

GC28-0628-3
File No. S370-36

Systems

**OS/VS2 System Programming
Library: Supervisor**

Release 3.8

IBM

Fourth Edition (August, 1979)

This is a major revision of, and obsoletes, GC28-0628-2 and Technical Newsletters GN28-2902 and GN28-2915 incorporating changes released in the following System Library Supplement:

MVS Processor Support 2 5752-864 GD23-0113 (dated December 29, 1978)

This edition, with Technical Newsletters GN28-4738, GN28-4740, and GN28-4978 applies to release 3.8 of OS/VS2 and to all subsequent releases of OS/VS2 until otherwise indicated in new editions or Technical Newsletters. Changes are continually made 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.

This material may contain reference to, or information about, IBM products (machines and programs), programming or services which are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Publications are not stocked at the address given below; requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Programming Systems Publications, Department D58, Building 706-2, PO Box 390, Poughkeepsie, N.Y. 12602. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Guide For Using This Publication

The following is a list of the requirements for using this publication.

- To use this publication you must have installed the following Selectable Unit:
MVS Processor Support 2 SU64
- The implied date of this publication, for the purpose of adding new Supplements/TNLs, is August 31, 1979. Always use the page with the latest date (shown in the slug at the top of the page) when adding pages from different Newsletters/Supplements.

Preface

This publication describes supervisor facilities that can be influenced by the system programmer.

Part I: Supervisor Services discusses supervisor functions restricted to system programmers. In this publication, a system programmer is defined as a programmer whose programs run supervisor state, system key 0-7 or access APF authorized libraries. The publication includes the macro instructions and parameters used to obtain the functions.

Part I is divided into seven topics. For convenience and compatibility, these services are grouped in the same manner as in *OS/VS2 Supervisor Services and Macro Instructions*, GC28-0683. Accordingly, some of the subjects discussed are merely continuations of previous coverage, but are extended to include discussions of the new functions. However, some subjects discussed under the topics are completely new, but are restricted in use to the system programmer.

Part II: Macro Instructions contains the formats and descriptions of the supervisor macro instructions. It provides system programmers with the information necessary to code the macro instructions.

Note: The system activity measurement facility (MF/1) and the dynamic support system (DSS) are not supported with MVS/System Extensions (program number 5740-XYN, -XYS).

Publications referenced:

OS/VS2 System Programming Library: Data Management, GC26-3830.
IBM System/370 Principles of Operation, GA22-7000.
OS/VS2 Supervisor Services and Macro Instructions, GC28-0683.
OS/VS2 System Logic Library, SY28-0713, SY28-0714, SY28-0715,
SY28-0716, SY28-0717, SY28-0718, SY28-0719. (7 volumes)
OS/VS - DOS/VS - VM/370 Assembler Language, GC33-4010.
OS/VS2 System Programming Library: System Generation Reference,
GC26-3792.
OS/VS2 System Programming Library: MVS Diagnostic Techniques,
GC28-0725.
OS/VS2 System Programming Library: Debugging Handbook, Volumes 1, 2,
and 3, GC28-0708-1, GC28-0709-1, and GC28-0710-0.
OS/VS2 System Programming Library: Initialization and Tuning Guide,
GC28-0681.
OS/VS2 Message Library: VS2 System Codes, GC38-1008
OS/VS2 MVS System Programming Library: Service Aids, GC28-0674
OS/VS2 System Programming Library: Job Management, GC28-0627
Operator's Library: OS/VS2 MVS System Commands, GC38-0229
OS/VS2 MVS Resource Access Control Facility (RACF): Installation Reference
Manual, SY28-0734

Contents

Summary of Amendments	xiii
Part I: Supervisor Services	1
Subtask Creation and Control	3
Creating a New Task	3
User Modifications	4
Operator Communication with a Problem Program	5
Providing an EXTRACT Answer Area	6
Program Management	9
Synchronous Exits	9
Multiprocessing Programming Considerations	9
Checkpoint/Restart	9
Re-entrant Modules	10
Priority	10
Resource Control	11
Resource Serialization	11
Serialization Requirements	11
Locking	12
Categories of Locks	12
Types of Locks	12
Classes of Locks	13
Locking Hierarchy	13
Locking Considerations	14
SETLOCK Macro Instruction	14
SETLOCK Macro Instruction (5740-XE1)	14.1
INTSECT Macro Instruction (5740-XE1)	15
The Must Complete Function	15
Characteristics of the Must Complete Function	16
Programming Notes	16
Shared Direct Access Storage Devices (Shared DASD)	17
Devices that Can be Shared	17
Volume/Device Status	17
System Configuration	18
Volume Handling	18
Macro Instructions Used with Shared DASD	18
Releasing Devices	18
Preventing Interlocks	18
Volume Assignment	19
Program Libraries	19
Finding the UCB Address	19
Authorized Program Facility (APF)	23
APF Authorization	23
Authorized Programs	23
Authorized Libraries	24
Using APF	25
Restricting SVC Routines	25
Restricting Load Module Access	25
Assigning Authorization	26
Authorization Requests Under Various Conditions	26
Guidelines for Using APF	27
Changing System Status	28
Generating an SVC	28
Generating Inline Code	28
Interprocessor Communications (IPC)	28
Service Classes	29
Status Conditions	30
Event Completion	31
Cross Memory POST	31
Bypassing the POST Routine	31
Waiting for Event Completion	32
Writing POST Exit Routines	32

