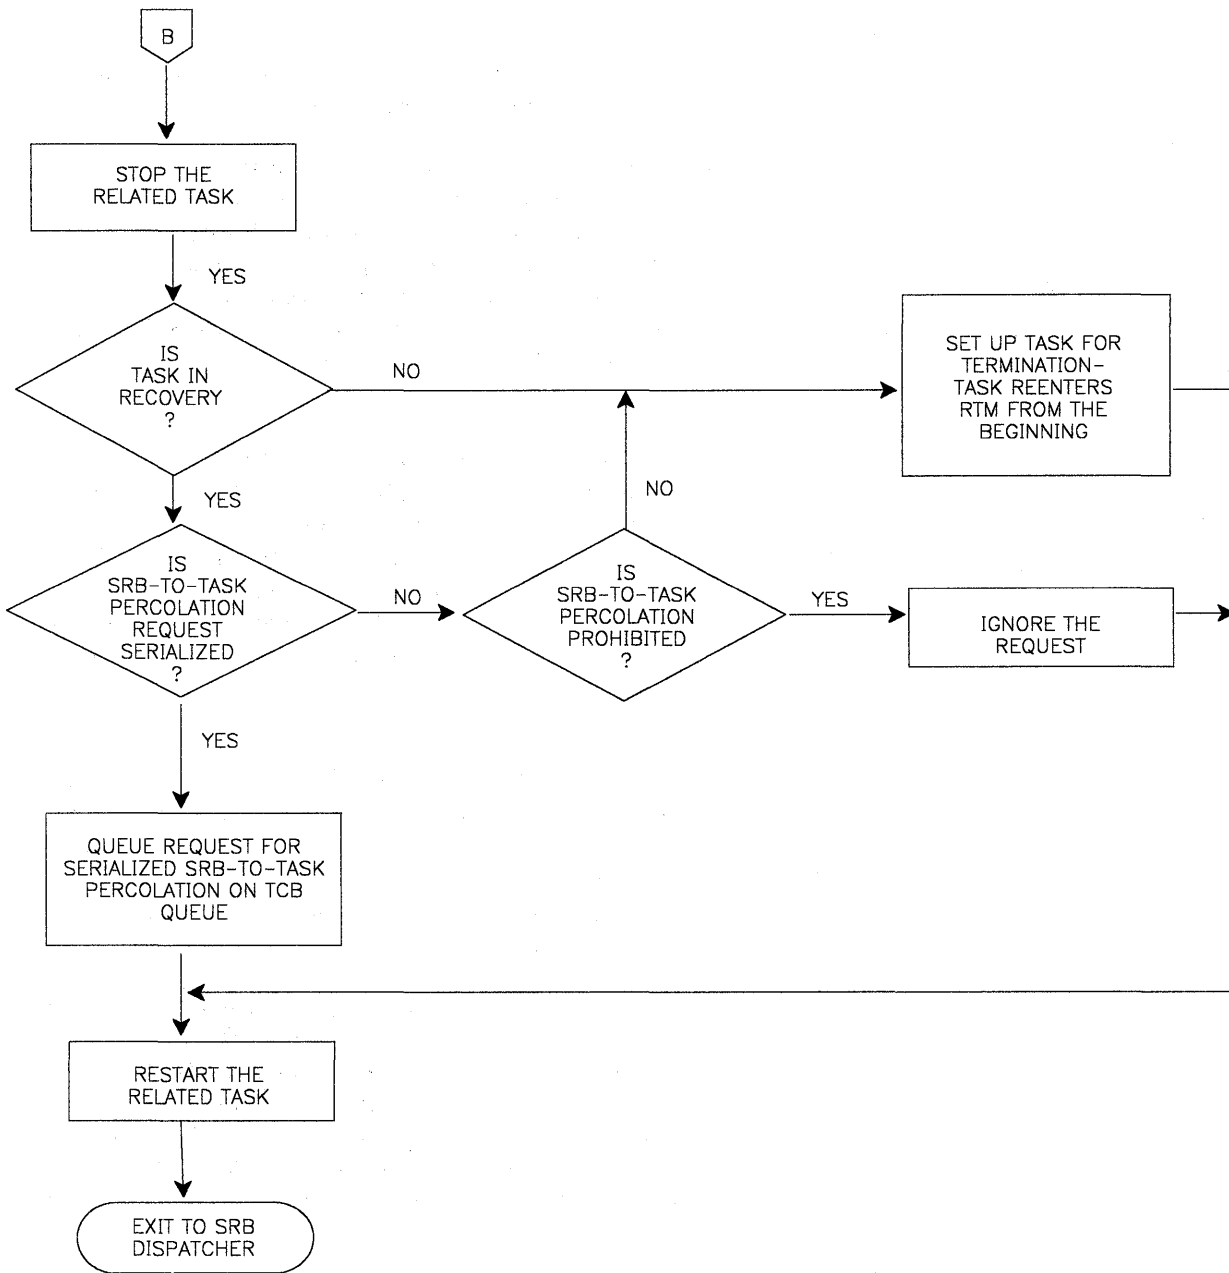


Figure 56 (Part 3 of 3). Routing Control to Recovery Routines

Percolation for the Same Unit of Work

Percolation for the same unit of work causes control to be given to one recovery routine after another for that same unit of work, which can be either a task or an SRB.

Percolation to an FRR always occurs from another FRR. The environment (except for the cross memory environment) in which a subsequent FRR gets control is the same as the one that existed when the first FRR was entered. The cross memory environment varies because FRRs can turn off the super bits. The lock status, which implies enablement or disablement, could also be different. If an FRR obtains locks that were not held when it was entered and then requests percolation, RTM frees those locks before giving control to the next FRR. Also, if the percolating FRR requested that RTM free any locks, RTM frees these locks before giving control to the next FRR.

Percolation to an ESTAE-type recovery routine can occur from either an FRR or another ESTAE-type recovery routine. The environment in which an ESTAE-type recovery routine gets control does not vary regardless of whether the percolation request came from an FRR or another ESTAE-type recovery routine. Note that a recovery routine established by the ESTAI parameter of the ATTACH macro instruction can choose either to percolate to a previous ESTAI routine (by setting a return code of 0 for RTM) or to bypass further ESTAI recovery routine processing and continue with termination (by setting a return code of 16 for RTM).

SRB-TO-TASK Percolation

When an SRB is scheduled and the fields SRBPASID and SRBPTCB are supplied, the task whose TCB address is in SRBPTCB and is executing in the address space whose ASID is in SRBPASID is defined as the SRB's related task. When an SRB with a related task terminates abnormally and the FRR for the SRB does not exist or does not request a retry, the error is percolated to the recovery for the related task. This percolation is called SRB-to-task percolation.

SRB-to-task percolation occurs if none of the FRRs established by the SRB retry or if the SRB does not have an FRR. Either case creates a request for RTM to perform SRB-to-task percolation. RTM ignores the request whenever the related task is terminated. RTM may ignore the request when the related task is already in recovery. If serialization is requested on the SETRP macro instruction in the FRR for that SRB, the percolation request is deferred. (See the SERIAL = YES parameter of the SETRP macro instruction in Volume 2 of this publication.) Serializing SRB-to-task percolation ensures that information about multiple SRB failures is not lost.

Note: SERIAL = YES should not be specified unless the task's recovery routine expects it.

RTM processes requests for non-serialized SRB-to-task percolation as follows:

- If the task is in recovery, RTM ignores the request.
- If the task is not in recovery, then RTM abnormally terminates the task and passes the information about the SRB's error to the task's recovery.

RTM processes requests for serialized SRB-to-task percolation as follows:

- If the task is already in recovery, RTM saves and queues the information about the SRB's error for processing later when the task recovers from the previous error.
- If the task is not in recovery, RTM abnormally terminates the task and passes the information about the SRB's error to the task's recovery.

When one of the task's recovery routines that is not a nested recovery routine (defined in "Recovery for Recovery" later in this chapter) requests a retry, RTM checks for queued requests for SRB-to-task percolation and takes the following actions before performing the retry:

1. If any requests are queued, RTM dequeues a request and again enters the recovery routine that requested the retry. RTM repeats this process as long as there are queued requests.
2. When the queue is empty or depleted, RTM honors the retry request and gives control to the retry routine.

Figure 56 shows this process.

The environment for a task recovery routine entered as a result of SRB-to-task percolation is the same as the environment described earlier under "Percolation for the Same Unit of Work."

However, the information in the SDWA describes the error that occurred in the SRB. Whenever serialized SRB-to-task percolation is requested and RTM must queue a request, RTM obtains an area from the user's private area to preserve the information about the SRB's error. If no space is available, RTM cannot preserve that information but still enters the task recovery for the request. If an SDWA is available, RTM sets the SDWARPIV bit to indicate that error-time information is not available. Also, RTM sets the SDWACOMU field to zeroes because RTM cannot preserve its contents to pass from the SRB's FRR to the task recovery routine.

Recovery for Recovery

In some situations, the function a recovery routine performs is so essential that you should establish a recovery routine to recover from errors in the recovery routine. Two examples of such situations are:

1. The availability of some resources can be so critical to continued system or subsystem operation that it might be necessary to establish a recovery routine for the recovery routine, thus ensuring the availability of the critical resources.
2. A recovery routine might perform a function that is, in effect, an extension of the main routine's processing. For example, a system service might elect to check a caller's parameter list for fetch or store protection. The service references the user's data in the user's key and, as a result of protection, suffers a program check. The recovery routine gets control and requests a retry in order to pass a particular return code to the main routine. If this recovery routine terminates abnormally and does not establish its own recovery, then the caller's recovery routine gets control, and the caller does not get an opportunity to check the return code that it was expecting.

You can establish an FRR from either another FRR or from an ESTAE-type recovery routine. You can also establish an ESTAE-type recovery routine from an ESTAE-type recovery routine.

However, do not establish an ESTAE-type recovery routine from an FRR because RTM gives control to all FRRs in a recovery path before giving control to any ESTAE-type recovery routines. Therefore, an ESTAE-type recovery routine established in an FRR might not get control in the proper sequence.

Any recovery routine established in a recovery routine is called a nested recovery routine (either a nested FRR or a nested ESTAE recovery routine). Nested ESTAE recovery routines can retry; the retry routine executes under the RB of the ESTAE-type recovery routine that established the nested recovery routine.

Nested FRRs, however, cannot retry. If you need to provide recovery for one FRR with a second FRR, and the second FRR must be able to request a retry, establish both FRRs in the main routine. For example, assume that you want to provide recovery for FRR A with another FRR, FRR B. In the main routine, issue the SETFRR macro instruction for FRR B first, followed by the SETFRR macro instruction that established FRR A. If the main routine terminates abnormally, FRR A gets control. If FRR A then encounters an error, FRR B gets control and, because it is not nested, it can request a retry. Note that RTM deletes FRR A when it gives control to FRR B. If FRR A is needed for the rest of the processing, the retry routine must establish it again. If FRR A is not established again, RTM passes any subsequent errors to FRR B.

All FRRs (nested or not) as well as the main routine, execute under the top RB while in task mode.

Recovery Routine Guidelines

The actions a recovery routine should take are highly dependent on the function of the main routine; the information presented here summarizes the major decisions -- whether or not you need a recovery routine and what kind of recovery routine you need -- and presents some considerations related to recovery routines that request retry and those that request a dump. "FRR Summary" identifies the information in this book that can help you to design an FRR; "ESTAE-Type Recovery Routine Summary" does the same for ESTAE-type recovery routines.

Deciding Whether Recovery Is Needed

The first decision you must make is whether or not the main routine requires a recovery routine. Usually, if a function acquires resources that might be requested by another function or that are not known to be related to the task, a recovery routine should be established to free the resources. An example of this type of resource is storage with a subpool that is not task-related (such as subpool 231). Another case that requires a recovery routine occurs when the main routine manipulates such resources as data areas, queues, and data sets that are used by more than one function. The recovery routine in this case should maintain integrity of the resource in case of failure. Recovery routines can also be used to:

- Intercept errors and perform clean-up processing
- Intercept expected program checks and perform the desired action
- Isolate an error to a particular section of processing and continue further processing if possible
- Intercept abends and provide tailored dumps

Deciding What Type of Recovery Routine to Establish

The second decision is what type of recovery routine to establish. If the function holds a lock, is physically disabled, or is an SRB, an FRR can intercept errors. If the function is running under a task and holds a lock during some portion of its processing, an ESTAE can intercept errors, but the lock is freed before the ESTAE routine gets control. An ESTAE-type recovery routine is useful when losing the locked status can be tolerated, such as when the lock is used only to protect a queue from change while it is being read. Also, if the function is running as an enabled unlocked SRB and there is no need to retry at some point in the SRB, an ESTAE routine associated with its related task could be used to intercept errors in the SRB.

If the function attaches any subtasks, it can also provide recovery for the subtask by specifying the ESTAI parameter on the ATTACH macro instruction.

Requesting a Retry

A recovery routine that requests a retry should not assume the registers in the SDWA are its own because there are problems, such as errors in called routines that have no recovery or errors in an asynchronous routine (such as an SRB or IRB) that can affect the register contents in the SDWA. The safest method to ensure a successful retry is for the main routine to save volatile information, such as register contents and addresses, in the parameter area passed to the recovery routine and for the recovery routine to use that information for retry. For example, the issuer of the macro instruction can save the base register and data register for the function in the parameter area. This information enables the recovery routine to reference the areas that belong to the function.

Requesting an ABEND Dump From an FRR

When writing an FRR, note that RTM places the SYSABEND/SYSUDUMP/SYSMDUMP dump options specified on the SETRP macro instruction into the SDWA. Dump options that an FRR specifies replace any dump options that an ABEND macro instruction or a previous recovery routine specified. Also, the CHNGDUMP operator command can add to or override the options. RTM takes one ABEND dump based on the accumulated options. RTM does not take the dump if a retry occurs before the error percolates to an ESTAE-type recovery routine. RTM takes the dump only if all FRRs percolate and no subsequent recovery routine suppresses the dump.

Requesting an ABEND Dump From an ESTAE-Type Recovery Routine

When writing an ESTAE-type recovery routine, note that RTM accumulates the SYSABEND/SYSUDUMP/SYSMDUMP dump options specified by means of the SETRP macro instruction and places them in the SDWA. During percolation, these options are merged with any dump options specified on an ABEND or CALLRTM macro instruction or by other recovery routines. Also, the CHNGDUMP operator command can add to or override the options. RTM takes one dump as specified by the accumulated options. If the recovery routine requests a retry, RTM takes the dump before the retry. If the recovery routine does not request a retry, RTM percolates through all recovery routines before taking the dump.

Requesting SVC Dumps and LOGREC Recording From Recovery Routines

By the time an ABEND dump is taken after percolation, valuable function-related information might have been cleaned up by the recovery routines that have already executed. Also, an ABEND dump is returned to the person who ran the abnormally terminating job. Usually the system programming staff wants to see the dumps requested by authorized programs immediately. A recovery routine that is an authorized program can issue the SDUMP macro instruction to request an SVC dump for system programmers to use. Being an authorized program is only one of the rules for requesting SVC dumps. Other restrictions are described in “Using the SDUMP Macro Instruction” later in this section.

Even when SVC dumps are requested, extra care must be taken to preserve valuable function-related information for routines that need to use the branch entry to SVC dump and that cannot wait by using the ECB option on the SDUMP macro instruction.

- For example, information in control blocks used for communication with other address spaces can be changed before a branch entry SVC dump is taken. The recovery routine should use footprint areas to save this information. The available footprint areas include the function’s work areas, the variable recording area (VRA) in SDWAVRA, the 4K SQA buffer provided for the SDUMP macro instruction, and any areas the SUMLIST or SUMLISTA parameter specified in an SDUMP macro instruction. Documentation describing how footprint areas are used is very helpful when analyzing the dump or LOGREC entries.
- If volatile information is required, the recovery routine should issue an SDUMP macro instruction after saving the volatile information. Volatile information can be saved by specifying SUMLISTA or SUMLIST = *list address*, BRANCH = YES, and SDATA = SUMDUMP in the SDUMP macro instruction. The areas indicated by the SUMLIST parameter are included in the summary dump. For branch entered SVC dumps, the summary dump is taken by default while the system is disabled. Enabled routines can request SUSPEND = YES on branch entry SVC dump requests in order to preserve pageable volatile data. Volatile data can also be saved by moving it into the 4K SQA buffer described in this section under “Using the SDUMP Macro Instruction,” however, this buffer can be used only if no other recovery routine is using it.

Other considerations related to dumping and LOGREC recording are:

- Dumps are usually not required to solve type x37 abends (caused when not enough space is allocated for a data set) and type 913 abends (caused when an operator or user does not supply the correct password). Also, when a prior task is abending and has already taken an ABEND or SVC dump (indicated by the SDWA bits SDWACTS and SDWAMABD being on or the SDWAEAS bit being on), another dump is not necessary. When several recovery routines are written for the same component, the recovery routines that are higher in the hierarchy can test the SDWAEAS bit to check if a lower routine has already taken an SVC dump. (Lower routines should set the SDWAEAS bit to indicate that an SVC dump has been taken.) If other routines have not issued an error message for an abend, the higher recovery routine should issue one.
- The default dump data set is the data set that the user specifies on the SYSABEND, SYSMDUMP, or SYSUDUMP DD statement. If the function executes in key 0, or if it has access to restricted data that should be kept secure in a dump, either the data set specified on the DD statement should be secure or the SDUMP macro instruction should be used to take the dump. The SDUMP macro instruction can include the DCB option for a secure data set.

- When using the SDUMP macro instruction, consider the following:
 - The SDUMP macro instruction should specify a title that summarizes the problem and function. The print dump service aid inserts the first 62 characters of the title on each output page. The title should include at least the name of the module that failed and the name of the recovery module, as in the following example:


```
OC4 ABEND IN OPEN, ERRMOD=IFGORROA,
ISSUER=IGG020FC, JOBN=C49JACIA, STEPN=GO,
SDWAVRA=8417F0
```
 - The SDUMP macro instruction should not specify all SDATA parameters unless all storage areas are necessary to correct the error. If a particular area is not required, omit the corresponding keyword or do not specify the SDATA option, if there is one. Whenever possible, tailor the SDUMP using the storage list options (STORAGE =, LIST =, SUMLSTA =, SUMLIST =, and LISTA =).
- When LOGREC recording is requested (the default for the first FRR if not specified on the SETRP macro instruction), the name of the module that failed and the name of the recovery module should be saved in the appropriate SDWA fields (SDWAMODN, SDWACSCT, and SDWAREXN). Additional functional information should be saved in the SDWACRC, SDWACID, SDWAMLVL, SDWASC, and SDWARRL fields, and in the SDWA variable recording area (SDWAVRA). For a way to map the contents of the VRA, see the VRADATA macro instruction in Volume 2.

FRR Summary

If you decide to use an FRR to provide recovery, review the following:

- The syntax of the SETFRR macro instruction in Volume 2.
- The information about FRRs presented earlier under “System Environment.”
- The register use information for FRRs presented earlier under “Register Interface.”
- The information about the SDWA presented earlier under “System Diagnostic Work Area (SDWA).” Note that the SDWA is fully described in the *Debugging Handbook*.
- If your FRR might request a resume or a retry, the information presented earlier under “Resume” and “Retry from an FRR.”
- The syntax of the SETRP macro instruction in Volume 2.
- The information presented under “Recovery Routine Guidelines.”

ESTAE-Type Recovery Routine Summary

If you decide to use an ESTAE-type recovery routine to provide recovery, review the following:

- If you are using ESTAE, the syntax of the ESTAE macro instruction in Volume 2.
- If you are using FESTAE, the information presented earlier under “Using the FESTAE Macro Instruction” and the syntax of the FESTAE macro instruction in Volume 2.

- If you are using the ESTAI parameter, the syntax of the ATTACH macro instruction in Volume 2.
- The register use information for ESTAE-type recovery routines presented earlier under “Register Interface.”
- The information about ESTAE-type recovery routines presented earlier under “System Environment.”
- The information about the SDWA presented earlier under “System Diagnostic Work Area (SDWA).” Remember that an SDWA might not always be available to your routine. The SDWA is fully described in *Debugging Handbook*.
- If your ESTAE-type recovery routine might request a retry, the information presented earlier under “Retry from an ESTAE-Type Recovery Routine.”
- The syntax of the SETRP macro instruction in Volume 2.
- If your main routine is a task related to an SRB, the information presented earlier under “SRB to Task Percolation.”
- The information presented earlier under “Recovery Routine Guidelines.”

STAE/STAI Exit Routines

The STAE macro instruction causes a recovery routine address to be made known to the control program. This recovery routine is associated with the task and the RB that issued STAE. Use of the STAI option on the ATTACH macro instruction also causes a recovery routine to be made known to the control program, but the routine is associated with the subtask created via ATTACH. Furthermore, STAI recovery routines are propagated to all lower-level subtasks of the subtask created with ATTACH that specified the STAI parameter.

If a task is scheduled for abnormal termination, the exit routine specified by the most recently issued STAE macro instruction gets control and executes under a program request block created by the SYNCH service routine. Only one STAE routine receives control. The STAE routine must specify, by a return code in register 15, whether a retry routine is to be scheduled. If no retry routine is to be scheduled (return code = 0) and this is a subtask with STAI recovery routines, the STAI recovery routine is given control. If there is no STAI recovery routine, abnormal termination continues.

If there is more than one STAI recovery routine existing for a task, the newest one receives control first. If it requests that termination continue (return code = 0), the next STAI routine receives control. This continues until either all STAI routines have received control and requested that the termination continue, a STAI routine requests retry (return code = 4 or 12), or a STAI routine requests that the termination continue but no further STAI routines receive control (return code = 16).

Programs running under a single TCB can issue more than one STAE macro instruction with the create (CT) parameter. Each issuance makes the previous STAE environment temporarily inactive. The environment becomes active when the current STAE environment is canceled.

A STAE environment is canceled when the RB that created it goes away (unless it issues XCTL and specified the XCTL = YES parameter on the STAE macro instruction), when the STAE macro instruction is issued with the CANCEL option, or when the STAE routine receives

control. If a STAE exit routine receives control and requests retry, the retry routine reissues the STAE macro instruction if it wants continued STAE protection.

A STAI environment is canceled if the task completes or if it requests that termination continue and no further STAI processing be done. In the later case, all STAI exits for the task are canceled.

Interface to a STAE/STAI Routine: Prior to entering a STAE/STAI recovery routine, the control program attempts to obtain and initialize a work area that contains information about the error. The first word of the SDWA contains the address of the parameter list specified on the STAE macro instruction or the STAI parameter or the ATTACH macro instruction.

Upon entry to the STAE routine, parameter registers are as follows:

If an SDWA was obtained:

Register	0	a code indicating the type of I/O processing performed:
	0	active I/O has been quiesced and is restorable.
	4	active I/O has been halted and is not restorable.
	8	no active I/O at ABEND time.
	16	active I/O, if any, was allowed to continue.
Register	1	address of the SDWA.
Register	13	save area address.
Register	14	return address.
Register	15	address of STAE exit routine.

If no SDWA was available:

Register	0	a code of 12 to indicate that no SDWA was obtained.
Register	1	ABEND completion code.
Register	2	address of user-supplied parameter list.
Register	13	unpredictable.
Register	14	return address.
Register	15	address of STAE exit routine.

When the STAE or STAI routine has completed, it should return to RTM via the contents of register 14. Register 15 should contain one of the following return codes:

Return Code	Action
0	Continue the termination. The next STAI, ESTAI, or ESTAE routine will be given control. No other STAE routines will receive control.
4,8,12	A retry routine is to be scheduled.
16	No further STAI/ESTAI processing is to occur. This code may only be issued by a STAI/ESTAI routine

For the following situations, STAE/STAI routines are not entered:

- If the abnormal termination is caused by an operator's CANCEL, job step timer expiration, or the detaching of an incomplete task without the STAE = YES option.
- If the failing task has been in a wait state for more than 30 minutes.
- If the STAE macro instruction was issued by a subtask and the attaching task abnormally terminates.
- If the recovery routine was specified for a subtask, via the STAI parameter of the ATTACH macro instruction, and the attaching task abnormally terminates.

- If a problem other than those above arises while RTM is preparing to give control to the STAE routine.
- If another task in the jobstep terminates without the step option.

STAE/STAI Retry Routines: If the STAE retry routine is scheduled, the system automatically cancels the active STAE environment; the preceding STAE environment, if one exists, then becomes the active one. Users wanting to maintain STAE protection during retry must reestablish an active STAE environment within the retry routine, or must issue multiple STAE requests prior to the time that the retry routine gains control.

Like the STAE/STAI exit routine, the STAE/STAI retry routine must be in storage when the exit routine determines that retry is to be attempted. If not already resident in your program, the retry routine may be brought into storage via the LOAD macro instruction by either the main program or exit routine.

If the STAE/STAI routine indicates that a retry routine has been provided (return code = 4, 8, or 12), register 0 must contain the address of the retry routine. The STAE environment that requested retry is canceled and the request block queue is purged up to, but not including, the RB of the program that issued the STAE macro instruction. This is done by pointing each RB old PSW to an SVC 3 (EXIT) instruction. In addition, open DCBs that can be associated with the purged RBs are closed and queued I/O requests associated with the DCBs being closed are deleted from the I/O restore chain.

The RB purge is an attempt to cancel the effects of partially executed programs that are at a lower level in the program hierarchy than the program under which the retry occurs. However, certain effects on the system are not canceled by this RB purge. Generally, these effects are TCB-related and are not identifiable at the RB level. Examples of these effects are as follows:

- Subtasks created by a program to be purged. Reason: subtasks cannot be associated with an RB; the structure is defined via TCBs.
- Resources allocated by the ENQ macro instruction. Reason: ENQ resources are associated with the TCB and are not identifiable at the RB level.
- DCBs that exist in dynamically acquired virtual storage. Reason: Only DCBs in the program, as defined by the RB via the CDE itself, are closed.

If there are quiesced restorable input/output operations, they can be restored, in the STAE retry routine, by using word 2 in the SDWA. Word 2 contains the pointer to the purged I/O request list (PIRL) passed as a parameter to SVC Restore. SVC Restore is used to have the system restore all I/O requests on the PIRL. (For additional information on SVC Restore, see *SPL: Data Management*.)

If an SDWA was obtained upon entry to the STAE/STAI retry routine, register contents are as follows:

Register	0	0
Register	1	Address of the SDWA.
Register	2-13	Unpredictable.
Register	14	Address of an SVC 3 EXIT instruction.
Register	15	Address of the STAE/STAI retry routine.

When the storage is no longer needed, the retry routine should use the FREEMAIN macro instruction to free the first 104 bytes of the work area. If the retry routine is in the user key, this storage should be freed from subpool 0 which is the default subpool for the FREEMAIN

macro instruction. If the retry routine is in the control program key, storage must be freed from subpool 250. The remainder of the work area's storage was freed by RTM during STAE/STAI processing.

If the ABEND/STAE interface routine was not able to obtain storage for the work area, register contents are as follows:

Register	0	12
Register	1	ABEND completion code.
Register	2	Address of PIRL or 0 if I/O is not restorable.
Register	14	Address of an SVC EXIT instruction
Register	15	Address of the STAE/STAI retry routine

The retry routine is entered in supervisor state if the RBOPSW of the retry RB is in supervisor state and the task was authorized at the time the STAE routine was established or at the time of the error. Otherwise, the retry routine is entered in problem program state.

The task is considered to be authorized at the time the STAE routine is established when at least one of the following is true:

1. The task is APF-authorized.
2. The requestor is in supervisor state.
3. The requestor has a PSW key less than 8.
4. The TCBPKF of the task is less than 8.
5. The PKM of the requestor allows keys less than 8.

The main routine is considered to be authorized at the time of the error when at least one of the following is true:

1. The task is APF-authorized.
2. The TCBPKF of the task in error is less than 8.
3. The bit TCBFSM of the task in error (in TCBFLGS3 of the TCB) is set to one, indicating that all RBs for that TCB execute in supervisor state.

The retry routine is entered with the same PSW key as the one in RBOPSW of the retry RB when one of the following is true:

1. The task was authorized at the time of the error as described above.
2. The RBOPSW of the retry RB has a key greater than or equal to 8 and is in problem program state, and the PKM of that RB (XSBKM field in the XSB control block for that RB) does not have authority to keys less than 8.

Otherwise, the PSW key of the retry routine is that of the TCBPKF.

Uses of Resource Managers

A resource manager routine gets control during normal and abnormal termination of a task or an address space. Task or address space termination is the process of removing a task or address space from the system, releasing the resources from the task or address space, and making the resources available for reuse. MVS/XA provides resource managers that are invoked to establish routines that “clean up” the queues and control blocks associated with the resources.

In general, MVS/XA does not provide a resource manager for a function that is totally self-contained. A self-contained function does not allocate any resources to the requestor or reclaims all allocated resources before the function returns control. Examples of self-contained functions are EXTRACT and TTIMER. MVS/XA does provide a resource manager for functions that are not self-contained -- that do allocate resources to the requestor. Such services fall into two categories: paired requests (such as OPEN/CLOSE and ENQ/DEQ) and requests that invoke an asynchronous process (such as STIMER and I/O requests).

The system-provided resource managers can, for example, perform clean-up work for a subsystem. The responsibilities of a resource manager include:

- For task termination, removing all traces of the fact that the TCB for the terminating task at one time was connected to, allocated to, or associated with the resource in question. The resource should be left in such a state that it can be reused by another task in the address space or in the system.
- For address space termination, releasing all system queue area and common storage area control blocks obtained for the use of the terminating address space. Also, any buffers, bit settings, pointers, and so on relating to the terminating address space should be reset to make the system appear as if the ASID or ASCB of the terminating address space never existed.

The resource manager is also responsible for establishing a recovery environment when first entered to protect itself against errors during its own processing. For SRBs, the clean-up routine issues the PURGEDQ macro instruction to ensure that all undispatched SRBs are removed from the SRB dispatching queue.

If an installation creates a function that is not totally self-contained, the installation should also provide a resource manager for that function. The installation-created resource manager should perform the same basic tasks for this new function that system resource managers perform for system functions. See *SPL: System Modifications* for information on how to write an installation resource manager.

Protecting the System

Protecting the system or maintaining system integrity is a major consideration in large systems. This chapter includes information concerning the following topics:

- System integrity
- Using the authorized program facility (APF)
- Using the resource access control facility (RACF)
- Protecting low storage

System Integrity

System integrity is defined as the ability of the system to protect itself against unauthorized user access to the extent that security controls cannot be compromised. That is, there is no way for an unauthorized program using any system interface to bypass store or fetch protection, bypass password checking, bypass RACF checking, or obtain control in an authorized state.

Note: An authorized program in MVS/XA is one that executes in a system key (keys 0-7), in supervisor state, or is authorized via the authorized program facility (APF).

Documentation on System Integrity

This section contains information about MVS/XA system integrity. The related topic of security in regard to the physical environment of a computing system is discussed in the following publications:

- *The Considerations of Physical Security in a Computer Environment*
- *Data Security Controls and Procedures--A Philosophy for DP Installations*
- *Security Assessment Questionnaire*

Installation Responsibility

To ensure that system integrity is effective and to avoid compromising any of the integrity controls provided in the system, the installation must assume responsibility for the following:

- Physical environment of the computing system.
- Adoption of certain procedures (for example, the password protection of appropriate system data sets) that are a necessary complement to the integrity support within the operating system itself.

- That its own modifications and additions to the control program do not introduce any integrity exposures. That is, all installation-written authorized code (for example, an installation SVC) must perform the same or an equivalent type of validity checking and control that the MVS/XA control program employs to maintain system integrity.

Elimination of Potential Integrity Exposures

MVS/XA system integrity support restricts only unauthorized problem programs. It is the responsibility of the installation to verify that any authorized programs added to the system control program will not introduce any integrity exposures. To do this effectively, an installation should consider these areas for potential integrity exposure:

- User-supplied addresses for user storage areas.
- User-supplied addresses for protected control blocks.
- Resource identification.
- SVC routines calling SVC routines.
- Control program and user data accessibility.
- Resource serialization. (See the section "Locking.")

Each of the following descriptions is a guideline to aid the installation in:

- Eliminating that area as a potential integrity exposure.
- Determining whether an impact on existing installation-written code might occur, especially where that code is dependent on the use of non-standard interfaces to the system control program.

There should be no impact on installation-written routines that use standard interfaces (problem program/system interface described in an SRL) because no standard interfaces for system integrity support have been removed from the MVS/XA system control program. However, some routines now require authorization for use.

User-Supplied Addresses for User Storage Areas

A potential integrity exposure exists whenever a routine having a system protection key (key 0-7) accepts a user-supplied address of an area to which a store or fetch is to be done. If the system routine does not adequately validate the user-supplied address to ensure that it is the address of an area accessible to the user for storing and fetching data, an integrity violation can occur when the system-key routine:

- Stores into (overlays) system code or data (for example, in the nucleus or the system queue area), or into another user's code or data.
- Moves data from a fetch-protected area that is not accessible to the user (for example, fetch-protected portion of the common service areas) to an area that is accessible to the user.

To eliminate this problem system-key routines should always verify that the entire area to be stored into, or fetched from, is accessible (for storing or fetching) to the user in question. The primary validation technique is the generally established MVS/XA convention that system-key routines obtain the protection key of the user before accessing the user-specified area of storage. For example, MVS/XA data management SVC routines (which generally execute in key 5) assume the user's key before modifying a data control block (DCB) or an I/O block (IOB).

User-Supplied Addresses for Protected Control Blocks

A potential integrity exposure exists whenever the control program (system key/privileged mode) accepts the address of a protected system control block from the user. For most system control blocks, this situation should not be permitted to exist. However, in certain cases it is necessary to allow the user to provide the address of a system control block that describes his allocation/access to a particular resource (for example, a data set), in order to identify that resource from a group of similar resources (for example, a user might have many data sets allocated). Inadequate validity checking in this situation can create an integrity exposure, because an unauthorized problem program could provide its own (counterfeit) control block in place of the system block and thereby gain the ability to:

- Access a resource in an uncontrolled manner (because the control block in this case would normally define the restrictions, such as read-only for a data set, on the user's allocation to the resource).
- Gain control in a privileged state (because such control blocks might contain the addresses of routines that run in privileged mode or with a system (0-7) key).
- Cause various other problems depending on exactly what data is in the control block involved.

To avoid this type of exposure, the control program must verify, for every such address accepted from a problem program, that the address is that of:

1. A protected control block created by the control program.
2. The correct type of control program block (for example, a TCB versus a DEB, or a QSAM DEB versus an ISAM DEB).
3. A control block created for use in connection with the user (job step) that supplied the address.

In MVS/XA, verification is generally accomplished by establishing a chain or table of the particular type of control block to be validated. This chain or table is located via a protected and jobstep-related control block that is known to be valid. Addresses that are not allowed to be supplied by the user, are located via a chain of protected control blocks that begins with a control block known to be valid or fixed at a known location at IPL time, such as the CVT. Therefore, a control block can only be entered in the chain/table by:

- An authorized program satisfying point 1.
- Definition, where the chain/table establishes the type of control block satisfying point 2.
- Definition, where each chain/table is located only through a jobstep-related control block satisfying point 3.

Note: This does not imply that a system routine must go back to the CVT or similar control block every time it wants to establish a valid chain. Typically, a control block address not too far down on such a chain is available and already validated in a register. For example, the first load of an SVC can receive control with a valid TCB address in a register.

Resource Identification

Resource identification is another area that can be subject to integrity exposures. Exposures can result if the control program does not maintain and use sufficient data to uniquely distinguish one resource from other similar resources. For example, a program must be identified by both name and library to distinguish it from other programs. The consequences of inadequate resource identification are problems such as the ability of an unauthorized program to create counterfeit control program code or data, or to cause varying types of integrity problems by intermixing incompatible pieces of control program code and/or data.

The general solution can only be stated as the reverse of the problem; that is, the control program must maintain and use sufficient (protected) data on any control program resource to distinguish between that resource and other control program or user resources. The following are examples of the controls that MVS/XA employs to comply with the requirement:

- In general, authorized program requests to load other authorized programs are satisfied only from authorized system libraries (see the topic “Control Program Extensions” described in this section.)
- MVS/XA takes explicit steps to ensure that routines loaded from authorized system libraries are used only for their intended purpose. This includes expanded validity checking to remove any potential for the unauthorized program to specify explicitly which of the authorized library routines are to gain control in any given situation.
- Sensitive system control blocks are validated as being the “correct” blocks to be used in any given control program operation. (See the topic “User-Supplied Addresses of Protected Control Blocks” described earlier in this section.)

SVC Routines Calling SVC Routines

A potential problem area exists whenever a problem program is allowed to use one SVC routine (routine A) to invoke a second SVC routine (routine B) that the problem program could have invoked directly. An integrity exposure occurs if:

- SVC routine B bypasses some or all validity checking based on the fact that it was called by SVC routine A (an authorized program) or
- User-supplied data passed to routine B by routine A either is not validity checked by routine A, or is exposed to user modification after it was validated by routine A.

These problems will not exist if the user calls SVC routine B directly, because the validity checking will be performed on the basis of the caller being an unauthorized program.

SVC routine A, which is aware that it has been called by an unauthorized program, must ensure that the proper validity checking is accomplished. However, it is usually not practical for SVC routine A to do the validity checking itself, because of the potential for user modification of the data before or during its use by SVC routine B. The general solution should be for SVC routine A to provide an interface to SVC routine B, informing routine B that the operation is being requested with user-supplied data in behalf of an unauthorized problem program (implying that normal validity checking should be performed).

In practice, in MVS/XA, most SVC B-type routines that could be subject to this problem use the key of their caller as a basis for determining whether or not to perform validity checking. Therefore, most SVC A-type MVS/XA routines have simply adopted the convention of assuming the key of their caller before calling the SVC B routine. (For additional information see the section "Writing SVC Routines" later in this book.)

Control Program and User Data Accessibility

Important in maintaining system integrity is the consideration of what system data is sensitive and must be protected from the user, and what data can be exposed to user manipulation. The implications of the exposure of the wrong type of data are obvious.

In general, it is necessary to store protect the following types of data:

- Code, and the location of code, that is to receive control in an authorized state.
- Work areas for such code, including areas where it saves the contents of registers.
- Control blocks that represent the allocation or use of system resources.

MVS/XA maintains such items in system storage, or in a separate address space in the case of some APF-authorized programs.

It might also be necessary to protect, for a limited period, certain data that is normally under the control of the user (for example, to prevent its modification during a critical operation). In this case MVS/XA provides fetch protection for such data if:

- The data consists of proprietary information (such as passwords).
- The control program cannot determine the nature of the contents of the data area.

Fetch Protection Provided for the PSA

The last 2K locations of the PSA (addresses 2048 through 4095) contain sensitive system data that must be protected. These locations are key 0 fetch protected. This means that only key 0 programs can fetch data from the last 2k of the PSA. Also the entire PSA of one CPU is key 0 fetch-protected from programs attempting to access the PSA while executing on another CPU.

Control Program Extensions

This potential problem area involves the somewhat hazy distinction that exists between the control program and certain types of problem programs. In most installations, there are problem state/user key (keys 8-15) programs that are actually extensions to the control program in that they are allowed (by means of various special SVCs, and so forth) to bypass normal system controls over access to system resources. For example, a special utility program that scans all the data on a pack might be able to avoid the normal system extent checking on a direct access volume.

If an installation has its own control program extensions and SVCs that allow the bypass of normal system security or integrity checks (for example, an SVC that returns control in key 0), and if such SVCs are not currently restricted from use by an unauthorized program, the APF facility should be used to restrict them and to authorize the control program extensions that use them.

Using the Authorized Program Facility (APF)

The authorized program facility (APF) is a facility that an installation manager uses to protect the system. In MVS, certain system functions, such as all or part of some SVCs, are sensitive; their use must be restricted to users who are authorized. An authorized program is one that executes in supervisor state, with a PSW key of 0-7, or with APF authorization. In addition, an authorized program should be thoroughly tested so that its use does not compromise the system.

The MVS/XA supervisor uses APF to protect the system as follows:

1. The supervisor limits the use of sensitive system SVC routines and optionally, sensitive user SVC routines, to authorized programs by issuing a TESTAUTH macro instruction before giving control to the SVC routine. TESTAUTH determines, among other things, the authorization status of the calling program. The installation can then use APF to prohibit unauthorized programs from using sensitive SVC routines.
2. The supervisor ensures that all modules in an authorized job step task are fetched only from authorized libraries. The supervisor thus prevents the unauthorized counterfeiting of any module in an authorized job step task's module flow.

APF Authorization

APF is a mechanism that allows a program to become authorized. APF authorization is established at the job step task level and depends on the authorization status of the first module of the job step task at the time the job step task is initiated. Only when the first module loaded meets both of the following conditions is the job step task marked APF-authorized:

- The module comes from an authorized library.
- The module was link-edited with the authorization code AC=1. This code is contained in a bit setting in the partitioned data set (PDS) directory entry for the module.

The authorization code is meaningful only when the load module resides in an authorized library and is executed as the first module of a job step task. Thus, even though a program is link-edited with AC=1, it can run as an authorized program only if it is loaded from an authorized library. When this occurs, the program manager then verifies that all subsequent modules for that program also come from authorized libraries; if they do not, a 306 abend results. Authorized libraries are marked at OPEN time by having a bit turned on in their DEBs.

The names of programs that are APF-authorized in the systems IBM provides can vary from one release to another. To determine which programs are APF-authorized in the current system, list the PDS directories of the authorized libraries (see "Authorized Libraries" later in this discussion) and check for modules that are marked authorized (AC=1). Only modules marked authorized and any modules they invoke from authorized libraries can ever run as authorized programs.

Authorized Programs

MVS/XA considers a program authorized when that program executes in any one of the following states:

- Supervisor state (bit 15 of the PSW is zero).
- A system key (bits 8-11 of the PSW are in the range 0-7).
- As part of an APF-authorized job step task (bit JSCBAUTH in the JSCB is 1).

However, MVS/XA sometimes distinguishes authorization between a program that runs either in supervisor state or system key and a program that is APF-authorized. For example, the use of certain keywords in some macro instructions is restricted to programs running in supervisor state or system key; programs that are APF-authorized but not in supervisor state or system key cannot use these keywords.

When MVS/XA attaches the first load of a job step task (identified via the JSTCB keyword in the ATTACH macro instruction), program management decides whether to mark the task authorized or to leave it unauthorized. If the first load module has AC=1 and comes from an authorized library, the program is considered APF-authorized, and the task is marked APF-authorized. If these conditions are not met, the task is not marked APF-authorized and cannot normally become so during the life of the job step. (See "Authorization Results Under Various Conditions" for exceptions to this rule.)

You should recognize the distinction between an authorized program and an authorized user. To use restricted functions in MVS/XA a program must be authorized; that is, in supervisor state, running with a system key, or part of an APF-authorized task or some combination of these that satisfies the restriction. Any user, however, can submit a job that executes an authorized program. To restrict a program to an individual user or a class of users, you can use existing data set security facilities to place the program in a library, other than LINKLIB, SVCLIB, or LPALIB, protected by RACF or by a password. If the program is an authorized program, the library must be an authorized library.

Note: An authorized program could also be restricted by defining it as a RACF resource and using the RACHECK macro instruction in the program to verify the user's authorization.

Authorized Libraries

APF-authorized programs must reside in authorized libraries. The authorized libraries in MVS/XA are:

- SYS1.LINKLIB
- SYS1.SVCLIB
- SYS1.LPALIB (only during an IPL; see note 2)
- Installation authorized libraries

To allow an installation to authorize libraries, MVS/XA provides member IEAAPF00 and supports optional members IEAAPFxx in SYS1.PARMLIB. During IPL, the system uses the contents of one member, specified at IPL, to build the APFTABLE, which contains the names of authorized libraries and the serial numbers of the volumes on which they reside. SYS1.LINKLIB and SYS1.SVCLIB are automatically placed in the first two entries of the APFTABLE; the remaining entries contain the names from member IEAAPFxx (where xx is the identifier of the member). The volume serial numbers in the entries prevent the system from obtaining data from a non-authorized library having the same name as an authorized library but residing on a different volume.

The LNKSTxx members of SYS1.PARMLIB denote libraries to be concatenated to SYS1.LINKLIB. If libraries in the LNKST concatenation are accessed through either JOBLIB or STEPLIB DD statements, MVS/XA does not consider them authorized unless the installation has also placed the names of the libraries in IEAAPFxx. As long as a load module is in an authorized library, an authorized program can load it. The installation is responsible for ensuring that duplicate module names are not permitted across authorized libraries to preclude access to an incorrect module resulting in possible integrity exposures.

For details concerning SYS1.PARMLIB members IEAAPFxx and LNKSTxx, see *Initialization and Tuning*.