Identifying and Deleting Exit Routines	33
Initializing Extended ECBs and ECB Extensions	33
POST Interface with Exit Routines	34
Re-entry to POST from a POST Exit	34
Example of Using a POST Exit Function	35
Branch Entry to the POST Macro Instruction	36
Branch Entry to the WAIT Macro Instruction	38
Suspension and Resumption of Request Block Tasks	38
Suspension and Resumption of Tasks (5740-XE1)	45
Using the SUSPEND Macro Instruction	38
Using the RESUME Macro Instruction	40
Using the TCTL Macro Instruction	40
Using the TCTL Macro Instruction (5740-XE1)	40.1
Using the BRANCH=YES Option of the CALLDISP Macro Instruction	41
System Integrity	41
System Integrity (5740-XE1)	48.1
Documentation on System Integrity	41
Documentation on System Integrity (5740-XE1)	48.1
Installation Responsibility	42
Elimination of Potential Integrity Exposures	42
User-Supplied Addresses for User Storage Areas	42
User-Supplied Addresses for Protected Control Blocks	43
Resource Identification	43
SVC Routines Calling SVC Routines	44
Control Program and User Data Accessibility	44
Control Program Extensions	45
Resource Access Control Facility (RACF)	45
RACDEF Macro Instruction	46
RACINIT Macro Instruction	46
RACHECK Macro Instruction	46
RACLIST Macro Instruction	46
Interruption, Recovery/Termination, and Dumping Services	47
SPIE Processing	47
Recovery/Termination	48
Invoking the Recovery/Termination Manager	48
CALLRTM	48
ABEND	50
Types of Recovery Routines	50
Functional Recovery Routines (FRRs)	50
Task Recovery Routines (STAE/STAI ESTAE/ESTAI)	50
Establishing Recovery Routines	50
Functional Recovery Routines	50
Task Recovery Routines	51
RTM/Recovery Routine Interface	54
Interface to Functional Recovery Routines	54
Interface to Task Recovery Routines	55
RTM/Retry Routine Interface	57
FRR-Requested Retry Routine	57
Task Recovery Retry Routines	57
Recovery Routine Guidelines	59
FRRs	60
Task Recovery	60
Clean-Up Routines	64
Support for Installation-Written Clean-Up Routines	64
Programming Considerations	65
Intercepting System Errors	65
Using the SLIP Command	66
SLIP Command Examples	66
Example 1: Obtaining a Dump with SQA Control Blocks	66
Example 2: Obtaining a Dump with Queue Elements and Control Blocks	67
Example 3: Replacing an SVC Dump with a Standalone Dump	67
Example 4: Setting SLIP Definitions for an Application Program	68
Example 5: Setting a SLIP Trap Using the Instruction Fetch PER Event	68
Example 6: Setting a SLIP Trap Using the Storage Alteration PER Event	68
Example 7: Setting a SLIP Trap Using the Storage Alteration PER Event and the ACTION=IGNORE Option	69
Example 8: Setting a SLIP Trap Using the Successful Branch PER Event	69
Dumping Virtual Storage	70
Using the SDUMP Macro Instruction	70
SQA Buffer	70

Using the CHNGDUMP Command	71
SNAP/ABEND User Exit	72
Correlating Diagnostic Material	73
Establishing a Timer Disabled Interrupt Exit	75
Using SETDIE	75
Recovery Considerations	76
DIE Characteristics	76
Task Queue Element Control	78
Virtual Storage Management	80
The BRANCH Parameter	81
The KEY Parameter	81
Real Storage Management	83
Fixing/Freeing Virtual Storage Contents	83
Completion Considerations	84
PGFIX/PGFREE Completion Considerations (5740-XE1)	84
PGFIXA/PGFREEA Completion Considerations (5740-XE1)	84
Virtual Subarea List (VSL)	84
Virtual Subarea List (VSL) (5740-XE1)	84.1
Reconfiguration Using Vary Storage	85
Multiprocessing Configuration Considerations	87
Miscellaneous Services	89
Writing Operator Messages	89
Routing the Message	89
Writing a Multiple-Line Message	90
Message Routing Exit Routines	90
Characteristics of MCS	90
Programming Conventions for WTO/WTOR Routines	91
Messages Not Using Routing Codes	92
Writing a WTO/WTOR Exit Routine	92
Adding a WTO/WTOR Exit Routine to the Control Program	92
Service Management	93
Scheduling Service Requests	94
Service Request Blocks (SRBs)	94
Priorities	95
Characteristics of Service Requests	95
Purging Service Requests	96
PURGEDQ Parameters	96
Asynchronous Exit Routines	97
Stage 1 Initialization	97
Stage 2 Scheduling	98
Stage 3 Execution	99
Execution and Termination Characteristics	99
Writing SVC Routines	100
Characteristics of SVC Routines	100
Programming Conventions for SVC Routines	100
Inserting SVC Routines Into the Control Program	103
Specifying SVC Routines	103
Inserting SVC Routines During the System Generation Process	104
Type 5 SVC Facility	104
Type 6 SVC Facility	105
Non-Preemptable SVCs	106
Subsystem SVC Screening	107
Missing Interruption Handler	108
Choosing a Time Interval	109
Modifying the CSECT IGFINTVL	109
Mass Storage System Missing Interrupts	109
Intercepting Hot I/O Interrupts	110
Adding Code to the Power Warning Feature Support	111
Writing Code for the Machine Check Handler Appendage	112
Coding Considerations	113
Writing Code for the Master Scheduler Initialization Routine	114
Coding Considerations	115
Control Track Record	116
Limiting User Region Size — IEALIMIT	117
Low Address Protection (5740-XE1)	119
Using the SRM Reporting Interface (5740-XE1)	120