Notes:

1. *If a JCL DD statement concatenates an authorized library in any order with an unauthorized library, the entire set of concatenated libraries is treated as unauthorized.*
2. *After IPL, SYS1.LPALIB is not an authorized library. SYS1.LPALIB is authorized only during NIP processing when the system builds the pageable link pack area (PLPA). All modules in PLPA are marked as coming from an authorized library. SYS1.LPALIB becomes an authorized library after IPL only when an installation places its name in the IEAAPFxx member of SYS1.PARMLIB, and there is no reason to do so.*

Mixing APF and Non-APF Libraries in LNKST

You can include the LNKAUTH parameter in the system parameter list, IEASYSxx, to indicate whether all data sets in the LNKST concatenation are to be treated as APF authorized (LNKAUTH=LNKST) or whether only those that are named in the APFTABLE are to be treated as APF authorized (LNKAUTH=APFTAB). This means that it is possible to mix APF authorized libraries and non-APF authorized libraries in the LNKST concatenation. The LNKAUTH parameter is also described in *Initialization and Tuning*.

Using APF

APF allows an MVS/XA installation to restrict use of SVC routines to authorized programs and to restrict access to load modules; that is, APF prevents authorized programs from accessing any load module that is not in an authorized library.

Restricting SVC Routines (TESTAUTH)

MVS/XA provides two methods to restrict the use of sensitive SVC routines:

1. You can specify the FC01 parameter on the SVCTABLE macro instruction during system generation to restrict sensitive SVCs to authorized callers. The SVC first-level interrupt handler (FLIH) ensures that only authorized programs can access routines so restricted. If an unauthorized program tries to access a restricted SVC, an 047 abend results. The TESTAUTH macro instruction in the SVC FLIH specifies testing for supervisor state, system key, or APF; any one of these authorizations satisfies the test and allows the caller to proceed.
2. You can use the TESTAUTH macro instruction to restrict an entire SVC or particular paths through an SVC when only a portion of the SVC's function is sensitive. Though you can specify any combination of system key, supervisor state, and APF authorization for TESTAUTH to test, you should specify only those conditions that you consider essential. If any of the conditions specified to TESTAUTH are present, TESTAUTH returns an

indication that the caller is authorized. For example, to validate the authorization status of programs requesting restricted functions, various system routines use TESTAUTH to make the following distinctions:

- a. The caller is executing in supervisor state, system key, or both.
- b. The caller is an APF-authorized task (the JSCBAUTH bit in the JSCB is on).
- c. Either a or b.

The TESTAUTH macro instruction, inserted at appropriate locations in an SVC routine, returns an authorized or unauthorized indication; the SVC routine can then take appropriate action based upon this return.

The TESTAUTH macro is not used to control the use of I/O appendages. I/O appendages are controlled by means of the IEAAPP00 member of SYS1.PARMLIB. (See the description of this member in *Initialization and Tuning*.)

Restricting Load Module Access

To authorize a program, the installation must first assign the authorization code to the first load module of the program. APF prevents authorized programs from accessing any load module that is not in an authorized library. If an authorized program tries to access a module that is not in an authorized library, a search is done to find a copy of the module in an authorized library. If a copy is found, then processing continues with that copy of the module. If there is not a copy of the module in any authorized library, a 306 abend results.

Assigning Authorization

An installation can assign load modules the APF-authorization code either through the PARM field on the link edit step or through a linkage editor control statement. The authorization code of a load module has meaning only when it resides on an APF-authorized library and when it is executed as the first program of a job step attach.

To assign an authorization code via JCL, code AC=1 in the operand field of the PARM parameter of the EXEC statement as follows:

```
//LKED EXEC PGM=HEWL,PARM='AC=1',...
```

If no authorization code is assigned in the linkage editor step, the default is non-authorization. The authorization code for a given output load module can be overridden with the SETCODE control statement.

The SETCODE statement establishes authorization for a specific output load module. If it is used, you must place it before the NAME statement for the load module. The format of the SETCODE statement is:

```
SETCODE AC(1)
```

If more than one SETCODE statement is assigned to a given output load module, the last statement found is used.

In the example in Figure 57, the SETCODE statement assigns an authorization code to the output load module MOD1.

```

//LKED      EXEC  PGM=HEWL
//SYSPRINT  DD    SYSOUT=A
//SYSUT1   DD    UNIT=SYSDA,SPACE=(TRK,(10,5))
//SYSLMOD  DD    DSNNAME=SYS1.LINKLIB,DISP=OLD
//SYSLIN   DD    DSNNAME=&&LOADSET,DISP=(OLD,PASS),
//          UNIT=SYSDA
//          DD    *
//          SETCODE AC(1)
//          NAME   MOD1(R)
/*

```

Figure 57. Assigning Authorization via SETCODE

No security or integrity exposure exists if a program is link-edited into an unauthorized library with authorization code AC=1. The job step task is not authorized when the first module of the job step task is loaded and no abend occurs. However, if the loaded module tries to execute functions or SVCs that require authorization, the program is abended.

Authorization Results Under Various Conditions

When a program issues an SVC or accesses a load module through a LINK, LOAD, or XCTL macro instruction, authorization is straight-forward; the only factors considered are whether the calling program is authorized and whether the called program is a restricted SVC or a load module in an authorized library. Figure 58 summarizes the authorization rules.

Rule	Abend Resulting From Violation
1. An unauthorized routine cannot call a restricted SVC.	047
2. A routine running in supervisor state, system key, or APF-authorized cannot call programs residing outside APF-authorized libraries.	306

Figure 58. Authorization Rules

The rules shown in Figure 58 are also true when the ATTACH macro instruction is used unless the RSAPF keyword is specified. An attaching task that specifies RSAPF= YES and is running in supervisor state or PSW key 0-7 can attach programs residing outside APF-authorized libraries if the following conditions are met:

- The caller is not running APF-authorized or the caller turns authorization (JSCBAUTH) off until program fetch obtains the subtask.
- The caller is attaching a subtask in problem state.
- The attached task's TCB key is 8-15 (non-system key).

The newly attached subtask does not run APF-authorized. If the attaching task is not in supervisor state or PSW key 0-7, the default, RSAPF= NO, is taken and a 306 abend might result.

However, if the subtask comes from an APF-authorized library and is link edited with the APF-authorized attribute, then the task executes with APF authorization.

Another factor to be considered when using ATTACH is the JSTCB keyword. A routine running in supervisor state or system key can attach a task that is allowed to become APF-authorized if the routine specifies JSTCB = YES.

Guidelines for Using APF

Installations using APF authorization must control which programs are stored in authorized libraries. If the first module in a program sequence is authorized, the system assumes that the flow of control to all subsequent modules is known and secure as long as these subsequent modules come from authorized libraries. To ensure that this assumption is valid, the installation should:

- Ensure that all programs that will run as authorized programs adhere to the installation's integrity guidelines.
- Protect authorized libraries through RACF or passwords to ensure that only selected users can store programs in these libraries.
- Ensure that no two load modules with the same name exist across the set of authorized libraries. Two modules with the same name could lead to accidental or deliberate mix-up in module flow, possibly introducing an integrity exposure.
- Link edit with the authorization code (AC = 1) only the first load module in a program sequence. Do not use the authorization code for subsequent load modules, thus ensuring that a user cannot call modules out of sequence, become APF-authorized, and thus possibly bypass validity checking or critical logic flow.
- Ensure that IEAAPFxx does not contain the names and volume serial numbers of data sets that no longer exist. If it does, a user could assign his own data sets with the same names on the same volumes and cause his own libraries to become authorized.

Resource Access Control Facility (RACF)

The Resource Access Control Facility (RACF) provides software access control measures that can be used to enhance data security in a computing system. RACF can be used in addition to any data security measure currently being used.

RACF provides the ability to specify access authorities under which the permanent DASD data sets, tape volumes, DASD volumes, terminals, and other resources are made available to the users of the system. RACF can protect VSAM, non-VSAM, cataloged, and uncataloged data sets can be protected by RACF.

When users, groups, DASD data sets, tape volumes, DASD volumes, terminals, and other resources are defined to RACF, RACF builds and stores their descriptions in profiles on the RACF data set. RACF uses these profiles for RACHECK authorization checking and RACINIT user identification and verification.

For more information on RACF, see *Resource Access Control Facility (RACF) - General Information Manual*.

Defining a Resource to RACF (RACDEF)

The RACDEF macro instruction can be used to define, modify, and delete resource profile (for example, a tape volume profile and a DASD data set profile) for RACF.

The resource manager responsible for establishing and maintaining the resources issues the RACDEF macro instruction to define or delete the resource profile.

Identifying a RACF-Defined User (RACINIT)

The RACINIT macro instruction can be used to determine if a userid is defined to RACF and if the user has supplied a valid password, group name, and operator identification. RACF builds an access environment element for the user if the userid, password, group name, and terminal id (for the terminal user) are accepted. The identification and verification in the case of a terminal or batch job user, is based on the information contained in the TSO LOGON or IMS /SIGN command or data specified in the JOB statement for the batch job. The access environment element identifies the scope of the user's authorization to be used during the current terminal session or batch job.

Checking RACF Authorization (RACHECK)

RACHECK processing determines if a user is authorized to obtain use of a resource (for example, DASD data set, tape volume, or DASD volume) protected by RACF. When a user requests access to a RACF-protected resource, acceptance of the request is based upon the identity of the user and whether the user has been permitted sufficient access authority to the resource.

RACF performs system authorization checking when a resource manager that controls a RACF-protected resource issues the RACHECK macro instruction before allowing a user access to the resource.

The system programmer using this macro instruction to check a user's authorization to a resource has available three parameters (CSA, LOG, and PROFILE) that are not available to the application programmer. These parameters permit the system programmer to specify that a profile is to be copied and maintained in main storage for the resource and that different types of access attempts are or are not to be recorded on the SMF data set.

Retrieving and Encrypting Data (RACXTRT)

The RACXTRT macro instruction can be used for either of two purposes. It can be used to retrieve certain specified fields from a RACF user profile or it can be used to encrypt certain clear-text (readable) data.

Building In-Storage Profiles (RACLIST)

The RACLIST macro instruction can be used to build in-storage profiles from RACF defined class resources. RACLIST processes only the resources described by class descriptors. Once profiles are brought into main storage by RACLIST, FRACHECK and RACHECK macros can be issued for the resources without requiring access to the RACF data set.

System Authorization Facility (SAF)

The System Authorization Facility (SAF) provides a system interface that conditionally directs control to the Resource Access Control Facility (RACF), if RACF is present, and/or a user-supplied processing routine when receiving a request from a resource manager. SAF does not require any other program product as a prerequisite, but overall system security functions are greatly enhanced and complemented by the concurrent use of RACF. The key element in SAF is the MVS router.

MVS Router

SAF provides an installation with centralized control over system security processing by using a system service called the MVS router. The MVS router provides a focal point and a common system interface for all products providing resource control. The resource managing components and subsystems call the MVS router as part of certain decision-making functions in their processing, such as access control checking and authorization-related checking. These functions are called "control points." This single SAF interface encourages the use of common control functions shared across products and across systems.

The router is always present whether or not RACF is present. If RACF is available in the system, the router passes control to the RACF routine (ICHRFR00) that invokes the appropriate RACF function based on the parameter information and the RACF router table (ICHRFR01), which associates router invocations with RACF functions. The RACF router table is described in *SPL: Resource Access Control Facility (RACF)*. Before it calls the RACF routine, the router calls an optional, user-supplied security processing exit if one has been installed. (See the following topic "MVS Router Exit.")

Control points that issue the RACROUTE macro instruction enter the MVS router in the same key and state as the RACROUTE issuer. Control points that continue to issue the RACF macro instructions go directly to RACF, bypassing the router.

For use on an MVS/XA system, the MVS router exit must be link-edited with AMODE(ANY) and RMODE(24.)

MVS Router Exit

The MVS router provides an optional installation exit that is invoked whether or not RACF is installed and active on the system. If RACF is not available, the router exit acts as an installation written security processing (or routing) routine. If RACF is available, the exit acts as a RACF preprocessing exit. The installation exit should have an AMODE of ANY and an RMODE of 24.

The only way to invoke the MVS router installation exit is by issuing the RACROUTE macro. The exit is entered via a branch and link macro and thus will execute in the same key and state as the issuer of the RACROUTE macro. The exit must be named ICHRTX00 and must be located in the link pack area (LPA). The router passes the parameter list to the installation exit. In addition, the exit receives the address of a 150-byte work area.

Control points that continue to use the RACF macro instructions do not invoke the installation exit.

On entry to the MVS router exit routine, register 1 contains the address of the following area:

Offset	Length	Description
0	4	Parameter list address - points to the MVS router parameter list
4	4	Work area address - points to a 150-byte work area that the exit can use

Return Codes

MVS Router Exit Codes

The exit routine returns one of the following return codes in register 15:

Hex (Decimal)	Meaning
0 (0)	The exit has completed successfully. Control proceeds to the RACF front end routine for further security processing and an invocation of RACF.
C8 (200)	The exit has completed successfully. The MVS router translates this return code to a router return code of 0 and returns control to the issuer of the RACROUTE macro, bypassing RACF processing. (See the note below.)
CC (204)	The exit has completed successfully. The MVS router translates this return code to a router return code of 4 and returns control to the issuer of the RACROUTE macro, bypassing RACF processing. (See the note below.)
D0 (208)	The exit has completed processing. The MVS router translates this return code to a router return code of 8 and returns control to the issuer of the RACROUTE macro, bypassing RACF processing. (See the note below.)
Other	If the exit routine sets any return code other than those described above, the MVS router returns control directly to the issuer of the RACROUTE macro and passes the untranslated code as the router return code. The exit routine should place the return and reason code information in the parameter list. Further RACF processing is bypassed.

Note: The installation exit routine is responsible for putting RACF compatible return and reason codes in the first two full words of the parameter list. If the exit routine does not issue a specific reason code, it should issue a zero reason code.

Simulating a Call to RACF

Normally, a caller, such as DFP, IMS, or JES invokes the MVS router and passes it class, requestor, and subsystem parameters via the RACROUTE exit parameter list. Using those parameters, the MVS router calls the router exit, which then returns to the router with a return code. If the return code is 0, as defined above, the router invokes RACF. RACF reports the results of that invocation to the router by entering return and reason codes in registers 15 and 0 respectively. The router converts the RACF return and reason codes to router return and reason codes and passes them to the caller. The router provides additional information to the caller by placing the unconverted RACF return and reason codes in the first and second words (respectively) of the ROUTER input parameter list.

Instead of invoking RACF processing, your installation may choose to have the MVS router exit respond to the caller's request. If that is the case, you must still provide the caller with the RACF return and reason codes that it expects to receive. To do so, you must set the router exit return code, as defined above, so that RACF is not invoked. However, you must still simulate the results of a RACF invocation by coding the exit so it places the RACF return and reason codes in the first and second fullwords (respectively) of the RACROUTE input parameter list.

MVS Router Parameter List

The MVS router parameter list (mapped by macro ICHSAFP) is generated when the RACROUTE macro is issued and describes the security processing request by providing the request type. If the router installation exit exists, the router passes the parameter list to this exit. If RACF is active, the router uses the request type information to invoke the appropriate RACF function.

Field Name	Offset	Length	Description
SAFPRRET	0 (0)	4	Return code - Defines the RACF or installation exit return code.
SAFPRREA	4 (4)	4	Reason code - Defines the RACF or installation exit reason code.
SAFPPLN	8 (8)	2	Length - Defines the length of the SAFP parameter list.
	10 (A)	2	Reserved.
SAFPREQT	12 (C)	2	Request type - A binary halfword corresponding to the request type on the RACROUTE macro. The request type and the associated request numbers are listed below. AUTH (RACHECK) - 1 (01) FASTAUTH (FRACHECK) - 2 (02) LIST (RACLIST) - 3 (03) DEFINE (RACDEF) - 4 (04) VERIFY (RACINIT) - 5 (05)
	14 (E)	2	Reserved.
SAFPREQR	16 (10)	4	Request name address - Points to an 8-byte character field containing the control point name.
SAFPSUBS	20 (14)	4	Subsystem name address - Points to an 8-byte character field containing the calling subsystem's name, version, and release level.
SAFPWA	24 (18)	4	SAF work area address - Points to a 512-byte work area for use by the MVS router and the RACF front end routine.
	28 (1C)	4	Reserved.
	32 (20)	4	Reserved.
SAFPRACT	36 (24)	4	Offset - Contains the (signed) offset from the start of the MVS router parameter list to the RACF parameter list.

Interface to the MVS Router (RACROUTE)

The RACROUTE macro instruction is the interface to the MVS router that provides a focal point and a common system interface for all products providing resource control. The MVS router first invokes an optional installation exit and then invokes RACF, if RACF is active and installed on the system.

The RACROUTE macro accepts all valid parameters for any of the RACF macros (RACDEF, RACINIT, RACHECK, RACLIST and FRACHECK) and internally issues the appropriate RACF macro to generate a RACF parameter list. When the RACROUTE macro internally invokes the RACF macros, RACROUTE verifies that only valid parameters have been coded and then passes the parameters to the MVS router. Existing control points that invoke RACF processing via the supervisor call interface can continue to do so or can replace the RACF supervisor calls with the RACROUTE macro.

See the RACROUTE macro in Volume 2 for a description of the return codes.

Changing System Status (MODESET)

The MODESET macro instruction alters selective fields of the program status word (PSW). You can code the standard form of MODESET in two separate ways: one form generates an SVC and the other form generates inline code.

Generating an SVC

This form of MODESET, which executes as APF-authorized, in supervisor state, or under protection key 0-7, changes the status of programs between supervisor state and problem program state, and key zero and non-key zero. The parameters that must be specified to perform the changes are MODE and KEY respectively.

The MODE parameter specifies whether bit 15 of the PSW is to be set to one or zero. When bit 15 is one, the processor is in the problem state. For problem state, the PKM is changed to reflect the PSW key. When bit 15 is zero, the processor is in the supervisor state.

The KEY parameter specifies whether bits 8-11 are to be set to zero or set to the value in the caller's TCB. Bits 8-11 form the processor protection key. The key is matched against a key in storage whenever information is stored, or whenever information is fetched from a location that is protected against fetching.

Generating Inline Code

This form of MODESET is used to ensure that storage areas and the control program functions they are associated with have the same protection key. The EXTKEY parameter of MODESET indicates the key to be set in the current PSW.

You can set the following keys:

- Scheduler
- Job entry subsystem
- Real storage management
- Virtual storage management
- System resource management
- Supervisor
- Data management
- Telecommunications access method
- Key of zero
- Key of TCB
- Key of caller of type 1 SVC issuing MODESET
- Key of caller of type 2, 3, or 4 SVC issuing MODESET

Other parameters of MODESET allow the original key to be saved and restored upon completion of the desired changes.

Protecting Low Storage (PROTPSA)

The low address protection facility provides protection against altering storage addresses in the range of 0-511. The facility is designed to prevent inadvertent program destruction of storage used by hardware to fetch new PSWs for interruption processing. Low address protection does not apply to the storing of status by the processor, such as old PSWs, logout data, and processor logouts, nor does it apply to the data address in channel commands words (CCWs) and indirect data address words (IDAWs).

The PROTPSA macro instruction is used to disable and enable low address protection. To use this macro, programs must execute in supervisor state with PSW protection key 0, must be physically disabled for I/O and external interrupts, and must not issue any SVCs. The program must not call or transfer control to another program while low address protection is disabled. The protection-disabled window, that part of the program that is executing with low address protection disabled, should be as small as possible.

To insure the proper functioning of the low address protection facility, the contents of control register 0 must be maintained. This requires any program that modifies control register 0, except those using the PROTPSA macro instruction, to place a copy of control register 0 into the field PSACROSV of the PSA.

Bit 3 of control register 0 is defined as the protection bit. When this bit is zero, low address protection is disabled and stores are permitted. If the bit is set to one, low address protection is enabled and stores are not permitted. When a store in the address range 0-511 is attempted and low address protection is enabled, the content of the storage area addressed by the instruction is not modified. The execution of the current instruction is terminated and a protection exception occurs.

Exit Routines

This chapter includes the following topics:

- Using asynchronous exit routines
- Establishing a timer disabled exit routine

Additional information on exit routines can be found in *SPL: User Exits*.

Using Asynchronous Exit Routines

An authorized user can request an asynchronous exit routine to execute on behalf of a specific task. Before execution, the exit routine must complete three system control stages, each stage carried out by an individual “exit effector” routine.

The stage 1 exit effector routine creates and initializes an interrupt request block (IRB) that identifies the user’s asynchronous exit routine to the system. Interface to the stage 1 exit effector is through the create interrupt request block (CIRB) macro instruction. Volume 2 describes the CIRB macro instruction.

The stage 2 exit effector routine schedules the user’s exit routine for execution. Input to the stage 2 exit effector is an interrupt queue element (IQE) initialized by the caller. The IQE identifies the task the exit routine is to execute under and the associated IRB, and it also contains information about the exit routine’s characteristics. At the conclusion of stage 2 exit effector processing, the user’s exit routine is logically ready for system dispatch and execution. The exit routine is executed only once each time the caller invokes the stage 2 exit effector for that exit routine. Reuse of previously defined control blocks is possible, however. Thus, two or more invocations of the exit routine can use one interface to the stage 1 exit effector.

The MVS/XA dispatcher invokes the stage 3 exit effector to queue the exit routine’s IRB to the specified task. The exit routine will then execute on the next dispatch of that task (provided no other IRBs have been scheduled for that task).

Stage 1 Initialization

The CIRB macro instruction provides the interface to the stage 1 exit effector and is normally a type 1 SVC interface. Also, a branch entry interface is available by specifying `BRANCH=YES` as one of the macro options. Detailed information about using the CIRB macro instruction is in Volume 2. When invoked, the stage 1 exit effector routine obtains the storage for the IRB and, optionally, the IQE and a problem program work area. Also, the stage 1 exit effector initializes those fields in the IRB necessary to control the execution of the exit routine. At completion of the stage 1 exit effector, the IRB address is returned to the caller in register 1. The IQE, if requested, occupies storage contiguous to the IRB and is pointed to by the

RBNEXAV word in the IRB. The problem program work area, if requested, is pointed to by the RBPPSAV1 word in the IRB. The data area configuration is shown in Figure 59.

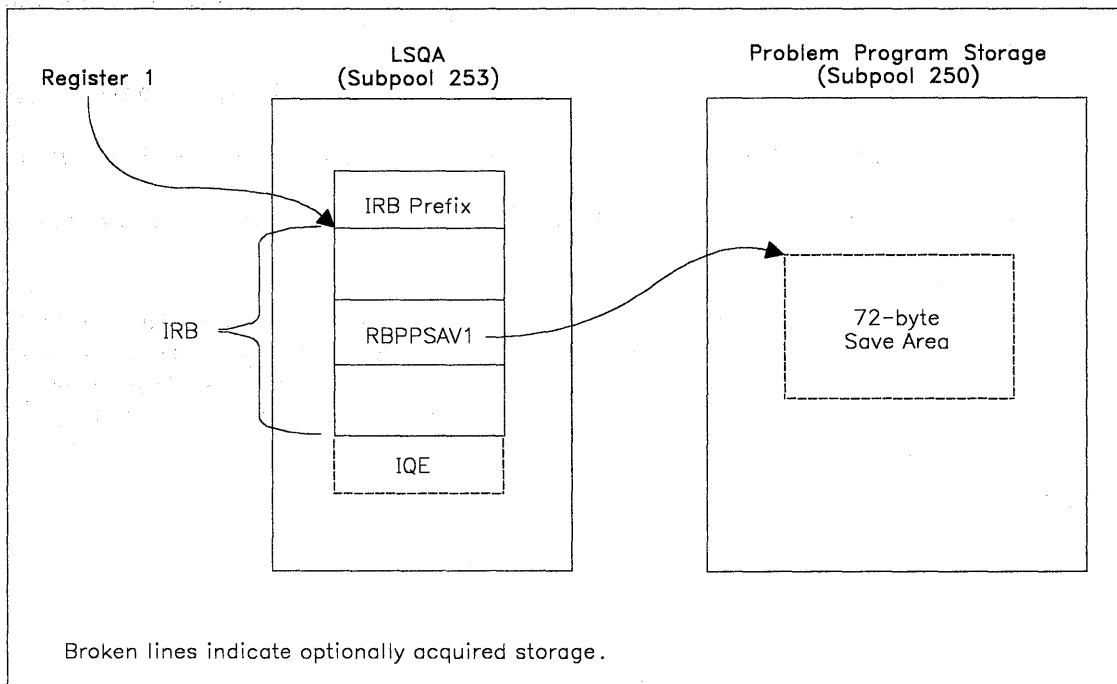


Figure 59. Asynchronous Exit Data Area Configuration

The fields in the IRB initialized by the stage 1 exit effector are:

- RBEP -- entry point address of the exit routine.
- RBSTAB -- flags indicating how the IRB and IQE are to be treated upon termination of the exit routine (defined according to the STAB and RETRN parameters of CIRB).
- RBIQETP -- flag indicating the type of queue element (RQE or IQE) associated with the exit request.

Note: Only the I/O supervisor uses RQEs.

- RBOPSW -- PSW to be loaded to initiate execution of the exit routine:
 - PSW is enabled for interrupts.
 - Protection key: 0 if KEY=SUPR specified on CIRB macro; TCB key (TCBPKF) of the caller if KEY=PP is specified on CIRB.
 - Mode: Supervisor state if MODE=SUPR on the CIRB macro instruction; problem program state if MODE=PP on CIRB.
- RBSIZE -- the size of the IRB (including the size of the IQE if the CIRB specification included the WKAREA parameter).
- RBNEXAV -- the address of the IQE if WKAREA was specified (occupies the first four bytes of the work area).
- RBPPSAV1 -- the address of the problem program save area if SVAREA was specified.

Stage 2 Scheduling

The user must initialize the IQE to define to the system the task under which the exit routine is to execute. The fields to be initialized are:

- IQEPARAM -- optional address of the parameter list to be passed to the exit routine.
- IQEIRB -- address of the IRB as returned in register 1 by the stage 1 exit effector routine.
- IQETCB -- address of the TCB for the task under which the user's exit routine is to execute.

When IRB/IQE initialization is complete, the user should invoke the stage 2 exit effector routine to queue the request (IQE) to the appropriate system asynchronous exit queue. The entry to stage 2 is by branch only, where the branch entry point address is found in the communications vector table (CVT) field CVT0EF00. The interface to the stage 2 exit effector is defined as follows:

Register	Contents
0	Irrelevant for scheduling an exit via IQE.
1	Twos-complement IQE address.
2-13	Irrelevant.
14	Return address.
15	Irrelevant.

Note: Upon return, registers 0, 2-9 and 11-14 are unchanged, register 1 contains a true (non-complemented) IQE address, and register 10 is destroyed.

The caller of stage 2 must:

- Hold the local lock.
- Be addressable in the address space in which the exit routine is to be dispatched.
- Be in supervisor state under protection key zero.

Stage 3 Execution

Once scheduled by stage 2, the user's exit routine is logically ready for dispatch. Stage 3, effectively a subroutine of the MVS/XA dispatcher, is called to queue the IRB associated with the user's exit to the task the IQE indicates. The user's exit executes as a result of the next dispatch of the task unless a subsequent IRB has been scheduled for the same task. In this case, the second exit routine might be executed first. Stage 3 execution depends on valid information being in the IQE fields. Because stage 3 performs no validity checks on IQE initialization, the user must ensure that the IQE fields are correctly initialized.

Execution and Termination Characteristics

The following characteristics of an asynchronous exit routine can influence the user's choice of CIRB options and should be considered:

- The exit routine executes as an IRB under the TCB defined by the IQE passed to the stage 3 exit effector routine.
- The exit routine executes enabled in the key and state requested by the CIRB macro instruction interface to the stage 1 exit effector routine.

- Register contents upon entry to the exit routine are:

Register	Contents
0	IQE address.
1	Parameter list address (IQEPARAM).
13	Problem program register save area address (if any).
14	Return address (CVTEXTIT).

- Upon termination of the asynchronous exit routine:
 - The IQE is returned to a “next available” queue anchored by the RBNEXAV field in the IRB if the user specified the WKAREA and RETRN = YES options on the CIRB macro instruction. This allows subsequent invocations of the exit routine without requiring the user to repeat requests for stage 1 processing (data area setup).
 - If the user specified the SVAREA and STAB = (DYN) options on the CIRB macro instruction, the problem program register save area is freed.
 - If the user specified the STAB = (DYN) option of the CIRB macro instruction, the IRB and IQE are freed.

Establishing a Timer Disabled Interrupt Exit

Timer supervision provides a function called set DIE that allows a user-written program to establish a disabled interrupt exit (DIE) routine. The DIE routine gains control asynchronously after a specified real time interval has elapsed.

The set DIE function is available only to programs executing in supervisor state with PSW key zero. The set DIE function allows users to initiate a real time interval by branching to the set DIE system service (IEAVRT02). When the time interval expires, the user’s DIE routine gains control as an extension of the timer second level interrupt handler (IEAVRTI0). It is also possible for a user to set a new time interval from the DIE routine.

Although a program can have an unlimited number of outstanding time intervals at one time, storage and system performance considerations may impose practical and reasonable limits.

Note: The time during which a DIE routine is executing is not charged to the job step time of the interrupted address space.

The caller of the set DIE service routine can be executing in either task control block (TCB) or service request block (SRB) mode, but must be in PSW key zero and supervisor state. The entry point to the set DIE service routine is in field TPCSDIE in the timer supervision work area mapped by macro IEAVVTPC. The address of this work area is in CVT field CVTTPC.

The caller of the set DIE service routine must provide the following input environment:

1. Register 1 must contain the address of a user-supplied timer queue element (TQE) whose fields are available from the IHATQE mapping macro. The IHATQE macro is available in macro library APVTMACS. This user TQE must:
 - Be a contiguous block of 128 bytes aligned on a double word boundary.
 - Reside in SQA

- Include the following field initialization:

TQEALD -- zero or a valid ASID, important in case of an address space failure (see "Obtaining and Freeing the TQE").

TQEVAL -- the desired real time interval (a 64 bit unsigned binary number with bit 51 = 1 microsecond).

TQEAMODE bit in TQEFLGS3 field -- set to 1 to indicate that the address of the user's DIE in TQEEXIT is pointer defined.

TQEEXIT -- address of the user's DIE. If the TQEAMODE bit in the TQEFLGS3 field is set to 1, the high-order bit of this field, TQEXMODE, must indicate the addressing mode of the user's DIE. If the user's DIE is to execute in 24-bit addressing mode, TQEXMODE=0; if the user's DIE is to execute in 31-bit addressing mode, TQEXMODE=1.

- Have all the other fields cleared to zero.
2. Registers 2 - 12 must be parameter registers whose input values will be restored in the same registers on entry to the DIE routine.
 3. Register 14 must contain the caller's return address.

Loss of the contents of register 1 and 11-13 occurs upon return from the set DIE service routine. Register 15 contains a return code as follows:

Code	Meaning
0	The TQE was successfully enqueued onto the system's real time queue.
4	Failure - needed clocks are unavailable.

Note: The set DIE service routine obtains the dispatcher lock if the caller does not already hold it. After completing its processing, the set DIE routine releases the lock if the caller did not previously hold it. The caller must not hold any lock higher in the locking hierarchy than the dispatcher lock.

The set DIE service routine does not establish its own recovery routine. Any system program calling the set DIE service routine should have its own FRR or ESTAE routine. A program check occurs in the set DIE service routine if the caller is not both in PSW key zero and in supervisor state.

The DIE routine executing out of the timer SLIH gains control under timer supervision's FRR on the current stack. The DIE itself can optionally establish its own FRR, which should terminate by percolation to let the timer supervision FRR gain control. When the timer supervision FRR gets control, it tries to repair any damage to the system's real time queue and then percolates.

DIE Characteristics

Entry to the DIE routine is in supervisor state, with PSW key zero, disabled, with no disabled global spin locks held. Register contents upon entry are as follows:

- Register 1 contains the address of the TQE. At this time the TQE is not enqueued upon the real time queue. Fields TQETCB and TQEASCB respectively contain a TCB address and an ASCB address, if previously set by the user on entry to the set DIE service routine.
- Register 2 - 12 are as they were upon entry to the set DIE service routine (or as changed by a previous DIE entry -- "DIE Execution").
- Register 14 contains the return address.
- Register 15 contains the entry point of the DIE routine.
- The contents of floating point registers are unpredictable.

While a system program has a TQE enqueued upon the real time queue, it must ensure that the associated DIE routine is available for the timer SLIH (second level interrupt handler) to access from any address space. Additionally, because the DIE is entered disabled, its code must be resident or fixed to avoid a page fault at entry.

Exit from the DIE Routine: must be to the address specified in register 14. This exit must also occur in supervisor state with PSW key zero, and disabled. The routine must release all locks it obtained but need not save the general purpose registers. Floating point registers, however, must have the same contents on exit as on entry to the DIE.

DIE Execution: must be like the execution of an interrupt handler because it executes as an extension of the timer SLIH. Specifically, the DIE routine executes under the following restrictions:

- The DIE must be capable of executing in any address space because the timer interruption may occur while any address space enabled for external interruptions is executing.
- The DIE cannot reference any private storage areas.
- The DIE must execute disabled. Hence, it cannot cause a page fault.
- The DIE cannot request a local lock or the CMS lock because these are suspend locks and might therefore already be in use. Furthermore, the DIE routine cannot assume whether or not these locks are held upon entry.
- The DIE cannot execute any SVCs.

The DIE routine may re-enqueue the TQE to set another real time interval by using the timer's TQE ENQUEUE routine (whose entry point is in CVT field CVTQTE00). The DIE routine must hold the dispatcher lock upon entry to the TQE ENQUEUE routine.

The input environment for the TQE ENQUEUE routine must be as follows:

- Supervisor state, key zero, and holding the dispatcher lock.
- Register 1 must contain the address of the TQE supplied to the DIE routine. Only the following TQE fields can be changed.

TQEVAL -- This field should contain the clock comparator value for the next interruption. This value is equivalent to the desired interval added to the value in TQEVAL when the DIE routine was entered. Alternatively, TQEVAL can be calculated by adding the desired interval to the current TOD clock reading (as obtained by a STCK instruction). The choice of which method to use is further discussed under "Clock Failure."

TQEAMODE bit in TQEFLGS3 field -- set to 1 to indicate that the address of the user's DIE in TQEEXIT is pointer defined.

TQEEXIT -- This field should contain the new address if a DIE routine address different from the current one is desired. Otherwise the field should remain unchanged. If the TQEAMODE bit in the TQEFLGS3 field is set to 1, the high-order bit of this field, TQEXMODE, must indicate the addressing mode of the user's DIE. If the user's DIE is to execute in 24-bit addressing mode, TQEXMODE=0; if the user's DIE is to execute in 31-bit addressing mode, TQEXMODE=1.

TQEDREGS -- If the parameter values in registers 2 - 12 are to be changed for the subsequent DIE routine entry, the new values should be set in this eleven word field.

- Register 2 must contain the caller's return address.

Upon return from the TQE ENQUEUE routine, all registers are as they were on entry except for registers 13 and 15.

Although the set DIE function is similar to the TQE ENQUEUE function, the routines differ in the following respects:

- Although TQE ENQUEUE expects an already established and fully initialized TQE as input, the set DIE service routine completes the user-supplied TQE (including important flag bits) to make it acceptable to timer supervision.
- For TQE ENQUEUE, TQEVAL in the TQE must be set to the clock comparator value for the next interruption. With the set DIE service routine, it must be set to the desired interval. The set DIE service routine then converts it to the proper clock comparator value.
- TQE ENQUEUE assumes that the clocks are functioning correctly. The set DIE service routine must use the clocks directly and therefore verifies (rather than assumes) that the clocks are functioning correctly. The set DIE service routine is therefore capable of advantageously using alternate clocks in a multiprocessing environment in which one or more clocks have failed.

Timer Queue Element Control

The major aspects of controlling the timer queue element (TQE) associated with the user's DIE routine are:

- Obtaining and freeing the TQE
- Serializing the use of each TQE
- Time-of-day clock failure
- Interval cancellation

Descriptions of each of these aspects follow.

Obtaining and Freeing the TQE: is your responsibility as user of the set DIE function because the TQE resides in SQA. Thus, you must explicitly free the TQE when it is no longer necessary and (with one exception) in error situations as well. Timer supervision frees a TQE for you for a failing address space only if the TQE is enqueued on the real time queue and has field TQEASID set to the ASID of the failing address space.

Before freeing the TQE, however, you must ensure that it is not currently on the real time queue. There are several ways to accomplish this:

- Always free the TQE in the DIE routine because it is never on the real time queue when the routine receives control.
- Before freeing the TQE, use timer supervision's TQE DEQUEUE routine. This routine either removes the TQE from the real time queue or, if the TQE is not on the queue, takes no action.

Notes:

1. *You must not alter the TQE (other than in the fields previously described).*
2. *The interface for the TQE DEQUEUE routine is described in the section "Interval Cancellation."*

Serializing the Use of Each TQE: is also your responsibility. Serialization includes the execution of the set DIE service routine, TQE ENQUEUE, and TQE DEQUEUE routines for a given TQE because these routines update the supplied TQE. Never update a TQE, however, while it is on the real time queue. Timer supervision serializes the use of the real time queue by means of the dispatcher lock.

Clock Failure: can keep a DIE routine from receiving control. If a clock required by a DIE routine's TQE fails while the TQE is on the real time queue, timer supervision leaves the TQE on the queue, thereby denying control to the DIE routine. To permit the DIE routine to receive control, a properly functioning TOD clock and clock comparator must be varied online. For this remedy to work, the DIE routine must be in resident or fixed storage as long as its TQE is on the real time queue. These storage locations make the DIE routine available to the timer SLIH from any address space.

When the DIE routine gains control under these circumstances, the clock comparator value in TQEVAl could be behind the TOD clock. If the DIE routine re-enqueues the TQE on each successive entry and adds a new interval to TQEVAl, then the DIE routine gains control each time, immediately upon enablement of the external interruptions. This sequence continues until the value in TQEVAl is equal to the TOD clock value. To avoid this synchronization loop,

the DIE routine can calculate the new TQEVAL as the sum of the new interval plus the current TOD clock value. This method, however, requires that the DIE routine contain error recovery code in case the STCK instruction fails due to a bad TOD clock in the executing processor.

Interval Cancellation: can occur by using timer supervision's TQE DEQUEUE routine. This routine removes a specific TQE from the real time queue and resets clocks if necessary. The entry point to the TQE DEQUEUE routine is in CVT field CVTQTD00. Entry to this routine must be by branch entry, in supervisor state, with PSW key zero, and with the dispatcher lock. The input environment is as follows:

- Register 1 must contain the address of the TQE to be dequeued.
- Register 2 must contain the caller's return address.

Upon return, all registers except 13 and 15 are the same as they were on entry.

User-Written SVC Routines

This chapter includes the following topics:

- Writing SVC routines
- Inserting SVC routines into the control program
- Subsystem SVC Screening

Writing SVC Routines

You can introduce user-written SVC routines into the control program whenever you IPL the system. When you write an SVC routine, you must follow the same programming conventions used by SVC routines supplied with MVS/XA. Five types of SVC routines are supplied with MVS/XA, and the programming conventions for each type are different.

SVC routines, including user-written ones, can either be part of the resident control program (the nucleus), or be part of the fixed or pageable link pack area. Types 1, 2, and 6 SVC routines become part of the resident control program, and types 3 and 4 go into the link pack area. Before IPLing the system, you must place your SVC routine in SYS1.NUCLEUS or SYS1.LPALIB. You must also create, before IPLing the system, an IEASVCxx member in SYS1.PARMLIB with SVC Parm statements that describe the characteristics of your SVC routine.

SVC routines receive control with PSW key zero and in supervisor state. They must be reenterable and, if you want to aid system facilities in recovering from machine malfunctions, they must also be refreshable.

If your routines must execute serially with respect to other parts of the control program, then you must use the same locking conventions as the control program. If you write two or more SVC routines that must serialize with each other, use the locking facilities or the ENQ and DEQ macro instructions.

When you insert an SVC routine into the control program, you specify which locks the routine will require. When an SVC routine receives control, it is normally enabled and it can be holding one or more locks. However, if you specified that the routine requires a disabled spin lock, the routine is disabled when it receives control. The routine is also entered in a disabled state if it is a type 6 SVC routine.

Type 6 SVC Routines

You must define your user-written SVC routine as being one of the five valid types, including type 6. The type 6 SVC routine performs functions similar to the type 1 SVC routine. However, because the instruction path lengths for receiving and releasing control are shorter, the type 6 routine offers performance advantages over the type 1. The type 6 SVC routine cannot require the LOCAL lock, as noted later.

The type 6 SVC also provides a more efficient way to change from TCB mode to SRB mode processing. The type 1 SVC must schedule an SRB, which then goes through queuing and dequeuing operations before it is eventually dispatched. The type 6 SVC, however, normally results in immediate scheduling and dispatching of the SRB.

Because a type 6 SVC routine executes under the control of the SVC first level interrupt handler (FLIH), it has the same limitations that apply to the FLIH. When a type 1 SVC routine exits, it always returns to the SVC FLIH. There are three exit options for a type 6 SVC:

- Return to the caller directly
- Return to the dispatcher
- Dispatch an SRB (service request block)

To exit from a type 6 SVC routine, either issue the T6EXIT macro or use the original contents of register 14 as a return address. The use of T6EXIT results in the following register conditions:

T6EXIT Option	Register 14	Register 0	Register 1	Register 15
CALLER	As on entry	Returned to caller	Returned to caller	Returned to caller
DISPATCH	CVTT6SVC	N/A	0	N/A
SRB	CVTT6SVC	N/A	SRB address	N/A
BR 14	As on entry	Returned to caller	Returned to caller	Returned to caller

If a type 6 SVC uses the RETURN = SRB exit option on the T6EXIT macro instruction, register 1 must point to an SRB. The SRBASCBC field must indicate the current address space.

The system neither acquires nor releases any locks for type 6 SVCs. Because a type 6 SVC executes disabled, it has exclusive use of the processor. Thus, the SVRB save areas are unavailable to a type 6 SVC routine, although the PSA areas can be used instead. When a type 6 SVC is executing, no other task-related activity can occur concurrently. To indicate this situation, the TCBACTIV flag is set. Type 6 SVC routines should be short enough to minimize any adverse effect on performance and they should provide for recovery by using the SETFRR macro instruction.

Non-Preemptable SVC routines

You can define a user-written SVC routine as non-preemptable for I/O interruptions. If a non-preemptable SVC routine sustains an I/O interrupt, the SVC, rather than the highest priority ready work, gets control when I/O processing is complete. The non-preemptable SVC cannot issue other SVCs and remain non-preemptable because the exit function always resets the non-preemptable indicator in the TCB associated with the SVC. This action causes the issuing SVC to lose its non-preemptable state. If a non-preemptable SVC issues a STAX DEFER = NO macro instruction, the SVC routine remains non-preemptable until it exits.

Programming Conventions for SVC Routines

Figure 60 summarizes the programming conventions for the five types of SVC routines. Details about many of the conventions are in the reference notes that follow the figure. The numbers in the right most column of the figure correspond to the reference notes. If a reference note for a convention does not pertain to a specific type of SVC routine, that type is indicated by an asterisk.