Part II: Reference — Macro Instructions	121
Macro Instruction Forms	123
Coding the Macro Instructions	123
Continuation Lines	125
ATTACH — Create a New Task	126
CALLDISP — Force Dispatcher Entry	130
CALLRTM — Call Recovery/Termination Manager	131
CHANGKEY — Change Virtual Storage Protection Key	133
CIRB — Create Interruption Request Block	135
DEQ — Release a Serially Reusable Resource	138
DEQ (List Form)	141
DEQ (Execute Form)	142
DSGNL — Issue Direct Signal	143
ENQ — Request Control of a Serially Reusable Resource	145
ENQ (List Form)	149
ENQ (Execute Form)	150
ESTAE — Extended STAE	152
ESTAE (List Form)	154
ESTAE (Execute Form)	155
EVENTS — Wait for Events	157
EXTRACT — Extract TCB Information	158
EXTRACT (List Form)	160
EXTRACT (Execute Form)	161
FESTAE — Fast Extended STAE	162
FREEMAIN — Free Virtual Storage	164
FREEMAIN (List Form)	167
FREEMAIN (Execute Form)	168
GETMAIN — Allocate Virtual Storage	169
GETMAIN (List Form)	171
GETMAIN (Execute Form)	172
INTSECT — Intersect With the Dispatcher (5740-XE1)	172.1
MODESET — Change System Status	173
MODESET (List Form)	176
MODESET (Execute Form)	177
NIL — Provide a Lock Via an AND IMMEDIATE (NI) Instruction	178
OIL — Provide a Lock Via an OR IMMEDIATE (OI) Instruction	180
PGFIX — Fix Virtual Storage Contents	182
PGFIXA — Fast Path to Fix Virtual Storage Contents (5740-XE1)	184.1
PGFREE — Free Virtual Storage Contents	185
PGFREEA — Fast Path to Free Virtual Storage Contents (5740-XE1)	186.1
POST — Signal Event Completion	187
POST (List Form)	189
POST (Execute Form)	190
PROTPSA — Disable, Enable Low Address Protection (5740-XE1)	190.1
RGEDQ — Purge SRB Activity	191
PURGEDQ (List Form)	192
PURGEDQ (Execute Form)	193
QEDIT — Command Input Buffer Manipulation	194
RACDEF — Define a Data Set to RACF	195
RACDEF (List Form)	198
RACDEF (Execute Form)	199
RACHECK — Check RACF Authorization	200
RACHECK (List Form)	202
RACHECK (Execute Form)	203
RACINIT — Identify a RACF Defined User	204
RACINIT (List Form)	208
RACINIT (Execute Form)	210
RACLIST — Build In-Storage Profiles	211
RACLIST (List Form)	214
RACLIST (Execute Form)	215
RESERVE — Reserve a Device (Shared DASD)	216
RESERVE (List Form)	220
RESERVE (Execute Form)	221
RESUME — Resume Execution of a Suspended Request Block Task	222
RISGNL — Issue Remote Immediate Signal	224
RPSGNL — Issue Remote Pendable Signal	226
SCHEDULE — Schedule System Services for Asynchronous Execution	228
SDUMP — Dump Virtual Storage	229
SDUMP (List Form)	236
SDUMP (Execute Form)	237
SETFRR — Set Up Functional Recovery Routines	240

SETLOCK – Control Access to Serially Reusable Resources	243
SETRP – Set Return Parameters	250
SPIE – Specify Program Interruption Exit	254
SPIE (List Form)	256
SPIE (Execute Form)	257
SPOST – Synchronize POST	258
STAE – Specify Task Asynchronous Exit	259
STAE (List Form)	262
STAE (Execute Form)	263
STATUS – Change Subtask Status	264
SUSPEND – Suspend Execution of a Request Block Task	267
SYNCH – Take a Synchronous Exit to a Processing Program	268
SYSEVENT – Notify Transaction Complete (5740-XE1)	269.0
TCTL – Transfer Control from an SRB Process	270
TESTAUTH – Test Authorization of Caller	271
T6EXIT – Type 6 Exit	273
WTO – Write to Operator	274
WTO (List Form)	277
WTOR – Write to Operator with Reply	278
WTOR (List Form)	280
Index	281

Figures

Figure 1.	EXTRACT ECB and CIB Pointers	5
Figure 2.	Command Input Buffer Contents	5
Figure 3.	EXTRACT Answer Area Fields	6
Figure 4.	Summary of Locking Characteristics	13
Figure 5.	Valid Volume Characteristic and Device Status Combinations	17
Figure 6.	Example of an Interlock Environment	19
Figure 7.	Example of Subroutine Issuing RESERVE and DEQ	22
Figure 8.	Assigning Authorization via SETCODE	26
Figure 9.	Authorization Rules	27
Figure 10.	Bypassing the POST Routine	32
Figure 11.	ECB Extension (ECBE)	33
Figure 12.	Extended ECB	34
Figure 13.	Data Areas POST Exit for Example	35
Figure 14.	POST Function and Branch Entry Points	36
Figure 15.	POST Branch Entry Input	37
Figure 16.	POST Branch Entry Output	37
Figure 17.	ESTAE Environment	54
Figure 18.	Virtual Subarea List Entries	85
Figure 19.	IPL Designation of Processor Storage Units	86
Figure 20.	Asynchronous Exit Data Area Configuration	98
Figure 21.	Programming Conventions for SVC Routines	101
Figure 22.	SRB Input Format	106
Figure 23.	Changing the Missing Interruption Handler Time Interval	109
Figure 24.	Logical Placement of Your Code in the Machine Check Handler Appendage	113
Figure 25.	Logical Placement of Your Code in the Master Scheduler Initialization Module	114
Figure 26.	Control Track Record	116
Figure 27.	Storage Assignments on MP Systems	117
Figure 28.	The Effects of IEALIMIT and REGION Values on Various GETMAINS	119
Figure 29.	Macro Instruction Coverage	122
Figure 30.	Sample Macro Instruction	124
Figure 31.	Continuation Coding	125
Figure 32.	Return Code Area Used by DEQ	139
Figure 33.	Return Code Area Used by ENQ	147
Figure 34.	Return Code Area Used by RESERVE	218
Figure 35.	MCSFLAG Fields	276

**Summary of Amendments
for GC28-0628-3
by TNL GN28-4978
OS/VS2 Release 3.8**

The following changes have been made to support Resource Access Facility (RACF) Version 1 Release 4.