Conventions	Type 1	Type 2	Type 3	Type 4	Type 6	Reference Code
Part of resident control program	Yes	Yes	No	No	Yes	
Size of routine	Any	Any	Any	Any	Any	
Reenterable routine	Yes	Yes	Yes	Yes	Yes	1
Refreshable routine	No	No	Yes	Yes	No	2
Locking requirements	Yes	No	No	No	No	3
Entry point	Must be on a halfword boundary and must be the first instruction to get control. Need not be the first byte of the module.					
Number of routine	Numbers assigned to your SVC routine should be in descending order from 255 through 200					
Name of routine	IGCnnn	IGCnnn	IGC00nnn	IGC00nnn	IGCnnn	4
Register contents at entry time	Registers 3, 4, 5, 6, 7, and 14 contain communication pointers; registers 0, 1, 13, and 15 are parameter registers					5
Supervisor request block (SVRB) size	No SVRB exists	224	224	224	No SVRB exists	6
May issue WAIT macro instruction	No	Yes	Yes	Yes	No	7
May suspend their caller	Yes	No	No	No	Yes	8
May issue XCTL macro instruction	No	Yes	Yes	Yes	No	9
May pass control to what other types of SVC routines	None	Any	Any	Any	None	10
Type of linkage with other SVC routines	Not Applicable	Issue supervisor call (SVC) instruction			Not Applicable	11
Exit from SVC routine	Branch using return register 14				T6 EXIT or BR 14	12
Method of abnormal	ABEND	ABEND			ABEND	
Recovery	FRR	ESTAE or FRR			FRR	13

Figure 60. Programming Conventions for SVC Routines

Reference Code	SVC Routine Types	Reference Notes
1	all	If your SVC routine is to be reenterable, you cannot use macro instructions whose expansions store information into an inline parameter list.
2	3,4	Types 3 and 4 in the pageable LPA must be refreshable. Types 3 and 4 in the fixed LPA must be reenterable, but not necessarily refreshable.
3	all	<p>The following conventions on locking requirements apply:</p> <ul style="list-style-type: none"> * Type 1 SVC routines always receive control with the LOCAL lock held and must not release the LOCAL lock. Additional locks may be requested prior to entry via the SVCTABLE macro instruction or may be requested dynamically within the SVC routine. * Types 2, 3, and 4 may also request locks via the SVCTABLE macro instruction or may obtain them dynamically. * Types 1 and 2 may request that any locks be held on entry. Types 3 and 4 may only request that the LOCAL or LOCAL and CMS lock be held. * If no locks are held or obtained, or only suspend locks (LOCAL and CMS) are held or obtained, the SVC routine executes in supervisor state key zero, enabled mode. * If disabled spin locks are held or obtained, the SVC routine executes in supervisor state, key zero, disabled mode. No SVCs may be issued. * SVCs may not take disabled page faults. Therefore, if a disabled spin lock is held, the SVC routines must ensure that any referenced pages are fixed. For types 3 and 4, all pages containing code must be fixed. * An FRR may be defined for any SVC routine that holds or obtains locks to provide for abnormal termination (see reference code 9). * Type 6 may not request any locks.
4	all	<p>You must use the following conventions when naming SVC routines:</p> <ul style="list-style-type: none"> * Types 1, 2, and 6 must be named IGCnnn; nnn is the decimal number of the SVC routine. You must specify this name in an ENTRY, CSECT, or START instruction. * Types 3 and 4 must be named IGC00nnn; nnn is the signed decimal number of the SVC routine. <p>The following conventions regarding type 3 and 4 SVCs are not enforced by SVC processing, but have traditionally been used to distinguish between the two types:</p> <ul style="list-style-type: none"> * A type 3 SVC identifies a function that is contained in a single load module. * A type 4 SVC identifies a function that loads additional modules. You can identify these loaded modules as IGC01nnn, IGC02nnn, ..., and IGC0xnnn. (IGC01nnn is the first module that IGC00nnn loads, IGC02nnn is the second module that IGC00nnn loads, and IGC0xnnn is the last module that IGC00nnn loads.)
5	all	<p>Before your SVC routine receives control, the contents of all registers are saved. In general, the location of the register save area is unknown to the routine that is called. When your SVC routine receives control, the status of the registers is as follows:</p> <ul style="list-style-type: none"> * Register 0 and 1 contain the same information as when the SVC routine was called. * Register 2 contains unpredictable information. * Register 3 contains the starting address of the communication vector table (CVT). * Register 4 contains the address of the task control block (TCB) of the task that called the SVC routine. * Register 5 contains the address of the supervisor request block (SVRB), if a type 2, 3, or 4 SVC routine is in control. If a type 1 or 6 SVC routine is in control, register 5 contains the address of the last active request block. * Register 6 contains the entry point address. * Register 7 contains the address of the address space control block (ASCB). * Registers 8 through 12 contain unpredictable information. * Register 13 contains the same information as when the SVC routine was called. * Register 14 contains the return address. * Register 15 contains the same information as when the SVC routine was called. <p>You must use register 0, 1, and 15 if you want to pass information to the calling program. The contents of registers 2 through 14 are restored when control is returned to the calling program.</p>

Reference Code	SVC Routine Types	Reference Notes
6	2,3,4	When a type 2, 3, or 4 SVC routine receives control, register 5 contains the address of the SVRB with in this 224-byte area. This SVRB contains a 48-byte "extended save area." In addition, an area is provided for a STAE control block (SCB); this SCB is used by the FESTAE macro instruction.
7	2,3,4	You can issue the WAIT macro instruction if you hold no locks. You can issue WAIT macro instructions that await either single or multiple-events. The event control block (ECB) for single-event waits on the ECB list and ECBs for multiple-event waits must be in virtual storage. Type 6 SVCs may not issue WAIT but may issue SUSPEND.
8	1,6	Both type 1 and 6 SVC routines can issue SUSPEND RB=CURRENT to suspend their callers.
9	2,3,4	When you issue an XCTL macro instruction in a routine under control of a type 2, 3, or 4 SVRB, the new load module must be located in the fixed or pageable link pack area. The contents of registers 2 through 13 are unchanged when control is passed to the load module; register 15 contains the entry point of the called load module.
10	all	No SVC routines except ABEND may be called if locks are held. ABEND may be called at any time.
11	all	No locks may be held. If locks are held, branch entry to SVCs is acceptable, or the locks may be freed, the SVC issued, and the locks reobtained.
12	all	Branch using return register 14 should be used. SVC routines that exit via BR 14 or T6EXIT must return control in the same state in which they received control, such as, key zero, supervisor state. Otherwise, if locks are held, SVC 3 results in abnormal termination. Note: To ensure that control is returned to the dispatcher, the SVC routine can load register 14 with the address in the CVTEXPI field of the CVT before issuing BR 14.
13	all	If an SVC routine is entered with a lock held or if an SVC routine obtains a lock, it should specify a functional recovery routine (FRR) for as long as the lock is held (see SETFRR macro instruction). The FRR receives control if an error occurs, and ensures the validity of the data being serialized by the lock; the FRR either recovers or releases the lock and continues with termination. If no FRR is specified, the recovery termination manager releases the lock and terminates the task. No cleanup of the data is performed. (Note that the lock is released before any STAI/ESTAI/ESTAE (or STAE) recovery routine is entered. If no locks are acquired for or by an SVC routine, then an ESTAE may be used to define your recovery processing (see ESTAE and SETRP macro instructions).

Inserting SVC Routines Into the Control Program

To supply user-written SVC routines to the system, you place descriptions of your user SVC routines in SYS1.PARMLIB, and you place the actual routines in SYS1.NUCLIB and SYS1.LPALIB. When the system is IPLed, the NIP program translates the SVC definitions that you placed on SYS1.PARMLIB into SVC table entries. Then NIP places the SVC table in the extended read-only nucleus.

When NIP processes SYS1.PARMLIB, it searches for member names that it uses to build the SVC table. These member names have the form, IEASVCxx, where xx is the field specified by the SVC= option in the IPL system parameters.

In the IEASVCxx members, you code SVC Parm statements. The SVC Parm statements describe the properties and attributes of individual SVC routines. Each SVC Parm statement that you code describes a single SVC routine; it generates one entry in the SVC table. Using the SVC Parm statement, you specify the SVC number, type, entry point name, lock requirements, authorization level, and whether or not the SVC is preemptable. See *MVS/XA SPL: Initialization and Tuning* for a description of the SVC= and SVC Parm statements.

The user SVC entries, which are represented by the SVC numbers 200-255, are the only ones you are allowed to define. You cannot modify SVCs that are in the range of 0-199. When you

define an SVC with an SVC Parm statement, you define its type as type 1 through type 6, excluding type 5. The system provides no SVC routines in the range 200-255. Therefore, unless the user defines some SVC routines in this range, execution of an SVC 200 through 255 will cause an abend.

Modifying the SVC Table at Execution Time (SVCUPDTE)

After the IPL, the SVC table can be dynamically modified by authorized users via the SVCUPDTE macro. For example, authorized subsystems such as VTAM can alter the SVC table when the subsystem starts and restore the table when the subsystem terminates. For additional flexibility, the EPNAME and EXTRACT parameters of the SVCUPDTE macro allow the authorized user to dynamically associate SVC numbers with entry points of SVC routines.

An SVC update recording table is maintained in parallel with the SVC table. This table provides a record of changes to the SVC table. Entries are created whenever a change is made to the SVC table with SVC Parm statements or the SVCUPDTE macro.

Intercepting an SVC Routine

When you execute an SVC instruction, the unique program to which control is passed is called the SVC routine. A common programming technique is to intercept an SVC routine by inserting another program in the path between the SVC instruction and the SVC routine. The inserted program is sometimes called a *front end* to the original SVC routine. After the front end program is inserted, the resulting body of code, including the front end program and the original SVC routine, is the new SVC routine.

Intercepting SVC routines can be recursive. Thus, if an SVC routine already has a front end, you can still add another front end onto it, and so on, indefinitely.

To intercept an SVC routine, you must obtain and save the address of the existing SVC routine for use by the front end program. To change the entry in the SVC table so it points to the front end program, you must use the REPLACE function of the SVCUPDTE macro.

In a user environment where the interception of SVC routines is recursive, it might be necessary to serialize the modification of the SVC table. To serialize, use the ENQ and DEQ macros to secure and hold the SYSZSVC TABLE resource while you are changing the SVC table.

Before you obtain the SVC table entry, use ENQ to secure this resource, and hold it until you have replaced the SVC table entry with the pointer to the front end routine. Then you can DEQ the resource. The major and minor names of this resource are, respectively, SYSZSVC and TABLE.

Subsystem SVC Screening

After you write an SVC routine and insert it into the system, the routine is generally available unless you take steps to regulate access to the routine. Subsystem SVC screening allows a system routine to define those SVCs that a specific task can validly issue. When SVC screening is active for a task, the SVC first level interrupt handler (FLIH) determines, for each SVC issued by that task, whether the task can request that SVC function. If the SVC request is invalid, the SVC FLIH gives control to a special error subroutine supplied by the routine that activated the screening function.

The subsystem, executing under PSW protection key zero, activates SVC screening by setting two fields in each TCB for which screening is desired. The two fields consist of a screen flag bit and a one-word field containing the address of the subsystem screen table, which provides the interface between the SVC FLIH and the subsystem subroutine. In addition to these fields, the subsystem may optionally set the TCBSVCSP bit to indicate that ATTACH processing is to pass the SVC screening information to the attached task. The important SVC screening fields in the TCB are:

- TCBSVCS - A flag bit. When set to one, it indicates that screening is in effect for this task.
- TCBSVCA2 - Address of the subsystem screen table.
- TCBSVCSP - Propagation bit. When set to one, it indicates that ATTACH processing should pass the SVC screening information in these three fields to the attached task.

When the screening facility detects an invalid SVC, it gives control to the specified error routine. The error routine receives control as an SVC and is subject to the same restrictions as SVC routines. Before giving control to the subroutine, the SVC FLIH provides the setup for the subroutine as defined by the subsystem SVC entry (SSTSVCN) in the subsystem screen table. This setup includes:

- Initializing the SVRB if the subroutine is to execute as a type 2, 3, or 4 SVC.
- Obtaining the LOCAL lock if the subroutine is to execute as a type 1 SVC.
- Acquiring all locks necessary for the subroutine's execution.

The subsystem that needs SVC screening obtains storage via GETMAIN for a 264 byte area called the subsystem screen table. To prevent a page fault, this area must come from the LSQA (subpool 253-255), the SQA (subpool 245), or must be in fixed storage. If the subsystem screen table is in fixed storage, the subsystem must ensure that the storage is protected from user modification. The subsystem screen table contains two areas as follows:

1) SSTVCN -- Subsystem SVC entry (8 bytes)

Bytes	Content
0-3	Entry point address of the subsystem subroutine that will get control whenever a task has issued an SVC against which there is a screening restriction.
4	X'00' -- means that the subroutine is to execute as a Type 1 SVC X'80' -- means that the subroutine is to execute as a Type 2 SVC X'C0' -- means that the subroutine is to execute as a Type 3 or Type 4 SVC X'20' -- means that the subroutine is to execute as a Type 6 SVC
5	Zero.
6-7	Locks to be held on entry to the subroutine. If the appropriate lock bit is one, the lock will be acquired by the SVC FLIH. The lock bits are:
	Bit Lock
	0 LOCAL
	1 CMS
	2 SRM
	3 SALLOC
	4 Dispatcher
	Bits 5-15 are always zero (off).

2) SSTMASK -- SVC screening mask (256 bytes)

Bytes	Content
8-263	Each byte corresponds to an SVC number in ascending order in the range 0-255. When the high order bit in a byte is one, the task may validly issue the respective SVC; when the bit is zero, there is a screening restriction that prohibits the task from issuing the SVC.

The subsystem must get and initialize the subsystem screen table properly. The subsystem must also free the table before terminating.

UCB Scan Services

The UCB scan routine (IOSVSUCB) allows you to scan each UCB in the system or in a specified device class. The device classes are: tape, communication, channel-to-channel adapter, direct access, display, unit record, and character reader. Using IOSVSUCB you can, for example, find the UCB currently associated with a particular VOLSER or find all tape devices currently allocated.

IOSVSUCB runs in the caller's key, state, and addressing mode. The caller can be in either task or SRB mode with the following restrictions:

- If in task mode, the caller must be enabled and must hold no locks.
- If in SRB mode, the caller cannot hold the UCB lock or any lock higher than that in the locking hierarchy.

Invoking IOSVSUCB

IOSVSUCB is more general than the IOSLOOK macro instruction, which requires the user to be in supervisor state and to provide the device address as input. See Volume 2 for a description of the IOSLOOK macro.

Each time that you invoke IOSVSUCB, you will obtain the address of the common segment of one unit control block (UCB). In order to scan several UCBs, you must invoke IOSVSUCB repeatedly, once for each UCB. IOSVSUCB keeps track of your position in the UCB chain by information that it stores in the 100-byte work area that you provide as input. To start your scan you must clear this work area to binary zero. The zeros indicate that IOSVSUCB is to start the scan at the first UCB in the system or device class. If you want to continue the scan to obtain the next UCB, you must not change the work area. See Figure 63 for an example of how to use IOSVSUCB.

Input to IOSVSUCB

To use IOSVSUCB, the caller must:

- Obtain a 100-byte work area that starts on a double-word boundary and clear the work area to binary zero each time that the scan is to start with the first UCB in the system or the first UCB in a device class.
- Build a parameter list as shown in Figure 61.

Address of the 100-byte work area provided by the caller.
Address of the byte containing the device class to which the search is being restricted. See the topic "Limiting the UCB Scan" for information on how to restrict the search to a specific device. If all UCBs are to be scanned, the byte pointed to must contain X'00'.
Address of the word in which IOSVSUCB is to return the UCB address. The high order bit of this field must be 1 to indicate it is the last word in the parameter list.

Figure 61. Parameter List for the UCB Scan Routine (IOSVSUCB)

- Set up the registers to contain the following information:

Register	Contents
1	Address of the parameter list
13	Address of caller's 18-word save area
14	Caller's return address
15	Entry point of the UCB scan routine (IOSVSUCB) (The CVTUCBSC field in the CVT contains the entry point address.)

Note: The data areas that the caller passes to IOSVSUCB must be addressable in the addressing mode of the caller. If the program runs in 31-bit addressing mode, the data areas can be anywhere. If the program runs in 24-bit addressing mode, the data areas must be below 16 megabytes.

Limiting the UCB Scan

If you want to limit the UCB scan to a specific device class, you must provide the address of a one-byte field containing the hexadecimal code for that class. These fields are defined in the UCBDVCLS (or UCBTBYT3) bit string in the UCB. Figure 62 lists the valid device class specifications with their UCB definitions. For example, to restrict the search to tapes, set the byte containing the device class equal to the constant UCB3TAPE. If you use the UCB definitions in your program, you must include the UCB mapping macro (IEFUCBOB). To scan all of the UCBs in the system, provide the address of a one-byte field containing X'00'.

UCB Definition	Device Class
UCB3TAPE	Tape
UCB3COMM	Communication
UCB3CTC	Channel-to-channel adapter
UCB3DACC	Direct access
UCB3DISP	Display
UCB3UREC	Unit record
UCB3CHAR	Character reader

Figure 62. Device Classes

Output from IOSVSUCB

When IOSVSUCB returns, register 15 contains one of the following return codes:

Return Code	Meaning
00	IOSVSUCB stored a UCB address in the location specified in the third word of the parameter list.
04	There are no more UCBs. IOSVSUCB set the 100-byte work area to binary zeros.

Notes:

- 1. A dynamic device reconfiguration (DDR) swap might occur during a scan. Because this type of swap results in the interchange of information in UCBs, it might cause a UCB address to be skipped or returned twice.*
- 2. Do not place any dependencies on the order in which the UCB addresses appear during a scan. The address of the UCB representing device 250, for example, might be returned before the one representing device 140.*
- 3. Devices with optional channels are associated with only one UCB. Therefore, IOSVSUCB returns only one UCB address for those devices. Devices with multiple exposures have one UCB associated with each exposure. Therefore, IOSVSUCB returns one UCB address for each exposure.*

	LTR	R15,R15	HAS A UCB BEEN RETURNED?	
	BNZ	NOMATCH	NO, AT END OF DEVICE CLASS AND NO MATCH FOUND	X
	L	R2,ADDRUCB	GET UCB ADDRESS THAT THE SCAN SERVICE RETURNED	X
	CLC	UCBVOLI,SRCHVOL	IS THIS THE VOLSER WE'RE LOOKING FOR?	X
	BNE	SEARCH	NO, CONTINUE SCAN OF UCBS.	
*			NOTE: THE WORK AREA MUST NOT BE CHANGED BETWEEN CALLS TO THE SCAN SERVICE ROUTINE	
	FOUND	DS	OH	
			
	NOMATCH	DS	OH	
			
	ENDIT	DS	OH	
		L	R13,MYSAVE+4	
		LM	R14,R12,12(R13)	RESTORE CALLER'S REGISTER 13 RESTORE REMAINDER OF CALLER'S REGISTERS
		BR	R14	
	EJECT			
	PARMLIST	DS	3F	PARMLIST MAPPING
		ORG	PARMLIST	
	PARMWA	DS	F	ADDRESS OF 100-BYTE WORK AREA
	PARMDEVT	DS	F	ADDRESS OF BYTE CONTAINING THE DEVICE TYPE TO BE SEARCHED
	PARMUCB	DS	F	ADDRESS OF WORD TO CONTAIN THE UCB ADDRESS
		SPACE		
	DEVCLASS	DS	CL1	BYTE CONTAINING DEVICE CLASS TO BE SEARCHED FOR
	ADDRUCB	DS	F	WORD IN WHICH UCB SCAN WILL PLACE THE ADDRESS OF THE UCBS. ALIGN ON DOUBLE-WORD BOUNDARY. (THE WORK AREA FOR SCAN SERVICE MUST BE ON A DOUBLE-WORD BOUNDARY.)
		DS	0D	WORK AREA
*				
	WORKAREA	DS	CL100	
	MYSAVE	DS	18F	
			
			
	DSECT			
	IEFUCBOB			UCB MACRO ID
	CVT	DSECT = YES		
	EJECT			
	END	FINDVOL		

Figure 63 (Part 2 of 2). Example of the UCB Scan Routine (IOSVSUCB)

Obtaining Information from the Input/Output Supervisor (IOS)

The IOSINFO macro instruction obtains the subchannel number for a specified unit control block (UCB) from the input/output supervisor (IOS) without being dependent on the location or format of the information as it is maintained by IOS. The macro returns the subsystem identification word (SID), which identifies the subchannel number of the UCB, in a user-specified location. The SID is a fullword value; it contains the subchannel number in its ending halfword. (The first halfword contains X'0001'.)

IOSINFO obtains the number of the subchannel that was associated with the UCB at NIP time. However, the subchannel and the UCB might become disassociated during system operation. Any disassociation of the UCB and the subchannel means the subchannel number in the SID might not be valid. Therefore, IOS returns information consistent with NIP time but does not guarantee that the subchannel will always be associated with the UCB.

If the UCB is disassociated from the subchannel at the time of the IOSINFO macro invocation, IOSINFO can detect the situation and notify the user via a return code. If the UCB is disassociated from the subchannel after the IOSINFO macro invocation, IOSINFO can not notify the caller. For detailed information on how IOSINFO deals with the disassociation of the subchannel from the UCB, see the instructions for coding the macro in Volume 2 of this book.

Dynamic Allocation

The allocation of resources performed in response to JCL at step allocation (or at logon, for time-sharing users) can be altered prior to step deallocation (or logoff) by invoking dynamic allocation functions. A job's device requirements might not be fully evident prior to execution; using dynamic allocation, the program can acquire resources as the need develops. Similarly, each job can use common resources more efficiently with dynamic allocation: the resources can be acquired just before use and/or released immediately after use.

You request dynamic allocation functions by invoking SVC 99.

“Dynamic Allocation” introduces the functions available through SVC 99, along with some concepts and processing features that are special to SVC 99.

In addition, this topic describes installation options you can use to control the processing of all SVC 99 requests.

“Requesting SVC 99 Functions” describes how to request SVC 99 functions, including details on coding the SVC 99 parameter list and information on the return codes issued by SVC 99 functions.

“Requesting SVC 99 Functions” also presents all the SVC 99 text unit keys, in numerical order by key within verb code groups, in reference format. In addition, “Requesting SVC 99 Functions” includes a detailed discussion of the processing involved in dname allocation, and an example of a dynamic allocation request.

The term “resource” means a dname-data set combination, with any attendant volumes and devices.

“Dynamic allocation functions” refers to the allocation of I/O resources *during program execution*, and all of the related functions of resource allocation performed by SVC 99.

To avoid confusion between the specific function of dynamically allocating a resource and the set of functions provided by the dynamic allocation routines, the first is called *dynamic allocation* and the second *SVC 99 functions*, throughout the remainder of the book.

Note: Throughout the remainder of the book, the word “deallocate” is used to denote the action, and the word “unallocated” is used to denote the state. In cases where the common usage is hard to change – for example, in the name of an SVC 99 function – without causing confusion, the words “unallocate/unallocation” have been retained.

Introduction to SVC 99 Functions

When you invoke SVC 99, you can request five different functions. They are:

- Dynamic allocation — acquiring a resource
- Dynamic unallocation — deallocating a resource
- Dynamic concatenation — associating acquired data sets
- Dynamic deconcatenation — separating associated data sets
- Dynamic information retrieval — obtaining certain data set information

A typical use for SVC 99 functions is in a program that needs temporary use of a volume for which there is heavy contention. In such a case, dynamic allocation and dynamic unallocation provide the means for a program to tie up the volume for only as long as necessary rather than for the total execution time of the program.

Another common use for SVC 99 functions is in a program whose need for I/O resources varies according to the input. Dynamic allocation and dynamic unallocation permit such programs to allocate and then free only the files necessary to process the input, so the specific resources supporting the required files can be in use for the minimum time.

You request SVC 99 functions via the DYNALLOC macro. Instead of specifying operands on the macro, you supply information in the SVC 99 parameter list.

You request a specific SVC 99 function via information in two fields of the parameter list:

- **Verb code** — a one-byte hexadecimal code that describes the function being requested.
- **Text Unit Key** — a two-byte hexadecimal value that describes the processing being requested from the SVC 99 function routine. A number of keys are available for each verb code.

For example, verb code 01 with the key of 2 requests that a new data set be allocated, while verb code 02 with key 7 requests that an existing data set be deallocated.

You use other fields in the SVC 99 parameter list to supply the information required to process your request. See Figure 65 for the entire SVC 99 parameter list.

Concepts Needed to Understand SVC 99 Processing

SVC 99 processing involves some features dictated by the environment in which dynamic allocation or deallocation takes place, and by the very fact that the SVC 99 function is performed dynamically — while a program is executing. Because SVC 99 was developed to speed up time-sharing operations, these features might appear to be relevant only to a program running in an interactive environment. However, any program that invokes SVC 99 makes use of the features.

An understanding of the features — and the concepts behind them — is essential to an understanding of the SVC 99 functions.

Processing Control Features

SVC 99 routines provide some controls designed specifically for the unpredictable interactive environment.

Time-sharing command processors use SVC 99 functions to dynamically allocate data sets required for their own processing (for example, work areas), in addition to the data sets the user requests via the time-sharing commands. Because the same command processor can be called again, and different command processors might need the same data sets, the command processors do not deallocate the data sets they allocate for their own use. This avoids allocation processing that would be required when a subsequent command processor requested the same data sets. However, keeping all data sets allocated until the end of a terminal session can also tie up resources that might no longer be needed. The following features help to avoid tying up resources that are not being used:

In-Use Bit and Attribute

An **in-use bit** for each data set is located in the data set association block (DSAB).

The in-use bit is turned on when a data set is dynamically allocated. In a time-sharing environment, the terminal monitor program (TMP) turns off the in-use bits of all data sets that were dynamically allocated by a command processor when that command processor completes execution. (Turning off the in-use bit does **not** deallocate the data set.) If a subsequent command processor dynamically requests a previously-allocated data set, the in-use bit is turned on again until the command processor completes execution and returns control to the TMP.

As part of the in-use feature, the system also keeps track of the data sets that have been “not-in-use” for the longest time.

Control Limit

The **control limit** limits the number of data sets that can be allocated but marked “not-in-use” (that is, the in-use bit is turned off).

This control limit is determined by the JCL parameter DYNAMNBR on the EXEC statement in the logon procedure and the number of DD statements in the logon procedure. If the control limit is exceeded when the SVC 99 routines receive a request for a new dynamic allocation, the routines automatically attempt to deallocate enough data sets to meet the control limit, starting with eligible data sets that have been not-in-use for the longest time.

If the control limit is still exceeded after all eligible resources have been deallocated, the request for a new allocation fails. In this case, you must explicitly request deallocation of an existing allocation before the new allocation can be satisfied.

Note: The control limit would not seem to apply to a program running in a batch environment. However, some utilities (IDCAMS, for example) may expect a control limit, and could cause an 043C error reason code from SVC 99 (see Figure 69). If a batch program requires a large number of allocations, you might need to include the DYNAMNBR parameter on the EXEC statement, specifying a large value (3273 is the system maximum).

Permanently Allocated Attribute

The **permanently allocated attribute** prevents the SVC 99 routines from automatically deallocating a particular data set to meet the control limit. The effect of this attribute is to determine a data set's eligibility for automatic deallocation.

The permanently allocated attribute is automatically assigned to data sets allocated through JCL and the ALLOCATE command. In addition, you can request (via the SVC 99 parameter list) that a data set be assigned this attribute when you dynamically allocate the data set.

Here is an example of the use of the permanently allocated attribute: The ATTRIBUTE command processor assigns the permanently allocated attribute to a dummy data set that is associated with other attributes specified on the command. The dummy data set must not be automatically deallocated if the control limit is exceeded; it must remain available, to be referenced by a USING keyword on an ALLOCATE command, until the user FREES the data set. Therefore, it needs the permanently allocated attribute.

Note: Because permanently-allocated resources are not automatically deallocated, and all resources allocated via the ALLOCATE command and JCL are permanently allocated, the control limit primarily limits the number of resources that a terminal user can have allocated at the same time. If no control limit is established, the limit is 1635.

Because a batch application is more predictable than a time-sharing terminal session, the programmer who is using SVC 99 functions in a batch application probably will not need to be concerned with the in-use bit, the control limit, or the permanently allocated attribute. For an exception, see the note under "Control Limit."

Convertible Attribute

Because a data set requested by a command processor or an application program might already be allocated, the dynamic allocation routines first check for an existing allocation that matches the current request. This avoids redundant allocation processing. In some cases, an existing allocation matches the current request except for some parameters. The dynamic allocation routines can change certain unmatching parameters of the existing allocation to meet the current request if the existing allocation has the **convertible attribute**. The convertible attribute allows the dynamic allocation routines to change the following parameters of the existing allocation:

- Ddname
- Member name
- Status
- Normal and conditional dispositions
- Space
- Deallocation at CLOSE
- Input only
- Output only
- DCB attributes
- Password

In addition, the routines can change the permanently allocated attribute.

The convertible attribute is automatically assigned to all data sets dynamically allocated without the permanently allocated attribute. You can, however, assign both the convertible attribute and the permanently allocated attribute to a resource: although you might want to

prevent a data set from being automatically deallocated, you might also want to allow some of its parameters to be changed to satisfy a new allocation.

Although the setting and use of the convertible attribute are transparent to a program running in a batch environment, the SVC 99 routines make use of this feature. All the resources dynamically allocated on behalf of a program running in a batch environment are automatically assigned the convertible attribute. When, for example, a program in your batch job dynamically allocates a data set for input only and later allocates the same data set for output only, the dynamic allocation routines will automatically avoid redundant allocation processing by using the first allocation and changing “input only” to “output only.”

Functions Available Through SVC 99

When you invoke SVC 99, you request that its routines perform one of five general functions. The following topics discuss those five, and some subdivisions among them, in general terms. “Requesting SVC 99 Functions” presents the individual text unit keys you use to request SVC 99 functions, grouped within the general divisions discussed here.

Dynamic Allocation

You can request one of two types of dynamic allocation: allocation by *dsname* or allocation by *ddname*.

Dsname Allocation

Dynamic allocation by dsname is equivalent to data set allocation during job step initiation; the SVC 99 parameter list is equivalent to a DD statement. In the parameter list, you request the allocation-by-*dsname* function by specifying verb code 01. You can request most of the JCL services that you can code in a DD statement – such as data set disposition, volume label information, expiration date, and SYSOUT destination – by specifying different text units in the parameter list. Figure 64 lists JCL DD statement facilities that *cannot* be used in dynamic allocation.

In addition, you can specify the following, which do not have a JCL equivalent, via SVC 99 text units:

- The password for a password-protected data set. If you specify the password in your program via SVC 99, the system need not prompt the operator.
- The permanently allocated attribute.
- The convertible attribute.
- Return of certain information.

Restricted DDnames	JOB CAT, STEP CAT, JOB LIB, and STEP LIB	
Keyword Parameters	CHKPT, DDNAME, DLM, and DSID	
Positional Parameters	*, DATA, and DYNAM	
Selected Subparameters of Keywords	Keyword	Subparameter Not Supported
	DCB	reference to ddname of a previous step CYLOFL NTM RKP
	DISP	PASS specification
	DSN	reference to ddname (as in *.ddname) ISAM area name
	SPACE	ABSTR specification
	UNIT	AFF
	VOLUME	RETAIN specification REF = ddname

Figure 64. JCL DD Statement Facilities not Supported by Dynamic Allocation

Consult the detailed description of each text unit key (see “Requesting SVC 99 Functions”) for the capabilities supported by the key. While a subparameter may be supported (for example, DCB = DSORG), all values of that subparameter may not be supported (for example, DCB = DSORG = IS).

See “Requesting SVC 99 Functions” for programming considerations for dynamic allocation by dsname.

Ddname Allocation

Ddname allocation allows you to reuse, by specifying only the associated ddname, a previously-allocated data set that was marked not-in-use. Ddname allocation processing sets the in-use bit on.

This type of dynamic allocation request is useful in time-sharing command processors for re-allocating groups of data sets that were allocated and concatenated by an earlier command processor but whose in-use bits were then turned off by the TMP.

In MVS/XA, the HELP command processor uses ddname allocation to allocate the SYSHELP data set. As a user, you may allocate, in your logon procedure, a group of data sets to be searched for HELP information. These data sets are concatenated with the ddname of SYSHELP. When the TMP receives control after the LOGON command processor completes execution, the in-use bits of all the SYSHELP data sets are off. Then, when you issue the HELP command, the HELP command processor invokes the allocation-by-ddname function, specifying SYSHELP as the ddname. As a result, all of the data sets concatenated to SYSHELP will have the in-use bit turned on. (If the system cannot locate the SYSHELP ddname, the HELP command processor uses the allocation-by-dsname function to allocate the SYS1.HELP data set; it will be associated with a system-generated ddname.)

You request dynamic allocation by ddname by specifying verb code 06 and putting the ddname to be allocated in the SVC 99 parameter list. You are, in effect, asking the SVC 99 routines to use a specific existing allocation to satisfy your request.

In order for the SVC 99 routines to satisfy your ddname dynamic allocation request, the existing allocation must not be in use. In addition, it must not have the convertible attribute; or it must be permanently concatenated. In other words, it must have properties that ensure that the ddname could not have been disassociated from the existing allocation. (See “The Permanently Concatenated Attribute” in the topic “Dynamic Concatenation” for a description of this attribute.)

If the existing allocation with the specified ddname does not meet these requirements, or if the ddname is not associated with any of your program’s existing allocations, the request is failed and an error reason code is returned in the SVC 99 parameter list.

If the existing allocation meets the requirements, its in-use bit is turned on and the request has been satisfied. If the existing allocation is a member of a concatenated group, all members of the group are assigned the in-use attribute, so the entire group has been allocated.

You may specify that an indication is to be returned if the existing allocation that satisfies the request is associated with a dummy data set.

Dynamic Unallocation

The SVC 99 dynamic unallocation routines deallocate resources when you request it via the appropriate verb code. Two functions are available through dynamic unallocation:

- **Releasing a data set**, specified by verb code 02 with key 7 in the SVC 99 parameter list.

Releasing a data set involves the following processes:

- Disassociating the ddname from the data set name, which allows the ddname to be used in subsequent dynamic allocations.
- Processing the data set disposition.
- Releasing the data set for use by other jobs except when:
 - A temporary VIO data set is released at the end of the last step in which it is referenced.
 - A temporary non-VIO data set is released at the end of the job.
- Freeing the unit(s) to which the data set was allocated.
- Releasing the volume(s) on which the data set was allocated.

- **Removing the in-use attribute**, specified by verb code 02 with text unit key 8 in the SVC 99 parameter list.

This function turns off the data set’s in-use bit. When this processing is completed, the data set is referred to as “not-in-use.” You can also request remove in-use processing based on task-id by specifying verb code 05; see “Removing the In-Use Attribute by Task-ID.”

If you code verb code 02 without specifying key 7 or key 8, the dynamic unallocation routines *remove the in-use attribute* from data sets allocated through JCL, the time-sharing ALLOCATE command, or dynamically with the permanently allocated option; the routines *release* data sets that were allocated dynamically without the permanently allocated option.

You use key 7 and key 8 to specify explicitly the type of processing you prefer. An explicit specification will be satisfied in all but one case: the routines will not remove the in-use attribute from a non-permanently allocated, non-*&dsname* data set with a disposition of DELETE. Such a resource cannot be used to satisfy a subsequent request, so it will be released.

You can request either dynamic unallocation function for a *dsname* or a *ddname*.

Dynamic Unallocation Processing: The following considerations apply to all dynamic unallocation requests:

- If a data set is open, is a member of an open concatenated group, or is a private catalog, it will not be deallocated.
- When a SYSOUT data set is released, it is immediately made available for output (unless you specify an overriding disposition of DELETE, in which case the data set is deleted).

The following considerations apply to unallocation requests specifying a *dsname*:

- If no *ddname* is specified and the *dsname* is associated with more than one *ddname*, all associated data sets are deallocated. If an error occurs while deallocating one *ddname*, processing continues for the others and an error code is returned in the SVC 99 parameter list. If errors occur for more than one *ddname*, the error code applies to the last *ddname* for which there was an error.
- If a member name is specified with the *dsname*, only those associations containing both the member name and *dsname* are deallocated. (Both the member name and *dsname* text unit keys must be coded for a valid request.)

The following considerations apply to unallocation requests specifying a *ddname*:

- Only the occurrence of the data set associated with the specified *ddname* is deallocated, even if that data set is associated with other *ddnames*.
- If a *dsname* or *dsname* and membername are specified in addition to the *ddname*, they must be associated with that *ddname* or the request fails.

Deallocating Concatenated Groups: If the specified resource is associated with a permanently concatenated group, the in-use attribute is removed from all members of the group, and the count of the number of resources held for reuse is increased by the number of members in the group. (An exception occurs when the concatenated group was generated by the system, as, for example, VSAM data sets spanning device types and GDG ALL groups. In these cases, the group is treated as a single resource.)

If the concatenated group does not have the permanently concatenated attribute, the group is deconcatenated and the member associated with the specified *dsname* is released. (The first member is released if the group's *ddname* is specified.)

If a concatenated group has the permanently concatenated attribute and you specify a *ddname* with a *dsname*, a VSAM *dsname*, or GDG ALL, the entire group is released. If you specify a *dsname* with a VSAM *dsname* or GDG ALL, the request for dynamic unallocation fails.

Changing Parameters at Dynamic Unallocation: You can include with your dynamic unallocation request text units to change a data set's parameters as it is being deallocated. If your request is in the form of verb code 02, key 8, the changes are honored when the data set is actually released, unless they have been overridden in the meantime. If your dynamic unallocation request is in the form of verb code 02, key 7, the changes take effect immediately.

The parameters that can be changed at deallocation are:

- Output class
- HOLD/NOHOLD parameters
- Remote work station destination
- Disposition

Allocation disposition cannot be overridden for passed data sets, VSAM data sets, or system-named data sets. For all other types of data sets, the disposition specified on an unallocation request overrides the disposition specified at allocation.

If you request that a data set be cataloged on an OS CVOL when it is deallocated, you must provide linkage to the CVOL. The system will not dynamically allocate the OS CVOL.

Members of partitioned data sets cannot be deleted with a disposition of DELETE; the entire data set is deleted. An overriding disposition of DELETE for data sets allocated as shared is invalid; the overriding disposition request is failed.

Removing the In-Use Attribute by Task-ID

In addition to requesting removal of the in-use attribute by specifying a ddname or dsname, you may request, via verb code 05, that the in-use attribute be removed based on task-ID. The attribute may be removed from all resources associated with a specified task, or all resources except those associated with the current task, its higher-level tasks, and the initiator. This function is used by, for example, the time-sharing terminal monitor program (TMP) to turn off the in-use bits of any data sets allocated by a command processor when the command processor completes execution.

Dynamic Concatenation

Dynamic concatenation logically connects allocated data sets into a concatenated group. You request this function by specifying verb code 03 in the SVC 99 parameter list. You can only identify data sets to be concatenated by their associated ddnames. These data sets must not be open; if they are, the request for dynamic concatenation fails.

The order in which you specify the ddnames is the order in which the SVC 99 routines will concatenate their associated data sets. The name associated with the concatenated group is the ddname that was specified first; the other ddnames are no longer associated with any data set.

If a ddname you specify is already associated with a concatenated group, that entire group will be included in the new concatenation.

After the request for dynamic concatenation is satisfied, all members of the dynamically-concatenated group are assigned the in-use attribute.

The Permanently Concatenated Attribute

You can request that a concatenated group created via SVC 99 be assigned the permanently concatenated attribute. A concatenated group defined via JCL is automatically assigned the permanently concatenated attribute, as is a concatenated group defined by the system via JCL or SVC 99. A GDG ALL request and a request for a VSAM data set that spans device types are examples of the latter situation.

A group with the permanently concatenated attribute has the following characteristics:

- The group cannot be dynamically deconcatenated into its member data sets.
- If a permanently concatenated group is dynamically concatenated with other data sets to form a new non-permanently concatenated group, the permanently concatenated group remains intact if the new group is dynamically deconcatenated.
- If the group is not a system-defined permanently concatenated group, it is automatically assigned the permanently allocated attribute.

Note: To dynamically release a non-system-defined permanently concatenated group, you specify the ddname, not the dsname, in the unallocation request.

Dynamic Deconcatenation

Dynamic deconcatenation logically disconnects the members of a concatenated group. You request dynamic deconcatenation by specifying verb code 04 in the SVC 99 parameter list. You identify the concatenated group to be deconcatenated by specifying the ddname of the group.

The request for dynamic deconcatenation fails if the concatenated group is open. A permanently concatenated group, or members of a concatenated group that are permanently concatenated, remain concatenated.

When a concatenated group is dynamically deconcatenated, the ddnames that were associated with the data sets before they were concatenated are restored unless this would result in duplicate ddnames. This situation could arise if a dynamic allocation with the ddname to be restored occurred after a dynamic concatenation. In this case, the deconcatenation request fails.

Dynamic deconcatenation has no effect on the in-use attributes associated with the members of the group.

Dynamic Information Retrieval

Dynamic information retrieval provides you with information about your current allocation environment. You request this function by specifying verb code 07 in the SVC 99 parameter list. You can request information about ddnames or dsnames.

In addition, you can ask for information about any or all of your currently-allocated requests by specifying a relative request number. For example, you could obtain information about all your allocation requests by successively asking for information about the 1st, 2nd,...nth allocation request. Code the DINRTLST text unit key (key 13) with this series of requests, to receive an indication of the *last* relative entry.

You can request the following kinds of information using SVC 99 verb code 07:

- Data set name
- Ddname
- Member name
- Data set organization
- Status
- Normal disposition
- Conditional disposition
- Attribute status, including the permanently allocated, in-use, permanently concatenated, convertible, and dynamically allocated attributes
- Data set types, including dummy, SYSIN, SYSOUT, and allocation of the user's terminal as an I/O device
- The number of resources held in anticipation of reuse that exceeds the control value; that is, the number of existing allocations that must be deallocated before a request for a new allocation can be satisfied
- Whether or not the allocation is the last relative request.

Installation Options For SVC 99 Functions

This section describes the values and options your installation might want to modify in order to control SVC 99 processing. The values and options discussed in the following topics are:

- Default values for space and unit information
- Default values for the TIOT
- Mounting volumes and bringing devices online
- Installation validation routines.

Space and Unit Defaults

This section describes how to change the allocation default values for space and unit. It also describes remote user workstation and TIOT defaulting.

Existing Default Values for Space: If no space information is specified on a request for a new direct access data set, and the request is eligible to MSS exclusively (with MSVGP specified), the SVC 99 routines use the MSVGP defaults.

If the request is not eligible to MSS exclusively, or MSVGP is not specified, the defaults are:

- Block length of 1000
- 10 primary blocks and 50 secondary blocks
- Release unused space (RLSE) specification

These space defaults are contained in the allocation default CSECT IEFAB445 (load module IEFAB445). The contents of the module are as follows:

Note: Beginning at offset zero; the hexadecimal numbers in parentheses show the values supplied by IBM.

- Three bytes for the primary quantity value (X'00000A')
- Three bytes for the secondary quantity value (X'000032')
- Three bytes for the average block length (X'0003E8')
- Three bytes for the number of directory blocks (X'000000')
- One byte of flags with the following bit meanings:
 - bit 0: TRK (0)
 - bit 1: CYL (0)
 - bit 2: blocklength (1)
 - bit 3: RLSE (1)
 - bit 4: CONTIG (0)
 - bit 5: MXIG (0)
 - bit 6: ALX (0)
 - bit 7: ROUND (0)

Existing Default Values for Unit: If a time-sharing user's dynamic allocation request does not include unit information, the SVC 99 routines obtain a unit description from the UADS entry.

If the user is not a time-sharing user, or if the UADS entry does not contain a unit description, 'SYSALLDA' (the system's esoteric name indicating all direct access devices) is the default. This default is contained in the allocation default CSECT IEFAB445, in the eight bytes beginning at offset 13 (decimal).

The unit description you supply in your dynamic allocation request can override the unit type for a cataloged data set. The unit description from the UADS, however, cannot override the unit information in the catalog.

Default Value for a Remote User Work Station: If a remote user work station – or destination – is not specified by a time-sharing user allocating a SYSOUT data set, the SVC 99 routines obtain a default value from the UADS entry.

Default Value for The TIOT: The size of the TIOT is determined by the DEFTIOTS field in the allocation default CSECT (IEFAB445). The IBM-supplied default for the size of the TIOT is 32K. An installation may change that to a larger (maximum 64K) value or a smaller (minimum 16K) value. This default is contained in the allocation default CSECT IEFAB445, in the eight bytes beginning at offset 23 (decimal).

The size of the TIOT controls how many DDs are allowed per jobstep. By specifying any integer from 16 to 64 as the value of the DEFTIOTS field, the user controls the DD allowance. The following table shows the relationship between the size of the TIOT and the maximum number of DDs allowed:

DEFTIOTS		Size of TIOT	Maximum number of DDs allowed
Dec	(Hex)		
16	10	16384 (16K)	816
24	18	24576 (24K)	1225
32	20	32768 (32K)	1635
40	28	40960 (40K)	2045
48	30	49152 (48K)	2454
56	38	57344 (56K)	2864
64	40	65536 (64K)	3273

Mounting Volumes and Bringing Devices Online

Dynamic allocation processing can bring devices online and have volumes mounted.

This function is optional for time-sharing users, because it is time-consuming and requires operator intervention; it is not always desirable in an interactive environment. If selected, the option is assigned via the UADS entries.

Other users of SVC 99, however, can always have volumes mounted and devices brought online. If you do not want this function performed during dynamic allocation, you can indicate in the SVC 99 parameter list that volumes are not to be mounted and that devices are not to be brought online for a request.

In addition, the operator may inform the dynamic allocation routines that a volume is not to be mounted or that a device is not to be brought online. If the operator thus prevents the mounting or bringing online of a volume or device, the allocation request fails.

If you allow volume mounting, the SVC 99 routines will wait for tape volumes to be mounted. Common allocation processing, by contrast, does not wait for tape volumes to be mounted. When a volume is mounted for a dynamic allocation request, OPEN verifies that it is the correct volume.

If the option to have volumes mounted and devices brought online is not in effect, tape and direct access devices that have an outstanding mount request, or that are not ready, are not eligible for use by dynamic allocation.

Installation Input Validation Routine for SVC 99

An exit (IEFDB401) from the allocation control routine provides for a user-written routine to validate or alter any request to SVC 99. The routine is entered for all system and user SVC 99 requests. You must code it so it does not interfere with system requests.