- The keywords UNIT, SPECIAL, OWNER, LEVEL, UACC, DATA, and AUDIT have been added to the RACDEF macro instruction.
- The DSTYPE and TYPE keywords of the RACDEF macro instruction have been extended to include DSTYPE=M and TYPE=CHGVOL.

- The keyword OWNER has been added to the RACHECK macro instruction.
- The DSTYPE keyword of the RACHECK macro instruction has been extended to include DSTYPE=M.

**Summary of Amendments
for GC28-0628-3
by TNL GN28-4738
OS/VS2 Release 3.8**

This technical newsletter updates the Resource Access Control Facility (RACF) macro instructions.

**Summary of Amendments
for GC28-0628-3
by TNL GN28-4740
OS/VS2 Release 3.8**

This technical newsletter adds the topic "Getting More Than One SYSMDUMP" to the Interruption, Recovery/Termination, and Dumping Services section in Part I.

Part I: Supervisor Services

The supervisor provides the resources that your programs need while assuring that as many of these resources as possible are being used at a given time. Well designed programs use system resources efficiently. Knowing the conventions and characteristics of the VS supervisor helps you design more efficient programs.

This section describes those supervisor services that should be restricted in use to systems programmers and installation-approved personnel. In most cases, the services correspond to macro instructions and parameters that are described in part II.

For convenience and compatibility, the services you can request from the supervisor are grouped in the same manner as in **OS/VS2 Supervisor Services and Macro Instructions**. The service groupings are as follows:

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.

Program Management: You can use the supervisor to aid communication between segments of a program. Save areas, addressability, and passage of control from one segment of a program to another are included in this topic.

Resource Control: Portions of some tasks depend on the completion of events in other tasks, thus requiring planned task synchronization. Planning is also required when more than one program uses a serially reusable resource.

Interruption, Termination, and Dumping Services: 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.

Virtual Storage Management: While virtual storage allows you to write large programs without the need for complex overlay structures, virtual storage must be obtained for your job step. Virtual storage is allocated by both explicit and implicit requests.

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.

In addition to the services outlined above, the supervisor provides the facilities for timing events, extended precision floating-point simulation, and operator communication with both the system and application programs.

Subtask Creation and Control

One task is created by the control program as a result of initiating execution of the job step (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 tasks are competing for control than the task in the job you are concerned with. When a wait condition occurs in one of your tasks, it is not necessarily a task from some other job that gets control; it may be one of your tasks, a portion of your job.

The general rule is that parallel execution of a job step (that is, more than one task in a job step) should be chosen only when a significant amount of overlap between two or more tasks can be achieved. The amount of time taken by the control program in establishing and controlling 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 **OS/VS2 Supervisor Services and Macro Instructions**. This chapter continues discussion in the following areas:

- Task creation (ATTACH macro instruction)
- Communication with a problem program (EXTRACT and QEDIT macro instructions)

Creating a New Task

The ATTACH macro instruction causes the control program to create a new task. The complete use of the macro instruction is described in **OS/VS2 Supervisor Services and Macro Instructions**.

The macro instruction has nine parameters that permit the authorized user (protection key 0-7 or supervisor state) greater flexibility in using the services of the macro instruction. If the parameters are not used by authorized tasks, the default values are 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 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 address of 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=YES -- reset the APF authorization of the step according to the attributes of the subtask.

User Modifications

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 first load of a job is attached as a job step task by the initiator. For such an attach, the program manager will not search the job library of the attaching task. When the job step task issues ATTACH, LOAD, LINK, or XCTL, the job library of the job step task may be searched for the load module being fetched.
Also, only under a job step task can a system program (system key or supervisor state) attach a load module from a nonsystem 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 before privileged instructions (for example, LPSW) can be executed. You can specify supervisor mode via this parameter or via the MODESET macro instruction.
- **SVAREA=NO** -- a save area is not needed for the new task.
The save area is obtained from the user's region. Since it may not always be desirable to have a save area (for example, the user's region may not be defined at the time of a system ATTACH), this parameter may be used to specify that no save area should 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, this allows the originating task to alter the new TCB. The new task will remain nondispatchable until TCBANDSP has been reset via the STATUS macro instruction with the RESET option. (*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.
- **RSAPP=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 "Resource Control" section.

Operator Communication with a Problem Program

The operator can pass information to a problem 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.

An EXTRACT macro instruction is issued 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 1.

```
EXTRACT answer area, FIELDS=COMM
```




Figure 1. EXTRACT ECB and CIB Pointers

The CIB contains the information specified on the STOP or the MODIFY command, as shown in Figure 2. 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.