Your validation routine can test and modify the SVC 99 input request, and it can indicate through a return code whether processing of the request is to continue. For example, the routine might perform the following functions:

- Control the amount of direct access space requested
- Check for authorization to use specified units
- Check for authorization to use certain data sets
- Check for authorization to hold certain resources for reuse

See *SPL: User Exits* for information about programming conventions that the input validation routine must observe.

Requesting SVC 99 Functions

To request an SVC 99 function, you must code the DYNALLOC macro (it has no operands) and supply the SVC 99 parameter list. The SVC 99 routines perform requested functions based on the information you provide in the parameter list, in the form of text unit keys.

This topic consists primarily of descriptions of the various text unit keys you use to request SVC 99 functions. It also includes a discussion of the parameter list itself, a detailed review of the process of dynamic allocation by dsname, the SVC 99 return codes, and some general considerations for using SVC 99. An example of a dynamic allocation request is provided at the end of this topic.

Programming Considerations When Using SVC 99

Before deciding to use any of the SVC 99 functions, you should consider the environment of the program that invokes SVC 99. Your program interacts with the job entry subsystem, with the MVS/XA initiator, with data management functions, and with the common allocation functions of the MVS/XA scheduler, in addition to the SVC 99 routines themselves. If the program is a system routine, there are cross-memory considerations, as well.

You must make sure your routine does not interfere with the normal functioning of the operating system. Conversely, you should be aware that operating system functions could affect the interface between your program and the SVC 99 routines; and that the SVC 99 routines could affect the functioning of your program.

Following are some of the things you need to consider when preparing to use SVC 99. Other considerations are included with the topics to which they apply.

- **Serialization of Resources**

Your program might serialize the same resources as SVC 99. The SVC 99 routines can serialize the following resources, depending on the path taken in SVC 99 processing.

Major Name	Minor Name
SYSDSN	data set name
SYSIEFSD	CHNGDEVS
SYSIEFSD	DDRDA
SYSIEFSD	DDRTPUR
SYSIEFSD	Q4
SYSZOPEN	data set name
SYSZPCCB	PCCB
SYSZTIOT	address of the DSAB QDB.asid
SYSZVMV	ucbaddr
SYSZVOLS	volume serial number

- **Other System Routines and SVC 99**

System routines invoked by various paths of SVC 99 processing might also serialize a system resource. Some of the system functions invoked by SVC 99 processing are LOCATE, OBTAIN, CATALOG, SCRATCH, DADSM Allocate, and MSS Interface (SVC 126).

- **No Cross Memory Support for SVC 99**

SVC 99 does not work in cross memory mode. For further information on the use of cross memory services, see "Synchronous Inter-Address Space Communication."

- **Enqueuing on the SYSZTIOT; Avoiding 138 Abends**

Avoid requesting SVC 99 functions in routines that run under the control of an interruption request block (IRB), especially when the program that might be interrupted issues OPEN, OPENJ, CLOSE, EOVS, or FEOVS, or any other SVC that enqueues on the SYSZTIOT. An SVC 99 request issued in such an environment can cause a 138 abend when SVC 99 tries to enqueue on the SYSZTIOT resource.

For the same reason, user exits for OPEN/CLOSE/EOVS, or for any other routines that enqueue on SYSZTIOT, should not issue SVC 99 requests.

- **Avoiding 0B0 Abends**

Programs that issue SVC 99 should not receive control during START (initialization) processing for LOGONS, MOUNTS, or started tasks.

Programs that get control during START processing (user exits, for example) should not issue LOCATE, OPEN, OBTAIN, CATALOG, SCRATCH, or DADSM Allocate for data sets that have not been preallocated to the program; to do so will cause an 0B0 abend.

Subsystems that receive control during step allocation as a result of the SUBSYS parameter should not issue SVC 99; to do so might cause an 0B0 abend.

- **Accessing ICF CATALOGS, CVOLs and VSAM Private Catalogs**

Programs that get control during 'START' (Initialization, User Exits, for example) should not issue LOCATE, OPEN, OBTAIN, CATALOG, SCRATCH, or DADSM allocate for data sets that have not been preallocated to the program; to do so could cause MSGIEC331I, RC4, RC84.

- **JES Consideration**

The program that requests SVC 99 functions must not receive control when the job entry subsystem is being started; an unending wait could result, causing the system to crash.

- **Effect of Outstanding STIMER**

The program that requests SVC 99 functions must not have an STIMER macro outstanding, because certain paths in SVC 99 issue another STIMER, which causes an overlay of the first STIMER. This situation causes the program that issued the SVC 99 request never to receive control because of the expiration of the timer.

- **Considerations for System Routines**

System routines that invoke SVC 99 functions should be aware that a non-zero return code might be returned. Because system routines cannot always diagnose non-zero return codes, the system routine should print an error message including the error code (S99ERROR) and information code (S99INFO) fields from the SVC 99 request block (S99RB).

- **Accessing CVOLs or VSAM Private Catalogs**

Routines should not allocate data sets that are cataloged in OS CVOLs or VSAM private catalogs to long-running tasks, because the private catalog or CVOL will be allocated and will remain allocated until the step terminates. This is especially important in installation exits for system tasks, because the private catalog or CVOL might be allocated to an initiator or subsystem such as JES2 or JES3.

Volumes that contain a CVOL or VSAM private catalog for data sets allocated to long-running steps should be assigned the permanently resident attribute.

- **Changes to the TIOT By SVC 99 Routines**

SVC 99 routines might cause changes to the task input/output table (TIOT). Depending on the function requested via SVC 99, an entry could be added, deleted, or reordered; you cannot assume a fixed order for TIOT entries.

You should make sure a problem program is aware of changes to the TIOT, especially if the EXTRACT macro is also being used within that program, or in the program that will gain control when SVC 99 processing is finished.

If you need to reference TIOT entries after SVC 99 is invoked, do so via the data set association block (DSAB) chain, which is pointed to by the DSAB queue descriptor block (DSAB QDB) field, JSCDSABQ, in the active JSCB.

Datasets that are cataloged in a private catalog or CVOL should not be dynamically allocated in the JES3 address space. If they are dynamically allocated, a system deadlock may occur involving SYSZTIOT, which is required for the SVC 99, and SYSZPCCB, which is required in the modification of the private catalog control block (PCCB).

SVC 99 Parameter List

When you code DYNALLOC, you must supply a parameter list. The SVC 99 parameter list includes a request block, text pointers, and text units.

The request block indicates the function you want the SVC 99 routines to perform; the text pointers contain the addresses of the text unit keys; and the text units contain the keys and parameters for the SVC 99 functions. Figure 65 illustrates the structure of the SVC 99 parameter list.

IBM supplies two macros, IEFZB4D0 and IEFZB4D2, to aid in constructing the SVC 99 parameter list. IEFZB4D0 provides symbolic names (DSECTs) for the positional information in the structure; IEFZB4D2 provides mnemonics for the text unit keyword values. The names in Figure 65 are those assigned by the macro IEFZB4D0.

On entry to SVC 99, register 1 must point to a pointer to the request block.

The pointers to the parameter list, the request block and the text units must be created in storage with the same key as the one the caller of SVC 99 is running in.

In addition, the request block and any information retrieval text units must be in non-store-protected storage. This is to prevent an 0C4 abend when SVC 99 copies the caller's parameters into its own work area and then restores the information retrieval text units into the caller's storage.

Note: SVC 99 only restores information retrieval text units originally specified by the caller and validated by the IEFDB401 exit routine (see *SPL: User Exits*). Text units added or modified by IEFDB401 are not copied back to the calling program's storage.

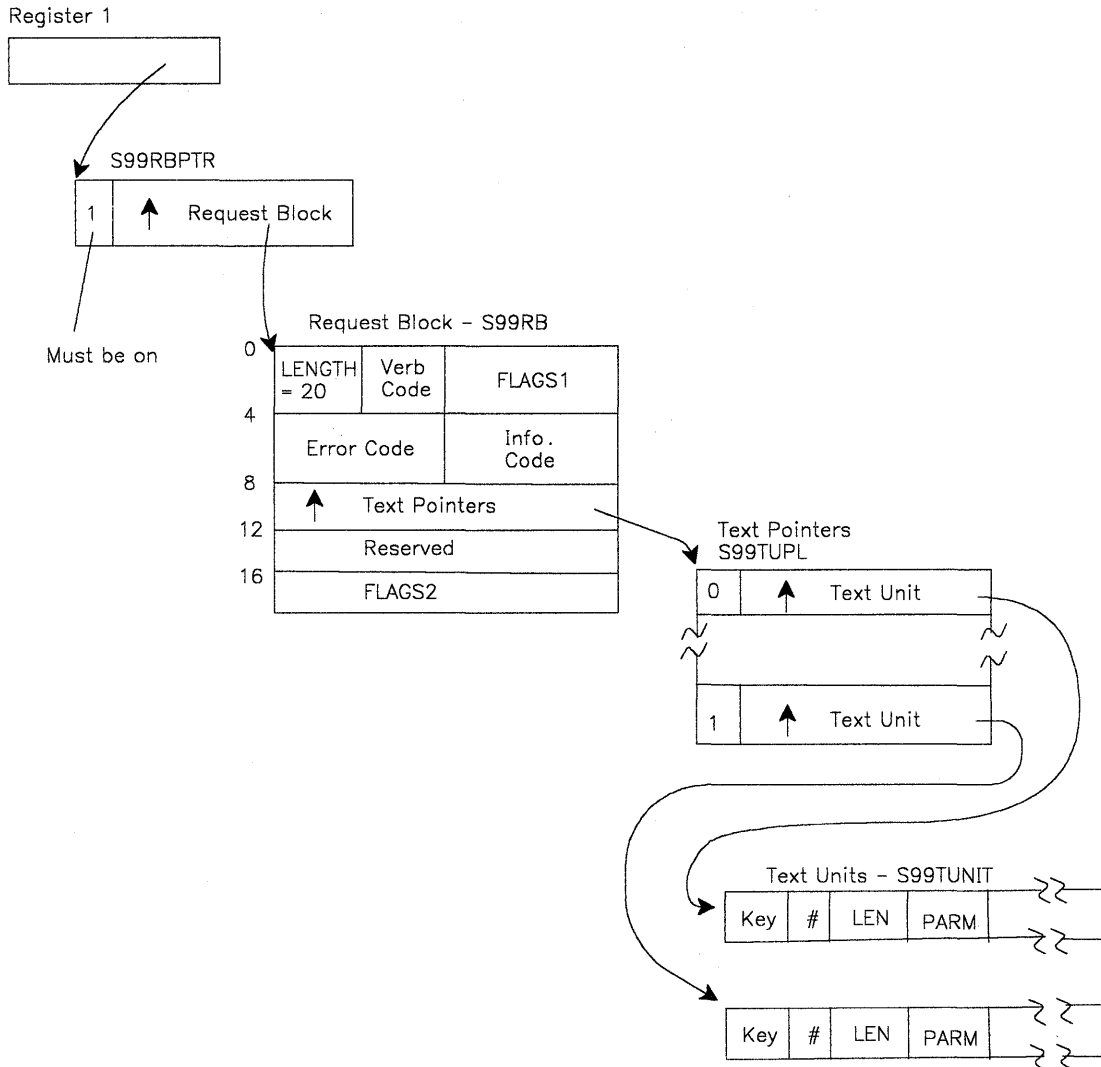


Figure 65. Structure of the SVC 99 Parameter List

Request Block Pointer

The request block pointer, DSECT name S99RBP, is a single fullword containing the address of the SVC 99 request block. IEFZB4D0 assigns the label S99RBPTR to the address. The high-order bit in this field must be set to one.

Request Block

The request block must begin on a fullword boundary. Mapping macro IEFZB4D0 assigns it a DSECT name of S99RB. It contains the following fields (the names in parentheses are those assigned by IEFZB4D0):

- **LENGTH (S99RBLN)** A one-byte field containing the length of the request block. The length is always 20 bytes.
- **VERB CODE (S99VERB)** A one-byte field that identifies the SVC 99 function to be performed. You may specify the following verb codes:

Verb Code	Name	Meaning
01	S99VRBAL	Request for dsname allocation
02	S99VRBUN	Request for deallocation (based on dsname or ddname)
03	S99VRBCC	Request for concatenation
04	S99VRBDC	Request for deconcatenation
05	S99VRBRI	Request for removing the in-use attribute based on task-ID
06	S99VRBDN	Request for ddname allocation
07	S99VRBIN	Request for information retrieval

- **FLAGS1 (S99FLAG1)** A two-byte field that instructs the system on how to satisfy dsname allocation requests. The meaning of the bits in the field are as follows:

Bit	Bit Name	Meaning When On
0	S99ONCNV	Only use an existing allocation that has the convertible attribute to satisfy the request.
1	S99NOCNV	Do not use an existing allocation to satisfy this request.
2	S99NOMNT	Do not mount volumes or consider offline devices. (This bit overrides S99MOUNT and S99OFFLN in FLAGS2.) If this bit is one and the request causes a private catalog to be allocated, mounting will not be allowed for that catalog.
3	S99JBSYS	Treat the data set as part of the job's normal output. The data set is not expected to be dynamically deallocated (spun off). This flag is used for SYSOUT data sets. If the data set is dynamically deallocated, it will be printed immediately, but paging space will not be released until the job ends.
4	S99CNENQ	Issue a conditional ENQ on the TIOT resource. If the TIOT is not available, an error code is returned to the user.
5-16		Reserved; must be zero

Note: The FLAGS1 indicators (except S99CNENQ) are used only for dsname allocation requests.

- **ERROR CODE (S99ERROR)** A two-byte field that SVC 99 uses to return error reason codes. See "SVC 99 Return Codes."

- **INFO CODE (S99INFO)** This two-byte field is used by SVC 99 to return information reason codes. See "SVC 99 Return Codes."
- **TEXT POINTERS ADDRESS (S99TXTPP)** A fullword field containing the address of a list of pointers to the text units.
- **RESERVED** An area, one fullword in length, that contains zeros.
- **FLAGS2 (S99FLAG2)** A four-byte field of indicators. These indicators may be set only by authorized programs. To be authorized, the requesting program must meet at least one of the following criteria:
 - It must have a system storage protection key (0-7).
 - It must be in supervisor state.
 - It must be APF authorized.

The meanings of the FLAGS2 bits are:

Bit	Bit Name	Meaning When On
0	S99WTVOL	Wait for volumes.
1	S99WTDSN	Wait for dsname. (See Note 2.)
2	S99NORES	Do not reserve data sets.
3	S99WTUNT	Wait for units. (See note 3.)
4*	S99OFFLN	Consider offline devices. The system ignores this bit if S99NOMNT in FLAGS1 is on. For a batch user, or if a time-sharing user has the mount attribute in his UADS entry, the system proceeds as if this bit were on, regardless of the setting.
5	S99TIONQ	TIOT ENQ already performed.
6	S99CATLG	Set special catalog data set indicators.
7*	S99MOUNT	Volumes may be mounted. The system ignores this bit if S99NOMNT in FLAGS1 is on. For a batch user, or if a time-sharing user has the mount attribute in her UADS entry, the system proceeds as if this bit were on, regardless of the setting.
8	S99UDEVT	Unitname parameter for DALUNIT is a device type. If you are using the output from the DEVTYPE macro, be sure the shared DASD bits are turned off.
9	S99PCINT	Allocate a private catalog on behalf of the initiator.
10-31		Reserved. Must be zero.

*These fields override the NOMOUNT option from the TSO user attribute data set (UADS).

Notes:

1. *The FLAGS2 indicators (except S99TIONQ) are used only for dsname allocation requests.*
2. *In a JES3 environment, authorizing a dynamic allocation request to wait for data set availability might cause a system interlock.*
3. *Authorizing a dynamic allocation request to wait for currently allocated devices may cause the job to hang if the only devices available are already allocated to this job.*

Text Pointers

The text pointer part of the parameter list is a variable-length list of fullwords containing pointers to the text units. You indicate the end of the list by setting the high-order bit of the last pointer to one. A fullword of zeros is ignored.

Mapping macro IEFZB4D0 assigns the DSECT name S99TUPL to the list, the label S99TUPTR to each pointer in the list, and label S99TUPLN to an equate that allows you to turn on the end-of-list indicator.

Text Units

Each SVC 99 text unit is a variable-length field (assigned the DSECT name S99TUNIT by macro IEFZB4D0) that contains the following subfields:

- **KEY (S99TUKEY)** A two-byte field containing a unique binary number that identifies the type of information to be found in the PARM subfield. For example, a key of '0004' for a dsname allocation request indicates that the value of the PARM subfield specifies data set status. SVC 99 ignores a KEY field of zero. See "Text Units by Function" for a description of the text units that can be coded for each SVC 99 function.
- **NUMBER (S99TUNUM)** A two-byte binary number specifying the number of length and parameter combinations in the text unit. If a key of zero is specified, S99TUNUM must also be zero.
- **COMBINATION (S99UENT)** The label for length and parameter combinations. IEFZB4D0 provides a separate DSECT (named S99TUFLD) for use when specifying multiple parameters in a single text unit. This DSECT places the length field at displacement 0 for the second and subsequent combinations:

S99TUFLD	Label for the DSECT
S99TULEN	Label for the length field
S99TUPRM	Label for the parameter
- **LENGTH (S99TULNG)** A two-byte binary number specifying the length of the following parameter field.
- **PARAM (S99TUPAR)** A variable-length field in which you put the parameter information identified by the value in the KEY field. See "Text Units by Function" for a description of the values you can code for each text unit key.

The following notes apply to the structure of the text units; you will find rules for coding specific text units in "Text Units by Function."

Notes:

1. *Special characters – of the type requiring apostrophes in JCL statements – are not valid in PARM values, except in the DALUSRID text unit.*
2. *Parameters whose values consist of alphameric and national characters may include trailing blanks.*
3. *The text units may be in any order.*

4. *Each function of SVC 99 has an associated set of text units, and each set is independent of any other. For example, the functions of allocation and unallocation may both use a KEY value of '0007', but that value does not necessarily have the same meaning for both functions.*

Detailed Review of Dsname Allocation Processing

The major function performed by the SVC 99 routines, and the function most often requested of them, is that of dynamically allocating a data set/resource according to its data set name (dsname). Following is a detailed discussion of the processing the SVC 99 routines perform in satisfying dsname allocation requests.

When you invoke SVC 99 to perform dsname dynamic allocation, an “allocation environment” already exists for your request. It consists of the allocation requests made via your JCL or internal dynamic allocation, that have not yet been deallocated. The system considers these resources to be **existing allocations**, and goes to them first to fill your SVC 99 requests.

For dsname allocation, the SVC 99 routines first check for environmental conflicts by noting the types of resources that are currently available to the task that includes your program. The routines then try to satisfy the request with an existing allocation that matches or can be made to match the request. (See the discussion of the convertible attribute in the “Dynamic Allocation” topic.) If the routines cannot make the match, they proceed with a new allocation. If an existing allocation can be used, much allocation processing is avoided.

Checking for Environmental Conflicts

The SVC 99 routines cannot satisfy a dsname allocation request that is in conflict with your existing allocation environment. Following is a list of the environmental conflicts that can cause your request to fail:

- The specified ddname is associated with an existing allocation that is in use.
- The specified ddname is associated with one of a group of concatenated data sets defined as permanently concatenated. (For a definition of permanently concatenated, see “The Permanently Concatenated Attribute.”)
- The request specifies a new non-temporary data set with the same dsname as that of an existing allocation. (This is not a conflict if the request specifies a different volume serial number.)
- A status of OLD or SHR is specified for a dsname associated with an existing allocation that is not permanently allocated, not in-use, and has a disposition of DELETE. (This is not a conflict if the request specifies a different volume serial number.)
- The specified ddname is associated with an existing allocation that does not have the convertible attribute or that does not fulfill the conditions listed under “Using an Existing Allocation.”

Using an Existing Allocation

If possible, the SVC 99 routines will use an existing allocation — an allocated resource marked not-in-use — to satisfy your dsname allocation request. Although some parameters can be changed if necessary, the request and the existing allocation must match according to several criteria before the allocation can be selected to satisfy your request.

In order to be satisfied by an existing allocation, your request must be one of the following:

- A request for an explicit data set name (dsname), or
- A request for the allocation of your terminal as an I/O device, or
- A request for a dummy data set

In order to be satisfied by an existing allocation, your request must **not** specify any of the following:

- Data set sequence number
- DCB reference
- Label type
- Parallel mounting
- Private volume
- Unit count
- Unit description (If the dsname is in the form “&dsname,” the unit name description is ignored.)
- Volume count
- Volume reference
- Volume sequence number

Note: MSVGP is ignored if an existing allocation is used.

In order to be used to satisfy your request, the data set that is the existing allocation must have the following properties:

- It must not be in use.
- It must not be a member of a concatenated group.
- It must have the same volume serial number as any explicitly specified in the request.
- It must have the permanently-allocated attribute, if its disposition is DELETE and the request specifies a status of MOD.
- It must not be a generation data group data set.

- It must either have the convertible attribute or, if the request is in a form other than “&dsname,” **all** of the following must be true:
 - The request does not specify a ddname; or the specified ddname matches the ddname associated with the existing allocation. A terminal request that does not specify a ddname cannot be satisfied by an existing allocation that does not have the convertible attribute.
 - For partitioned data sets, the member name specified in the request is the same as the member name associated with the existing allocation; or a member name is neither specified in the request nor associated with the existing allocation.
 - The request does not specify input only, output only, or any DCB parameters.
 - If a status of MOD is specified in the request, MOD is also associated with the existing allocation; or it is neither specified in the request nor associated with the existing allocation.
 - The request does not specify that the convertible attribute be assigned to the allocation.
 - The request does not specify that only existing allocations with the convertible attribute may be used.

If the request specifies dsname in the form “&dsname,” only the first item need be true.

Even with all the restrictions listed here, more than one existing allocation could match your dsname request. Then, if you specified a ddname and one of the matching existing allocations is associated with that ddname, that is the allocation that the SVC 99 routines select to satisfy your request.

If you did not specify a ddname, the SVC 99 routines select the matching existing allocation whose in-use bit was most recently turned off. (Data sets allocated via JCL are considered to have had their in-use attributes removed at step allocation.)

It could happen that an existing allocation does not match your request even though it is associated with the same ddname you specify. Since the ddname is going to be associated with the resource that is allocated to your program, the system gives the rejected allocation a new ddname, of the form ‘SYS’ followed by five digits. The association of a system-generated ddname with an existing allocation cannot occur in the following cases:

- The existing allocation is in use.
- The existing allocation is open.
- The existing allocation does not have the convertible attribute.
- The existing allocation is associated with a permanently concatenated group that does not represent an entire generation data set group or a multi-device-type VSAM data set.

Changing the Parameters of an Existing Allocation

When the dynamic allocation routines use an existing allocation to satisfy a dsname allocation request, some of the parameters of the existing allocation might have to be changed to reflect the parameters specified in the request. Only existing allocations that were dynamically allocated, with the convertible attribute, can have their parameters changed. Resources allocated via JCL or the TSO ALLOCATE command cannot have their parameters changed (with the exception of status and disposition specified via JCL), but they may be used if no changes are necessary.

Note that, if your request does not specify the permanently allocated attribute, the allocated resource is automatically assigned the convertible attribute. An allocation request may specify both the permanently allocated and the convertible attributes in its parameter list.

The following parameters are eligible for change by the SVC 99 routines:

- Ddname
- Member name
- Status
- Normal disposition
- Conditional disposition
- Space
- Unallocation at CLOSE
- Input only
- Output only
- DCB attributes
- Password
- Permanently allocated attribute

No other parameters may be changed.

Notes:

1. *You cannot change an exclusive status to shared status. For example, you cannot change OLD to SHR. However, it is possible to change SHR to OLD if no other jobs are enqueued on the requested data set, or the next job enqueued on the data set has requested exclusive use of the data set.*
2. *You cannot change the parameters on an explicitly-referenced OUTPUT JCL statement (DALOUTPT).*

Using a New Allocation

The dynamic allocation routines attempt a **new allocation** when they cannot satisfy your request with an existing allocation. New allocations cannot be processed by the dynamic allocation routines while a job step holds (for possible reuse) more dynamically allocated resources than permitted. The number of allocated resources permitted is determined in two ways:

- The maximum number of concurrent allocation requests allowed per job step or time-sharing terminal session is 1635. This includes both dynamic requests and requests made via JCL.

- The control limit set by the DYNAMNBR parameter on the JCL EXEC statement plus the number of DD statements limits the number of resources that can be held for reuse (that is, resources that are allocated but whose in-use bits are off).

When the maximum number of resources have been allocated and you request additional allocations, the dynamic allocation routines automatically attempt to deallocate enough resources to meet the control limit.

Automatic Unallocation of Resources Held for Re-use

The only resources eligible for automatic deallocation are those that were allocated dynamically without the permanently allocated attribute and whose in-use bit has been turned off. (Resources allocated through JCL and through the time-sharing ALLOCATE command are not eligible because they automatically have the permanently allocated attribute.)

When many resources are eligible for automatic deallocation, the dynamic allocation routines choose those that have been designated as not-in-use for the longest time. These are deallocated and the new allocation is processed.

If the control value is still exceeded after all eligible resources have been deallocated, the request for a new allocation fails. In this case, you must explicitly request deallocation of an existing allocation before the new allocation can be performed.

Considerations When Requesting Dsname Allocation

- If you do not specify a ddname, the system generates one. The ddname created consists of the characters 'SYS' followed by five digits.
- You may specify passwords as part of a dynamic allocation request to bypass prompting the operator.
- If you allocate a data set with a status of MOD but do not specify any volume information, and the data set cannot be found in the catalog, it is treated as a new data set.
- If you specify a normal disposition of CATLG for a new direct access data set, the system catalogs the data set when it is allocated rather than when it is deallocated. If the data set cannot be cataloged, then no allocation will take place; if the data set cannot be allocated, it will not be cataloged.
- Rather than wait for another user to release a data set, volume, or device in order to obtain use of it, the dynamic allocation routines fail a request by an unauthorized program. If an authorized program specifically requests a wait, the routines will wait.
- You can request that the ddname, data set name, and volume serial number assigned by the allocation routines be returned in the SVC 99 parameter list.

- You can also request that the data set organization (DSORG) of the allocated data set be returned in the SVC 99 parameter list. The SVC 99 routines return whatever you specify as the DSORG, if any. If you do not specify a DSORG on the allocation request, the system assigns and returns data set organizations according to the following defaults:
 - If the allocation request is for a terminal as an I/O device or for a SYSOUT data set, 'PS' (physical sequential) is returned as a default value.
 - If the allocation request is for a tape data set, 'PS' is returned as a default value.
 - If the allocation request is for a NEW direct access data set, 'PO' (partitioned organization) is returned if you specified a directory space quantity; otherwise, the data set is assigned the DSORG of 'PS'.
 - If the allocation request is for an existing direct access data set, the data set organization obtained from the data set control block (DSCB) is returned. If the DSORG cannot be obtained from the DSCB, the allocation request is failed.
 - For other types of allocation requests where you do not specify a DSORG for the data set, the system returns zeros in the SVC 99 parameter list field.
- For time-sharing users allocating new data sets using the TSO ALLOCATE command, DSORG is defaulted to partitioned organization (PO) if a directory quantity is specified, or to physical sequential (PS) otherwise.
- You cannot create ISAM or VSAM data sets through dynamic allocation.
- If you request an allocation by dsname and the dsname is already allocated but not available, the SVC 99 routines copy the unit and volume information from the existing allocation and associate the information with your request.
- An allocation of a GDG data set will refer to the same data set for the life of the job (or TSO logon session), even if another generation is added during the job.
- The dynamic allocation routines will not use passed data set information to retrieve volume information.

SVC 99 Return Codes

Note: The data area labels used in this topic are assigned by macros IEFZB4D0 and IEFZB4D2.

When the SVC 99 routines return control to your program, register 15 contains a return code. Depending on the return code, the S99ERROR and S99INFO fields in the input request block (S99RB) may also contain error and information reason codes. The return codes that can appear in register 15 are shown in Figure 66.

Code	Meaning
0	Successful completion; there will also be an information reason code if a non-terminating error occurred during request processing.
4	An error resulted from the current environment, the unavailability of a system resource, or a system routine failure; there will also be an error reason code.
8	The installation validation routine denied this request. (See "Installation Input Validation Routine" for additional information.)
12	The error is due to an invalid parameter list; there will also be an error reason code from class 3. (Class 3 reason codes are listed in Figure 67.)

Note: The DAIRFAIL TSO service routine can be used to issue write-to-programmer or TSO PUTLINE failure messages for both DAIR and SVC 99 error codes. Refer to *TSO Guide to Writing a Terminal Monitor Program or a Command Processor* for information on using DAIRFAIL.

Figure 66. SVC 99 Return Codes

The next two topics describe the information and error reason codes that the SVC 99 routines return in the SVC 99 request block. The last part of this chapter contains the descriptions of the SVC 99 text units, and an example of a dynamic allocation request.

Information Reason Codes

When the SVC 99 routines encounter a non-terminating error during processing, an information reason code appears in the request block field labelled S99INFO. The possible codes and their meanings are:

Code	Meaning
0004	Reserved
0008	Overriding disposition ignored for one of the following reasons: <ul style="list-style-type: none">• Data set was originally allocated with a disposition of PASS• Data set is a non-subsystem data set that has a system-generated name; you cannot override disposition on this type of data set• Data set is a VSAM data set

In these cases, the data set is deallocated using the disposition specified when the request was allocated.

000C-001C Reserved

002n The data set was successfully deallocated but processing of the requested CATLG or UNCATLG disposition was unsuccessful. The digit "n" is a code representing the reason for the failure. The possible codes and their meanings are:

Code Meaning

- 1 A control volume is required; a utility program must be used to catalog the data set.
- 2 The data set to be cataloged is already cataloged; or the data set to be uncataloged could not be located; or no change was made to the volume serial list of a data set with a disposition of CATLG.
- 3 The specified index does not exist.
- 4 The data set could not be cataloged because the space was not available in the catalog.
- 5 Not enough storage was available to perform the specified cataloging.
- 6 The data set to be cataloged in a generation index is improperly named.
- 7 The data set to be cataloged has not been opened; no density information is available (for dual density tape requests only).
- 8 Reserved
- 9 An uncorrectable I/O error occurred in reading or writing the catalog.

003n The data set was successfully deallocated but processing of the requested DELETE disposition was unsuccessful. The digit "n" is a code representing the reason for the failure. The possible codes and their meanings are:

Code Meaning

- 1 The expiration date has not occurred.
- 2 Reserved
- 3 Reserved
- 4 No device was available for mounting the volume during deletion.
- 5 Not enough storage was available to perform the specified deletion.
- 6 Either no volumes were mounted or volumes that were mounted could not be demounted to permit the remaining volumes to be mounted.
- 8 The SCRATCH routine returned an error code. If the user's JOB statement requested allocation/termination messages, message IEF283I appears in the SYSOUT listing. This message lists the volume serial numbers of the data sets that were not deleted; following each number is a code that explains why each data set was not deleted.

Error Reason Codes

When the SVC 99 routines return a nonzero return code in register 15, the request block field labelled S99ERROR contains a code that explains the reason for the return code.

Error reason codes are divided into the following classes:

Class	Description
1	Reserved
2	Unavailable system resource
3	Invalid parameter list
4	Environmental error
5	Reserved
6	Reserved
7	System routine error

The error reason codes are shown in Figure 67 through Figure 70. The class designations listed here appear as the second digit of the reason code.

Note: The explanations of the codes in these figures are followed, in parentheses, by an indication of the kind of request associated with the code.

CLASS 2 CODES (UNAVAILABLE SYSTEM RESOURCE)

Hex Code	(Decimal)	Meaning
0204	(516)	Real storage unavailable. (dsname allocation)
0208	(520)	Reserved.
020C	(524)	Request for exclusive use of a shared data set cannot be honored. (dsname allocation) ¹
0210	(528)	Requested data set unavailable. The data set is allocated to another job and its usage attribute conflicts with this request. (dsname allocation) ¹
0214	(532)	Unit(s) not available; or, if allocating an internal reader, all defined internal readers are already allocated. (dsname allocation) ¹
0218	(536)	Specified volume or an acceptable volume is not mounted, and user does not have volume mounting authorization. (dsname allocation) ¹
021C	(540)	Unit name specified is undefined. (dsname allocation)
0220	(544)	Requested volume not available. (dsname allocation) ²
0224	(548)	Eligible device types do not contain enough units. (dsname allocation) ¹
0228	(552)	Specified volume or unit in use by system. (dsname allocation)
022C	(556)	Volume mounted on ineligible permanently resident or reserved unit. (dsname allocation)
0230	(560)	Permanently resident or reserved volume on required unit. (dsname allocation)
0234	(564)	More than one device required for a request specifying a specific unit. (dsname allocation)
0238	(568)	Space unavailable in task input output table (TIOT). (dsname allocation, concatenation)
023C	(572)	Required catalog not mounted, and user does not have volume mounting authorization. (dsname allocation)
0240	(576)	Requested device is a console. (dsname allocation)
0244	(580)	Telecommunication device not accessible. (dsname allocation)
0248	(584)	MSS virtual volume cannot be mounted. (dsname allocation)
024C	(588)	Operating-system-managed resource was unavailable to the subsystem. (dsname allocation) ³
0250	(592)	Subsystem resource not available. (dsname allocation) ³
0254	(596)	The TIOT resource is currently unavailable and the user requested conditional ENQ on the resource. (all SVC 99 functions)
0258	(600)	There was not a sufficient number of non-restricted units to satisfy the request, or JES3 selected a JES3-managed restricted unit to satisfy the request. Restricted units are described in <i>SPL: System Generation Reference</i> . (dsname allocation)
025C	(604)	Requested device is boxed and cannot be accessed, as a result of an I/O error condition or the operator issuing a VARY X, OFFLINE, FORCE command. (dsname allocation)
0260	(608)	Unit does not meet specified status requirements. (dsname allocation)
0264	(612)	Invalid request due to current unit status. (dsname allocation)
0268	(616)	Tape device is broken. (dsname allocation)

¹ The conditions that cause these return codes are detected by MVS/XA or JES3.

² For MSS requests, the MSSC reason code for this failing job step is contained in message IEF710I on the hardcopy log. An explanation of the MSSC reason code is contained in *Mass Storage System Extensions Messages*. For non-MSS requests, this code is accompanied by message IEF485I. It may result from a JES3 failure because of a busy or unavailable situation.

³ The information reason code contains a subsystem-defined value to further describe the error. This value is documented in publications associated with the particular subsystem.

Figure 67. Class 2 Error Reason Codes (Unavailable System Resource)

CLASS 3 CODES

(INVALID PARAMETER LIST)

Hex Code	(Decimal)	Meaning
0304-0338	(772-824)	Assigned by DAIR. (See <i>TSO Guide to Writing a Terminal Monitor Program or a Command Processor</i> .)
033C-0354	(828-852)	Reserved.
0358	(856)	Overriding disposition of DELETE invalid for data set allocated as SHR. (unallocation) ¹
035C	(860)	Invalid PARM specified in text unit. (all SVC 99 functions) ²
0360	(864)	Invalid KEY specified in text unit. (all SVC 99 functions) ²
0364	(868)	JOBLIB/STEPLIB/JOBCAT/STEP CAT specified as ddname, or associated with specified dsname. (dsname allocation, ddname allocation, unallocation, concatenation, deconcatenation) ¹
0368	(872)	Authorized function requested by unauthorized user. (all SVC 99 functions)
036C	(876)	Invalid parameter list format. (all SVC 99 functions)
0370	(880)	Reserved.
0374	(884)	Invalid # specified in text unit. (all SVC 99 functions) ²
0378	(888)	Duplicate KEY specified in text unit. (all SVC 99 functions) ²
037C	(892)	Invalid LEN specified in text unit. (all SVC 99 functions) ²
0380	(896)	Mutually exclusive KEY specified. Two keys that cannot be used together were used in the request. (dsname allocation, unallocation, information retrieval, remove-in-use processing) ²
0384	(900)	Mutually inclusive KEY not specified. One key was used; two should have been used. (dsname allocation, unallocation) ²
0388	(904)	Required key not specified. (ddname allocation, information retrieval, concatenation, deconcatenation, remove-in-use processing, unallocation)
038C	(908)	Duplicate ddnames specified. (concatenation)
0390	(912)	GDG group name specified with relative generation number exceeds 35 characters. (dsname allocation)
0394	(916)	Status and relative generation number are incompatible. (dsname allocation)
0398	(920)	Volume sequence number exceeds the number of volumes. (dsname allocation)
039C	(924)	Device type and volume are incompatible. (dsname allocation)
03A0	(928)	Subsystem detected an invalid parameter. (dsname allocation) ³
03A4	(932)	Unable to protect data set/volume because of conflicting keyword specification.

¹ The information reason code field contains 0004 if the requested function was performed, although an error occurred, as the error reason code indicates.

² The information reason code contains the value of the key that caused the error.

³ The information reason code field contains a subsystem-defined value to further describe the error. This value is documented in publications associated with the particular subsystem.

Figure 68. Class 3 Error Reason Codes (Invalid Parameter List)

CLASS 4 CODES (ENVIRONMENTAL ERROR)		
Hex Code	(Decimal)	Meaning
0404-040C	(1028-1036)	Reserved.
0410	(1040)	Specified ddname unavailable. (dsname allocation, ddname allocation)
0414-041C	(1044-1052)	Reserved.
0420	(1056)	Specified ddname or dsname associated with an open data set. (ddname allocation, concatenation, deconcatenation, unallocation, dsname allocation) ¹
0424	(1060)	Deconcatenation would result in duplicate ddnames (deconcatenation). ¹
0428-0430	(1064-1072)	Reserved.
0434	(1076)	Ddname specified in ddname allocation request is associated with a convertible or non-permanently allocated resource. (ddname allocation)
0438	(1080)	Specified ddname not found. (information retrieval, ddname allocation, concatenation, deconcatenation, unallocation)
043C	(1084)	The system could not deallocate enough of the resources being held in anticipation of reuse to meet the control limit. (dsname allocation)
0440	(1088)	Specified dsname not found. (information retrieval, unallocation)
0444	(1092)	Relative entry number specified in information retrieval request not found. (information retrieval)
0448	(1096)	Request for a new data set failed; the data set already exists. (dsname allocation)
044C	(1100)	Request was made for a data set that has a disposition of delete; this request cannot be honored because the data set may be deleted at any time. (dsname allocation)
0450	(1104)	Request would cause the limit of 1635 concurrent allocations to be exceeded. (dsname allocation)
0454	(1108)	Ddname in DCB reference not found. (dsname allocation)
0458	(1112)	Dsname in DCB reference or volume reference is a GDG group name. (dsname allocation)
045C	(1116)	Specified dsname to be deallocated is a member of a permanently-concatenated group. (unallocation) ¹
0460	(1120)	Specified dsname or member to be deallocated is not associated with specified ddname. (unallocation)
0464	(1124)	Specified dsname to be deallocated is a private catalog. (unallocation) ¹

- ¹ The information reason code field contains 0004 if the requested function was performed, although an error occurred as the error reason code indicates.
- ² The MSSC reason code for this failing job step is contained in message IEF710I on the hardcopy log. An explanation of the MSSC reason code is contained in *Mass Storage System Extensions Messages*.
- ³ This code corresponds to MSSC reason code X'007', which is explained in *Mass Storage System Extensions Messages*.
- ⁴ This code corresponds to MSSC reason code X'207', which is explained in *Mass Storage System Extensions Messages*.
- ⁵ The information reason code contains a subsystem-defined value to further describe the error. This value is documented in publications associated with the particular subsystem.
- ⁶ The information reason code field contains the MSSC reason code. An explanation of the MSSC reason code is contained in *Mass Storage System Extensions Messages*.

Figure 69 (Part 1 of 2). Class 4 Error Reason Codes (Environmental Error)

CLASS 4 CODES, continued

Hex Code	(Decimal)	Meaning
0468	(1128)	Error while allocating or opening a private catalog. (dsname allocation)
046C	(1132)	Remote work station not defined to job entry subsystem. (dsname allocation, unallocation)
0470	(1136)	User unauthorized for subsystem request. (dsname allocation)
0474	(1140)	Error while attempting to select optimum device. (dsname allocation).
0478	(1144)	Unable to process job entry subsystem request. (dsname allocation, unallocation)
047C	(1148)	Unable to establish ESTAE environment. (all SVC 99 functions)
0480	(1152)	The number of units needed to satisfy the request exceeds the limit. (dsname allocation)
0484	(1156)	Request denied by operator. (dsname allocation)
0488	(1160)	GDG pattern DSCB not mounted. (dsname allocation)
048C	(1164)	GDG pattern DSCB not found. (dsname allocation)
0490	(1168)	Error changing allocation assignments. (dsname allocation)
0494	(1172)	Error processing OS CVOL. (dsname allocation)
0498	(1176)	MSS virtual volume not accessible. (dsname allocation) ²
049C	(1180)	MSS virtual volume not defined. (dsname allocation) ³
04A0	(1184)	Specified MSVGP name not defined. (dsname allocation) ⁴
04A4	(1188)	Subsystem request in error. (dsname allocation) ⁵
04A8	(1192)	Subsystem does not support allocation via key DALSSNM. (dsname allocation)
04AC	(1196)	Subsystem is not operational.
04B0	(1200)	Subsystem does not exist.
04B4	(1204)	Protect request not processed; RACF not in system or not active.
04B8	(1208)	MSS not initialized for allocation. (dsname allocation)
04BC	(1212)	MSS volume select error. (dsname allocation) ⁶
04C0	(1216)	Protect request failed; user not defined to RACF. (dsname allocation)
04C4	(1220)	The last request was for a VOL = REF to a dsname or DCB = dsname that exceeded the maximum allowable dsname referbacks. (A maximum of 972 backward references are allowed if the data set names are 44 characters in length.)
04C8	(1224)	A non-zero return code was set in register 15 from either common allocation or JFCB housekeeping; however, the SIOT reason code (SIOTRSNC) was not set. This problem might result from installation modification of the eligible device table (EDT). (dsname allocation)
04CC	(1228)	Invalid reference to an OUTPUT JCL statement. (dsname allocation)

- ¹ The information reason code field contains 0004 if the requested function was performed, although an error occurred, as the error reason code indicates.
- ² The MSSC reason code for this failing job step is contained in message IEF710I on the hardcopy log. An explanation of the MSSC reason code is contained in *Mass Storage System Extensions Messages*.
- ³ This code corresponds to MSSC reason code X'007', which is explained in *Mass Storage System Extensions Messages*.
- ⁴ This code corresponds to MSSC reason code X'207', which is explained in *Mass Storage System Extensions Messages*.
- ⁵ The information reason code contains a subsystem-defined value to further describe the error. This value is documented in publications associated with the particular subsystem.
- ⁶ The information reason code field contains the MSSC reason code. An explanation of the MSSC reason code is contained in *Mass Storage System Extensions Messages*.

Figure 69 (Part 2 of 2). Class 4 Error Reason Codes (Environmental Error)

CLASS 7 CODES**(SYSTEM ROUTINE ERROR)**

Note: The failing system routine returns the code represented by "zz."

Hex Code	Meaning
17zz	LOCATE error. (Note: Hexadecimal '08', '18', and '2C' are the only expected LOCATE return codes. 'FF' is returned as the value of zz if an unexpected return code is returned by LOCATE.) (dsname allocation) ¹
27zz	Reserved
37zz	Reserved
47zz	DADSM allocate error. (dsname allocation) ²
57zz	CATALOG error. (dsname allocation) ¹
67zz	OBTAIN error. (dsname allocation, information retrieval) ¹
7700	Subsystem error. (dsname allocation) ³
7704	A subsystem interface system error occurred while processing key DALSSNM.
8700	Scheduler JCL facility error.

¹ LOCATE and CATALOG return codes can be found in *MVS/XA CVOL Processor Logic Manual, LY26-3895*. OBTAIN and DADSM return codes can be found in *MVS/XA DADSM Diagnosis Reference Manual, SY26-3904*

² The information reason code field might contain a value that further describes the error. An explanation of this value is contained in the allocation message corresponding to the error code. See the section on reason codes in the allocation/unallocation component of *System Logic Library* for the message number associated with the error and *Message Library: System Messages* for the message itself.

³ The information reason code field contains a subsystem-defined value to further describe the error. This value is documented in publications associated with the particular subsystem.

Figure 70. Class 7 Error Reason Codes (System Routine Error)

SVC 99 Text Units, by Function

The following pages contain descriptions of each of the text units you can use in the SVC 99 parameter list that accompanies the DYNALLOC macro. The text units are arranged according to the functions they request, in ascending order of their KEY values. See Figure 65 for a general description of the text unit and text unit keys.