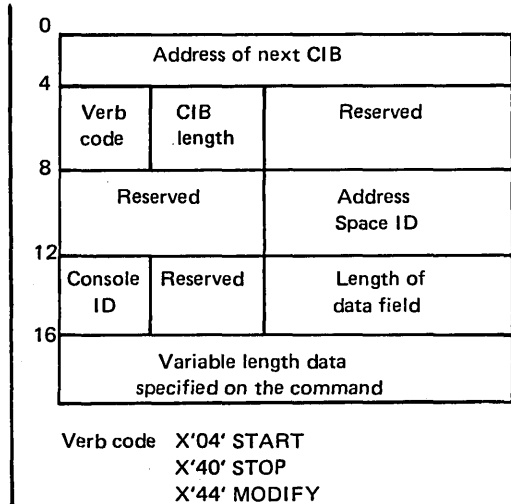


Figure 2. Command Input Buffer Contents

If the address of the START CIB is present, the QEDIT macro instruction should be used 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 1.
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. STOP commands, however, 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, the user must establish addressability via IEZCOM based on the address returned by the EXTRACT.

For the duration of the job, the communications ECB may be waited on or checked at any time to see if a command has been entered for the program. The verb code in the CIB should be examined to determine whether a STOP or a MODIFY command has been entered. After the data in the CIB has been processed, a QEDIT macro instruction should be issued to free the CIB.

The communications ECB is cleared 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.

Providing an EXTRACT Answer Area

The EXTRACT macro instruction is used to provide TCB information for either the active task or one of its subtasks. The information from the requested field is returned in the relative order shown in Figure 3. If the information from a field is not requested, the associated fullword is omitted.

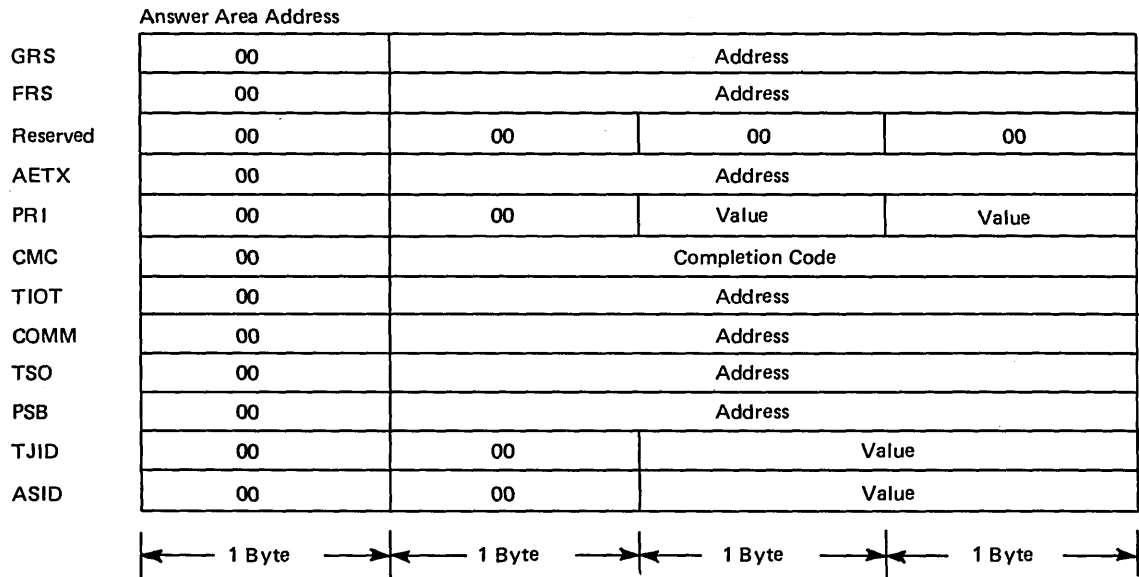


Figure 3. EXTRACT Answer Area Fields

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.

Addresses are always returned in the low-order three bytes of the fullword, and the high-order byte is set to zero. Fields for which no addresses or values are specified in the task control block are set to zero.

For example, if **FIELDS=(TIOT,GRS,PRI,TSO,PSB,TJID)** is coded, a 6-fullword answer area is required, and the extracted information appears in the same relative order as shown in Figure 3. (That is, **GRS** is returned in the first word, **PRI** in the second word, **TIOT** in the third word, etc.)

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

Program Management

The supervisor can be used to aid communication between segments of a program. The descriptions of load module structures, facilities for passing control between programs, and the use of associated macro instructions are available to all users and are described in *OS/VS2 Supervisor Services and Macro Instructions*.

Following is a description of synchronous exits and the SYNCH macro instruction associated with it. The facility should be used only by system programmers or other installation-approved personnel.

Synchronous Exits

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 supervisor state) where the processing program is expected to return control to the supervisor state control program. The program to which control is given must be in virtual storage when the macro instruction is issued. 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 use of the SYNCH macro instruction is similar to that of the BALR instruction in that register 15 may be used for the entry name address.

Registers are not saved or restored by SYNCH processing when control is returned to the caller unless RESTORE=YES is specified. IF RESTORE=NO is coded or defaulted to, 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, to preserve system integrity, the authorized program must either save the registers in a protected save area and then restore them, or validate the contents of the returned registers.

As an example of the use of the SYNCH macro instruction, label processing as the result of an OPEN macro instruction may be carried out to a point at which a user's processing program indicates that private processing is desired (or necessary). 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.

Multiprocessing Programming Considerations

Many of the following programming considerations are true in a multi-tasking environment (such as OS/MVT) as well as in a multiprocessing system. However, because of their increased importance in a multiprocessing environment, they should be closely reviewed.

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.

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. This is important since in a multiprocessing system:

- 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 the other processor.
- Recovery routines can execute on either 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-entrancy of macros is to use the LIST and EXECUTE forms of the macro instructions with a dynamically acquired parameter list.