You request a particular SVC 99 function by coding the appropriate verb code in the request block of the SVC 99 parameter list. The text units are grouped within verb codes; the largest group (verb code 01) is further divided into three subgroups. The verb codes and the functions they represent are listed below:

Verb Code	SVC 99 Function
X'01'	Dsname allocation
X'02'	Unallocation
X'03'	Concatenation
X'04'	Deconcatenation
X'05'	Remove-in-use processing based on task-ID
X'06'	Ddname allocation
X'07'	Information retrieval

Figure 71 through Figure 79 present the SVC 99 text units in list form, introductory to the descriptions of the text units in each verb code group. The mnemonics given for the text units are those assigned by mapping macro IEFZB4D2.

A suggested approach to setting up your *dsname* text unit keys is to code the applicable JCL DD statement and then look up the text unit keys you need, by their mnemonics, in Figure 71. The descriptions of the text units are in order of the key numbers, for easy reference.

Note that the values you specify in the text units are in binary (hexadecimal) and EBCDIC.

Hex Text Unit Key	IEFZB4D2 Mnemonic	Dsname Allocation Function
0001	DALDDNAM	Associates a ddname with an allocation request.
0002	DALDSNAM	Names the data set to be allocated.
0003	DALMEMBR	Specifies data set number or relative generation number.
0004	DALSTATS	Specifies the data set status.
0005	DALNDISP	Specifies the data set's normal disposition.
0006	DALCDISP	Specifies the data set's conditional disposition.
0005	DALNDISP	Specifies the data set's normal disposition.
0007	DALTRK	Specifies the space allocation in tracks.
0008	DALCYL	Specifies the space allocation in cylinders.
0009	DALBLKLN	Specifies the average data block length.
000A	DALPRIME	Specifies a primary space quantity.
000B	DALSECND	Specifies a secondary space quantity.
000C	DALDIR	Specifies the number of PDS directory blocks.
000D	DALRLSE	Deletes unused space at data set closure.
000E	DALSPFRM	Ensures a specific allocated space format.
000F	DALROUND	Specifies space allocation in whole cylinders.
0010	DALVLSER	Specifies volume serial numbers.
0011	DALPRIVT	Specifies the private volume use attribute.
0012	DALVSEQ	Specifies the volume sequence number processing.
0013	DALVLCNT	Specifies the data set's volume count.
0014	DALVLRDS	Specifies volume reference to a cataloged data set.
0015	DALUNIT	Describes the unit specification.
0016	DALUNCNT	Specifies the number of devices to be allocated.
0017	DALPARAL	Specifies parallel mounting for a data set's volumes.
0018	DALYSOU	Specifies the SYSOUT data set and defines its class.
0019	DALSPGNM	Specifies the SYSOUT program name.
001A	DALFMNO	Specifies the SYSOUT form number.
001B	DALOUTLM	Limits the SYSOUT data set's logical record count.
001C	DALCLOSE	Frees a data set at closure.
001D	DALCOPYS	Specifies the SYSOUT listing copies count.
001E	DALLABEL	Specifies the type of volume label.

Figure 71 (Part 1 of 2). Verb Code 01 (Dsname Allocation) - Text Unit Keys, Mnemonics, and functions

Hex Text Unit Key	IEFZB4D2 Mnemonic	Dsname Allocation Function
001F	DALDSSEQ	Specifies a tape data set's relative position.
0020	DALPASPR	Password protects the created data set.
0021	DALINOUT	Specifies "input only" or "output only" data set processing.
0022	DALEXPDT	Specifies the data set's expiration date.
0023	DALRETPD	Specifies the data set's retention period.
0024	DALDUMMY	Allocates a dummy data set.
0025	DALFCBIM	Identifies the forms control buffer image.
0026	DALFCBAV	Requests operator verification of the image display or forms alignment.
0027	DALQNAME	Names a TPROCESS macro, and a TCAM procedure.
0028	DALTERM	Specifies a time sharing terminal as an I/O device.
0029	DALUCS	Specifies a universal character set.
002A	DALUFOLD	Specifies "fold mode" for loading the requested print chain or train.
002B	DALUVERFY	Requests operator verification of the correct print chain or train mounting.
002C	DALDCBDS	Specifies the retrieval of DCB information from a cataloged data set's label.
002D	DALDCBDD	Specifies the retrieval of DCB information from a ddname-related, currently allocated data set.
0058	DALSUSER	Specifies remote work station routing for the SYSOUT data set.
0059	DALSHOLD	Specifies hold queue routing for the SYSOUT data set.
005E	DALMSVGP	Specifies a group of MSS virtual volumes.
005F	DALSSNM	Requests allocation of a subsystem data set.
0060	DALSSPRM	Specifies subsystem-defined parameters for use with key DALSSNM.
0061	DALPROT	Requests that the direct access data set or the tape volume be RACF-protected.
0063	DALUSRID	Specifies a user ID to which the SYSOUT data set is to be routed.
0064	DALBURST	Specifies which stacker of the 3800 Printing Subsystem is to receive the paper output.
0065	DALCHARS	Specifies the name or names of character arrangement tables for printing a data set on the 3800.
0066	DALCOPYG	Specifies how copies are to be grouped if printing is done on a 3800.
0067	DALFFORM	Specifies the forms overlay to be used on the 3800 Printing Subsystem.
0068	DALFCNT	Specifies the number of copies on which the forms overlay is to be printed.
0069	DALMMOD	Specifies the name of the copy modification module to be loaded into the 3800 Printing Subsystem.
006A	DALMTRC	Specifies the table reference character that corresponds to a character arrangement table used for printing the copy modification data.
006C	DALDEFER	Specifies that the system should allocate a device to the data set, but defer mounting the volume(s) until the data set is opened.
8002	DALOUTPT	Refers to a specific OUTPUT JCL statement.

Figure 71 (Part 2 of 2). Verb Code 01 (Dsname Allocation) - Text Unit Keys, Mnemonics, and functions

Dsname Allocation Text Units

Most of the information that can be specified on a JCL DD statement can also be specified in text units for the dsname allocation function (verb code X'01'). These text units are described on the following pages and listed in Figure 71. The text units that represent DCB attributes are described under "DCB Attribute Text Units" and listed in Figure 72. The meaning of the parameters is the same as when specified on a DD statement as described in *JCL*.

For dsname allocation text units that do not have a JCL equivalent, see Figure 73 and the topic "Non-JCL Dsname Allocation Functions."

Ddname Specification - Key = X'0001'

DALDDNAM specifies a ddname to be associated with a dsname allocation request. When you code this key, # must be one, LEN is the length of the parameter field, and PARM contains the ddname.

Example: to specify the ddname DD1, code

KEY	#	LEN	PARM
0001	0001	0003	C4 C4 F1

Dsname Specification - Key = X'0002'

DALDSNAM specifies the name of the data set to be allocated. Dynamic allocation does not support backward references. See Figure 64. The QNAME (DALQNAME) and IPLTXTID (DALIPLTX) keys are mutually exclusive with DALDSNAM. When you code this key, # must be one, LEN is the length of the dsname, and PARM contains the dsname.

Example: to specify the dsname MYDATA, code

KEY	#	LEN	PARM
0002	0001	0006	D4 E8 C4 C1 E3 C1

Example: to specify the temporary dsname &LOAD, code

KEY	#	LEN	PARM
0002	0001	0005	50 D3 D6 C1 C4

Example: to specify the dsname A.B, code

KEY	#	LEN	PARM
0002	0001	0003	C1 4B C2

Member Name Specification - Key = X'0003'

DALMEMBR specifies that a particular member of a data set is to be allocated, rather than the entire data set. A relative generation group number may be specified as the member name. When you specify DALMEMBR, you must also specify the dsname key (DALDSNAM). The QNAME (DALQNAME) and IPLTXTID (DALIPLTX) keys are mutually exclusive with DALMEMBR. When you code this key, # must be one, LEN is the actual length of the member name, and PARM contains the member name.

Example: to specify the member name MEM1, code

KEY	#	LEN	PARM
0003	0001	0004	D4 C5 D4 F1

Example: to specify the relative generation number +1, code

KEY	#	LEN	PARM
0003	0001	0002	4E F1

Data Set Status Specification - Key = X'0004'

DALSTATS specifies the data set status desired. It is mutually exclusive with the SYSOUT key (DALSYSOU). When you code DALSTATS, # and LEN must be one, and PARM contains one of the following values:

X'01'	if OLD is desired
X'02'	if MOD is desired
X'04'	if NEW is desired
X'08'	if SHR is desired

Example: to specify a status of NEW, code

Key	#	LEN	PARM
0004	0001	0001	04

Do not code MOD for temporary data sets dynamically allocated as &&dsname.

Data Set Normal Disposition Specification - Key = X'0005'

DALNDISP specifies the normal data set disposition desired. It is mutually exclusive with the SYSOUT key (DALSYSOU). When you code DALNDISP, # and LEN must be one, and PARM contains one of the following values:

X'01'	if UNCATLG is desired
X'02'	if CATLG is desired
X'04'	if DELETE is desired
X'08'	if KEEP is desired

Example: to specify a normal disposition of DELETE, code

KEY	#	LEN	PARM
0005	0001	0001	04

Data Set Conditional Disposition Specification - Key = X'0006'

DALCDISP specifies the conditional data set disposition desired. It is mutually exclusive with the SYSOUT key (DALSYSOU). The values for #, LEN, and PARM are the same as for normal disposition.

Example: to specify a conditional disposition of DELETE, code

KEY	#	LEN	PARM
0006	0001	0001	04

Track Space Type (TRK) Specification - Key = X'0007'

DALTRK specifies that space is to be allocated in tracks. The primary quantity space key (DALPRIME) or the secondary quantity space key (DALSECND) must also be specified when you code DALTRK. The text unit keys that define space in terms of cylinders (DALCYL, DALROUND) or blocks (DALBLKLN) are mutually exclusive with DALTRK. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify a space request in tracks, code

KEY	#	LEN	PARM
0007	0000	-	-

Cylinder Space Type (CYL) Specification - Key = X'0008'

DALCYL specifies that space is to be allocated in cylinders. The primary quantity space key (DALPRIME) or secondary quantity space key (DALSECND) must also be specified when you code this key. The text unit keys that define space in terms of tracks (DALTRK) or blocks (DALBLKLN) are mutually exclusive with DALCYL. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify a space request in cylinders, code

KEY	#	LEN	PARM
0008	0000	-	-

Block Length Specification - Key = X'0009'

DALBLKLN specifies the average data block length to be used by the system in computing the amount of space to allocate. The primary quantity space key (DALPRIME) or the secondary quantity space key (DALSECND) must also be specified when you code this key. The text unit keys that request space in terms of tracks (DALTRK) or cylinders (DALCYL, DALROUND) are mutually exclusive with DALBLKLN. When you code this key, # must be one, LEN must be three, and PARM contains the average data block length. The maximum PARM value is X'00FFFF' (65,535).

Example: to specify an average data block length of 80, code

KEY	#	LEN	PARM
0009	0001	0003	00 00 50

Primary Space Quantity Specification - Key = X'000A'

DALPRIME specifies a primary space quantity. You must also code one of the space type keys (DALBLKLN, DALCYL, DALTRK) when you specify DALPRIME. When you code this key, # must be one, LEN must be three, and PARM contains the primary quantity value.

Example: to specify a primary quantity of 20, code

KEY	#	LEN	PARM
000A	0001	0003	00 00 14

Secondary Space Quantity Specification - Key = X'000B'

DALSECND specifies a secondary space quantity. You must also code one of the space type keys (DALBLKLN, DALCYL, DALTRK) when you specify DALSECND. When you code this key, # must be one, LEN must be three, and PARM contains the secondary quantity value.

Example: to specify a secondary space quantity of 10, code

KEY	#	LEN	PARM
000B	0001	0003	00 00 0A

Directory Block Specification - Key = X'000C'

DALDIR specifies the number of blocks to be contained in the directory of a partitioned data set. You must also specify a space type key (DALBLKLN, DALCYL, or DALTRK) and the primary quantity key (DALPRIME) when coding DALDIR. When you code this key, # must be one, LEN must be three, and PARM contains the number of directory blocks.

Example: to specify two directory blocks, code

KEY	#	LEN	PARM
000C	0001	0003	00 00 02

Unused Space Release (RLSE) Specification - Key = X'000D'

DALRLSE specifies that unused space is to be deleted when the data set is closed. When you code this key, # must be zero; LEN and PARM are not coded.

Example: to specify the release of unused space, code

KEY	#	LEN	PARM
000D	0000	-	-

Format of Allocated Space Specification - Key = X'000E'

DALSPFRM specifies a particular format of allocated space. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'02'	if different areas of contiguous space are to be allocated (ALX)
X'04'	if maximum contiguous space is required (MXIG)
X'08'	if space must be contiguous (CONTIG)

Example: to specify contiguous space format, code

KEY	#	LEN	PARM
000E	0001	0001	08

Whole Cylinder Allocation (ROUND) Specification - Key = X'000F'

DALROUND specifies that allocated space is to be equal to one or more whole cylinders when requested in units of blocks. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify allocation of whole cylinders, code

KEY	#	LEN	PARM
000F	0000	-	-

Volume Serial Specification - Key = X'0010'

DALVLSER specifies volume serial numbers. It is mutually exclusive with the SYSOUT (DALSYSOU) and volume reference (DALVLRDS) keys. When you code DALVLSER, # contains the number of volume serials being specified, LEN contains the length of the immediately following volume serial, and PARM contains the volume serial.

Example: to specify the volume serials 231400 and 231401, code

KEY	#	LEN	PARM	LEN	PARM
0010	0002	0006	F2 F3 F1 F4 F0 F0	0006	F2 F3 F1 F4 F0 F1

Private Volume Specification - Key = X'0011'

DALPRIVT specifies that the volume(s) allocated are to be assigned the volume use attribute of private. This key is mutually exclusive with the SYSOUT key (DALSYSOU). When you code DALPRIVT, # must be zero; LEN and PARM are not specified.

Example: to specify the private volume attribute, code

KEY	#	LEN	PARM
0011	0000	-	-

Volume Sequence Number Specification - Key = X'0012'

DALVLSEQ specifies which volume, of a multi-volume data set, processing is to begin with. This key is mutually exclusive with the SYSOUT key (DALSYSOU). When you code DALVLSEQ, # must be one, LEN must be two, and PARM contains the volume sequence number. The maximum PARM value is X'00FF' (255).

Example: to specify a volume sequence number of two, code

KEY	#	LEN	PARM
0012	0001	0002	0002

Volume Count Specification - Key = X'0013'

DALVLCNT specifies the maximum number of volumes an output data set may require. This key is mutually exclusive with the SYSOUT key (DALSYSOU). When you code DALVLCNT, # and LEN must be one, and PARM contains the volume count.

Example: to specify a volume count of 10, code

KEY	#	LEN	PARM
0013	0001	0001	0A

Volume Reference to a Dsname Specification - Key = X'0014'

DALVLRDS indicates that the system is to obtain volume serial information from the specified cataloged data set. This key is mutually exclusive with the SYSOUT (DALSYSOU) and volume serial (DALVLSER) keys. (You cannot use a volume reference to a ddname for dynamic allocation.) When you code this key, # must be one, LEN is the actual length of the dsname, and PARM contains the dsname (a name of all blanks is invalid).

Example: to specify volume reference to the data set DSN1, code

KEY	#	LEN	PARM
0014	0001	0004	C4 E2 D5 F1

Unit Description Specification - Key = X'0015'

DALUNIT specifies a unit as a group (esoteric) name, a device type (generic), or a specific unit address (in EBCDIC). When you code DALUNIT, # must be one, LEN is the actual length of the unit description, and PARM contains the unit description.

Example: to specify the group name SYSDA, code

KEY	#	LEN	PARM
0015	0001	0005	E2 E8 E2 C4 C1

Example: to specify the device type 3330, code

KEY	#	LEN	PARM
0015	0001	0004	F3 F3 F3 F0

Example: to specify the unit address 230, code

KEY	#	LEN	PARM
0015	0001	0003	F2 F3 F0

Unit Count Specification - Key = X'0016'

DALUNCNT specifies the number of devices to be allocated. It is mutually exclusive with the parallel mount key (DALPARAL). When you code DALUNCNT, # and LEN must be one, and PARM contains the unit count. The maximum PARM value is X'3B' (59).

Example: to specify a unit count of ten, code

KEY	#	LEN	PARM
0016	0001	0001	0A

Parallel Mount Specification - Key = X'0017'

DALPARAL specifies that each volume of a data set is to be mounted on a separate device. It is mutually exclusive with the unit count key (DALUNCNT). When you code DALPARAL, # must be zero; LEN and PARM are not specified.

Example: to specify parallel mount, code

KEY	#	LEN	PARM
0017	0000	-	-

SYSOUT Specification - Key = X'0018'

DALSYSOU specifies that a system output data set is to be allocated and defines the output class of the data set. When you code this key and want a class other than the default, # and LEN must be one, and PARM contains the output class. To obtain the class from the OUTPUT DD statement, if specified, code zero in the # field; LEN and PARM are not specified. If no OUTPUT DD statement is found, the default message class is used. DALSYSOU is mutually exclusive with the following text unit keys:

DALSTATS, DALNDISP and DALCDISP
DALVLSER, DALPRIVT, DALVLSER, DALVLCNT, and DALVLRDS
DALQNAME
DALSSNM, DALSSPRM, and DALSSATT

Example: to specify a SYSOUT data set in class A, code

KEY	#	LEN	PARM
0018	0001	0001	C1

Example: to specify a SYSOUT data set and to default the class, code

KEY	#	LEN	PARM
0018	0000	-	-

SYSOUT Program Name Specification - Key = X'0019'

DALSPGNM specifies the SYSOUT program name. The SYSOUT key (DALSYSOU) must also be specified when you code DALSPGNM. The subsystem name request (DALSSNM), subsystem parameter (DALSSPRM), and SYSOUT userid (DALUSRID) keys are mutually exclusive with DALSPGNM. When you code this key, # must be one, LEN is the actual length of the name, and PARM contains the program name.

Example: to specify the program name MYWRITER, code

KEY	#	LEN	PARM
0019	0001	0008	D4 E8 E6 D9 C9 E3 C5 D9

SYSOUT Form Number Specification - Key = X'001A'

DALFMNO specifies the SYSOUT form number. The SYSOUT (DALSYSOU) key must also be specified when you code DALFMNO. The subsystem name request (DALSSNM) and subsystem parameter (DALSSPRM) keys are mutually exclusive with DALFMNO. When you code this key, # must be one, LEN is the actual length of the form number, and PARM contains the form number.

Example: to specify the form number 1234, code

KEY	#	LEN	PARM
001A	0001	0004	F1 F2 F3 F4

SYSOUT Output Limit Specification - Key = X'001B'

DALOUTLM specifies the number of logical records in a SYSOUT data set. The SYSOUT key (DALSYSOU) must also be specified when you code DALOUTLM. When you code this key, # must be one, LEN must be three, and PARM contains the output limit.

Example: to specify an output limit of 1000, code

KEY	#	LEN	PARM
001B	0001	0003	00 03 E8

Unallocation at CLOSE Specification - Key X'001C'

DALCLOSE requests unallocation when a DCB is closed rather than at step unallocation. When you code DALCLOSE, # must be zero; LEN and PARM are not specified.

Example: to specify unallocation at CLOSE, code

KEY	#	LEN	PARM
001C	0000	-	-

SYSOUT Copies Specification - Key = X'001D'

DALCOPYS requests up to 255 hardcopy listings of a particular SYSOUT data set. The SYSOUT key (DALSYSOU) must also be specified when you code DALCOPYS. When you code this key, # and LEN must be one, and PARM contains the number of copies being requested.

Example: to specify a request for 25 copies, code

KEY	#	LEN	PARM
001D	0001	0001	19

Label Type Specification - Key = X'001E'

DALLABEL specifies the type of label associated with a volume. It is mutually exclusive with the SYSOUT (DALSYSOU) key. When you code DALLABEL, # and LEN must be one, and PARM contains one of the following values:

X'01'	if the volume has no label (NL)
X'02'	if the volume has an IBM standard label (SL)
X'04'	if the volume has a non-standard label (NSL)
X'0A'	if the volume has both an IBM standard label and a user label (SUL)
X'10'	if label processing is to be bypassed (BLP)
X'21'	if the system is to check for and bypass a leading tape mark on DOS unlabeled tape (LTM)
X'40'	if the volume has an American National Standard label (AL)
X'48'	if the volume has an American National Standard label and an American National Standard user label (AUL)

Example: to specify no labels, code

KEY	#	LEN	PARM
001E	0001	0001	01

Note: If your installation has not specified the BLP feature in the JES2 reader cataloged procedure, specifying BLP has the same effect as specifying NL.

Data Set Sequence Number Specification - Key = X'001F'

DALDSSEQ specifies the relative position of a data set on a tape volume (data set sequence number). It is mutually exclusive with the SYSOUT (DALSYSOU) key. When you code DALDSSEQ, # must be one, LEN must be two, and PARM contains the sequence number. The maximum PARM value is X'270F' (9999).

Example: to specify a data set sequence number of 2, code

KEY	#	LEN	PARM
001F	0001	0002	00 02

Password Protection Specification - Key = X'0020'

DALPASPR specifies that the data set being created is to be password protected. It is mutually exclusive with the SYSOUT (DALSYSOU) key. When you code DALPASPR, # and LEN must be one, and PARM contains one of the following values:

X'10'	if the data set should not be read, changed, extended, or deleted without the password.
X'30'	if the data set should not be changed, extended, or deleted without the password. Reading is permitted.

Example: to specify complete password protection, code

KEY	#	LEN	PARM
0020	0001	0001	10

Input Only or Output Only Specification - Key = X'0021'

DALINOUT specifies that the data set is to be processed for input only or output only. In the case of BDAM and BSAM data sets, this key overrides OPEN macro options (INOUT, UPDAT, OUTIN, OUTINX) the same way the JCL LABEL parameter options IN and OUT do. See *JCL* for details.

DALINOUT is mutually exclusive with the SYSOUT (DALSYSOU) key. When you code DALINOUT, # and LEN must be one, and PARM contains one of the following values:

X'40'	if output only is to be requested.
X'80'	if input only is to be requested.

Example: to specify processing for input only, code

KEY	#	LEN	PARM
0021	0001	0001	80

Expiration Date Specification - Key = X'0022'

DALEXPDT specifies the date when the data set can be deleted or overwritten by another data set. This key is mutually exclusive with the retention period (DALRETPD) and SYSOUT (DALSYSOU) keys. When you code DALEXPDT, # must be one, LEN must be five, and PARM contains five digits – a two-digit year value and a three-digit day value (yyddd).

Example: to specify an expiration date of January 1, 1985 (85001), code

KEY	#	LEN	PARM
0022	0001	0005	F8 F5 F0 F0 F1

Retention Period Specification - Key = X'0023'

DALRETPD specifies the number of days that must pass before the data set can be deleted or overwritten by another data set. It is mutually exclusive with the expiration date (DALEXPDT) and SYSOUT (DALSYSOU) keys. When you code DALRETPD, # must be one, LEN must be two, and PARM contains the retention period. The maximum PARM value is X'270F' (9999).

Example: to specify a retention period of 10 days, code

KEY	#	LEN	PARM
0023	0001	0002	000A

Dummy Data Set Specification - Key = X'0024'

DALDUMMY requests that a dummy data set be allocated. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to request allocation of a dummy data set, code

KEY	#	LEN	PARM
0024	0000	-	-

Forms Control Buffer (FCB) Image Identification Specification - Key = X'0025'

DALFCBIM specifies the code that identifies the image to be loaded into the forms control buffer (FCB). It is mutually exclusive with the DCB INTVL (DALINTVL) and FRID (DALFRID) keys. When you code DALFCBIM, # must be one, LEN contains the length of the image-id (maximum of 4), and PARM contains the image-id.

Example: to specify the image-id STD1, code

KEY	#	LEN	PARM
0025	0001	0004	E2 E3 C4 F1

Form Alignment and Image Verification Specification - Key = X'0026'

DALFCBAV requests that the operator be prompted to check the alignment of the printer forms before the data set is printed, or to visually verify the image displayed on the printer as the desired one. The FCB image-id (DALFCBIM) key must also be coded when DALFCBAV is specified. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'04' if verification is requested (VERIFY).
X'08' if alignment is requested (ALIGN).

Example: to specify verification, code

KEY	#	LEN	PARM
0026	0001	0001	04

QNAME Specification - Key = X'0027'

DALQNAME specifies the name of a TPROCESS macro and an optional qualifier designating the particular TCAM job or started task associated with the process name. The dsname (DALDSNAM), member name (DALMEMBR), IPLTXTID (DALIPLTX), and SYSOUT (DALSYSOU) keys are mutually exclusive with DALQNAME. The DCB BLKSIZE (DALBLKSZ), BUFL (DALBUFL), LRECL (DALLRECL), OPTCD (DALOPTCD) and RECFM (DALRECFM) keys (see "DCB Attribute Text Units") are meaningful with DALQNAME.

When you code this key, # must be one, LEN is the length of the entire process name (maximum of eight characters for each name, plus a period if you are coding two names; total of 17), and PARM contains the process name itself.

Example: to specify the process name TP1, code

KEY	#	LEN	PARM
0027	0001	0003	E3 D7 F1

Example: to specify the process name TPR.TCAM2, code

KEY	#	LEN	PARM
0027	0001	0009	E3 D7 D9 4B E3 C3 C1 D4 F2

Terminal Specification - Key = X'0028'

DALTERM specifies that a time-sharing terminal is to be used as an I/O device. In a batch environment, the specification is not used, but is checked for syntax. In a time-sharing environment, all other specifications except DCB specifications are ignored when DALTERM is coded. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify a terminal allocation, code

KEY	#	LEN	PARM
0028	0000	-	-

Universal Character Set (UCS) Specification - Key = X'0029'

DALUCS identifies a special character set to be used for printing a data set. The DCB INTVL (DALINTVL) and RESERVE (DALRSRVF and DALRSRVS) keys (see "DCB Attribute Text Units") are mutually exclusive with DALUCS. When you code this key, # must be one, LEN is the length of the character set name code (maximum is four) and PARM contains the character set code.

Example: to specify the character set code AN, code

KEY	#	LEN	PARM
0029	0001	0002	C1 D5

Fold Mode Specification - Key = X'002A'

DALUFOLD specifies that the chain or train corresponding to the desired character set is to be loaded in the fold mode. You must also specify the universal character set key (DALUCS) when you code DALUFOLD. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify fold mode, code

KEY	#	LEN	PARM
002A	0000	-	-

Character Set Image Verification Specification - Key = X'002B'

DALUVERFY requests that the operator be prompted to verify that the correct chain or train is mounted before the data set is printed. You must also specify the universal character set key (DALUCS) when you code DALUVERFY. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify character set image verification, code

KEY	#	LEN	PARM
002B	0000	-	-

DCB Reference to a Dsname Specification - Key = X'002C'

DALDCBDS specifies that DCB information is to be retrieved from the data set label of a cataloged data set. This data set must reside on a direct access volume and the volume must currently be mounted.

The DSORG, RECFM, OPTCD, BLKSIZE, LRECL, RKP, and KEYLEN DCB attributes, and the volume sequence number and expiration date are copied from the data set label. If text units for those parameters are coded in addition to this key, the text unit specifications override the parameters copied from the data set label.

DALDCBDS is mutually exclusive with DCB reference to a ddname (DALDCBDD). When DALDCBDS is specified, # must be one, LEN is the length of the dsname, and PARM contains the data set same. (A dsname of all blanks is invalid.)

Example: to specify DCB reference to the dsname ABC, code

KEY	#	LEN	PARM
002C	0001	0003	C1 C2 C3

DCB Reference to a Ddname Specification - Key = X'002D'

DALDCBDD specifies that DCB information is to be retrieved from the currently allocated data set associated with the specified ddname. For time-sharing users, the expiration date and INPUT/OUTPUT ONLY specifications are also retrieved. This key is mutually exclusive with DCB reference to a dsname (the DALDCBDS key). Any DCB attributes, expiration date (DALEXP), and INPUT/OUTPUT ONLY (DALINOUT) keys specified in addition to this key override the corresponding DCB parameters associated with the ddname.

When you code DALDCBDD, # must be one, LEN is the length of the ddname, and PARM contains the ddname.

Example: to specify DCB reference to the ddname DD1, code

KEY	#	LEN	PARM
002D	0001	0003	C4 C4 F1

SYSOUT Remote Work Station Specification - Key = X'0058'

DALSUSER requests that the SYSOUT data set being allocated be routed to a remote work station when it is deallocated. When coded in conjunction with the user ID key (DALUSRID), this key represents the node to which the user ID is assigned. The SYSOUT key (DALSYSOU) is required with this key. When you code DALSUSER, # must be one, LEN is the length of the work station name (maximum of 8), and PARM contains the work station name.

Example: to specify the work station USER01, code

KEY	#	LEN	PARM
0058	0001	0006	E4 E2 C5 D9 F0 F1

SYSOUT Hold Queue Specification - Key = X'0059'

DALSHOLD requests that the SYSOUT data set being allocated be placed on the hold queue when it is deallocated. The SYSOUT key (DALSYSOU) must also be specified when DALSHOLD is specified. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify hold, code

KEY	#	LEN	PARM
0059	0000	-	-

MSVGP Specification - Key = X'005E'

DALMSVGP specifies a group of MSS virtual volumes. It is mutually exclusive with the SYSOUT (DALSYSOU), QNAME (DALQNAME), and volume serial (DALVLSER) text unit keys. When you code DALMSVGP, # must be one, LEN is the actual length of the MSVGP name, and PARM contains the group name.

Example: to specify a MSS volume group of SYSGROUP, code

KEY	#	LEN	PARM
005E	0001	0008	E2 E8 E2 C7 D9 D6 E4 D7

Subsystem Name Request Specification - Key = X'005F'

DALSSNM specifies a subsystem data set. You must specify the name of the subsystem that is to process the request for allocation unless you want the request processed by the default subsystem.

- When you code DALSSNM to request a subsystem other than the default subsystem, # must be one, LEN specifies the length of the subsystem name (maximum of four) and PARM contains the subsystem name (one to four characters).

The first character of the subsystem name must be either alphabetic or national and the remaining characters must be either alphanumeric or national. See *JCL* for a list of the alphanumeric and national character sets.

- When you code DALSSNM to request the default subsystem, # must be zero; LEN and PARM are not be specified.

DALSSNM is mutually exclusive with the SYSOUT (DALSYSOU), SYSOUT program name (DALSPGNM) and SYSOUT form number (DALFMNO) keys.

Your installation's system programming staff can identify the subsystems at your installation that support DALSSNM requests.

Example 1: to request subsystem SUB1, code

KEY	#	LEN	PARM
005F	0001	0004	E2 E4 C2 F1

Example 2: to request the default subsystem, code

KEY	#	LEN	PARM
005F	0000	-	-

Subsystem Parameter Specification - Key = X'0060'

DALSSPRM specifies parameters that will be processed by a subsystem. When coding DALSSPRM, you must also specify the subsystem name (DALSSNM) key. DALSSPRM is mutually exclusive with the SYSOUT (DALSYSOU), SYSOUT program name (DALSPGNM), and SYSOUT form number (DALFMNO) keys.

When you code this key, # contains the number of LEN and PARM combinations that are present (maximum of 254), LEN specifies the length of the immediately-following parameter (value range from 0 to 67), and PARM contains the parameter to be passed to the subsystem. When you code a LEN value of 0, do not code a PARM value.

Example: to specify two parameters, PARM1 and PARAMETER2, code:

KEY	#	LEN	PARM
0060	0002	0005	D7 C1 D9 D4 F1
		000A	D7 C1 D9 C1 D4 C5 E3 C5 D9 F2

Note: For additional information about subsystem data sets and subsystem parameters, refer to the documentation for the particular subsystem.

PROTECT Specification - Key = X'0061'

DALPROT requests that the specified direct access data set or tape volume be RACF-protected when defined (DASD) or used (tape). It is mutually exclusive with the SYSOUT (DALSYSOU), FCB (DALFCBIM), QNAME (DALQNAME), terminal (DALTERM) and UCS (DALUCS) keys.

When you code DALPROT, # must be zero; LEN and PARM are not specified. See *JCL* for additional information about specifying the PROTECT function.

Example: to specify PROTECT, code

KEY	#	LEN	PARM
0061	0000	-	-

SYSOUT User ID Specification - Key = X'0063'

DALUSRID requests that the SYSOUT data set being allocated be routed to the specified user ID at a remote location. The SYSOUT (DALSYSOU) and SYSOUT remote work station (DALUSER) keys are required with this key. The SYSOUT program name key (DALSPGNM) is mutually exclusive with DALUSRID.

When you code this key, # must be one, LEN is the length of the user ID (maximum of 8), and PARM contains the user ID itself. The user ID may be any EBCDIC characters, including special characters.

Example: to send the Class A SYSOUT data set to user ID D58-VWM at remote work station (node) DALLAS, code

KEY	#	LEN	PARM
0063	0001	0007	C4 F5 F8 60 E5 E6 D4
0018	0001	0001	C1
0058	0001	0006	C4 C1 D3 D3 C1 E2

Burst Specification - Key = X'0064'

DALBURST specifies which stacker of the 3800 Printing Subsystem is to receive the paper output.

When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'02' for burster-trimmer-stacker
X'04' for continuous form stacking

Example: to specify continuous form stacking, code

KEY	#	LEN	PARM
0064	0001	0001	04

Character Arrangement Table Specification - Key = X'0065'

DALCHARS specifies the name or names of character arrangement tables for printing a data set on the 3800 Printing Subsystem.

When you code this key, # contains the number of character arrangement tables being specified, LEN contains the length of the immediately-following character arrangement table, and PARM contains the name of the character arrangement table.

Example: to specify the character arrangement tables GS10 and GS12, code

KEY	#	LEN	PARM	LEN	PARM
0065	0002	0004	C7E2F1F0	0004	C7E2F1F2

Copy Groups Specification - Key = X'0066'

DALCOPYG specifies how multiple copies of 3800 output are to be grouped. The copies specification (DALCOPYS) key is required with this key.

When you code DALCOPYG, # contains the number of group values being specified, LEN must be one, and PARM contains the number of copies of each page that are to be grouped together.

Example: to indicate that six copies of the data set are to be printed in three groups; and that the first group is to contain one copy of each page, the second group is to contain three copies of each page, and the third group is to contain two copies of each page, code

KEY	#	LEN	PARM	LEN	PARM	LEN	PARM
001D	0001	0001	06				
0066	0003	0001	01	0001	03	0001	02

Flash Forms Overlay Specification - Key = X'0067'

DALFFORM specifies the forms overlay to be used on the 3800 Printing Subsystem.

When you code this key, # must be one, LEN contains the length of the immediately-following form name, and PARM contains the name of the forms overlay frame that the operator is to insert into the printer before printing begins.

Example: to specify the forms overlay frame named ABCD, code

KEY	#	LEN	PARM
0067	0001	0004	C1C2C3C4

Flash Forms Overlay Count Specification - Key = X'0068'

DALFCNT specifies the number of copies on which the forms overlay is to be printed. When specifying DALFCNT, you must also specify the flash forms overlay (DALFFORM) key.

When you code DALFCNT, # and LEN must be one, and PARM contains the number of copies.

Example: to specify that the first five copies are to be flashed with the forms overlay, code

KEY	#	LEN	PARM
0068	0001	0001	05

Copy Modification Module Specification - Key = X'0069'

DALMMOD specifies the name of the copy modification module to be loaded into the 3800 Printing Subsystem.

When you code DALMMOD, # must be one, LEN contains the length of the immediately-following module name and PARM contains the name of the copy modification module.

Example: to specify that the data in the copy modification module named A is to replace the variable data in the data set, code

KEY	#	LEN	PARM
0069	0001	0001	C1

Copy Module Table Reference Specification - Key = X'006A'

DALMTRC specifies the table reference character that corresponds to a character arrangement table specified on the DALCHARS text unit key, and used for printing the copy modification data. When specifying DALMTRC, you must also specify the copy modification module specification (DALMMOD) key.

When you code this key, # and LEN must be one and PARM contains one of the following values:

X'00' for the first character arrangement table specified on the DALCHARS text unit
X'01' for the second character arrangement table specified
X'02' for the third character arrangement table specified
X'03' for the fourth character arrangement table specified

Example: to indicate that the first character arrangement table specified on the DALCHARS key is to be used, code

KEY	#	LEN	PARM
006A	0001	0001	00

DEFER Specification - Key = X'006C'

DALDEFER specifies that the system should allocate a device to the data set, but the volume(s) on which the data set resides should not be mounted until the data set is opened.

When you code DALDEFER, # must be zero; LEN and PARM are not specified. See *JCL* for the rules regarding the use of DEFER.

Example: to specify a request for deferred mounting of a volume or volumes, code

KEY	#	LEN	PARM
006C	0000	-	-

OUTPUT Statement Reference - Key = X'8002'

DALOUTPT explicitly associates a SYSOUT data set with the OUTPUT JCL statement specified in the PARM field. The SYSOUT (DALSYSOU) text unit key is required with this key.

When you code DALOUTPT, # is a number ranging from one to 128 (X'0080'), LEN is the length of the PARM field, (maximum of X'001A'), and PARM contains the name of the OUTPUT statement in one of the following forms:

name
stepname.name
stepname.procstepname.name

Example: to reference an OUTPUT JCL statement named OUT1 in the job step named STEP1, and another named OUTX in the current step, code

KEY	#	LEN	PARM
8002	0002	000A	E2 E3 C5 D7 F1 4B D6 E4 E3 F1
		0004	D6 E4 E3 E7

See *JCL* for programming considerations related to the use of OUTPUT JCL statements.

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
002E	DALBFALN	Specifies buffer alignment.
002F	DALBFTEK	Specifies the buffering technique.
0030	DALBLKSZ	Specifies blocksize.
0031	DALBUFIN	Specifies the receiving buffer count.
0032	DALBUFL	Specifies the buffer length.
0033	DALBUFMX	Specifies the buffer count per line.
0034	DALBUFNO	Specifies the buffer count per DCB.
0035	DALBUFOF	Specifies the buffer offset.
0036	DALBUFOU	Specifies the sending buffer count.
0037	DALBUFRQ	Specifies the buffer count per GET macro.
0038	DALBUFSZ	Specifies the line group buffer size.
0039	DALCODE	Specifies the data's paper tape code.
003A	DALCPRI	Specifies the relative sending and receiving priority.
003B	DALDEN	Specifies the magnetic tape density.
003C	DALDSORG	Specifies the data set organization.
003D	DALEROPT	Specifies reading and writing error options.
003E	DALGNCP	Specifies the GAM-I/O count per WAIT macro.
003F	DALINTVL	Specifies the line polling interval per group.
0040	DALKYLEN	Specifies the data set key lengths.
0041	DALLIMCT	Specifies the search limit.
0042	DALLRECL	Specifies the logical record length.
0043	DALMODE	Specifies card punch/reader operational mode.
0044	DALNCP	Specifies the READ/WRITE count per CHECK.
0045	DALOPTCD	Specifies the control program's operational services.
0046	DALPCIR	Specifies the relationship of the receiving PCI to the allocation and freeing of buffers.
0047	DALPCIS	Specifies the relationship of the sending PCI to the allocation and freeing of buffers.
0048	DALPRTSP	Specifies printer line spacing.
0049	DALRECFM	Specifies the record format.
004A	DALRSRVF	Specifies the first buffer's reserve byte count for insertion of data.
004B	DALRSRVS	Specifies the secondary buffer's reserve byte count for insertion of data.
004C	DALSOWA	Specifies the user's telecommunications input work areas size.
004D	DALSTACK	Specifies the card punch's stacker bin.
004E	DALTHRSRSH	Specifies the use percentage of nonreusable direct access message queue records per flush closedown.
004F	DALTRTCH	Specifies the 7-track tape recording technique.
0051	DALIPLTX	Specifies a TCAM network control program name.
0054	DALDIAGN	Requests OPEN/CLOSE/EOV diagnostic trace option.
005A	DALFUNC	Specifies the type of data set to be opened for the 3525 Card-Read-Punch-Print.
005B	DALFRID	Specifies input to the 3886 Character Reader.

Figure 72. Verb Code 01 (DCB Attributes) – Text Unit Keys, Mnemonics, and Functions

DCB Attribute Text Units

Use verb code 01 and the text unit keys listed in Figure 72 and described on the following pages to specify the DCB attributes of the data set being dynamically allocated. These attributes are described in *JCL* under the DCB parameter, and in *Data Management Macro Instructions*.

BFALN Specification - Key = X'0002E'

DALBFALN specifies the buffer alignment. It is mutually exclusive with the GAM-I/O count key (DALGNCP). When you code DALBFALN, # and LEN must be one, and PARM contains one of the following values:

X'01' for fullword not a doubleword boundary (F)
X'02' for doubleword boundary (D)

Example: to specify doubleword boundary, code

KEY	#	LEN	PARM
002E	0001	0001	02

BFTEK Specification - Key = X'002F'

DALBFTEK specifies the buffering technique to be used. It is mutually exclusive with the GAM-I/O count key (DALGNCP). When you code DALBFTEK, # and LEN must be one, and PARM contains one of the following values:

X'08' for dynamic buffering (D)
X'10' for exchange buffering (E)
X'20' for record buffering (R)
X'40' for simple buffering (S)
X'60' for record area buffering (A)

Example: to specify exchange buffering, code

KEY	#	LEN	PARM
002F	0001	0001	10

BLKSIZE Specification - Key = X'0030'

DALBLKSZ specifies the block size. It is mutually exclusive with the buffer size key (DALBUFSZ). When you code DALBLKSZ, # must be one, LEN must be two, and PARM contains the block size. The maximum PARM value is X'7FF8' (32,760).

Example: to specify a block size of 80, code

KEY	#	LEN	PARM
0030	0001	0002	00 50

BUFIN Specification - Key = X'0031'

DALBUFIN specifies the number of buffers to be initially assigned for receiving operations for each line in the line group. It is mutually exclusive with the buffer number (DALBUFNO) and buffer request (DALBUFRQ) keys. When you code DALBUFIN, # and LEN must be one, and PARM contains the number of buffers. The maximum PARM value is X'0F' (15).

Example: to specify 2 buffers, code

KEY	#	LEN	PARM
0031	0001	0001	02

BUFL Specification - Key = X'0032'

DALBUFL specifies the buffer length. When you code this key, # must be one, LEN must be two, and PARM contains the buffer length. The maximum PARM value is X'7FF8' (32,760).

Example: to specify a buffer length of 80, code

KEY	#	LEN	PARM
0032	0001	0002	00 50

BUFMAX Specification - Key = X'0033'

DALBUFMX specifies the maximum number of buffers to be allocated to a line at one time. It is mutually exclusive with the NCP key (DALNCP). When you code DALBUFMX, # and LEN must be one, and PARM contains the number of buffers. The maximum PARM value is X'0F' (15).

Example: to specify 4 buffers, code

KEY	#	LEN	PARM
0033	0001	0001	04

BUFNO Specification - Key = X'0034'

DALBUFNO specifies the number of buffers to be assigned to the data control block. It is mutually exclusive with the BUFIN (DALBUFIN), BUFOUT (DALBUFOU), and BUFRQ (DALBUFRQ) keys. When you code DALBUFNO, # and LEN must be one, and PARM contains the number of buffers.

Example: to specify 2 buffers, code

KEY	#	LEN	PARM
0034	0001	0001	02

BUFFOFF Specification - Key = X'0035'

DALBUFOF specifies the buffer offset. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'80'	the block prefix is four bytes long and contains the block length (L)
X'nn'	the length of the block prefix (maximum of X'63' (99))

Example: to specify an offset of 16, code

KEY	#	LEN	PARM
0035	0001	0001	10

BUFOUT Specification - Key = X'0036'

DALBUFOU specifies the number of buffers to be assigned initially for sending operations for each line in the group. It is mutually exclusive with the BUFNO (DALBUFNO) and BUFRQ (DALBUFRQ) keys. When you code DALBUFOU, # and LEN must be one, and PARM contains the number of buffers. The maximum PARM value is X'0F' (15).

Example: to specify 4 buffers, code

KEY	#	LEN	PARM
0036	0001	0001	04

BUFRQ Specification - Key = X'0037'

DALBUFRQ specifies the number of buffers to be requested in advance for the GET macro instruction. It is mutually exclusive with the BUFNO (DALBUFNO), BUFIN (DALBUFIN), and BUFOUT (DALBUFOU) keys. When you code DALBUFRQ, # and LEN must be one, and PARM contains the number of buffers.

Example: to specify 4 buffers, code

KEY	#	LEN	PARM
0037	0001	0001	04

BUFSZ Specification - Key = X'0038'

DALBUFSZ specifies the length in bytes of each of the buffers to be used for all lines in a particular line group. It is mutually exclusive with the blocksize key (DALBLKSZ). When you code DALBUFSZ, # must be one, LEN must be two, and PARM contains the buffer length.

Example: to specify a buffer length of 80, code

KEY	#	LEN	PARM
0038	0001	0002	00 50

CODE Specification - Key = X'0039'

DALCODE specifies the paper tape code in which the data is punched. It is mutually exclusive with the key length (DALKYLEN), MODE (DALMODE), printer spacing (DALPRTSP), STACK (DALSTACK), and TRTCH (DALTRTCH) keys. When you code DALCODE, # and LEN must be one, and PARM contains one of the following values:

X'02'	for Teletype 5-track (T)
X'04'	for USASCII 8-track (A)
X'08'	for National Cash Register 8-track (C)
X'10'	for Burroughs 7-track (B)
X'20'	for Friden 8-track (F)
X'40'	for IBM BCD 8-track (I)
X'80'	for no conversion (N)

Example: to specify USASCII, code

KEY	#	LEN	PARM
0039	0001	0001	04

CPRI Specification - Key = X'003A'

DALCPRI specifies the relative priority to be given to sending and receiving operations. It is mutually exclusive with the THRESH key (DALTHRS). When you code DALCPRI, # and LEN must be one, and PARM contains one of the following values:

X'01' for send priority (S)
X'02' for equal priority (E)
X'04' for receiving priority (R)

Example: to specify equal priority, code

KEY	#	LEN	PARM
003A	0001	0001	02

DEN Specification - Key = X'003B'

DALDEN specifies the magnetic tape density. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'03' for 200 bpi 7-track (0)
X'43' for 556 bpi 7-track (1)
X'83' for 800 bpi 7-track, 800 bpi 9 - track (2)
X'C3' for 1600 bpi 9-track (3)
X'D3' for 6250 bpi 9-track (4)

Example: to specify 1600 bpi 9 - track, code

KEY	##	LEN	PARM
003B	0001	0001	C3

DSORG Specifications - Key = X'003C'

DALDSORG specifies the data set organization. When you code this key, # must be one, LEN must be two, and PARM contains one of the following values:

X'0004' for TCAM 3705
X'0008' for VSAM
X'0020' for TCAM message queue (TQ)
X'0040' for TCAM line group (TX)
X'0080' for graphics (GS)
X'0200' for partitioned organization (PO)
X'0300' for partitioned organization unmovable (POU)
X'0400' for government of message transfer to or from a telecommunications message processing queue (MQ)
X'0800' for direct access message queue (CQ)
X'1000' for communication line group (CX)
X'2000' for direct access (DA)
X'2100' for direct access unmovable (DAU)
X'4000' for physical sequential (PS)
X'4100' for physical sequential unmovable (PSU)

Example: to specify Partitioned Organization, code

KEY	#	LEN	PARM
003C	0001	0002	02 00

EROPT Specification - Key = X'003D'

DALEROPT specifies the option to be executed if an error occurs in writing or reading a record. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'10'	for online BSAM testing (T)
X'20'	to cause abnormal end of task (ABE)
X'40'	to skip the block causing the error (SKP)
X'80'	to accept the block causing the error (ACC)

Example: to specify the SKP error option, code

KEY	#	LEN	PARM
003D	0001	0001	40

GNCP Specification - Key = X'003E'

DALGNCP specifies the maximum number of GAM input/output macros that will be issued before a WAIT macro is issued. It is mutually exclusive with the BFTEK (DALBFTK) and BFALN (DALBFAL) keys. When you code DALGNCP, # and LEN must be one, and PARM contains the GNCP value. The maximum PARM value is X'63' (99).

Example: to specify a GNCP value of four, code

KEY	#	LEN	PARM
003E	0001	0001	04

INTVL Specification - Key = X'003F'

DALINTVL specifies the polling interval for the lines in the line group. This key is mutually exclusive with the UCS (DALUCS) and FCB (DALFCB) keys. When you code this key, # and LEN must be one, and PARM contains the INTVL value.

Example: to specify an INTVL value of 10, code

KEY	#	LEN	PARM
003F	0001	0001	0A

KEYLEN Specification - Key = X'0040'

DALKYLEN specifies the length, in bytes, of the keys used in the data set. It is mutually exclusive with the CODE (DALCODE), MODE (DALMODE), PRTSP (DALPRTSP), STACK (DALSTACK), and TRTCH (DALTRTCH) keys. When you code this key, # and LEN must be one, and PARM contains the key length.

Example: to specify a key length of eight, code

KEY	#	LEN	PARM
0040	0001	0001	08

LIMCT Specification - Key = X'0041'

DALLIMCT specifies the search limit. When you code this key, # must be one, LEN must be three, and PARM contains the search limit value. The maximum PARM value is X'007FF8' (32,760).

Example: to specify a search limit of 1000, code

KEY	#	LEN	PARM
0041	0001	0003	0003E8

LRECL Specification - Key = X'0042'

DALLRECL specifies the actual or maximum length, in bytes, of a logical record. When you code this key, # must be one, LEN must be two, and PARM contains one of the following values:

X'8000' for variable length spanned records processed under QSAM and BSAM, the logical records exceed 32,756 bytes (X)

X'nnnn' the logical record length. The maximum value for nnnn is X'7FF8' (32,760).

Example: to specify a logical record length of 80, code

KEY	#	LEN	PARM
0042	0001	0002	0050

MODE Specification - Key = X'0043'

DALMODE specifies the mode of operation for a card reader or punch. It is mutually exclusive with the CODE (DALCODE), KEYLEN (DALKYLEN), PRTSP (DALPRTSP), and TRTCH (DALTRTCH) keys. When you code DALMODE, # and LEN must be one, and PARM contains one of the following values:

X'40' for EBCDIC mode (E)
X'50' for EBCDIC, read column eliminate mode (ER)
X'60' for EBCDIC, optical mark read mode (EO)
X'80' for card image mode (C)
X'90' for card image, read column eliminate mode (CR)
X'A0' for card image, optical mark read mode (CO)

Example: to specify EBCDIC mode, code

KEY	#	LEN	PARM
0043	0001	0001	40

NCP Specification - Key = X'0044'

DALNCP specifies the maximum number of READ or WRITE macros issued before a CHECK macro is issued. It is mutually exclusive with the BUFMAX (DALBUFMAX) key. When you code DALNCP, # and LEN must be one, and PARM contains the NCP value. The maximum PARM value is X'63' (99).

Example: to specify an NCP value of two, code

KEY	#	LEN	PARM
0044	0001	0001	02

OPTCD Specification - Key = X'0045'

DALOPTCD specifies optional services to be performed by the control program. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'01'	for relative block addressing (R), or to select character arrangement tables for the 3800 printer (J)
X'02'	for user totaling facility (T)
X'04'	for reduced tape error recovery or direct DASD search (Z)
X'08'	for direct addressing (A), or for translation of ASCII to or from EBCDIC (Q)
X'10'	for feedback (F), or for hopper-empty exit (H), or for online correction for optical readers (O)
X'20'	for chained scheduling or TCAM segment identification (C), or for extended search (E)
X'40'	for disregarding end-of-file recognition for tape (B), or for allowance of data checks caused by an invalid character, or for handling a TCAM work unit as a message (U)
X'80'	for write validity check, or to place TCAM message source in an eight-byte field in the workarea (W)

Note: When you are specifying more than one OPTCD value, PARM contains the **sum** of the values.

For more information regarding the OPTCD specification key, see *MVS/Extended Architecture JCL* and *Data Management Macros*.

Example: to specify OPTCD value U, code

KEY	#	LEN	PARM
0045	0001	0001	40

Example: to specify OPTCD values U and C, code

KEY	#	LEN	PARM
0045	0001	0001	60

Receiving PCI Specification - Key = X'0046'

DALPCIR specifies the relationship of program-controlled interrupts (PCI) during receiving operations to the allocation and freeing of buffers. When you code DALPCIR, # and LEN must be one, and PARM contains one of the following values:

X'02'	for a PCI and no new buffer allocated (R)
X'08'	for no PCIs (N)
X'20'	for a PCI and new buffer allocated (A)
X'80'	for a PCI, new buffer allocated, and the first buffer remains allocated (X)

Example: to specify no PCIs during receiving operations, code

KEY	#	LEN	PARM
0046	0001	0001	08

Sending PCI Specification - Key = X'0047'

DALPCIS specifies the relationship of PCIs during sending operations to the allocation and freeing of buffers. When this key is specified, # and LEN contain one, and PARM contains:

X'01' for a PCI and no new buffer allocated (R)
X'04' for no PCIs (N)
X'10' for a PCI and a new buffer allocated (A)
X'40' for a PCI, new buffer allocated, and first buffer remains allocated (X)

Example: to specify no PCIs during sending operations, code

KEY	#	LEN	PARM
0047	0001	0001	04

PRTSP Specification - Key = X'0048'

DALPRTSP specifies printer line spacing. It is mutually exclusive with the CODE (DALCODE), KEYLEN (DALKYLEN), MODE (DALMODE), STACK (DALSTACK), and TRTCH (DALTRTCH) keys. When you code DALPRTSP, # and LEN must be one, and PARM contains one of the following values:

X'01' for no spacing (0)
X'09' for one-line spacing (1)
X'11' for two-line spacing (2)
X'19' for three-line spacing (3)

Example: to specify no spacing, code

KEY	#	LEN	PARM
0048	0001	0001	01

RECFM Specification - Key = X'0049'

DALRECFM specifies the record format. When you code this key, # and LEN must be one, and PARM contains one of the following values:

X'02' for machine code printer control characters in record (M), or for complete QTAM record (R)
X'04' for ASA printer control characters in record (A), or for complete QTAM message (G)
X'08' for standard fixed records, spanned variable records, or segment of QTAM message (S)
X'10' for blocked records (B)
X'20' for variable ASCII records (D), or for track overflow (T)
X'40' for variable records (V)
X'80' for fixed records (F)
X'C0' for undefined records (U)

Note: When you code combinations of RECFM values, PARM contains the **sum** of the values.

Example: to specify fixed records, code

KEY	#	LEN	PARM
0049	0001	0001	80

Example: to specify variable blocked (VB) records, code

KEY	#	LEN	PARM
0049	0001	0001	50

First Buffer Reserve Specification - Key = X'004A'

DALRSRVF specifies the number of bytes to be reserved in the first buffer for insertion of data by the DATETIME and SEQUENCE macros. The UCS (DALUCS) key is mutually exclusive with DALRSRVF. When you code this key, # and LEN must be one, and PARM contains the number of bytes to reserve.

Example: to reserve 8 bytes in the first buffer, code

KEY	#	LEN	PARM
004A	0001	0001	08

Secondary Buffer Reserve Specification - Key = X'004B'

DALRSRVS specifies the number of bytes to be reserved in buffers other than the first for insertion of data by the DATETIME and SEQUENCE macros. The UCS (DALUCS) key is mutually exclusive with DALRSRVS. When you code this key, # and LEN must be one, and PARM contains the number of bytes to reserve.

Example: to reserve 8 bytes in secondary buffers, code

KEY	#	LEN	PARM
004B	0001	0001	08

SOWA Specification - Key = X'004C'

DALSOWA specifies the size, in bytes, of the user-provided input work areas for telecommunication jobs. When you code this key, # must be one, LEN must be two, and PARM contains the number of bytes. The maximum PARM value is X'7FF8' (32,760).

Example to specify a 256-byte work area, code

KEY	#	LEN	PARM
004C	0001	0002	0100

STACK Specification - Key = '004D'

DALSTACK specifies the stacker bin to receive cards. The CODE (DALCODE), KEYLEN (DALKEYLEN), PRTSP (DALPRTSP), and TRTCH (DALTRTCH) keys are mutually exclusive with DALSTACK. When you code this key, # and LEN are one, and PARM contains one of the following values:

X'01'	for bin 1 (1)
X'02'	for bin 2 (2)

Example: to specify stacker 2, code

KEY	#	LEN	PARM
004D	0001	0001	02

THRESH Specification - Key = X'004E'

DALTHRSRSH specifies the percentage of nonreusable disk message queue records to be used before a flush closedown occurs. The CPRI (DALCPRI) key is mutually exclusive with DALTHRSRSH. When you code this key, # and LEN must be one, and PARM contains the percentage. The maximum PARM value is X'64' (100).

Example: to specify a THRESH percentage of 99, code

KEY	#	LEN	PARM
004E	0001	0001	63

TRTCH Specification - Key = X'004F'

DALTRTCH specifies the recording technique for 7-track tape. It is mutually exclusive with the CODE (DALCODE), KEYLEN (DALKEYLEN), MODE (DALMODE), PRTSP (DALPRTSP), and STACK (DALSTACK) keys. When you code DALTRTCH, # and LEN must be one, and PARM contains one of the following values:

X'13'	for data conversion (C)
X'23'	for even parity (E)
X'2B'	for even parity and BCD/EBCDIC translation (ET)
X'3B'	for BCD/EBCDIC translation (T)

Example: to specify even parity, code

KEY	#	LEN	PARM
004F	0001	0001	23

IPLTXTID Specification - Key = X'0051'

DALIPLTX specifies the name of a TCAM network control program. It is mutually exclusive with the DSNAME (DALDSNAM), MEMBER NAME (DALMEMBR), and QNAME (DALQNAME) keys. When you code DALIPLTX, # must be one, LEN is the length of the name (maximum of 8), and PARM contains the name.

Example: to specify an IPLTXTID value of PGM, code

KEY	#	LEN	PARM
0051	0001	0003	D7 C7 D4

Diagnostic Trace Specification (DIAGNS = TRACE) - Key = X'0054'

DALDIAGN requests the OPEN/CLOSE/EOV trace option, which gives a module-by-module trace of OPEN/CLOSE/EOV's work area and the user's DCB. When you code DALDIAGN, # must be zero; LEN and PARM are not specified.

Note: GTF must be active in the system while the job that requested the trace is running.

Example: to specify the diagnostic trace specification, code

KEY	#	LEN	PARM
0054	0000	-	-

FUNC = Specification - Key = X'005A'

DALFUNC can be used with BSAM and QSAM; it specifies the type of data set to be opened for the 3525 Card Read-Punch-Print. When you code DALFUNC, # and LEN must be one, and PARM contains one of the following values:

X'10'	for W
X'12'	for WT
X'14'	for WX
X'16'	for WXT
X'20'	for P
X'30'	for PW
X'34'	for PWX
X'36'	for PWXT
X'40'	for R
X'50'	for RW
X'52'	for RWT
X'54'	for RWX
X'56'	for RWXT
X'60'	for RP
X'68'	for RPD
X'70'	for RPW
X'74'	for RPWX
X'76'	for RPWXT
X'78'	for RPWD
X'80'	for I

Where:

D	is data protection for a punch data set
I	is interpret punch data set
P	is punch
R	is read
T	is two line printer
W	is print
X	is printer

Notes:

- 1. In the absence of this information, the system assumes P.*
- 2. D, X, and T cannot be coded alone.*
- 3. If you specify D as part of a value, you must also specify the FCB image-id key (DALFCBIM), giving the image identifier for the data protection image.*

Example: to specify FUNC=RPWD, code

KEY	#	LEN	PARM
005A	0001	0001	78

FRID = Specification - Key = X'005B'

DALFRID specifies the last four characters of a SYS1.IMAGELIB member name to be used in the interpretation of documents for input to the IBM 3886 character reader. The FCB (DALFCBIM) key is mutually exclusive with DALFRID.

When you code DALFRID, # must be one, LEN is the number of characters specified, and PARM contains the characters of the IMAGELIB member name. The characters must be alphanumeric or national. If the length of the member name is four or less, code the entire name.

Example: to specify the last four characters of member name SHARK1, code

KEY	#	LEN	PARM
005B	0001	0004	C1 D9 D2 F1

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0050	DALPASSW	Specifies the password for a protected data set.
0052	DALPERMA	Specifies the permanently allocated attribute.
0053	DALCNVRT	Specifies the convertible attribute.
0055	DALRTDDN	Requests the return of the associated ddname.
0056	DALRTDSN	Requests the return of the allocated data set's name.
0057	DALRTORG	Requests the return of data set organization.
005C	DALSSREQ	Specifies allocation of a subsystem data set.
005D	DALRTVOL	Requests the return of the volume serial number.
0062	DALSSATT	Specifies allocation of a subsystem data set to SYSIN.

Figure 73. Verb Code 01 (Non-JCL Dsname Functions) -- Text Unit Keys, Mnemonics, and Functions

Non-JCL Dynamic Allocation Functions

The keys listed in Figure 73 and described on the following pages do not have JCL equivalents; they have meaning only to the SVC 99 routines in performing dynamic allocation by dsname (verb code 01).

Password Specification - Key = X'0050'

DALPASSW specifies the password for a password-protected data set. The dsname key (DALDSNAM) is required with this key. When you code DALPASSW, # must be one, LEN contains the length of the password, and PARM contains the password.

Example: to specify the password, MYKEY, code

KEY	#	LEN	PARM
0050	0001	0005	D4 E8 D2 C5 E8

Permanently Allocated Attribute Specification - Key = X'0052'

DALPERMA specifies that the permanently allocated attribute is to be assigned to this allocation. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify assignment of the permanently allocated attribute, code

KEY	#	LEN	PARM
0052	0000	-	-

Convertible Attribute Specification - Key = X'0053'

DALCNVRT specifies that the convertible attribute is to be assigned to this allocation.

Note: This specification is the default if the permanently allocated attribute key (DALPERMA) is not coded.

When you code DALCNVRT, # must be zero; LEN and PARM are not specified.

Example: to specify assignment of the convertible attribute, code

KEY	#	LEN	PARM
0053	0000	-	-

Ddname Return Specification - Key = X'0055'

DALRTDDN requests that the ddname associated with the allocation be returned to the caller of SVC 99. When you code DALRTDDN, # must be one, LEN must be eight, and PARM is an eight-byte field. The SVC 99 routines place the allocated ddname in PARM and update LEN to the length of the ddname.

Example: to request that the allocated ddname be returned, code

KEY	#	LEN	PARM
0055	0001	0008	-----

This specification would be updated upon the assignment of the ddname DD1 as follows:

KEY	#	LEN	PARM
0055	0001	0003	C4 C4 F1-----

Dsname Return Specification - Key = X'0056'

DALRTDSN requests that the dsname that is allocated be returned to the caller of SVC 99. When you code DALRTDSN, # must be one, LEN must be forty-four, and PARM is a forty-four byte field. The SVC 99 routines place the allocated dsname in PARM and update LEN to the length of the dsname.

Example: to request that the allocated dsname be returned, code

KEY	#	LEN	PARM
0056	0001	002C	-----...--

This specification would be updated for the allocation of the dsname ABC as follows:

KEY	#	LEN	PARM
0056	0001	0003	C1 C2 C3-----...--

DSORG Return Specification - Key = X'0057'

DALRTORG requests that the data set organization of the allocated data set be returned to the caller of SVC 99. When you code DALRTORG, # must be one, LEN must be two, and PARM is a two-byte field. The SVC 99 routines put one of the following values into PARM:

X'0000'	if the dynamic allocation routines cannot determine the DSORG
X'0004'	if TR
X'0008'	if VSAM
X'0020'	if TQ
X'0040'	if TX
X'0080'	if GS
X'0200'	if PO
X'0300'	if POU
X'0400'	if MQ
X'0800'	if CQ
X'1000'	if CX
X'2000'	if DA
X'2100'	if DAU
X'4000'	if PS
X'4100'	if PSU
X'8000'	if IS
X'8100'	if ISU

Example: to specify that the DSORG be returned, code

KEY	#	LEN	PARM
0057	0001	0002	--

This specification would be updated for a DSORG of PS as follows:

KEY	#	LEN	PARM
0057	0001	0002	4000

Subsystem Request Specification - key = X'005C'

DALSSREQ requests that a subsystem data set be allocated and, optionally, specifies the name of the subsystem for which the data set is to be allocated.

When you code DALSSREQ without specifying a subsystem name, # must be zero and LEN and PARM are not specified. The data set is then allocated to the primary subsystem.

When you code the subsystem name in the DALSSREQ key, # must be one, LEN is the length of the name (maximum of 4), and PARM contains the subsystem name.

Note: To specify DALSSREQ, your program must be APF authorized, in supervisor state, or running in a system protection key.

Example 1: to request a subsystem data set for the primary subsystem, code:

KEY	#	LEN	PARM
005C	0000	-	-

Example 2: to request a subsystem data set for JES2, code:

KEY	#	LEN	PARM
005C	0001	0004	D1 C5 E2 F2

Volume Serial Return Specification - Key = X'005D'

DALRTVOL requests that the volume serial number associated with the allocated data set be returned. Only the first volume serial of a multiple-volume data set is returned, and the volume sequence number, if any, is ignored.

When you code DALRTVOL, # must be one, LEN must be six, and PARM is a six-byte field.

If the allocated volume serial is available at the completion of allocation, the SVC 99 routines put the number in PARM. If the volume serial is not available at the completion of allocation, the SVC 99 routines set LEN to zero.

The volume serial will not be available at the completion of allocation if either of the following is true:

- No volume serial is allocated to the data set (a VIO or job entry subsystem data set)
- The request results in the allocation of a new data set on magnetic tape without a specific volume serial having been assigned.

Example: to specify that the allocated volume serial be returned, code

KEY	#	LEN	PARM
005D	0001	0006	-----

This specification would be updated for the allocation of data set ABC on volume 123456 as follows:

KEY	#	LEN	PARM
005D	0001	0006	F1 F2 F3 F4 F5 F6

Subsystem Request Type Specification - Key = X'0062'

DALSSATT specifies that the subsystem data set being requested is to be allocated to SYSIN. The subsystem request key (DALSSREQ) is required with this key.

Note: To specify DALSSATT, you must be APF authorized, or in supervisor state, or running in a system protection key.

When you code DALSSATT, # and LEN must be one, and PARM contains X'80', for SYSIN data set.

Example: to specify a subsystem SYSIN data set, code:

KEY	#	LEN	PARM
0062	0001	0001	80

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0001	DUNDDNAM	Specifies the ddname of the resource to be deallocated.
0002	DUNDSNAM	Specifies the data set to be deallocated.
0003	DUNMEMBR	Specifies the PDS member to be deallocated.
0005	DUNOVDSP	Specifies an overriding disposition.
0007	DUNUNALC	Specifies deallocation even if the resource has the permanently allocated attributed.
0008	DUNREMOV	Specifies removal of the "in-use" attribute, even if the resource has the permanently allocated attribute.
000A	DUNOVSNH	Specifies "nohold" status for a deallocated SYSOUT data set.
0018	DUNOVCLS	Specifies an overriding SYSOUT class.
0058	DUNOVSSUS	Specifies an overriding remote workstation.
0059	DUNOVSHQ	Puts the SYSOUT data set on the hold queue and overrides previous "nohold" specifications.

Figure 74. Verb Code 02 (Dynamic Unallocation) – Text Unit Keys, Mnemonics, and Functions

Dynamic Unallocation Text Units

Use verb code 02 and the text unit keys listed in Figure 74 and described on the following pages to request dynamic unallocation processing by the SVC 99 routines.

Ddname Specification - Key = X'0001'

DUNDDNAM specifies the ddname of the resource to be deallocated. When you code this key, # must be one, LEN is the length of the ddname, and PARM contains the ddname.

Example: to specify the ddname DD1, code

```
KEY      #      LEN      PARM
0001    0001    0003    C4 C4 F1
```

Dsname Specification - Key = X'0002'

DUNDSNAM specifies the data set name to be deallocated. When you code this key, # must be one, LEN contains the length of the dsname, and PARM contains the dsname.

Example: to specify the dsname MYDATA, code

```
KEY      #      LEN      PARM
0002    0001    0006    D4 E8 C4 C1 E3 C1
```

Member name Specification - Key = X'0003'

DUNMEMBR specifies that a particular member of the data set is to be deallocated. The dsname unallocation key (DUNDSNAM) is required with this key. When you code DUNMEMBR, # must be one, LEN is the length of the member name, and PARM contains the member name.

Example: to specify the member name MEM1, code

```
KEY      #      LEN      PARM
0003    0001    0004    D4 C5 D4 F1
```


Overriding Disposition Specification - Key = X'0005'

DUNOVDSP specifies a disposition that overrides the disposition assigned to a data set when it was allocated. When you code DUNOVDSP, # and LEN must be one, and PARM contains one of the following values:

X'01' for an overriding disposition of UNCATLG
X'02' for an overriding disposition of CATLG
X'04' for an overriding disposition of DELETE
X'08' for an overriding disposition of KEEP

Example: to specify an overriding disposition of CATLG, code

KEY	#	LEN	PARM
0005	0001	0001	02

Note: The SVC 99 routines ignore this key if any of the following are true:

- The overriding disposition was DELETE and the data set was originally allocated as SHARE.
- The data set was originally allocated with a disposition of PASS.
- The data set is a VSAM data set.
- The data set is a non-subsystem data set that has a system-generated name.

When the SVC 99 routines must ignore a DUNOVDSP request, they still perform the deallocation processing, but use the disposition from the original allocation request.

Unalloc Option Specification - Key = X'0007'

DUNUNALC specifies that the resource is to be deallocated even if it has the permanently allocated attribute. The remove option key (DUNREMOV) is mutually exclusive with DUNUNALC. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify the unalloc option, code

KEY	#	LEN	PARM
0007	0000	-	-

Remove Option Specification - Key = X'0008'

DUNREMOV specifies that the in-use attribute is to be removed even if the resource does not have the permanently allocated attribute. The unalloc option key (DUNUNALC) is mutually exclusive with DUNREMOV. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify the remove option, code

KEY	#	LEN	PARM
0008	0000	-	-

Overriding SYSOUT Nohold Specification - Key = X'000A'

DUNOVSNH specifies that the SYSOUT data set being deallocated is not to be placed on the hold queue. This specification overrides the HOLD/NOHOLD specification assigned when the data set was allocated.

This key is ignored if the data set is not a SYSOUT data set. The overriding hold key (DUNOVSHQ) is mutually exclusive with DUNOVSNH. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify nohold, code

KEY	#	LEN	PARM
000A	0000	-	-

Overriding SYSOUT Class Specification - Key = X'0018'

DUNOVCLS specifies a SYSOUT class that overrides the class assigned when the SYSOUT data set was allocated. This key is ignored if the resource is not a SYSOUT data set. When you code DUNOVCLS, # and LEN must be one, and PARM contains the overriding class.

Example: to specify an overriding class of C, code

KEY	#	LEN	PARM
0018	0001	0001	C3

Overriding SYSOUT Remote Workstation Specification - Key = X'0058'

DUNOVSSUS specifies that the SYSOUT data set being deallocated is to be routed to a remote user. This specification overrides the remote workstation specification assigned when the data set was allocated. DUNOVSSUS is ignored if the data set is not a SYSOUT data set.

When you code DUNOVSSUS, # must be one, LEN is the length of the remote workstation name (maximum of 8), and PARM contains the remote workstation (user) name.

Example: to specify the remote work station USER01, code

KEY	#	LEN	PARM
0058	0001	0006	E4 E2 C5 D9 F0 F1

Overriding SYSOUT Hold Queue Specification - Key = X'0059'

DUNOVSHQ specifies that the SYSOUT data set being deallocated is to be placed on the hold queue. This specification overrides the HOLD/NOHOLD specification assigned when the data set was allocated. This key is ignored if the data set is not a SYSOUT data set. The overriding nohold key (DUNOVSNH) is mutually exclusive with this key.

When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify hold, code

KEY	#	LEN	PARM
0059	0000	-	-

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0001	DCCDDNAM	Specifies the ddnames to be concatenated.
0004	DCCPERMC	Specifies the permanently concatenated attribute.

Figure 75. Verb Code 03 (Dynamic Concatenation) – Text Unit Keys, Mnemonics, and Functions

Dynamic Concatenation Text Units

Use verb code 03 and the text units listed in Figure 75 and described in the following paragraphs to request dynamic concatenation processing by the SVC 99 routines.

Ddname Specification - Key = X'0001'

DCCDDNAM specifies the ddnames that are associated with the data sets to be concatenated. When you code DCCDDNAM, # is the number of ddnames being specified (a minimum of two), LEN is the length of the immediately following ddname, and PARM contains the ddname.

Example: to specify concatenation of SYSLIB to MYLIB, code

KEY	#	LEN	PARM	LEN	PARM
0001	0002	0005	D4E8D3C9C2	0006	E2E8E2D3C9C2

Permanently Concatenated Attribute Specification - Key = X'0004'

DCCPERMC specifies that the concatenated group be assigned the permanently concatenated attribute. When you code this key, # must be zero; LEN and PARM are not specified.

Example: to specify assignment of the permanently concatenated attribute, code

KEY	#	LEN	PARM
0004	0000	-	-

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0001	DDCDDNAM	Specifies the ddname of the group to be deconcatenated.

Figure 76. Verb Code 04 (Dynamic Deconcatenation) – Text Unit Key, Mnemonic, and Function

Dynamic Deconcatenation Text Unit

Use verb code 04 and the following text unit to request dynamic deconcatenation processing by the SVC 99 routines

Ddname Specification - Key = X'0001'

DDCDDNAM specifies the ddname of the concatenated group that is to be deconcatenated. DDCDDNAM is required for dynamic deconcatenation.

When you code DDCDDNAM, # must be one, LEN is the length of the ddname and PARM contains the ddname.

Example: to request the deconcatenation of the group of data sets associated with the ddname DD1, code

```
KEY      #      LEN      PARM
0001     0001   0003     C4 C4 F1
```

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0001	DRITCBAD	Removes the "in-use" attribute from all resources associated with the specified TCB address.
0002	DRICURNT	Removes the "in-use" attribute from all resources but those of the current task and its higher-level tasks.

Figure 77. Verb Code 05 (Remove-In-Use Processing Based on Task-ID) – Text Unit Keys, Mnemonics, and Functions

Text Units for Removing the In-Use Attribute Based on Task-ID

Use verb code 05 and the text units in Figure 77 and described as follows to request that the SVC 99 routines turn off the in-use bits for resources based on task-ID.

TCB Address Specification - Key = X'0001'

DRITCBAD specifies that the in-use attribute is to be removed from all resources associated with the specified TCB address. The current task option key (DRICURNT) is mutually exclusive with this key.

When you code DRITCBAD, # must be one, LEN must be four, and PARM contains the TCB address.

Example: to specify the TCB address 22AC0, code

```
KEY      #      LEN      PARM
0001     0001   0004     00022AC0
```

Current Task Option Specification - Key = X'0002'

DRICURNT specifies that the in-use attribute is to be removed from all resources except those associated with the current task, its direct ancestors, and the initiator. This key is mutually exclusive with the TCB address key (DRITCBAD). When you code DRICURNT, # must be zero; LEN and PARM are not specified.

Example: to specify the current task option, code

```
KEY      #      LEN      PARM
0002    0000    -        -
```

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0001	DDNDDNAM	Specifies the ddname to be allocated.
0002	DDNRTDUM	Requests a dummy data set indication.

Figure 78. Verb Code 06 (Ddname Allocation) – Text Unit Keys, Mnemonics, and Functions

Ddname Allocation Text Units

Use verb code 06 and the text units listed in Figure 78 and described as follows to request ddname allocation processing. In ddname allocation, you are specifying that the SVC 99 routines are to use a particular existing allocation to satisfy your allocation request. You identify the data set you wish by specifying the ddname associated with it.

Ddname Specification - Key = X'0001'

DDNDDNAM specifies the ddname of the resource to be allocated. It is required for dynamic allocation by ddname.

When you code DDNDDNAM, # must be one, LEN contains the length of the ddname, and PARM contains the ddname.

Example: to specify the ddname SYSLIB, code

```
KEY      #      LEN      PARM
0001    0001    0006    E2 E8 E2 D3 C9 C2
```

Return DUMMY Indication Specification - Key = X'0002'

Code DDNRTDUM to request the return of an indication if the ddname specified in DDNDDNAM is associated with a dummy data set. When you code DDNRTDUM, # and LEN must be one, and PARM is a one-byte field. The SVC 99 routines set PARM as follows:

```
X'80'   if the ddname is associated with a dummy data set
X'00'   otherwise
```

Example: to specify that the DUMMY indication be returned, code

```
KEY      #      LEN      PARM
0002    0001    0001    -
```

Hex Text Unit Key	IEFZB4D2 Mnemonic	SVC 99 Function
0001	DINDDNAM	Specifies the ddname identifier of the requested information.
0002	DINDSNAM	Specifies the data set for which the information is requested.
0004	DINRTDDN	Requests the return of the associated ddname.
0005	DINRTDSN	Requests the return of the data set name.
0006	DINRTMEM	Requests the return of the PDS member name.
0007	DINRTSTA	Requests the return of the data set's status.
0008	DINRTNDP	Requests the return of the data set's normal disposition.
0009	DINRTCDP	Requests the return of the data set's conditional disposition.
000A	DINRTORG	Requests the return of the data set's organization.
000B	DINRTLIM	Requests the number of resources that must be deallocated before making a new allocation.
000C	DINRTATT	Requests the return of special attribute indications.
000D	DINRTLST	Requests the return of a last relative entry indication.
000E	DINRTTYP	Requests the return of the data set's type (terminal or dummy).
000F	DINRELNO	Specifies the desired allocation information retrieval by relative request number.

Figure 79. Verb Code 07 (Dynamic Information Retrieval) – Text Unit Keys, Mnemonics, and Functions

Dynamic Information Retrieval Text Units

Use verb code 07 and the text units listed in Figure 79 and described as follows to request that the SVC 99 routines return certain information about the allocated resources.

Ddname Specification - Key = X'0001'

DINDDNAM specifies the ddname associated with the allocation you are requesting information about. It is mutually exclusive with the dsname (DINDSNAM) and relative entry (DINRELNO) keys. When you code DINDDNAM, # must be one, LEN is the length of the ddname, and PARM contains the ddname.

Example: to specify the ddname DD1, code

```
KEY    #    LEN    PARM
0001  0001  0003  C4 C4 F1
```

Dsname Specification - Key = X'0002'

DINDSNAM specifies the dsname of the allocated resource you are requesting information about. It is mutually exclusive with the ddname (DINDDNAM) and relative entry (DINRELNO) keys. When you code DINDSNAM, # must be one, LEN is the length of the dsname, and PARM contains the dsname.

Example: to specify the dsname MYDATA, code

KEY	#	LEN	PARM
0002	0001	0006	D4 E8 C4 C1 E3 C1

Return Ddname Specification - Key = X'0004'

Code DINRTDDN to request the return of the ddname associated with the specified allocation. When you code this key, # must be one, LEN must be eight, and PARM is an eight-byte field. Upon return to your program, PARM will contain the requested ddname, and LEN will be set to its length.

Example: to request the return of the ddname, code

KEY	#	LEN	PARM
0004	0001	0008	-----

Return Dsname Specification - Key = X'0005'

Code DINRTDSN to request the return of the dsname of the specified allocation. When you code this key, # must be one, LEN must be forty-four, and PARM is a forty-four byte field. Upon return to your program, PARM will contain the dsname and LEN will be set to its length.

Example: to request that the dsname be returned, code

KEY	#	LEN	PARM
0005	0001	002C	-----

Return Member Name Specification - Key = X'0006'

Code DINRTMEM to request the return of the member name associated with the specified allocation. When you code this key, # must be one, LEN must be eight, and PARM is an eight-byte field. Upon return to your program, PARM will contain the member name and LEN will be set to its length (or to zero, if none).

Example: to request that the member name be returned, code

KEY	#	LEN	PARM
0006	0001	0008	-----

Return Status Specification - Key = X'0007'

Code DINRTSTA to request the return of the data set status of the specified allocation. When you code this key, # and LEN must be one, and PARM is a one-byte field. Upon return to your program, the PARM field will contain one of the following values:

X'01'	for OLD
X'02'	for MOD
X'04'	for NEW
X'08'	for SHR

Example: to request that the status be returned, code

KEY	#	LEN	PARM
0007	0001	0001	-

Return Normal Disposition Specification - Key = X'0008'

Code DINRTNDP to request the return of the normal disposition of the specified resource. When you code this key, # and LEN must be one. PARM is a one-byte field. Upon return to your program, PARM will contain one of the following values:

X'01' for UNCATLG
X'02' for CATLG
X'04' for DELETE
X'08' for KEEP
X'10' for PASS

Example: to request that the normal disposition be returned, code

KEY	#	LEN	PARM
0008	0001	0001	-

Return Conditional Disposition Specification - Key = X'0009'

Code DINRTCDP to request the return of the conditional disposition of the specified resource. The values for #, LEN and PARM are the same as for the return normal disposition key (DINRTNDP).

Example: to request that the conditional disposition be returned, code

KEY	#	LEN	PARM
0009	0001	0001	-

Return Data Set Organization Specification Key = X'000A'

Code DINRTORG to request the return of the data set organization (DSORG) of the specified resource. When you code this key, # must be one, LEN must be two, and PARM is a two-byte field. Upon return to your program, PARM will contain one of the following:

X'0000' if undetermined
X'0004' if TR
X'0008' for VSAM
X'0020' if TQ
X'0040' if TX
X'0080' for GS
X'0200' for PO
X'0300' for POU
X'0400' for MQ
X'0800' for CQ
X'1000' for CX
X'2000' for DA
X'2100' for DAU
X'4000' for PS
X'4100' for PSU
X'8000' for IS
X'8100' for ISU

Example: to request that the data set organization be returned, code

KEY	#	LEN	PARM
000A	0001	0002	--

Return Limit Specification - Key = X'000B'

Code DINRTLIM to request the return of the number of resources that must be deallocated before a new allocation can be made. When you code this key, # must be one, LEN must be two, and PARM is a two-byte field. Upon return to your program, PARM is set to the number of resources to be deallocated.

Example: to request that the number of not-in-use data sets over the control limit be returned, code

KEY	#	LEN	PARM
000B	0001	0002	--

If three data sets must be deallocated, the SVC 99 routines return DINRTLIM as follows:

KEY	#	LEN	PARM
000B	0001	0002	0003

Return Dynamic Allocation Attribute Specification - Key = X'000C'

Code DINRTATT to request indications of the attributes assigned to the specified resource. When you code this key, # and LEN must be one, and PARM is a one-byte field. Upon return to your program, PARM is set as follows:

Bit 0	on, if permanently concatenated
Bit 1	on, if in use
Bit 2	on, if permanently allocated
Bit 3	on, if convertible
Bit 4	on, if dynamically allocated
Bits 5-7	reserved

Example: to request return of the data set attributes, code

KEY	#	LEN	PARM
000C	0001	0001	--

If the allocation has the in-use and permanently allocated attributes, PARM contains the following on return:

KEY	#	LEN	PARM
000C	0001	0001	60

Return Last Entry Specification - Key = X'000D'

Code DINRTLST to determine if the relative entry number you specify is the last relative entry. When you code DINRTLST, # and LEN must be one, and PARM is a one-byte field. Upon return to your program, PARM contains one of the following values:

X'80'	if last relative entry
X'00'	otherwise

Example: to request the return of the last entry, code

KEY	#	LEN	PARM
000D	0001	0001	-

Return Data Set Type Specification - Key = X'000E'

Code DINRTTYP to determine the type of the specified data set. When you code this key, # and LEN must be one, and PARM is a one-byte field. Upon return to your program, PARM contains one of the following values:

X'80'	if a DUMMY data set
X'40'	if a terminal allocation
X'20'	if a SYSIN data set
X'10'	if a SYSOUT data set
X'00'	otherwise

KEY	#	LEN	PARM
000E	0001	0001	-

Relative Request Number Specification - Key = X'000F'

DINRELNO specifies the relative request number of the allocation you are requesting information about. It is mutually exclusive with the ddname (DINDDNAM) and dsname (DINDSNAM) keys. When you code DINRELNO, # must be one, LEN must be two, and PARM contains the relative number.

Example: to specify that information is to be returned about your tenth SVC 99 request, code

KEY	#	LEN	PARM
000F	0001	0002	000A

Example of a Dynamic Allocation Request

The assembler language example in Figure 80 is a dynamic allocation request allocating SYS1.LINKLIB with a status of SHARE. It also requests that the SVC 99 routines return the ddname associated with SYS1.LINKLIB.

Figure 81 shows the parameter list that results from the code in Figure 80.

```

DYN CSECT
  USING *,15
  STM 14,12,12(13)
  BALR 12,0
BEGIN DS OH
  USING BEGIN,12
  LA 0,50
  GETMAIN R,LV=(0)
  LR 8,1
  USING S99RBP,8
  LA 4,S99RBPTR+4
  USING S99RB,4
  ST 4,S99RBPTR
  OI S99RBPTR,S99RBPND
  XC S99RB(RBLN),S99RB
  MVI S99RBLN,RBLN
  MVI S99VERB,S99VRBAL
  LA 5,S99RB+RBLN
  USING S99TUPL,5
  ST 5,S99TXTTP
  LA 6,DSNTU
  ST 6,S99TUPTR
  LA 5,S99TUPL+4
  LA 6,STATUSTU
  ST 6,S99TUPTR
  LA 6,S99TUPL+8
  USING S99TUNIT,6
  LA 5,S99TUPL+4
  ST 6,S99TUPTR
  OI S99TUPTR,S99TUPLN
  MVC S99TUNIT(14),RETDDN
  LR 1,8
  DYNALLOC
  LM 14,12,12(13)
  BR 14
  RETURN TO CALLER.
RBLN EQU (S99RBEND-S99RB)
DSNTU DC AL2(DALDSNAM)
      DC X'0001'
      DC X'000C'
      DC C'SYS1.LINKLIB'
STATUSTU DC AL2(DALSTATS)
        DC X'0001'
        DC X'0001'
        DC X'08'
RETDDN DC AL2(DALRTDDN)
       DC X'0001'
       DC X'0008'
       DS CL8
       IEFZB4D0
       IEFZB4D2
DYN CSECT
  END

```

AMOUNT OF STORAGE REQUIRED FOR THIS REQUEST.
GET THE STORAGE NECESSARY FOR THE REQUEST.
SAVE THE ADDRESS OF THE RETURNED STORAGE.
ESTABLISH ADDRESSABILITY FOR S99RBP DSECT.
POINT FOUR BYTES BEYOND START OF S99RBPTR.
ESTABLISH ADDRESSABILITY FOR RB DSECT.
MAKE 'RBPTR' POINT TO RB.
TURN ON THE HIGH-ORDER BIT IN RBPTR.
ZERO OUT 'RB' ENTIRELY.
PUT THE LENGTH OF 'RB' IN ITS LENGTH FIELD.
SET THE VERB CODE FIELD TO ALLOCATION FUNCTION.
POINT PAST 'RB' TO START OF TUP LIST.
ESTABLISH ADDRESSABILITY FOR TEXT UNIT PTRS.
STORE ADDRESS OF TUP LIST IN THE RB.
GET ADDRESS OF FIRST TEXT UNIT
AND STORE IN TUP LIST.
GET ADDRESS OF NEXT TUP LIST ENTRY.
GET ADDRESS OF SECOND TEXT UNIT
AND STORE IN TUP LIST.
POINT PAST END OF TUP LIST.
ESTABLISH ADDRESSABILITY TO TEXT UNIT.
GET ADDRESS OF NEXT TUP LIST ENTRY.
STORE ADDRESS OF TEXT UNIT IN TUP LIST.
TURN ON HIGH-ORDER BIT IN LAST TUP LIST ENTRY.
MOVE RETURN DDNAME TEXT UNIT TO PARM AREA.
PUT ADDRESS OF REQUEST BLOCK POINTER IN REG 1.
INVOKE SVC 99 TO PROCESS THE REQUEST.

Figure 80. Example of a Dynamic Allocation Request

Note the following concepts that the example illustrates:

- You need to request storage via the GETMAIN macro for the request block and the DALRTDDN text unit, because the SVC 99 routines modify them. The DALDSNAM and DSLSTATS text units can be in virtual storage.

In the example, the GETMAIN request is for 50 bytes, derived as follows:

Bytes	Purpose
4	Pointer to the request block.
20	Request block space.
12	Four bytes each for three text unit pointers.
14	Text unit space for the requested return of the ddname.

- IEFZB4D0 provides DSECTs that map the parameter list structure.
- The example uses IEFZB4D2 mnemonics in the text unit keys.

Figure 81 shows the parameter list that results from the code in Figure 80. It is the SVC 99 request block structure needed to allocate data set SYS1.LINKLIB with a disposition of SHARE, and to return the ddname assigned by SVC 99.