Priority

Programs that use priority or precedence as a serialization mechanism are sensitive to changes in the dispatching algorithms used by the system. For example, the CHAP macro instruction does not ensure that tasks are dispatched in the expected order, due to dispatching on two processors. Also, the PRIORITY and DPRTY JCL parameters can no longer be used to accomplish serialization. First, the system resources manager might allow a task or job with a lower dispatching priority to execute before to a task with a higher priority. Second, since tasks are executed on both processors, tasks of different priority might be executed on both processors simultaneously.

Note: Tasks might not be redispached on the processor on which they were previously executing. Therefore, storage from 0-4K should not be used because redispach on a different processor results in different data being referenced.

For further multiprocessing considerations see the section "Locking".

Resource Control

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 time span. 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)
- Authorized program facility (TESTAUTH macro instruction)
- System status (MODESET macro instruction)
- Interprocessor communications (DSGNL, RISGNL, and RPSGNL macro instructions)
- Event completion (POST, SPOST, and EVENTS macro instructions)

Resource Serialization

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 period of time. 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.

Adequate control of resources is even more significant in a multiprocessing environment. Some uniprocessor serialization techniques (such as disablement) are not effective in a multiprocessing environment because of the possibility of multiple tasks using the same resource while running concurrently on different processors.

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 serialization is explicit; i.e. using one of the following methods: ENQ/DEQ; WAIT/POST/EVENTS; RB precedence at the TCB level; locking at the TCB level; CS (compare and swap instruction); and TS (test and set instruction). Such forms of serialization are required in the following cases:

- Scanning of the CIB chain. Manipulation of the CIB chain should be done by means of the QEDIT macro instruction.
- The use of data in subpools shared between tasks.
- Data referenced by more than one task. (For example, attached tasks can execute at the same time as the attaching task on different processors.)
- References to system control block fields that dynamically change after IPL. The serialization technique in this case must match that used by the system. Also, bits within a byte all require the same serialization technique, if the byte might be updated simultaneously by both processors.
- The access 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.
- Any common data references 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, make use of 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

Two categories of locks exist:

- 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 since it relates to the entire system. Also, a system-related GETMAIN for a global subpool, or a global ENQ, 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 queues and control blocks associated with that address space can be manipulated only by the owner of the lock. (For example, an address space-related GETMAIN for a user subpool requires a local lock.)

All of the locks described below, with the exception of the local lock, 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 lock, on the other hand, does not serialize across address spaces, but serializes functions executing within the address space.

Types of Locks

Two types of locks exist. The type determines what happens when a function on one processor makes an unconditional request for a lock that is held on the other processor. The types are:

- Spin locks -- prevent the requesting function on one processor from doing any work until the lock is freed on the other processor. The lock manager enters a loop that keeps testing the lock until it is released on the other processor. As soon as the lock is free, the lock manager spinning on the other processor can obtain the lock for the requesting function. As long as a spin 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.
- Suspend locks -- prevent the requesting program from doing work until the lock is available, but allow the processor to continue doing other work. The request is queued and other work may be dispatched on that processor. Upon release of the lock, the highest priority queued requestor is given control of the lock.

All of the locks described below, with the exception of the local and cross memory services (CMS) lock, are spin locks. The LOCAL and CMS locks run enabled, and are suspend locks; that is: they 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 lock contains the ID of the processor on which its owner is dispatched. The CMS lock contains the ASCB address under which the owner of the lock was executing. Special IDs are placed in the local lockword whenever the owner of either the CMS lock or local lock is not currently executing on a processor due to either an interruption or suspension.

The CMS lock was provided as an enabled global lock for the following reasons:

- Since disabled page faults are not allowed in the system, some global functions could use a lock which did 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 locks were left as disabled spin locks because the functions which run under the locks are of short duration. 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) are needed to 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

Two classes of locks exist:

- Single locks -- Only one lock exists at a given level of the locking hierarchy.
- Multiple locks (commonly referred to as class locks) -- More than one lock exists at a given level of the locking hierarchy.

Figure 4 summarizes the characteristics of the locks.

lock	global	local	spin	suspend	single	multiple
DISP	X		X		X	
ASM	X		X			X
SALLO	X		X		X	
IOSYNCH	X		X		X	
IOSCAT	X		X		X	
IOSUCB	X		X			X
IOSLCH	X		X			X
SRM	X		X		X	
CMS	X			X	X	
LOCAL		X		X		X

Figure 4. Summary of Locking Characteristics

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 DISP lock and requesting the SRM lock.
- Function B holding the SRM lock and requesting the DISP lock.

Deadlocks are prevented by requiring that a function on a processor can request only those locks that are higher in the hierarchy than the locks it currently holds. The hierarchy is shown in Figure 4, with DISP being the highest lock.

As noted above, some locks are single system locks (for example, DISP), and some locks are multiple locks in which there is more than one lock within the lock type (for example, IOSUCB).

For those global lock types that have more than one lock, only one lock of each type may be held. For example, if an IOSUCB lock is held on a CPU, a different IOSUCB lock may not be requested on the CPU.

The local lock must be held by the caller when requesting the CMS lock. Also the local lock cannot be released while holding a CMS lock. It is not necessary to hold all locks in the hierarchy up to the highest lock needed. The only locks that should be held are those you need.

Locking Considerations

In MVS, the local lock is provided to reduce disablement in the control program with a form of logical disabling that works across both processors. Locking, however, is not completely equivalent to disablement in VS2 Release 1. The significant differences are:

- A locked routine is not allowed to issue an SVC or to invoke a routine that would issue an SVC on the routine's behalf.
- In VS2 Release 1, all user-written functions that disable using the MODESET macro instruction, the SSM (set system mask) instruction, or some other method, should be changed to use the locking function. The SYSMASK and ENABLE parameters of MODESET are no longer supported. The SSM instruction causes a program check.
- In VS2 Release 1, disablement prevented storage from being freed or paged out while the storage was being referenced. In MVS, it is necessary to hold the local lock to prevent a FREEMAIN from being executed on the other processor even though a global spin lock is held on one processor. To prevent page stealing, it is necessary to either fix the pages or hold the SALLOC lock.