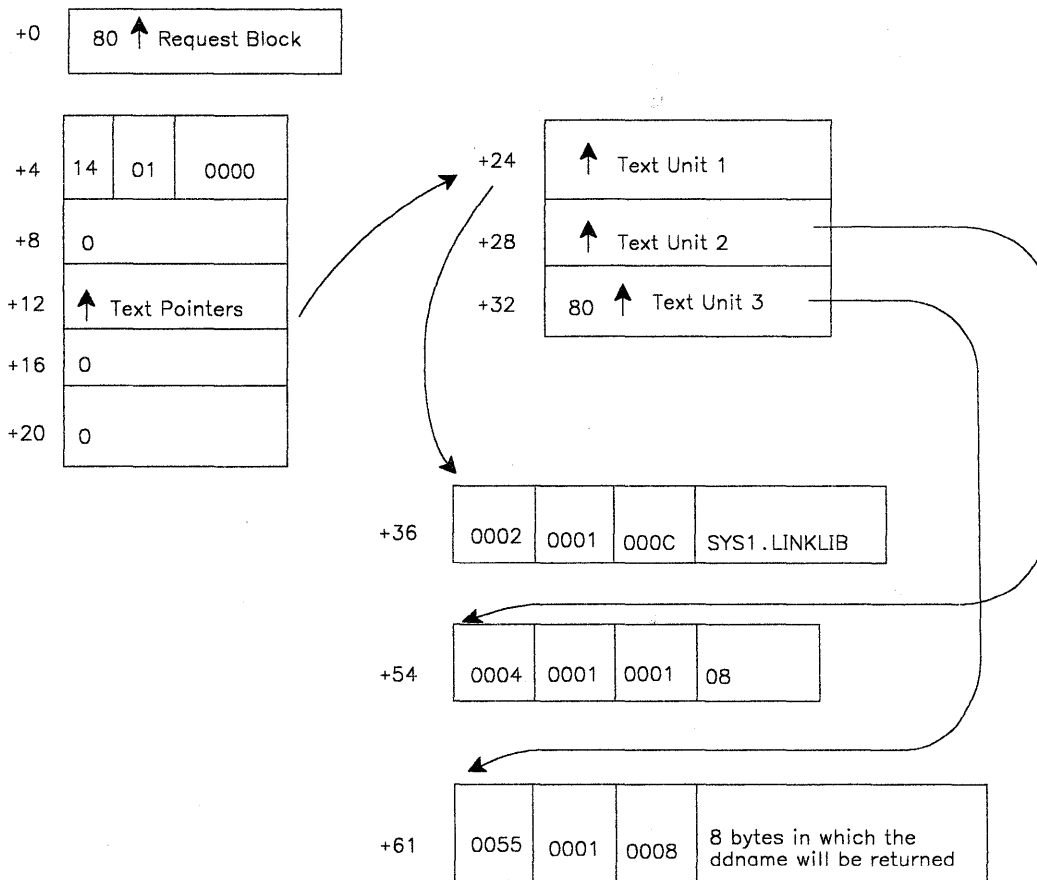


Figure 81. Parameter List Resulting From Dynamic Allocation Example

Index

A

- A = SVCD
 - specified with SLIP 1-147
- A = TRDUMP
 - specified with SLIP 1-147
- ABDUMP symptom area 1-159
- ABEND
 - completion code 1-165, 1-166
 - occurred because of DETACH 1-171
 - SVC 1-83
- abend dump
 - from an ETAE routine 1-184
 - from an FRR 1-184
- ABEND macro instruction 1-141, 1-149, 1-184
 - function 1-144
- abends 1-185
- abnormal
 - address space termination 1-143
 - condition 1-141
 - program termination 1-141
 - termination 1-158, 1-191
 - continue 1-142
 - of attaching task 1-168
- access
 - data from other address spaces 1-81
 - environment element 1-204
- access environment element 1-204
- accounting methods
 - affected by cross memory 1-83
- active
 - addressing bind
 - required by cross memory 1-82
 - binds 1-100
- address
 - range 0-511 1-209
- address space
 - authorization 1-85
 - current 1-82
 - DUMPSRV 1-157
 - home 1-81
 - options of SDUMP macro instruction 1-151
 - primary 1-82
 - secondary 1-82
 - swapped out 1-129
 - switch 1-91
 - termination 1-100, 1-141, 1-191
 - resource manager duties 1-191
 - tracing 1-148
- address space tracing
 - performing 1-148
- addressable
 - address spaces 1-80
 - virtual storage 1-127
- addressing
 - bind 1-82
 - environment
 - normal 1-162
 - restricted 1-163
 - mode 1-13
 - for retry 1-177
 - addressing mode 1-161
 - ADYSETxx 1-157
 - AKM
 - definition 1-86
 - in entry table 1-88
 - altering the dispatching queue 1-25
 - alternate clocks 1-217
 - alternate method
 - for indicating event completion 1-45
 - AMODE 1-13
 - assembler definition 1-13
 - contained in STKE 1-90
 - values for 1-14
 - answer area
 - for EXTRACT 1-33
 - APF
 - authorization 1-4, 1-198
 - guidelines 1-203
 - restricting load module access 1-201
 - using 1-200
 - APF-authorized 1-190
 - how to become 1-198
 - how to find out which programs are 1-198
 - APFTABLE 1-199
 - ASCB 1-155
 - ASID 1-79, 1-151, 1-155
 - in entry table 1-88
 - of home address space 1-81
 - option of SDUMP 1-152
 - parameter of CALLRTM macro instruction 1-143
 - record 1-155
 - translation exception 1-84
 - ASIDLST 1-151
 - ASIDTCB 1-79
 - acceptable values for 1-80
 - ASM
 - functions 1-127
 - recovery 1-144
 - ASMGL lock 1-21
 - assume key of caller 1-194
 - asynchronous
 - dump 1-152
 - exit routine
 - characteristics 1-211
 - register contents 1-214
 - inter-address space communication 1-75
 - AT
 - set 1-89
 - ATSET macro instruction 1-86, 1-89
 - example 1-94, 1-95
 - ATTACH macro

- ESTAI parameter 1-174
- ATTACH macro instruction
 - authorization 1-202
 - changing the defaults 1-4
 - defaults 1-3
 - ESTAI parameter 1-159, 1-168
 - function 1-3
 - STAI option 1-187
 - STAI parameter 1-159
- authorization
 - address space 1-85
 - APF 1-198
 - assigning 1-201
 - code 1-198, 1-201
 - assigned via JCL 1-201
 - assigned via SETCODE 1-201
 - default 1-201
 - cross memory 1-88
 - index 1-89, 1-162
 - for retry 1-176
 - how to reserve 1-93
 - macro instructions 1-89
 - program 1-86
 - requirements for STAE routines 1-190
 - results under various conditions 1-202
 - rules 1-202
 - table 1-85
- authorized
 - libraries 1-199
 - programs 1-199
 - user 1-199
 - definition of 1-3
- auxiliary storage management
 - global lock 1-21
- avoiding duplicate data
 - in a summary dump 1-154
- AX
 - determine 1-86
 - extract 1-89
 - for retry 1-173
 - free 1-89
 - how to reserve 1-93
 - how to set 1-93
 - initial value of 1-86
 - owner of 1-93
 - reserve 1-89
 - set 1-89
- AXEXT macro instruction 1-86
- AXFRE macro instruction 1-86, 1-89
 - example 1-96
- AXRES macro instruction 1-83, 1-86, 1-89, 1-100
 - example 1-93
- AXSET macro instruction 1-86, 1-89
 - example 1-93, 1-96, 1-98

B

- BALR 1-148
- BASR 1-148
- BASSM 1-148
- BDAM data set
 - how to reserve 1-34
- bind
 - active 1-82
 - addressing 1-82
- BLSABDPL macro instruction 1-150
- BLSQMDEF macro instruction 1-149
- BLSQMFLD macro instruction 1-149
- BLSRESSY macro instruction 1-149
- branch and link register 1-148
- branch and save and set mode 1-148
- branch and save register 1-148
- branch entry
 - to PGSER routine 1-132
 - to SDUMP 1-151
 - to stage 1 1-211
- BRANCH option
 - of SDUMP macro instruction 1-155
- BRANCH options
 - of SDUMP macro instruction 1-151
- branch tracing 1-148
 - performing 1-148
- BRANCH = SPECIAL option of PGSER 1-128
- BRANCH = YES option of CALLDISP 1-50
- buffer
 - for internal START command 1-6
- bypassing POST 1-37

C

- CALLDISP macro instruction
 - BRANCH = YES option 1-50
 - considerations for use 1-50
 - function 1-50
 - options 1-50
- CALLRTM macro instruction 1-141, 1-184
 - ASID parameter 1-143
 - function 1-142
 - restrictions 1-142
 - TCB parameter 1-143
 - TYPE = ABTERM 1-143
 - TYPE = MEMTERM 1-143
 - work area 1-143
- CCWs
 - protection of 1-209
- CDS instruction 1-17
- cell pool services 1-105
 - create 1-105
 - delete 1-105
 - free 1-105

- obtain 1-105
- changing
 - contents of registers for retry 1-167
 - system status 1-208
 - the PKM value 1-86
- changing the parameters at dynamic
 - unallocation 1-243
- channel
 - command words
 - protection of 1-209
- CHAP macro instruction
 - function 1-12
- characteristics
 - of a non-space switch PC routine 1-99
 - of a space switch PC routine 1-100
 - of valid storage subpools 1-107
- characters on a MCS console 1-70
- check stop 1-69
- checking
 - PER traps 1-146
- checkpoint/restart
 - restricted in cross memory 1-84
 - restrictions 1-15
 - using 1-15
- CHNGDUMP command 1-141, 1-149, 1-157, 1-184
- CIB
 - address exact 1-7, 1-8
 - contents 1-7
 - counter 1-8
 - free 1-8
 - verb code 1-7
- CIRB macro instruction 1-211
 - BRANCH = YES option 1-211
- clean up
 - cross memory 1-98
 - processing 1-191
 - queues 1-191
 - service available to all address spaces 1-98
- clock
 - comparator 1-218
 - failure 1-218
 - functioning of 1-218
 - resetting 1-219
- clocks
 - alternate 1-217
- CML
 - ASID 1-155
- CML lock 1-21, 1-22, 1-82
 - considerations 1-24
 - use 1-164
- CMS lock 1-20, 1-21, 1-224
- CMSEQDQ lock 1-20
- CMSSMF lock 1-20
- code protection 1-81
- collecting information
 - about resources 1-51
- command
 - MODIFY 1-7
 - START 1-1
 - STOP 1-7
- common
 - storage 1-107
 - storage for cross memory 1-82
- commonly addressable storage 1-75
 - avoidance of 1-81
 - for PC routine 1-94
 - when FRR must reside in 1-164
- communication
 - area 1-166
 - asynchronous 1-74
 - inter-address space 1-74
 - summary of 1-1
 - synchronous 1-74
 - with a problem program 1-7
- communications ECB 1-8
- compare
 - and swap instruction 1-17
 - double and swap 1-137
 - double and swap instruction 1-17
- completion of an event 1-47
- concatenated groups, dynamically deallocating 1-242
- concatenation
 - of authorized and unauthorized libraries 1-200
- concatenation of data sets by SVC 99
 - See dynamic concatenation
- concepts of SVC 99 processing 1-236
- connect ET 1-89
- connecting entry table to linkage table 1-94
- considerations for dsname dynamic allocation 1-260
- continue
 - with abnormal termination 1-142
- control
 - program extensions 1-197
 - register 0 1-209
 - routing of recovery routines 1-178
- control features of SVC 99 processing
 - See processing control features for SVC 99
- control limit, for SVC 99 1-237
- control limit, in dynamic allocation 1-237
- control program extensions 1-197
- convertible attribute 1-238
 - See also processing control features for SVC 99
- convertible attribute for dynamic allocation
 - See processing control features for SVC 99
- counterfeiting a module
 - preventing 1-198
- CPOOL macro instruction
 - function 1-105
- CPU
 - lock 1-19, 1-20
 - reset 1-68
 - work area save (WSA) vector table 1-155
- create
 - a new task 1-3
 - an address space 1-81
 - ET 1-89
 - IRB 1-211
- critical
 - address space resources 1-165
 - resource recovery 1-165
 - system resources 1-163

- cross memory
 - access 1-91
 - authorization 1-85
 - benefits of 1-80
 - clean up 1-91
 - considerations for recovery routines 1-100, 1-160
 - data access 1-81
 - data movement 1-81
 - environment 1-81
 - environment for retry 1-173
 - establishing environment 1-91
 - example
 - set up 1-92
 - examples 1-92
 - facilities available 1-84
 - general considerations 1-83
 - initialize the structure 1-91
 - instructions 1-82
 - linkage 1-87
 - linkage conventions 1-89
 - local lock 1-21
 - macro instructions 1-82
 - mode 1-82
 - mode for POST 1-42
 - mode for retry 1-176
 - post 1-37, 1-42
 - program sharing 1-81
 - provide service for 1-91
 - remove access 1-91
 - restrictions 1-83
 - restrictions for FESTAE 1-170
 - services lock
 - CMS 1-22
 - CMSEQDQ 1-22
 - CMSSMF 1-22
 - requesting 1-24
 - state 1-162
 - structures 1-85
 - terminology 1-81
- CS instruction 1-17
- CSA
 - subpools in 1-107
- current
 - address space 1-82
- custom-built delete functions for messages 1-74
- CVT
 - CVTSDBF field of 1-152
 - mapping macro 1-46, 1-49
 - use of 1-92, 1-213, 1-214, 1-216, 1-225
- CVTEXT 1-214
- CVTEXP1 1-225
- CVTQTD00 1-219
- CVTQTE00 1-216
- CVTSDBF 1-152
- CVTTPC 1-214
- CVTVWAIT 1-44
- CVT0EF00 1-213

D

- DAE
 - See dump analysis and elimination
- DALBFALN text unit 1-291
- DALBFTEK text unit 1-291
- DALBLKLN text unit 1-275
- DALBLKSZ text unit 1-291
- DALBUFIN text unit 1-292
- DALBUFL text unit 1-292
- DALBUFMX text unit 1-292
- DALBUFNO text unit 1-292
- DALBUFOF text unit 1-292
- DALBUFOU text unit 1-293
- DALBUFRQ text unit 1-293
- DALBUFSS text unit 1-293
- DALBURST text unit 1-287
- DALCDISP text unit 1-274
- DALCHARS text unit 1-287
- DALCLOSE text unit 1-279
- DALCNVRT text unit 1-303
- DALCODE text unit 1-293
- DALCOPYG text unit 1-287
- DALCOPYS text unit 1-280
- DALCPRI text unit 1-294
- DALCYL text unit 1-275
- DALDCBDD text unit 1-284
- DALDCBDS text unit 1-284
- DALDDNAM text unit 1-273
- DALDEFER text unit 1-289
- DALDEN text unit 1-294
- DALDIAGN text unit 1-300
- DALDIR text unit 1-276
- DALDSNAM text unit 1-273
- DALDSORG text unit 1-294
- DALDSSEQ text unit 1-280
- DALDUMMY text unit 1-281
- DALEROPT text unit 1-295
- DALEXPDT text unit 1-281
- DALFCBAV text unit 1-282
- DALFCBIM text unit 1-282
- DALFCNT text unit 1-288
- DALFFORM text unit 1-288
- DALFRID text unit 1-302
- DALFUNC text unit 1-301
- DALGNCP text unit 1-295
- DALINOUT text unit 1-281
- DALINTVL text unit 1-295
- DALIPLTX text unit 1-300
- DALKYLEN text unit 1-295
- DALLABEL text unit 1-280
- DALLIMCT text unit 1-296
- DALLRECL text unit 1-296
- DALMEMBR text unit 1-273
- DALMMOD text unit 1-288
- DALMODE text unit 1-296
- DALMSVGP text unit 1-285
- DALMTRC text unit 1-288

- DALNCP text unit 1-296
- DALNDISP text unit 1-274
- DALOPTCD text unit 1-297
- DALOUTLM text unit 1-279
- DALOUTPT text unit 1-289
- DALPARAL text unit 1-278
- DALPASPR text unit 1-280
- DALPASSW text unit 1-302
- DALPCIR text unit 1-297
- DALPCIS text unit 1-298
- DALPERMA text unit 1-303
- DALPRIME text unit 1-275
- DALPRIVT text unit 1-277
- DALPROT text unit 1-286
- DALPRTSP text unit 1-298
- DALQNAME text unit 1-282
- DALRECFM text unit 1-298
- DALRETPD text unit 1-281
- DALRLSE text unit 1-276
- DALROUND text unit 1-276
- DALRSRVF text unit 1-299
- DALRSRVS text unit 1-299
- DALRTDDN text unit 1-303
- DALRTDSN text unit 1-303
- DALRTORG text unit 1-304
- DALRTVOL text unit 1-305
- DALSECND text unit 1-275
- DALSFMNO text unit 1-279
- DALSHOLD text unit 1-285
- DALSOWA text unit 1-299
- DALSPFRM 1-276
- DALSPGNM text unit 1-279
- DALSSATT text unit 1-305
- DALSSNM text unit 1-285
- DALSSPRM text unit 1-286
- DALSSREQ text unit 1-304
- DALSTACK text unit 1-299
- DALSTATS text unit 1-274
- DALSUSER text unit 1-284
- DALYSOUC text unit 1-278
- DALTERM text unit 1-283
- DALTHRSH text unit 1-300
- DALTRK text unit 1-274
- DALTRTCH text unit 1-300
- DALUCS text unit 1-283
- DALUFOLD text unit 1-283
- DALUNCNT text unit 1-278
- DALUNIT text unit 1-278
- DALUSRID text unit 1-286
- DALUVERFY text unit 1-283
- DALVLCNT text unit 1-277
- DALVLRDS text unit 1-277
- DALVLSEQ text unit 1-277
- DALVLSER text unit 1-276
- DASD
 - allocation and management 1-127
- DAT
 - error 1-164
 - turned off 1-135
 - turned on 1-135
- DAT-OFF indexes 1-138
- DAT-OFF routines 1-135
 - restrictions 1-138
 - writing 1-138
- DAT-ON nucleus
 - obtaining information 1-139
- data
 - access 1-81
 - area for asynchronous exit routines 1-212
 - areas for POST 1-41
 - movement 1-81
 - protection 1-81
 - sets that cannot be shared 1-29
- data security 1-203
- DATOFF macro instruction
 - function 1-136
- DCB
 - address 1-33
 - contents 1-33
 - for a secure data set 1-185
 - in dynamically acquired virtual storage 1-174
- DCB attribute text units 1-291
 - See also SVC 99 text units
- DCBDEBAD 1-34
- DCCDDNAM text unit 1-309
- DCCPERMC text unit 1-309
- DD name 1-33
- DDCDDNAM text unit 1-310
- Ddname Allocation 1-240
- ddname allocation text units 1-311
 - See also SVC 99 text units
- DDNDDNAM text unit 1-311
- DDNRTDUM text unit 1-311
- DDs allowed per jobstep
 - default number of 1-245
 - modifying number of 1-245
- deadlock 1-23
 - example 1-23
 - preventing 1-23
- deadlocked address space 1-144
- deallocation by SVC 99
 - See dynamic unallocation
- DEB
 - bit for an authorized library 1-198
 - use 1-33
- DEBDVMOD 1-34
- declared storage for cross memory examples 1-92
- deconcatenating data sets via SVC 99
 - See dynamic deconcatenation
- default
 - dump data set 1-185
 - subpool 1-189
- deleting messages already written 1-73
- DEQ macro instruction
 - example 1-36
 - function 1-26
 - use of 1-51
- destroy
 - entry table 1-96
 - ET 1-89
- determine the AX of an address space 1-86

- device classes for UCB scan 1-230
- devices
 - releasing 1-31
 - reserving 1-30
 - sharing 1-28
 - that can be shared 1-28
- diagnostic
 - data collecting 1-146
 - information in cross memory 1-89
- DIE
 - restrictions on execution 1-216
- DIE routine
 - characteristics 1-216
 - execution 1-216
 - exit from 1-216
 - recovery for 1-219
 - register contents on entry 1-216
- DINDDNAM text unit 1-312
- DINDSNAM text unit 1-312
- DINRELNO text unit 1-316
- DINRTATT text unit 1-315
- DINRTCDP text unit 1-314
- DINRTDDN text unit 1-313
- DINRTDSN text unit 1-313
- DINRTLIM text unit 1-315
- DINRTLST text unit 1-315
- DINRTMEM text unit 1-313
- DINRTNDP text unit 1-314
- DINRTORG text unit 1-314
- DINRTSTA text unit 1-313
- DINRTTYP text unit 1-316
- direct
 - service class 1-67
- disable
 - low address protection 1-209
- disabled
 - locks 1-21
 - page fault 1-224
 - spin locks
 - held by SVC routine 1-224
- disabled/enabled state
 - for obtain 1-25
 - for release 1-25
- disablement 1-161
 - for retry 1-173, 1-175
 - legal 1-20
 - system recognized 1-20
- disconnect
 - entry table from linkage table 1-95
 - ET 1-89
- DISP
 - JCL parameter 1-158
 - lock 1-22
- dispatchable work queue 1-76
- dispatcher 1-213
 - in cross memory 1-84
 - lock 1-216
 - pass control to 1-50
- dispatching queue
 - altering 1-25
- DISPLAY command
 - used with SLIP 1-147
- DISPLAY DUMP command 1-149, 1-157
 - timestamp 1-157
- documentation
 - on system integrity 1-193
- DOM macro instruction
 - function 1-73
- double failure 1-77
- DRICURNT text unit 1-311
- DRITCBAD text unit 1-310
- DSNAME 1-158
- dsname allocation 1-239
 - coding considerations 1-260-1-261
 - in addition to JCL services 1-239
 - JCL services not available through dynamic allocation 1-240
 - processing details 1-256, 1-261
- dsname allocation processing 1-256
 - changing the parameters of an existing allocation 1-259
 - checking for environmental conflicts 1-256
 - criteria for using an existing allocation 1-257
 - required of the existing allocation 1-257
 - required of the request 1-257
 - using a new allocation 1-259
 - automatic deallocation of resources held for re-use 1-260
- dsname allocation text units 1-273
 - See also SVC 99 text units
- dsname processing by SVC 99
 - See dsname allocation processing
- dump
 - cross memory memory information needed 1-90
 - data sets 1-157
 - of virtual storage 1-149
 - options 1-171
 - overriding options 1-184
 - scheduled 1-151
 - summary 1-154
 - SVC 1-151
 - synchronous 1-152
 - tailoring 1-186
 - title 1-186
 - virtual storage 1-141
 - when RTM takes 1-184
- dump analysis and elimination (DAE) 1-147, 1-156
 - providing information for 1-159
 - status and error flags 1-167
 - symptom count 1-167
- DUMP command 1-149, 1-157
- dump suppression 1-156
 - bypassing 1-147
- DUMPDS command 1-149, 1-157
- DUMPSRV address space 1-157
- DUNDDNAM text unit 1-306
- DUNDSNAM text unit 1-306
- DUNMEMBR text unit 1-306
- DUNOVCLS text unit 1-308
- DUNOVDSP text unit 1-307

- DUNOVSHQ text unit 1-308
- DUNOVSNH text unit 1-308
- DUNOVSSUS text unit 1-308
- DUNREMOV text unit 1-307
- DUNUNALC text unit 1-307
- duration of fix 1-128
- DYNALLOC macro 1-236
- dynamic allocation 1-235, 1-239
 - ddname allocation 1-240-1-241
 - dsname allocation 1-239-1-261
 - installation options 1-245
 - summary of 1-2
- dynamic allocation by dsname, processing details
 - See dsname allocation processing
- dynamic allocation error reason codes
 - See SVC 99 error reason codes
- dynamic allocation information reason codes
 - See SVC 99 information reason codes
- dynamic allocation input validation
 - See installation input validation routine for SVC 99
- dynamic allocation installation options
 - See installation options for SVC 99 functions
- dynamic allocation parameter list
 - See SVC 99 parameter list
- dynamic allocation programming considerations
 - See programming considerations when using SVC 99
- dynamic allocation request block
 - See SVC 99 parameter list
- dynamic allocation request block pointer
 - See SVC 99 parameter list
- dynamic allocation return codes
 - See SVC 99 return codes
- dynamic allocation text pointers
 - See SVC 99 parameter list
- dynamic allocation text units
 - See SVC 99 parameter list
 - See SVC 99 text units
- dynamic concatenation 1-243
 - permanently concatenated attribute 1-244
 - characteristics 1-244
- dynamic concatenation text units 1-309
 - See also SVC 99 text units
- dynamic deconcatenation 1-244
- dynamic deconcatenation text unit 1-310
- dynamic information retrieval 1-244
 - kinds of information retrieved 1-245
- dynamic information retrieval text units 1-312
 - See also SVC 99 text units
- dynamic unallocation 1-241
 - changing parameters 1-243
 - of concatenated groups 1-242
 - processing considerations 1-242-1-243
 - removing in-use bit based on task id 1-243
- dynamic unallocation processing 1-242
- dynamic unallocation text units 1-306
 - See also SVC 99 text units

E

- EBCDIC characters 1-70
- EC PSW 1-166
- ECB
 - extended 1-40
 - not posted 1-129
 - posted by a system routine 1-42
 - posting 1-7
 - supplied with a page-fix request 1-127
 - supplied with a page-load request 1-127
 - target for cross memory post 1-37
 - used with page fix 1-129
- ECBE
 - content 1-39
 - used to identify a subsystem exit 1-42
- EKM
 - definition 1-87
 - in entry table 1-88
- ELSQA
 - subpools in 1-107
- emergency
 - signal 1-68
 - signal function 1-68
- enable
 - low address protection 1-209
- enabled
 - locks 1-21
 - summary dump 1-154
 - task mode 1-98
 - unlocked SRB 1-184
- end address
 - for paging services 1-127
- ENQ macro instruction 1-27
 - function 1-26
 - use of 1-51
- ENQ/DEQ 1-17
- ENQ/DEQ cross memory services lock 1-22
- entry index 1-88
- entry points
 - for cross memory services 1-92
 - to routines in DAT-OFF nucleus 1-135
- entry table
 - constructed via ETCRE macro instruction 1-88
 - contents 1-88
 - descriptions for examples 1-93
 - function of 1-92
 - owner of 1-92
 - second level table 1-87
 - set up 1-92
 - use 1-88
- environment
 - for a task recovery 1-182
 - for resource manager 1-191
 - system 1-160
- EOF 1-158
- EPA (external parameter area)
 - locate mode 1-125

- move mode 1-122, 1-123
- EPAR instruction 1-83
- EQT
 - use of 1-42
 - USERINFO field in 1-42
- equipment check 1-69
- EREP indicator 1-167
- ERRET routine 1-37
- error
 - associated with ABEND 1-166
 - documentation 1-142
 - during swap-in 1-144
 - during swap-out 1-144
 - exit routine 1-144
 - message for abend
 - who issues 1-185
 - on a higher task 1-171
 - recovery
 - for DIE routine 1-219
- error reason codes for dynamic allocation
 - See SVC 99 error reason codes
- ESAR instruction 1-83
- ESCA
 - subpools in 1-107
- ESPIE macro instruction 1-141, 1-144
 - AMODE of callers 1-144
- ESQA
 - subpools in 1-107
- establish
 - a cross memory environment 1-91
 - access 1-97
 - access for cross memory 1-91
 - addressability for QEDIT 1-8
- ESTAE
 - environment 1-168
 - exit 1-168
 - nested routine 1-183
 - recovery for SVCs 1-225
 - routine
 - when not to establish 1-183
- ESTAE macro instruction 1-141, 1-160, 1-168
 - TOKEN parameter 1-169
- ESTAE-type recovery
 - routines 1-173
 - deletion of 1-168
 - special considerations 1-170
 - summary 1-186
 - use of 1-142
- ESTAI parameter 1-168
 - of ATTACH 1-141
 - of ATTACH macro instruction 1-160
 - propagation of recovery routine 1-168
- ET
 - connect 1-89
 - create 1-89
 - destroy 1-89
 - disconnect 1-89
- ETCON macro instruction 1-89
 - example 1-94, 1-97
- ETCRE macro instruction 1-89, 1-98, 1-100
 - example 1-97
 - use of 1-92
- ETDES macro instruction 1-89
 - example 1-96, 1-98
 - PURGE= YES option 1-98
- ETDIS macro instruction 1-89
 - example 1-95
- EUT FRRs 1-160
- event
 - completion 1-45, 1-129
 - indicating completion of 1-47
 - waiting for completion 1-38
- EVENTS macro instruction 1-38
- EVENTS table 1-38
- EX
 - portion of PC number 1-89
 - use 1-88
- example
 - authorization assigned via SETCODE 1-202
 - bypassing the POST routine 1-38
 - cross memory set up 1-92
 - dump title 1-186
 - of deadlock 1-23
 - of subroutine issuing RESERVE and DEQ 1-36
 - of SUSPEND macro instruction 1-46
 - using POST exit function 1-41
- exclusive OR 1-137
- EXECUTE form of a macro instruction 1-16
- execution
 - and termination of asynchronous exits 1-213
 - time
 - in cross memory 1-84
- EXIT 1-190
 - for MVS router 1-205
- exit routines 1-211
 - asynchronous 1-211
 - deleting 1-39
 - error 1-144
 - identifying 1-39
 - POST 1-38
 - POST interface 1-40
 - SPIE 1-144
 - SPIE/ESPIE 1-145
 - STAE/STAI 1-188
 - summary of 1-2
 - timer disabled 1-214
- explicit
 - purging 1-130
 - tracing 1-148
- extended
 - area 1-107
 - control (EC) PSW 1-166
- external
 - call 1-68
 - call function 1-69
 - call pending 1-69
- extract
 - AX 1-89
 - information from the resource queues 1-51
 - primary ASN (EPAR) instruction 1-83

- secondary ASN (ESAR) instruction 1-83
- EXTRACT macro instruction
 - example 1-9
 - function 1-7, 1-31
 - use of 1-31
- EXTRACT option of PCLINK 1-90

F

- fast
 - ESTAE 1-168
 - path to fix virtual storage 1-128
 - path to free virtual storage 1-128
- features unique to dynamic allocation
 - See concepts of SVC 99 processing
- FESTAE macro instruction 1-141, 1-160, 1-168
 - cross memory restrictions 1-170
 - parameter area 1-170
 - restrictions 1-170
 - using 1-170
- fetch protection 1-193
 - of PSA 1-197
- first
 - level table for cross memory linkage 1-87
- fix
 - virtual storage 1-128
 - fast path 1-128
- FIX option of PGSER 1-128
- fixed frames
 - responsibility for freeing 1-130
- footprint areas 1-185
- FORCE command 1-144
- forced
 - stop 1-144
- FRACHECK macro instruction 1-204
- free
 - an AX 1-89
 - an AX value 1-86
 - AX 1-96
 - fixed frames 1-130
 - linkage 1-96
 - LX 1-89
 - real storage 1-128
 - the CIB 1-8
 - virtual storage 1-128
 - fast path 1-128
- FREE option of PGSER 1-129
- FREEMAIN macro instruction 1-106
 - BRANCH parameter 1-104
 - function 1-104
 - KEY parameter 1-105
 - used by retry routines 1-170
- FRR
 - establishing 1-182
 - global 1-163
 - recovery for SVCs 1-225
 - required option 1-78
 - summary 1-186

- FRR stack 1-161
 - normal 1-160
 - super 1-161, 1-163
- FRRs
 - nested 1-183
- FSTOP 1-144
- FULLXM
 - MODE parameter of SETFRR 1-162
- function codes, SWA manager
 - locate mode request 1-124
 - move mode request 1-121
- functions available through SVC 99 1-239

G

- general cross memory services lock 1-22
- GETMAIN macro instruction 1-106
 - BRANCH parameter 1-104
 - function 1-104
 - KEY parameter 1-105
- getting control as result of PC instruction 1-89
- global
 - dispatcher lock 1-22
 - intersect 1-25
 - locks 1-18
 - priority 1-78
 - resource serialization 1-31, 1-51
 - services available to all users 1-96
 - subpool 1-18
- GLOBAL MODE parameter of SETFRR 1-163
- global resource serialization
 - limiting requests 1-27
 - request 1-27
- GQSCAN macro instruction 1-27
 - function 1-51
 - results 1-53, 1-54
 - TOKEN parameter 1-51
- GTF trace record 1-146
- guidelines
 - for recovery routines 1-183
 - for using APF 1-203
- GVCTOL 1-28
- GVTREQ 1-27

H

- HASID 1-81
- hexadecimal representation of characters 1-70
- hierarchy of locks 1-23
- hierarchy of recovery routines 1-185
- high private storage 1-107
- home
 - address space 1-81
 - as entry table owner 1-92
 - as LX owner 1-92

owner of AX 1-93
address space and locking 1-21
mode 1-82
MODE parameter of SETFRR 1-162

I

I/O
error on page-in 1-141
interruptions 1-75
IAC instruction 1-83
ICHRFR00 1-205
ICHRFR01 1-205
ICHRTX00 1-205
ICHSAFP 1-207
IDAWs
protection of 1-209
identifier to PURGEDQ 1-77
identifying messages to be deleted 1-73
IEAAPFxx 1-200
IEAAPF00 1-199
IEAAP00 1-201
IEASYSxx 1-200
IEAVCDS
register contents 1-137
IEAVEUR1 1-135
IEAVEUR2 1-135
IEAVEUR3 1-135
IEAVEUR4 1-135
IEAVMVC0 1-135
IEAVMVKY 1-135
IEAVVTCP mapping macro 1-214
IEAVWAIT 1-44
IEAVXC0 1-135
IEA0PT0E 1-39
IEA0PT01 1-42
IEA0PT03 1-41
IEECVXIT 1-158
IEE331A 1-67
IEFAB445 1-246
IEFQMREQ macro 1-121
IEFZB4D0 1-251, 1-253, 1-255, 1-262
IEFZB4D2 1-251, 1-262
IEZCOM mapping macro 1-8
IEZMGCR mapping macro 1-6
IEZWPL mapping macro 1-71
IGC00nnn 1-224
IHAETD mapping macro 1-92
IHAPSA mapping macro 1-49, 1-50, 1-170
IHARB mapping macro 1-170
IHASCB mapping macro 1-170
IHASDWA mapping macro 1-166
IHASVT mapping macro 1-50
IHATQE mapping macro 1-214
IHSA 1-155
IKJTCB mapping macro 1-170
IKT001D 1-144
IKT010D 1-144
in-use bit and attribute, for SVC 99 1-237
See also processing control features for SVC 99
removing, via dynamic unallocation 1-241
incorrect state 1-69
independently dispatchable units of work 1-76
index
used with DATOFF 1-135
indicating event completion 1-47
indirect data address words
protection of 1-209
INDMVCLK 1-135
register contents 1-136
INDMVCL0 1-135
register contents 1-136
INDUSR1 1-135
INDUSR2 1-135
INDUSR3 1-135
INDUSR4 1-135
INDXC0 1-135
register contents 1-137
information
preserving for dumps 1-185
information reason codes for dynamic allocation
See SVC 99 information reason codes
initial
CPU reset 1-68
input
to paging services 1-130
to set DIE 1-214
insert address space control (IAC) instruction 1-83
inserting
SVC routines into the control program 1-225
installation
responsibility 1-193
installation input validation routine for SVC 99 1-248,
1-252
installation options for SVC 99 functions 1-245
installation-written input validation routine 1-248
mounting volumes and bringing devices
online 1-247
space and unit default values 1-245
instruction
address
in entry table 1-88
fetch 1-146
integrity
elimination of potential exposures 1-194
exposures
control program extensions 1-197
resource identification 1-196
sensitive system data 1-197
SVC routines calling SVC routines 1-196
user supplied addresses 1-195
system 1-193
inter-address space communication
asynchronous 1-75
intercept
errors 1-184
expected program checks 1-183
system errors 1-141, 1-146

- interface
 - to stage 2 1-213
- interlock 1-23, 1-31
 - example 1-23, 1-32
 - task 1-32
- interlocks
 - preventing 1-31
- interrupt handlers
 - in cross memory 1-84
- interruption
 - handler 1-75
 - types 1-145
- intersect
 - global 1-25
 - local 1-25
- interval cancellation 1-218, 1-219
- introduction to dynamic allocation
 - See introduction to SVC 99 functions
- introduction to SVC 99 functions 1-236
 - DYNALLOC macro 1-236
 - parameter list 1-236
 - See also SVC 99 parameter list
- INTSECT macro instruction
 - function 1-25
- invalid
 - function 1-69
 - parameter 1-69
- invoke
 - recovery termination 1-141
 - the emergency signal function 1-68
 - the external-call function 1-68
- IOS
 - obtaining information from 1-234
 - synchronization lock 1-22
 - unit control block lock 1-22
- IOSINFO macro instruction 1-234
- IOSUCB lock 1-22
- IOSVSUCB
 - input 1-229
 - output from
 - affected by DDR swap 1-231
 - for devices with multiple exposures 1-231
 - for devices with optional channels 1-231
 - parameter list 1-229
 - register contents 1-230
 - restrictions on use 1-229
 - return codes 1-231
- IOSYNCH lock 1-22
- IPC
 - function 1-67
 - service classes 1-67
- IPCS
 - dumping service 1-149
- IQE
 - function 1-213
 - initialization 1-213
- IQEIRB 1-213
- IQEPARAM 1-213
- IQETCB 1-213
- IRB

- address 1-211
- create 1-211
- errors in 1-184
- initialization 1-211, 1-212
- ISAM data set
 - how to reserve 1-34
- isolate
 - an error 1-183
 - data 1-81
- issuing RESERVE and DEQ 1-36

J

- JES3 1-254
 - class 2 reason code from SVC 99 1-265
 - notes on dynamic allocation (SVC 99) 1-254, 1-264
- JFCB
 - format of 1-118
- job
 - library 1-4
- job entry subsystem
 - and dynamic allocation (SVC 99) 1-249, 1-250, 1-267, 1-304
- job step
 - advantage of creating 1-3
 - control block 1-5
 - task 1-3
 - task owning entry tables 1-100
 - timer
 - expiration 1-171
- jobname 1-156
- JSCB 1-4
- JSCBAUTH 1-201

K

- key of caller
 - assume 1-194

L

- label
 - embedding lines 1-72
- latent parameter
 - address
 - in entry table 1-88
 - list address
 - in STKE 1-88
- LCCA 1-155
- length
 - of parameter list for SDUMP 1-150
 - of VRA 1-167

- libraries
 - installation authorized 1-199
 - program 1-32
 - SYS1.LINKLIB 1-199
 - SYS1.LPALIB 1-199
 - SYS1.SVCLIB 1-199
- library
 - search 1-4
- LIFO order 1-171
- limiting extent of message deletion 1-74
- link pack area 1-225
- linkage
 - and entry tables
 - for a global service 1-97
 - conventions
 - for PC recovery 1-89
 - in cross memory 1-100
 - index 1-87
 - format of 1-95
 - macro instructions 1-89
 - tables 1-87
 - connection to entry tables 1-95
 - contents 1-88
 - first level 1-87
 - use 1-88
- linking to routines in DAT-OFF nucleus 1-135
- LINKLIB
 - in dynamic allocation example 1-316
- LIST form of a macro instruction 1-16
- LISTA 1-151, 1-154
- LNKAUTH=APFTAB 1-200
- LNKAUTH=LNKLST 1-200
- LNKLST
 - mixing APF and non-APF libraries in 1-200
- LNKLST concatenation 1-200
- load
 - module 1-4
 - virtual storage 1-128
- LOAD macro instruction
 - function 1-14
 - to bring in retry routine 1-189
- local
 - intersect 1-25
 - level lock
 - obtaining more than one 1-24
 - lock 1-22, 1-213
 - not needed with PGSER 1-132
 - locks 1-18
 - lockword 1-20
 - priority 1-78
- LOCAL lock 1-22, 1-24
 - event completion 1-38
 - required option 1-78
 - with local intersect 1-25
- LOCAL MODE parameter of SETFRR 1-164
- locate mode request to SWA manager 1-124
 - EPA size and contents 1-125
 - input 1-124
 - output 1-124
 - summary 1-125
- lock
 - manager 1-19
 - obtaining more than one local 1-24
 - locked status
 - and recovery 1-184
 - locking 1-18, 1-161
 - at the TCB level 1-17
 - categories of locks 1-18
 - conventions
 - for SVCs 1-224
 - hierarchy 1-23
 - interface 1-79
 - restrictions 1-23
 - summary 1-19
 - locks
 - classes of 1-21
 - conditionally requested 1-21
 - CPU lock disablement 1-20
 - enabled 1-21
 - for retry 1-173, 1-175
 - global 1-18
 - in MVS/XA 1-21
 - local 1-18
 - multiple 1-21
 - obtaining 1-24
 - releasing 1-24
 - requests for shared/exclusive 1-19
 - shared/exclusive 1-19
 - single 1-21
 - spin 1-19
 - suspend 1-20
 - testing 1-24
 - types 1-19, 1-20
 - unconditionally requested 1-21
 - lockword 1-19
 - CPU 1-20
 - local 1-20
 - using the same 1-23
 - logoff
 - system initiated 1-171
 - LOGREC
 - recording 1-186
 - recording considerations 1-185
 - recording from recovery routines 1-185
 - LONG=Y option
 - of fix function 1-128
 - low
 - private storage 1-107
 - storage protection 1-209
 - LSQA
 - subpools in 1-107
 - LX
 - format of 1-95
 - free 1-89
 - how to reserve 1-92
 - portion of PC number 1-89
 - reserve 1-89
 - use 1-87
 - LXFRE macro instruction 1-89
 - example 1-96

LXRES macro instruction 1-89, 1-100
example 1-92, 1-97
with SYSTEM = YES option 1-97

M

machine check 1-141
macro instructions
EXECUTE form 1-16
LIST form 1-16
used for cross memory authorization 1-89
used for cross memory linkage 1-89
used with shared DASD 1-30
making services available
to all address spaces 1-96
to selected address spaces 1-92
managing SWA control blocks 1-117
mapping macros
CVT 1-46, 1-49
ICHSAFP 1-207
IEAVVTPC 1-214
IEZCOM 1-8
IEZMGCR 1-6
IEZWPL 1-71
IHAETD 1-92
IHAPSA 1-49, 1-50, 1-170
IHARB 1-170
IHASCB 1-170
IHASDWA 1-166
IHASVT 1-50
IHATQE 1-214
IKJTCB 1-170
master
scheduler address space 1-24
memory create 1-143
MEMTERM
routines that can invoke 1-143
message deletion
limiting extent of 1-74
messages
deleting 1-73
deletion 1-73
routing 1-70
writing 1-70
MGCR macro instruction
example 1-6
function 1-6
used to issue an internal START 1-6
migration
of code and data 1-81
mode
addressing 1-13
cross memory 1-82
for asynchronous exits 1-212
home 1-82
of set DIE caller 1-214
primary 1-82
residency 1-13

secondary 1-82
MODESET macro instruction 1-208
function 1-4
inline code 1-208
keys that you can set 1-208
SVC form 1-208
use by retry routine 1-170
MODESET SVC 1-86
MODIFY command 1-7
modifying the SVC table at execution time 1-226
modules
re-entrant 1-16
mounting and demounting with shared DASD 1-30
mounting volumes and bringing devices online 1-247
move
character long 1-136
character long in user key 1-136
data between address spaces 1-81
to primary (MVCP) instruction 1-83
to secondary (MVCS) instruction 1-83
with key (MVCK) instruction 1-83
move mode request to SWA manager 1-121
EPA format 1-123
EPA size and contents 1-122
input 1-122
output 1-122
summary 1-123
movement of virtual storage pages 1-127
MSSF failure 1-69
multiple
line messages
embedding label lines 1-72
virtual address spaces 1-74
multiple-event wait 1-225
must-complete function 1-26
characteristics 1-26
MVCK instruction 1-83
use 1-86
MVCL function 1-136
MVCP instruction 1-83
use 1-99
MVCS instruction 1-83
use 1-99
MVS router 1-205
exit 1-205
parameter list 1-207
MVS router exit 1-205
MVS router exit routine
return codes 1-206

N

naming conventions
for SVC routines 1-224
national characters 1-70
nested
ESTAE routines 1-183

FRRs 1-183
recovery routine 1-183
recovery routines 1-183

new

task
save area for 1-4
time interval
setting 1-214

next available queue for IQEs 1-214
non-JCL dynamic allocation functions 1-302

See also SVC 99 text units

non-pageable storage 1-127
non-preemptable routines 1-78
non-preemptable SVC routines 1-222
non-space switch service
definition 1-91

non-swappable address space 1-24
nondispatchability bit 1-5
nonquiesceable level 1-78

normal

FRR stack 1-160
program termination 1-141
system events 1-148
termination 1-191

NOSUM

option of SDUMP macro instruction 1-154
parameter of CHNGDUMP command 1-154

not-in-use attribute, for SVC 99 1-237, 1-240, 1-241

nucleus 1-135

DAT-OFF 1-135
linking to routines in DAT-OFF 1-135
summary of 1-2

NUCLKUP macro instruction

function 1-139

O

obtain

a global spin lock 1-25
a suspend lock 1-25
information about CSECTs in DAT-OFF
nucleus 1-139
information from IOS 1-234

opened data sets

finding the UCB address 1-34

operator

cancel 1-171
intervening 1-69
reply of FSTOP 1-144

options

override 1-184

P

page faults 1-78, 1-145
avoiding 1-216

page fix
reverse 1-129

page free
use of ECB 1-130

page out
virtual storage 1-128

paging services 1-128
branch entry
cross memory mode 1-132
non-cross memory mode 1-134
completion considerations 1-129
differences between MVS/370 and MVS/XA 1-127
input 1-130

paging supervisor 1-143

paired resources requests 1-191

parallelism 1-75

parameter

area
for recovery routines 1-172
of FESTAE 1-170

list

for ESTAE routine 1-166
length for SDUMP 1-150

registers

contained in STKE 1-90

parameter list for dynamic allocation

See SVC 99 parameter list

PASID 1-82, 1-90, 1-164

passing control to another address space 1-80

password

incorrect supplied 1-185
protection 1-193

PC 1-92

PC command 1-148

PC instruction 1-82, 1-87, 1-94
functions performed before issuing 1-90
issued from primary mode 1-98