The locks provided in MVS are:

- DISP (Global dispatcher lock) -- serializes all functions associated with the dispatching process.
- ASM (auxiliary storage manager lock) -- serializes the use of global ASM control blocks.
- SALLOC (real storage manager and virtual storage manager space allocation lock) -- serializes the global functions of real storage management and virtual storage management.
- IOSYNCH (IOS synchronization lock) -- serializes global IOS functions.
- IOSCAT (IOS channel availability table lock) -- serializes access and updates to the channel availability table.
- IOSUCB (IOS unit control block lock) -- serializes access and updates to the unit control blocks. There is one lock for each UCB.
- IOSLCH (IOS logical channel queue lock) -- serializes access and updates to the IOS logical channel queues. There is one lock for each channel queue.
- SRM (system resources manager lock) -- serializes use of the SRM control algorithms and associated data.
- Cross memory services lock (CMS) -- serializes on more than one address space where this serialization is not provided by one or more of the other global locks.
- 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.

SETLOCK Macro Instruction

The SETLOCK macro instruction is used to obtain, release, or test a specified lock or set of locks (using the OBTAIN, RELEASE, and TEST parameters). To use SETLOCK, you must be executing in supervisor state with protection key 0. Users of SETLOCK may also be

executing in SRB mode, as an extension of the interrupt handlers, or as a system service such as the MVS dispatcher.

Disabled/Enabled State for OBTAIN: When a global spin lock is successfully obtained, control is returned to the SETLOCK caller with the processor in a physically disabled state (except for machine check interruptions).

When a LOCAL lock or CMS lock is successfully obtained, control is returned to the SETLOCK caller with the processor in the state that existed before to the SETLOCK request. This is also true for unsuccessful attempts to obtain any of the locks where control is returned to the caller.

If a disabled caller unconditionally requests the LOCAL or CMS lock and the lock is not immediately available, the caller is abnormally terminated.

Disabled/Enabled State for RELEASE: When a global spin lock is successfully released, control is returned to the caller with the processor in a physically enabled state when no other global spin locks are held by that processor; control is returned in a physically disabled state when other global spin locks are held. Control is also returned in the disabled state when any disabled supervisor indicators are on when the RELEASE request is made or when the DISABLED parameter was specified.

When the LOCAL lock or CMS lock is released, control is returned with the processor in the state that existed before to the RELEASE request.

For all locks, if the RELEASE operation is unsuccessful and no ABEND condition exists, control is returned with the processor in the physically enabled or disabled state that existed before to the RELEASE request.

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

The Must Complete Function

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 system routines must complete their operations on the resource. Otherwise, the resource may be left incomplete or may 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 can place other tasks in a wait 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.

As an example, in multitask environments the integrity of the volume table of contents (VTOC) must be preserved during updating so that all future users may have access to the latest, correct version of the VTOC. ENQ on the VTOC and use the must complete function (to suspend processing of other tasks) when updating a VTOC.

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 placed in a wait state.

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 which owns a must-complete resource results in the abnormal termination of the entire job step. The programmer receives a message stating that the failure occurred in the step must complete status.

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 may occur if a routine issues an ENQ macro instruction while in the must complete mode. The requested resource may already be queued on by a task made nondispatchable due to the must complete request already made. Since the resource cannot be released, job step tasks are abnormally terminated, and system tasks are retried.
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 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=STEPItem 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 (Item b). Issue a DEQ macro instruction to terminate the must complete mode as soon as processing is finished.
6. The STATUS macro instruction changes the dispatchability status of tasks for users in supervisor state, with a protection key of 0-7, or APF-authorized. STATUS can also change the must complete status of a task. In all cases, the task remains enqueued.

Tasks set nondispatchable by the corresponding ENQ macro instruction are made dispatchable and asynchronous exits from the requesting task are enabled.

Shared Direct Access Storage Devices (Shared DASD)

The Shared DASD facility allows computing 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 may be desirable to achieve better performance.

Devices that Can be Shared

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

- IBM 2314 Direct Access Storage Facility equipped with the two-channel switch -- IBM 2314 Disk Storage Module.
- IBM 2314 Direct Access Storage Facility combined with the IBM 2844 Auxiliary Storage Control -- IBM Disk Storage Module. Device reservation and release are supported by this combination with or without the presence of the two-channel switch. Two channels -- one from System A and one from System B -- may be connected to the combination. In addition, the two-channel switch may be installed in either or both of the control units, thus permitting as many as four systems to share the devices.
- 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.

Alternate channels to a device from any one system may only be specified for the IBM 2314 Direct Access Storage Facility, or the IBM 3330, 3340/3344, 3350 Series Storage Unit.

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 5 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
Offline	Removable, reserved, or permanently resident

Figure 5. Valid Volume Characteristic and Device Status Combinations

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

PASSWORD	SYS1.NUCLEUS
SYS1.LOGREC	SYS1.SVCLIB
SYS1.LPALIB	VSAM page spaces
SYS1.MANX	Master catalog
SYS1.MANY	

Volume Handling

Volume handling on the Shared DASD option must be clearly defined since 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 effected on all systems before a device may 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 sharable devices offline prior to issuing START commands and then following parallel mount procedures.

Macro Instructions Used with Shared DASD

The RESERVE macro instruction is used to reserve a device for use by a particular system; it must be issued by each task needing device reservation. The EXTRACT macro instruction is used to obtain the address of the task input/output table (TIOT) from which the UCB address can be obtained. The topic "Finding the UCB Address" explains this procedure.