PC number 1-87, 1-89

constructing 1-93

contained in SFT 1-89

contents 1-88

indexing linkage and entry tables 1-88

making available 1-94

PC routine

active binds 1-100

characteristics of a non-space switch routine 1-99

characteristics of a space switch routine 1-100

data 1-99

designing 1-98

linkage conventions 1-100

non-space switch

mode of 1-98

preserve and restore PC linkage information 1-99

purpose of 1-98

- recovery considerations 1-100
- return to caller via PT 1-99
- secondary mode 1-99
- space switch
 - additional requirements 1-99
- space switch and non-space switch
 - requirements for 1-98
- space switch or non-space switch 1-99
- space-switch
 - mode of 1-98
 - use of checkpoint/restart not allowed 1-99
 - where loaded 1-100
- PC/AUTH services 1-92
- PC/PT linkage conventions 1-91
- PCCA 1-155
- PCLINK EXTRACT
 - use of 1-90
- PCLINK macro instruction
 - EXTRACT option 1-90
 - function 1-90
 - restrictions on 1-90
 - STACK option 1-90
 - standard method for saving status 1-100
 - UNSTACK THRU option 1-91
- PCLINK STACK 1-90, 1-155
- PER
 - events 1-146
 - hardware 1-146
- percolate
 - definition 1-177
- percolation 1-177
 - communication field 1-167
 - definition 1-142
 - for the same unit of work 1-181
 - SRB-to-task 1-181
 - to an ESTAE routine 1-181
 - to an ESTAI routine 1-181
 - to an FRR 1-181
- permanently allocated attribute 1-238
 - See also processing control features for SVC 99
 - changing 1-238
- permanently concatenated attribute 1-244
- PGFIX macro instruction
 - function 1-127, 1-128
- PGFIXA macro instruction
 - function 1-127, 1-128
- PGFREE macro instruction
 - function 1-128
- PGFREEA macro instruction
 - function 1-128
- PGLOAD macro instruction
 - function 1-128
- PGOUT macro instruction
 - function 1-128
- PGRLSE macro instruction
 - function 1-128
- PGSER
 - BRANCH entry 1-132
- PGSER macro instruction
 - BRANCH=SPECIAL 1-129
 - BRANCH=SPECIAL option 1-130
 - FIX option 1-129, 1-130
 - FREE option 1-128
 - function 1-127, 1-128
 - local lock 1-132
- PIRL 1-174
- PKM
 - changing value 1-86
 - checked for problem program 1-86
 - use 1-86
- PLPA
 - modules located in 1-200
- pointer defined address 1-77
 - definition of 1-77
- pointer to the SVC 99 request block
 - See SVC 99 parameter list
- POST
 - branch entry points and function 1-43
 - bypass 1-37
 - cross memory mode 1-42
 - data areas 1-41
 - entry points 1-41, 1-42
 - exit function 1-38
 - exit routines 1-38
 - input for branch entry 1-43
 - interface with exit routines 1-40
 - output for branch entry 1-44
 - re-entry 1-41
 - save area recursion with 1-41
 - service routine
 - branch entry 1-42
 - 702 abend 1-41
- POST macro instruction
 - function 1-37
- POST-without-ECB 1-44
- PRB 1-174
- precedence 1-12
- prefix format, SWA control blocks 1-119
- preventing
 - deadlock 1-23
 - interlock 1-23
- primary
 - address space 1-82
 - and CML lock 1-24
 - mode 1-82
- PRIMARY MODE parameter of SETFRR 1-162
- printable characters 1-70
- priority 1-12
 - association 1-76
 - dispatching 1-76
 - of SRB 1-76
- private
 - storage areas 1-216
- problem
 - program state 1-208
 - for retry 1-190
 - state bit in entry table 1-88
- process program interruptions 1-141
- processing control features for SVC 99 1-237
 - control limit 1-237
 - control limit on allocated data sets 1-237

- convertible attribute 1-238
 - parameters that can change 1-238
- in-use bit and attribute 1-237
- permanently allocated attribute 1-238
- TMP actions 1-237
- processor
 - affinity 1-77
 - lock 1-22
 - protection key 1-208
- profile 1-204
- profiles
 - in-storage 1-203
- program
 - authorization 1-86
 - call 1-148
 - call (PC) instruction 1-82
 - call instruction 1-87
 - call/program transfer sequence
 - consistent use of 1-89
 - check 1-141
 - interruption
 - processing 1-144
 - interruptions
 - processing 1-141
 - libraries 1-32
 - management 1-13
 - summary of 1-1
 - manager 1-4
 - request block 1-187
 - for STAE recovery 1-187
 - sharing 1-81
 - termination 1-141
 - transfer 1-148
 - transfer (PT) instruction 1-82
 - transfer instruction 1-87
- program termination
 - normal and abnormal
 - summary of 1-2
- programming considerations for SVC routines 1-223
- programming considerations when using SVC 99 1-249
 - accessing CVOLs or VSAM Private Catalogs 1-251
 - avoiding 0B0 abends 1-250
 - changes to the TIOT by SVC 99 routines 1-251
 - considerations for system routines 1-251
 - cross memory considerations 1-250
 - enqueueing on the SYSZTIOT 1-250
 - JES consideration 1-250
 - other system routines and SVC 99 1-250
 - outstanding STIMER, effect of 1-250
 - serialization of resources 1-249
- propagation
 - of STAI routines 1-187
- protecting
 - low storage 1-209
 - system data sets 1-193
 - the system 1-193
 - summary of 1-2

- protection
 - key 1-4
 - of low address range 1-209
- protection-disabled window 1-209
- PROTPSA macro instruction 1-209
- provide
 - non-space switch services 1-98
 - recovery routines 1-141
 - service 1-88
- PSA 1-155
 - fetch protection of 1-197
- PSAAOLD 1-81
 - and locking 1-24
- PSACROSV 1-209
- PSASEL 1-155
- PSASUPER 1-161
 - area of PSA 1-161
- PSL
 - contents of 1-131
- PSW 1-156
 - changing fields in 1-208
 - extended control (EC) 1-166
 - key mask of 1-86
 - of caller
 - in STKE 1-90
- PSW key 1-162
 - for retry 1-173, 1-176
 - set using the SPKA instruction 1-86
- PSW key mask
 - for retry 1-176
 - of caller
 - in STKE 1-90
- PT command 1-148
- PT instruction 1-82, 1-85, 1-87, 1-94
- PTRACE macro instruction 1-141, 1-148
 - function 1-148
 - restricted 1-148
- purged I/O
 - list 1-174
 - request list 1-174
- PURGEDQ macro instruction 1-191
 - function 1-76, 1-79
 - parameters 1-79
- PURGEDQ SVC 1-37
- purging service requests 1-79

Q

- QEDIT macro instruction 1-17
 - example 1-9
 - function 1-17
- queued access methods
 - finding the UCB address 1-33
- quiesced restorable I/O operations 1-174

R

- RACDEF macro instruction 1-204
 - use of 1-203
- RACF 1-203
 - building in-storage profiles 1-204
 - checking authorization 1-204
 - defining a resource 1-204
 - function 1-203
 - identifying a RACF user 1-204
- RACHECK macro instruction 1-204
 - use of 1-204
- RACINIT macro instruction 1-204
 - use of 1-204
- RACLIST macro instruction 1-204
- RACROUTE macro instruction 1-207
- RACXTRT macro instruction 1-204
- RB 1-156
 - purging of queue 1-174
 - queue
 - for STAE, how purged 1-189
 - things not cancelled by purge 1-174
 - RBWCF field 1-45
- RBEP 1-212
- RBFEPARM 1-170
- RBIQETP 1-212
- RBNEXAV 1-212
- RBOPSW 1-175, 1-176, 1-212
- RBPPSAV1 1-212
- RBSIZE 1-212
- RBSTAB 1-212
- RBWCF 1-45
- re-entrant modules
 - using 1-16
- real storage management
 - address space lock 1-22
 - common lock 1-22
 - cross memory lock 1-22
 - global lock 1-21
 - lock 1-22
 - steal lock 1-22
- real storage requirements
 - possible increase in cross memory 1-83
- receiver check 1-69
- recovery
 - an extension of main routine 1-182
 - considerations
 - for PC routine 1-100
 - environment for resource managers 1-191
 - for recovery routines 1-182
 - for SRBs 1-142
 - for subtasks 1-184
 - for tasks 1-142
 - routines 1-141, 1-158
 - cancellation 1-177
 - decisions 1-172
 - ESTAE 1-142, 1-173
 - ESTAE-type 1-168
 - ESTAE-type summary 1-186
 - for locked disabled SRB mode routine 1-159
 - FRR 1-142, 1-159
 - FRR summary 1-186
 - functions 1-142, 1-158, 1-183
 - guidelines 1-183
 - hierarchy of 1-185
 - major decisions regarding 1-183
 - nested 1-183
 - order 1-159
 - parameter area 1-172
 - propagation of 1-187
 - provide 1-142
 - recovery 1-182
 - requirements 1-160
 - restrictions 1-160
 - resume 1-172
 - retry 1-172
 - retry from an ESTAE-type 1-173
 - retry from an FRR 1-173
 - routing control to 1-178
 - selecting 1-160
 - types 1-142, 1-158
 - when required 1-183
 - SRB to task percolation 1-181
 - subtask 1-187
 - termination
 - invoke 1-141
 - reenterable SVCs 1-224
 - refreshable SVCs 1-224
 - region control task 1-144
 - regions
 - V = R 1-116
 - registers
 - at time or error 1-166
 - contents for SVC routines 1-224
 - for retry 1-173, 1-174
 - for STAE/STAI retry routine 1-189
 - in SDWA 1-184
 - on entry to ESTAE routine 1-165
 - on entry to FRR 1-165
 - on entry to STAE routine 1-188
 - update 1-175
 - upon entry to exit routine 1-216
 - release
 - virtual storage 1-129
 - RELEASE option of fix function 1-129
 - RELEASE option of free function 1-129
 - releasing devices 1-31
 - remote
 - service class 1-67
 - remove
 - access for all users 1-95
 - access for cross memory 1-91
 - PT authority 1-95
 - SSAR authority 1-95
 - removing in-use attribute based on task id
 - See dynamic unallocation
 - removing the in-use attribute based on task-ID 1-310
 - See also SVC 99 text units
 - reporting system characteristics 1-51

- summary of 1-1
- request block
 - resumption 1-45
 - suspension 1-45
- requesting
 - a retry 1-142, 1-184
- requesting SVC 99 functions 1-249
 - text units for parameter list 1-270
 - See also SVC 99 text units
- reserve
 - a device 1-30
 - AX 1-86, 1-89
 - LX 1-89
- RESERVE macro instruction 1-27
 - example 1-36
 - finding the UCB address 1-33
 - function 1-30
 - use of 1-51
- reset
 - AX 1-96
- reset must-complete 1-26
- residency mode 1-13
- resource
 - access control facility 1-203
 - accounting 1-75
 - management 1-100
 - in cross memory 1-83
 - management for cross memory 1-83
 - manager
 - environment 1-191
 - purpose 1-142
 - recovery environment for 1-191
 - responsibilities 1-191
 - routine 1-191
 - routine, for critical resources 1-163
 - system 1-191
 - system provided 1-191
 - uses 1-191
 - manager termination routine 1-77
 - profiles 1-204
 - scope of 1-51
 - serialized by local or global locks 1-100
- response time
 - transaction 1-55
- restart 1-68
 - with RESERVE 1-30
- RESTART key 1-172
- restricting
 - load module access 1-201
 - SVC routines 1-200
 - unauthorized users 1-198
- restrictions
 - for DAT-OFF routines 1-138
- resume 1-172
 - recovery routine 1-172
- RESUME macro instruction
 - ASYN option 1-48
 - caller in SRB mode 1-48
 - function 1-47
 - issued in cross memory mode 1-47
 - MODE option 1-48
 - options 1-48
- RESUME parameter 1-172
- retrieve
 - address and addressing mode of a nucleus
 - CSECT 1-139
 - name and entry point address of a nucleus
 - CSECT 1-139
- retrieving data set information via SVC 99
 - See dynamic information retrieval
- retry 1-158
 - address 1-173
 - from a recovery routine 1-184
 - from an ESTAE-type recovery routine 1-173
 - from an FRR 1-173
 - requesting 1-142, 1-184
 - routine 1-172
 - cross memory environment 1-176
 - freeing SDWA 1-170
 - when not permitted 1-172
- return
 - address
 - from PCLINK service, in STKE 1-90
 - codes
 - for STAE 1-188
 - from recovery routine 1-165
- return codes
 - from MVS router exit routine 1-206
- return codes from SVC 99
 - See SVC 99 return codes
- review of dsname processing
 - See dsname allocation processing
- RIB
 - format 1-52
 - used with GQSCAN 1-52
- RIBE
 - format 1-52
- RISGNL macro instruction 1-68
- RMF
 - transaction activity report 1-55
 - used to report SRM data 1-55
 - workload activity report 1-55
- RMODE 1-13
 - assembler definition 1-13
 - values for 1-14
- RMTR
 - interface to 1-80
 - parameter 1-79
- router exit 1-205
 - register 1's content on entry 1-206
- routing the message 1-70
- RPSGNL macro instruction 1-68
- RQE 1-212
- RSM
 - functions 1-127
 - lock 1-22
 - paging services 1-128
 - summary of 1-2
- RSMAD lock 1-22
- RSMCM lock 1-22

RSMGL lock 1-21
 RSMST lock 1-22
 RSMXM lock 1-22
 RTM
 invoking 1-142
 receives control 1-141
 summary of services 1-141
 use of 1-142
 RTM2 SVRB 1-174
 RTM2 work areas 1-156
 RTM2WA 1-156

S

SAC instruction 1-83
 SAF 1-205
 key element in 1-205
 SALLOC lock 1-22
 SASID 1-82, 1-91, 1-163
 save
 and restore registers with SYNCH 1-15
 save and restore status
 in cross memory 1-89
 save area
 address contained in STKE 1-90
 caller's in cross memory 1-90
 for CALLRTM 1-143
 new task 1-3
 POST routine 1-41
 standard in cross memory 1-100
 scanning the CIB chain 1-17
 SCB
 area provided for 1-225
 SCHEDULE macro instruction 1-76
 FRR required option 1-78
 function 1-75
 LOCAL lock required option 1-78
 scheduled dump 1-154
 definition 1-154
 when produced 1-154
 with an ECB 1-154
 scheduler work area 1-5
 scheduling an SRB 1-75
 scope
 ALL 1-53
 GLOBAL 1-53, 1-54
 LOCAL 1-53, 1-54
 STEP 1-51, 1-53
 SYSTEM 1-51, 1-53
 SYSTEMS 1-51, 1-53
 SDUMP macro instruction 1-141, 1-157, 1-185
 address space options 1-151
 BRANCH option 1-155
 BRANCH options 1-151
 considerations 1-186
 ECB option 1-185
 fails to dump 1-152
 in a reentrant program 1-150

 options 1-150
 parameter list length 1-150
 SQA buffer option 1-152
 SUSPEND option 1-155
 TYPE=FAILRC parameter 1-152
 use 1-149
 using 1-150
 SDWA 1-159
 availability for retry 1-173, 1-182
 freeing 1-170
 key fields and meanings 1-166
 SDWASDRC byte in 1-153
 use of 1-166
 SDWA extensions 1-159
 SDWACID 1-186
 SDWACLUP 1-172
 SDWACLUP bit 1-173
 SDWACMPC 1-166
 SDWACOMU 1-167, 1-182
 SDWACRC 1-166, 1-186
 SDWAC SCT 1-186
 SDWACTS 1-185
 SDWADAET 1-167
 SDWAEAS 1-185
 SDWAEBC 1-167
 SDWAEC1 1-166
 SDWAEC2 1-166
 SDWAFAIN 1-167
 SDWAGLBL 1-163
 SDWAGRSV 1-166
 SDWAHEX 1-167
 SDWALCL 1-164
 SDWALNTH 1-167
 SDWAMABD 1-185
 SDWAMLVL 1-186
 SDWAMODN 1-186
 SDWAOCUR 1-167
 SDWAPARM 1-166
 SDWAREXN 1-186
 SDWARPIV 1-182
 SDWARRL 1-186
 SDWASC 1-186
 SDWASPID 1-167
 SDWASRSV 1-167, 1-173
 SDWAURAL 1-167
 SDWAVRA 1-159, 1-185, 1-186
 SDWAVRAL 1-167
 second level table for cross memory linkage 1-87
 secondary
 access key 1-86
 address space 1-82
 and the CML lock 1-24
 mode 1-82
 segment and page faults
 in cross memory 1-84
 self-contained function 1-191
 sense 1-67
 sensitive functions
 protecting 1-198
 serialization 1-1, 1-12, 1-17, 1-26

- avoidance of 1-75
- methods 1-17
- of SRB to task percolation 1-181
- of task execution 1-26
- of the use of a task 1-48
- requirements 1-17
- summary of 1-1
- techniques
 - for the same volume 1-32
 - when needed 1-17
- serially reusable resource 1-18
- service
 - classes
 - direct 1-67
 - of IPC 1-67
 - remote 1-67
 - management facilities 1-76
 - purging of requests 1-79
- set
 - address space control (SAC) instruction 1-83
 - an address space's AT 1-86
 - an address space's AX 1-86
 - AT 1-89
 - AX 1-89, 1-93
 - prefix 1-68
 - secondary ASID 1-148
 - secondary ASN (SSAR) instruction 1-83
- SET DAE command 1-149
- SET DAE=xx command 1-157
- set DIE
 - function 1-214
 - input 1-214
 - mode of caller 1-214
 - restrictions on caller 1-216
 - return codes 1-215
- set must-complete 1-26
 - with RESERVE 1-30
- set up
 - connect entry table to all address spaces 1-92
 - cross memory environment 1-91
- SETFRR macro instruction 1-141, 1-159, 1-162, 1-183
 - MODE=FULLXM 1-163
 - MODE=GLOBAL 1-163
 - MODE=HOME 1-162
 - MODE=LOCAL 1-164
 - MODE=PRIMARY 1-163
 - options for cross memory 1-100
- SETLOCK macro instruction 1-79
 - function 1-24
- SETRP macro instruction 1-161, 1-166, 1-184
- setting up the buffer for MGCR 1-6
- SFT
 - contents 1-89
 - use of 1-89, 1-96
- shared DASD 1-28
 - equipment needed for 1-29
 - macros used with 1-30
 - use of 1-28
- shared direct access storage devices 1-28
- sharing
 - data 1-28
 - the same control program 1-67
- SIGP
 - order codes 1-67, 1-68
 - CPU reset 1-68
 - emergency signal 1-68
 - external call 1-69
 - initial CPU reset 1-68
 - restart 1-68
 - sense 1-67
 - set prefix 1-68
 - start 1-67
 - stop 1-67
 - stop and store status 1-68
 - store status at address 1-68
- single-event wait 1-225
- SLIP
 - command 1-141, 1-147, 1-157
 - restricted 1-147
 - establish traps 1-147
 - modify traps 1-147
 - obtaining an SVC dump 1-147
 - traps 1-147
 - displayed 1-147
 - non-PER 1-146
 - PER 1-146
 - use of 1-146
 - using 1-147
- SLIP command
 - ACTION keywords 1-147
- SMF cross memory services lock 1-22
- SNAP macro instruction 1-149
- space
 - not enough allocated 1-185
 - space allocation lock 1-22
 - space and unit defaults for dynamic allocation 1-245
 - space switch entry tables
 - ownership of 1-83
- SPIE environment
 - cancelling 1-144
 - deleted 1-147
- SPIE macro instruction 1-141
 - function 1-144
 - issued by problem program 1-145
- SPIE/ESPIE environment 1-144
- spin locks 1-19
- SPOST macro instruction
 - function 1-37
- SQA
 - buffer dumped by SDUMP 1-152
 - buffer option of SDUMP macro instruction 1-152
 - subpools in 1-107
- SRB
 - characteristics 1-78
 - cleanup 1-49
 - dispatching queue 1-191
 - errors in 1-184
 - priorities 1-78
 - required information for SCHEDULE 1-77
 - to task percolation 1-181

- transferring control 1-49
 - use of 1-77
- SRB's related task
 - definition of 1-181
- SRBASCB 1-77
- SRBCPAFF 1-77
- SRBEP 1-77
- SRBPARM 1-77
- SRBPASID 1-77, 1-181
- SRBPKF 1-77, 1-86
- SRBPTCB 1-77, 1-181
- SRBRMTR 1-77
- SRM
 - lock 1-22
 - reporting interface 1-55
- SSAR instruction 1-83, 1-85, 1-94, 1-148
- SSBLOCK 1-92
- SSL
 - contents of 1-131
- SSRB 1-156
- STACK option of PCLINK 1-90
- STAE
 - routines
 - authorization requirements for 1-190
 - STAE environment
 - when canceled 1-187
 - STAE macro instruction 1-187
 - STAE protection
 - continued 1-188
 - STAE/STAI
 - exit routines 1-187
 - retry routines 1-189
 - routine 1-188
 - interface to 1-188
 - when not entered 1-188
- stage 1
 - exit effector 1-211
 - interface to 1-211
 - initialization 1-211
- stage 2
 - exit effector 1-211
 - interface to 1-213
 - restrictions on caller 1-213
 - scheduling 1-213
- stage 3 1-213
 - execution 1-213
 - exit effector 1-213
- STAI parameter
 - of ATTACH 1-159
- standard
 - EBCDIC characters 1-70
 - interface 1-194
- start 1-67
- START CIB 1-7
- START command
 - issuing an internal 1-6
- started program 1-6
- state changes 1-75
- status conditions
 - check stop 1-69
 - equipment check 1-69
 - external call pending 1-69
 - incorrect state 1-69
 - invalid function 1-69
 - invalid parameter 1-69
 - MSSF failure 1-69
 - operator intervening 1-69
 - receiver check 1-69
 - stopped 1-69
- status indicators 1-69
- STATUS macro instruction 1-5
- STCK instruction 1-219
- stepname 1-156
- STKE
 - contents of 1-90
 - created by 1-90
- stop 1-67
- stop and store status 1-68
- STOP command 1-7
- stopped 1-69
- storage
 - alteration 1-146
 - auxiliary 1-127
 - dumping virtual 1-149
 - in cross memory 1-84
 - isolation 1-81
 - keys 1-107
 - non-pageable 1-127
 - protection
 - types of data needing 1-197
 - real frames 1-103
 - subpools 1-106
 - virtual 1-127
 - above 16 megabytes 1-127
- store status at address 1-68
- SUBPLST 1-151
- subpools 1-107
 - characteristics of 1-107
 - default 1-189
 - fetch protected 1-107
 - fixed 1-107
 - for STAE retry processing 1-190
 - global 1-18
 - ID 1-167
 - shared between tasks 1-17
 - storage key of 1-107
 - task related 1-84
 - type of storage 1-107
 - using 1-106
 - where backed 1-107
 - 236 and 237 1-5
- subsystem
 - how term is used in cross memory 1-91
 - identifiers 1-55
- subtasks
 - creating additional 1-3
 - creation and control
 - summary of 1-1
 - propagation of recovery routine to 1-168
 - recovery for 1-184, 1-187
- successful branch 1-146

SUMLIST 1-155, 1-156
SUMLSTA 1-155, 1-156
summary
 ESTAE-type recovery routine 1-186
 FRR 1-186
 of authorization rules 1-202
 of facilities available in cross memory 1-84
 of locking characteristics 1-19
 of macro instructions for cross memory authorization 1-89
 of macro instructions for cross memory linkage 1-89
 of PC/PT linkage conventions 1-91
summary dump 1-154
 avoid duplicate data 1-154
 disabled 1-155
 enabled 1-156
 parameters used 1-155
 suspend 1-155
super FRR stack 1-163, 1-164
supervisor
 assisted linkage 1-174
 control FRR 1-143
 state 1-208
 state for retry 1-190
supervisor/problem program state 1-161
 for retry 1-175
SUSMDUMP
 more than one 1-158
suspend
 count field 1-45
 locks 1-20, 1-216
 summary dump 1-155
SUSPEND and RESUME
 proper order 1-47
SUSPEND macro instruction
 considerations for use 1-46
 examples 1-46
 function 1-45
 used in cross memory mode 1-46
 used with RESUME 1-47
SUSPEND option of SDUMP 1-155
SUSPEND RB= CURRENT scenario 1-46
SUSPEND RB= PREVIOUS scenario 1-46
SUSPEND/RESUME 1-17
suspended FRR 1-165
SVC
 dump
 during SLIP processing 1-147
 from recovery routines 1-185
 initial status 1-153
 obtain 1-153
 obtaining 1-153
 ESTAE recovery for routines 1-225
 first-level interrupt handler 1-200
 FLIH 1-200
 FRR recovery for routines 1-225
 invalid use 1-141
 locking conventions for routines 1-224
 naming conventions for routines 1-224
 register contents 1-224
 restore 1-189
 SVC dumps 1-156
 SVC routines, user written
 exiting from 1-222
 SVC types 1-5 1-221
 SYS1.LPALIB 1-221
 SYS1.NUCLIB 1-221
 SYS1.PARMLIB 1-221
 TCBACTIV flag 1-222
 type 6 SVCs 1-222
 T6EXIT macro 1-222
 SVC routines, user-written
 IEASVC dataset member 1-225
 IEASYS dataset member 1-225
 inserting into control program at IPL time 1-225
 non-preemptable 1-222
 programming conventions 1-223
 screening access to 1-227
 STAX macro instruction 1-222
 SVC= statement 1-225
 SVC Parm statement 1-225
 SVCUPDTE macro instruction 1-226
 SYS1.PARMLIB 1-225
 writing 1-221
 SVC 13 1-160
 SVC 3 1-174, 1-189
 SVC 34 1-144
 SVC 51 1-154
 SVC 60 1-170
 SVC 99 error reason codes 1-264
 See also SVC 99 return codes
 classes of 1-264
 for an invalid parameter list 1-266
 for an unavailable system resource 1-265
 for environmental errors 1-267
 for system routine errors 1-269
 SVC 99 information reason codes 1-262
 See also SVC 99 return codes
 SVC 99 introduction
 See introduction to SVC 99 functions
 SVC 99 parameter list 1-251
 mapping macros for 1-251
 request block 1-253
 DSECT S99RB 1-253
 request block pointer 1-253
 DSECT S99RBP 1-253
 structure of 1-252
 text pointers 1-255
 text units 1-255
 notes on structure 1-255
 SVC 99 request block
 See SVC 99 parameter list
 SVC 99 request block pointer
 See SVC 99 parameter list
 SVC 99 return codes 1-262
 error reason codes 1-264
 information reason codes 1-262
 SVC 99 text unit pointer list
 See SVC 99 parameter list

SVC 99 text units 1-270
 See also SVC 99 parameter list
 dynamic concatenation text units 1-309
 ddname specification - DCCDDNAM 1-309
 permanently concatenated attribute -
 DCCPERMC 1-309
 dynamic deconcatenation text unit 1-310
 ddname specification - DDCDDNAM 1-310
 dynamic information retrieval 1-312, 1-316
 ddname specification 1-312
 dsname specification - DINDSNAM 1-312
 relative request number - DINRELNO 1-316
 return conditional disposition -
 DINRTCDP 1-314
 return control limit - DINRTLIM 1-315
 return data set type - DINRTTYP 1-316
 return ddname - DINRTDDN 1-313
 return dsname - DINDTDSN 1-313
 return DSORG - DINRTORG 1-314
 return dynamic allocation attributes -
 DINRTATT 1-315
 return last relative entry - DINRTLST 1-315
 return member name - DINRTMEM 1-313
 return normal disposition - DINRTNDP 1-314
 return status - DINRTSTA 1-313
 for ddname allocation 1-311
 ddname specification - DDNDDNAM 1-311
 return DUMMY indication -
 DDNRTDUM 1-311
 for dsname allocation 1-273
 align form or verify FCB image -
 DALFCBAV 1-282
 allocated space format - DALSPFRM 1-276
 block length - DALBLKLN 1-275
 burst specification - DALBURST 1-287
 character arrangement table specification -
 DALCHARS 1-287
 conditional disposition - DALCDISP 1-274
 copy groups specification - DALCOPYG 1-287
 copy modification module specification -
 DALMMOD 1-288
 copy module table reference character -
 DALMTRC 1-288
 cylinder space - DALCYL 1-275
 data set sequence number - DALDSSEQ 1-280
 data set status - DALSTATS 1-274
 DCB ddname reference - DALDCBDD 1-284
 DCB dsname reference - DALDCBDS 1-284
 ddname specification - DALDDNAM 1-273
 defer mounting - DALDEFER 1-289
 directory blocks - DALDIR 1-276
 dsname specification - DALDSNAM 1-273
 dummy data set - DALDUMMY 1-281
 expiration date (DALEXPDT) 1-281
 FCB image identification - DALDCBIM 1-282
 flash forms overlay count - DALFCNT 1-288
 flash forms overlay specification -
 DALFFORM 1-288
 fold mode - DALUFOLD 1-283
 input or output only - DALINOUT 1-281
 label type - DALLABEL 1-280
 member name specification -
 DALMEMBR 1-273
 MSVGP specification - DALMSVGP 1-285
 normal disposition - DALNDISP 1-274
 OUTPUT statement reference -
 DALOUTPT 1-289
 parallel mount - DALPARAL 1-278
 password protection - DALPASPR 1-280
 primary space quantity - DALPRIME 1-275
 private volume - DALPRIVT 1-277
 QNAME specification - DALQNAME 1-282
 RACF protection - DALPROT 1-286
 release unused space - DALRLSE 1-276
 retention period - DALRETPD 1-281
 round 1-276
 secondary space quantity - DALSECND 1-275
 subsystem name request - DALSSNM 1-285
 subsystem parameters - DALSSPRM 1-286
 SYSOUT copies - DALCOPYS 1-280
 SYSOUT form number - DALSFMNO 1-279
 SYSOUT hold queue - DALSHOLD 1-285
 SYSOUT output limit - DALOUTLM 1-279
 SYSOUT program name - DALSPGNM 1-279
 SYSOUT remote user - DALUSER 1-284
 SYSOUT specification - DALSYSOU 1-278
 SYSOUT User ID specification -
 DALUSRID 1-286
 terminal is an I/O device - DALTERM 1-283
 track space - DALTRK 1-274
 Unallocate at CLOSE - DALCLOSE 1-279
 unit count - DALUNCNT 1-278
 unit specification - DALUNIT 1-278
 universal character set - DALUCS 1-283
 verify character set image -
 DALUVFRY 1-283
 volume count - DALVLCNT 1-277
 volume reference - DALVLRDS 1-277
 volume sequence number - DALVLSEQ 1-277
 volume serial numbers - DALVLSER 1-276
 for dsname allocation (DCB) 1-291, 1-302
 blocksize - DALBLKSZ 1-291
 buffer alignment - DALBFALN 1-291
 buffer count per DCB - DALBUFNO 1-292
 buffer length - DALBUFL 1-292
 buffer offset - DALBUFOF 1-292
 buffer size per line group - DALBUFSZ 1-293
 buffering technique - DALBFTEK 1-291
 card reader/punch mode - DALMODE 1-296
 data set key length - DALKYLEN 1-295
 data set organization - DALDSORG 1-294
 DIAGNS = TRACE specification -
 DALDIAGN 1-300
 error option - DALEROPT 1-295
 first buffer reserve specification -
 DALRSRVF 1-299
 FRID = specification - DALFRID 1-302
 FUNC = specification - DALFUNC 1-301
 GET macro buffer request -
 DALBUFRQ 1-293
 GNCP specification - DALGNCP 1-295

IPLTXTID specification - DALIPLTX 1-300
 logical record length - DALLRECL 1-296
 maximum buffer numbers per line -
 DALBUFMX 1-292
 optional CP services - DALOPTCD 1-297
 polling interval - DALINTVL 1-295
 printer line spacing - DALPRTSP 1-298
 punch paper tape code - DALCODE 1-293
 READ/WRITE maximum - DALNCP 1-296
 receiving buffer count - DALBUFIN 1-292
 receiving PCI specification - DALPCIR 1-297
 record format - DALRECFM 1-298
 search limit - DALLIMCT 1-296
 secondary buffer reserve specification -
 DALRSRVS 1-299
 sending buffer count - DALBUFOU 1-293
 sending PCI specification - DALPCIS 1-298
 sending/receiving priority - DALCPRI 1-294
 size-of-work-area specification -
 DALSOWA 1-299
 STACK specification - DALSTACK 1-299
 tape density - DALDEN 1-294
 THRESH specification - DALTHRS 1-300
 TRTCH specification - DALTRTCH 1-300
 for dsname allocation (non-JCL) 1-302
 convertible attribute - DALCNVRT 1-303
 password specification - DALPASSW 1-302
 permanently allocated attribute -
 DALPERMA 1-303
 return ddname - DALRTDDN 1-303
 return dsname - DALRTDSN 1-303
 return DSORG - DALRTORG 1-304
 return volume serial - DALRTVOL 1-305
 subsystem request - DALSSATT 1-305
 subsystem request - DALSSREQ 1-304
 for dynamic unallocation 1-306, 1-309
 ddname unallocation - DUNDDNAM 1-306
 dsname unallocation - DUNDSNAM 1-306
 member name specification 1-306
 override SYSOUT class - DUNOVCLS 1-308
 override SYSOUT hold - DUNOVSHQ 1-308
 override SYSOUT nohold -
 DUNOVSNH 1-308
 override SYSOUT remote workstation -
 DUNOVSNH 1-308
 overriding disposition - DUNOVDSP 1-307
 remove in-use option - DUNREMOV 1-307
 unalloc option - DUNUNALC 1-307
 removing the in-use attribute 1-310
 current task option - DRICURNT 1-311
 TCB address specification - DRITCBAD 1-310
 SVC 99 text units for retrieving information
 See SVC 99 text units
 SVCs 1-216
 cross memory restrictions 1-78
 needing more than one recovery routine 1-170
 SVCTABLE macro instruction 1-200, 1-224
 SVCUPDTE macro instruction 1-226
 SVRB 1-170, 1-224
 SWA 1-5
 SWA control blocks
 above 16 megabytes 1-119
 assigning 1-117
 below 16 megabytes 1-119
 error processing 1-125
 format of 1-118
 locate mode processing 1-124
 move mode processing 1-121
 prefix format 1-119
 SWA manager processing 1-120
 SWA manager 1-120
 error processing 1-125
 locate mode 1-124
 move mode 1-121
 swapped-out 1-127
 address spaces in cross memory 1-84
 SWAREQ macro instruction 1-124
 symptom strings 1-159
 SYMREC (symptom recording)
 description of 1-56
 SYNCH macro instruction
 function 1-15
 saving and restoring registers 1-15
 SYNCH service routine 1-168, 1-187
 synchronization loop 1-218
 synchronous
 dump 1-154
 exits 1-15
 interrupts 1-45
 SYSABEND 1-184
 SYSDUMPs 1-156
 SYSEVENT macro instruction 1-55
 SYSDUMP 1-158, 1-184
 SYSDUMP DD statement 1-158
 system
 changing status 1-208
 configuration 1-29
 control stages 1-211
 data
 protecting 1-197
 data sets
 protecting 1-193
 environment 1-160
 addressing mode 1-161
 authorization index 1-162
 cross memory state 1-162
 definition 1-160
 disablement 1-161
 identify 1-161
 locking 1-161
 normal addressing environment 1-162
 PSW key 1-162
 PSW key mask 1-162
 register interface 1-165
 restricted addressing environment 1-163
 supervisor/problem program state 1-161
 errors 1-141
 interception of 1-141, 1-146
 function table 1-89
 integrity 1-193

- level 1-78
- program
 - definition 1-4
 - recovery area 1-159
 - resources manager lock 1-22
 - trace facilities 1-148
 - trace table 1-148
 - tracing in cross memory 1-85
- system authorization checking 1-204
- system authorization facility 1-205
- system initiated
 - logoff 1-171
- system log
 - writing message to 1-74
- system trace facilities 1-141
- system-initialized cancellation 1-144
- SYSUDUMP 1-184
- SYSZTIOT
 - See programming considerations when using SVC 99
- SYS1.MACLIB 1-8
- SYS1.PARMLIB 1-199
 - ADYSETxx 1-157
 - libraries concatenated to 1-200
 - LNKLSTxx 1-200

T

- tailored dumps 1-183
- task
 - changing the priority 1-12
 - creation 1-3
 - identifier 1-5
 - rules for creating 1-3
 - serializing the execution of 1-26
 - termination 1-141, 1-191
 - preventing 1-26
 - resource manager duties 1-191
- TCAS 1-144
- TCB 1-156
 - information for RESERVE 1-33
 - parameter of CALLRTM macro instruction 1-143
 - providing information for 1-11
- TCBFSM 1-176, 1-190
- TCBJSCB 1-5
- TCBPKF 1-86, 1-145, 1-176
- TCBTID 1-5
- TCTL macro instruction
 - considerations for use 1-49
 - function 1-49
- terminal monitor program (TMP) 1-237
- termination
 - abnormal 1-182, 1-191
 - address space 1-141
 - continuing with 1-187
 - normal 1-191
 - of job step task 1-100
 - task 1-141
- test and set instruction 1-17

- TESTAUTH macro instruction
 - SVC routines 1-200
- text unit keys
 - See SVC 99 text units
- text units for dynamic allocation
 - See SVC 99 text units
- timer
 - disabled interrupt exit 1-214
 - interruption 1-216
 - supervision 1-214
- TIOT
 - address of 1-33
 - default size for 1-245
 - modifying size of 1-245
 - obtaining the address of 1-31
- TOD
 - clock 1-219
 - current reading 1-217
- token
 - for internal START 1-6
 - for MGCR 1-6
 - issuing an internal 1-6
 - used with DOM macro instruction 1-73
 - used with the GQSCAN macro instruction 1-52
- TOKEN parameter
 - of DOM 1-73
 - of ESTAE 1-169
 - relation to XCTL 1-169
- TPCSDIE 1-214
- TQE
 - address of 1-217
 - controlling 1-218
 - freed when address space fails 1-218
 - freeing 1-218
 - obtain 1-218
 - serialization 1-218
- TQE DEQUEUE routine
 - function 1-218
 - input 1-219
- TQE ENQUEUE routine 1-216
 - differences from set DIE 1-217
 - input environment 1-217
 - register contents on exit 1-217
 - registers 1-217
- TQE AID 1-215
- TQE ASCB 1-216
- TQE DREGS 1-217
- TQE EXIT 1-215, 1-217
- TQE TCB 1-216
- TQE VAL 1-217
- TRACE command 1-148
- TRACE lock 1-22
- tracing
 - address space 1-148
 - branch 1-148
 - explicit 1-148
- transaction
 - activity report 1-55
 - of an interactive system 1-55
 - response time 1-55

- transfer pages 1-127
- transferring control
 - for SRB processing 1-49
- TS instruction 1-17
- type 1 or 6 SVC issuing CALLDISP 1-50
- type 1 SVC 1-38
- type 1 SVC interface 1-211
- type 1 SVC issuing SUSPEND 1-38
- type 17 interruption 1-145
- type 2 3 or 4 SVC 1-170
- type 2 3 or 4 SVC issuing SUSPEND 1-46
- type 6 SVC issuing SUSPEND 1-47
- type 6 SVC routines 1-222
- TYPE = XMEM 1-151
- TYPE = XMEME 1-151
- T6EXIT 1-47

U

- U.S. national characters 1-70
- UCB address
 - finding 1-33
 - finding via the DEB 1-33
 - of reserved device
 - finding 1-33
- UCB scan services
 - by device class 1-230
 - for all UCBs 1-230
 - limiting the scan 1-230
 - purpose 1-229
 - summary of 1-2
 - using IOSVSUCB 1-229
- unallocation by SVC 99
 - See dynamic unallocation
- uncorrectable translation errors 1-143
- unlocked task recovery routines 1-160
- unopened data sets
 - finding the UCB address 1-33
- unprivileged macro instructions 1-149
- UNSTACK THRU option of PCLINK 1-91
- use attribute of permanently allocated
 - See permanently allocated attribute
- user defined functions 1-135
- user-supplied addresses
 - for protected control blocks 1-195
 - for user storage areas 1-194
- user-written SVC routines
 - summary of 1-2
- userid
 - defined to RACF 1-204
- USERINFO field in EQT 1-41
- using
 - storage subpools 1-106

V

- valid volume characteristics 1-29
- validate
 - user-supplied addresses 1-194
- validating input for dynamic allocation functions
 - See installation input validation routine for SVC 99
- validation
 - primary technique 1-194
- verb codes
 - MODIFY 1-7
 - START 1-7
 - STOP 1-7
- verify control blocks 1-195
- virtual
 - addressing 1-74
 - equal real regions 1-127
 - storage 1-135
 - above 16 megabytes 1-127
 - allocated 1-111
 - allocating 1-104
 - dump 1-141
 - dumping 1-149
 - fix 1-128
 - fixing 1-128
 - free 1-115, 1-128
 - freeing 1-128
 - load 1-128
 - map 1-108
 - map of DAT-ON nucleus 1-108
 - obtaining information about 1-109
 - page out 1-128
 - release 1-128
 - unallocated 1-116
 - storage management
 - common area subpools lock 1-21
 - common VSM work area lock 1-22
- virtual page 1-127
- volatile information
 - saving 1-185
- volume
 - assignment 1-31
- volume and device status 1-29
- volume handling
 - rules with shared DASD 1-30
- VRA
 - data to be printed in EBCDIC 1-167
 - data to be printed in hexadecimal 1-167
 - length 1-167
 - use of 1-186
- VRADAE indicator 1-156
- VRADATA macro instruction 1-159, 1-186
- VSL
 - contents of 1-130
- VSM
 - services 1-103
 - summary of 1-2
- VSMFIX lock 1-21

VSMLIST macro instruction 1-103
 use 1-109
VSMLIST work area
 description of 1-110
 using 1-109
VSMLOC macro instruction 1-103
VSMPAG lock 1-22
VSMREGN macro instruction 1-103

W

WAIT
 count field 1-45
 entry point 1-44
 functions 1-44
 service routine
 branch entry 1-44
WAIT macro instruction
 function 1-38
wait time for job step exceeded 1-171
WAIT-without-ECB 1-44
WAIT/POST/EVENTS 1-17
waiting
 for an event to complete 1-45
 for event completion 1-38
work area
 for CALLRTM 1-143
 for recovery routines 1-188
workload activity report 1-55
writing
 operator messages 1-70
 user-written SVC routines 1-221
WSA vector table 1-155

WTL macro instruction
 function 1-74
WTO macro instruction
 function 1-70
WTOR macro instruction
 function 1-70

X

XC function 1-137
XCTL macro instruction 1-32
XSB 1-156, 1-176, 1-190
XSBKM 1-176, 1-190
x37 abend 1-185

Numerics

047 abend 1-200
052 abend 1-100
070 abend 1-48
070 system completion code 1-49
2305 fixed head storage facility 1-28
2835 storage control unit 1-28
30E abend 1-145
306 abend 1-201, 1-202
3830 storage control unit 1-28
3880 storage control unit 1-29
46D-18 abend completion-reason code 1-145
702 abend 1-41
913 abend 1-185



Printed in U.S.A.

GC28-1150-3

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. 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.

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

What is your occupation? _____

How do you use this publication? _____

Number of latest Newsletter associated with this publication: _____

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Cut or Fold Along Line

Fold and tape

Please Do Not Staple

Fold and tape




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 390
Poughkeepsie, New York 12602



Fold and tape

Please Do Not Staple

Fold and tape

Printed in U.S.A.



GC28-1150-3

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. 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.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

What is your occupation? _____

How do you use this publication? _____

Number of latest Newsletter associated with this publication: _____

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Cut or Fold Along Line

Fold and tape

Please Do Not Staple

Fold and tape




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 390
Poughkeepsie, New York 12602



Fold and tape

Please Do Not Staple

Fold and tape

Printed in U.S.A.



GC28-1150-3

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. 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.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

What is your occupation? _____

How do you use this publication? _____

Number of latest Newsletter associated with this publication: _____

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Cut or Fold Along Line

Fold and tape

Please Do Not Staple

Fold and tape




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.



POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 390
Poughkeepsie, New York 12602



Fold and tape

Please Do Not Staple

Fold and tape

Printed in U.S.A.



GC28-1150-03