Notes on RESERVE: The Set-Must-Complete (SMC) parameter available with the ENQ macro instruction may also be used with RESERVE.

If a 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.

Releasing Devices

The DEQ macro instruction is used in conjunction with RESERVE just as it is used with ENQ. It must describe the same resource 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 which 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 occur on the system.

Volume Assignment

Since 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. For example, as shown in Figure 6, data sets A and B reside on device C, and data sets D and E reside on device F. Task X in system X reserves device C in order to use data set A; task Y in system Y tries to reserve device F in order to use data set D. Now task X in system X tries to reserve device F in order to use data set E and task Y in system Y tries to reserve device C 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 would 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.

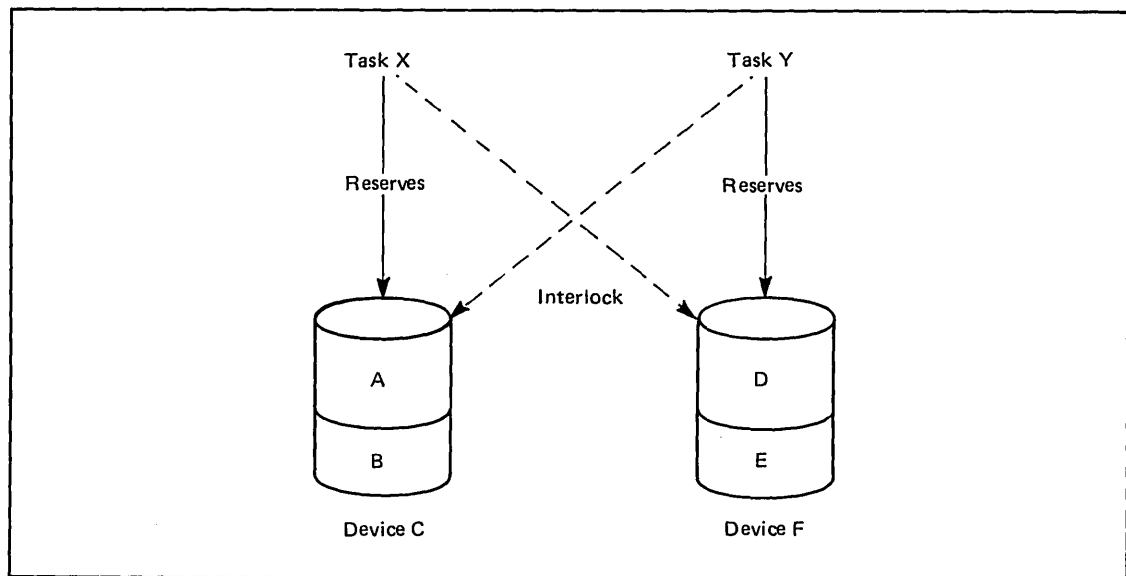


Figure 6. Example of an Interlock Environment

Program Libraries

When assigning program libraries to shared volumes, precaution must be taken to avoid interlock. For example, LINKLIB for system A resides on volume X, while LINKLIB for system B resides on volume Y. Task A in system A invokes a direct access device space management function for volume Y, resulting in that device being reserved. Task B in system B invokes a similar function for volume X, reserving that device. However, each load module transfers to another load module via XCTL. Since the LINKLIB for each system resides on a volume reserved by the other system, the XCTL macro instruction cannot complete the operation. An interlock occurs; since no access to LINKLIB is possible, both systems will eventually enter an enabled wait state.

Finding the UCB Address

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

Providing the Unit Control Block Address to RESERVE: The EXTRACT macro instruction is used 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. Prior to issuing an EXTRACT macro instruction, the user sets up an answer area in main storage which is to receive the requested information. One full word is required for each item to be provided by the control program. If the user wishes to obtain the TIOT address, he must specify `FIELDS=TIOT` in the EXTRACT macro instruction.

The address of the TIOT is then returned by the control program, right adjusted, in the full word answer area.

The TIOT is constructed by job management routines and 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, the user must locate the DD entry in the TIOT corresponding to the DD name of the data set for which he intends to issue the RESERVE macro instruction.

The UCB address can also be obtained via the data extent block (DEB) and the data control block (DCB). The DCB is the block within which data pertinent to the current use of the data set is stored. The address of the DEB is contained at offset 44 decimal after the DCB has been opened. 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 concerning the physical characteristics of the data set and other information that is used by the control program. 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, the user must locate the extent entry in the DEB for which he intends to issue the RESERVE macro instruction. (In disk addresses of the form MBBCCHHR, the M indicates the extent number starting with 0).

Procedures for Finding the UCB Address of a Reserved 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.

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 the user may 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 MBBCCCHR.
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 (MBBCCCHR) 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.

RES and DEQ Subroutines: The assembler language subroutine in Figure 7 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 DDCARD 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=pointer address as a parameter. If the DEQ macro is to be issued, a fullword of binary zeros must be placed in the DDNAME field before control is passed.

```

RESDEQ      CSECT
            SAVE    (14,12),T      SAVE REGISTERS
            BALR    2,0              SET UP ADDRESSABILITY
            USING   *,2
            ST      13,SAVE+4
            LA      11,SAVE          ADDRESS OF MY SAVE AREA IS
                                      STORED IN THIRD WORD OF CALLER'S
            ST      11,8(13)        SAVE AREA
            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 LENGTH
            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 RES/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
            BCR    15,14
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 7. Example of Subroutine Issuing RESERVE and DEQ

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 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 after 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 or selectable unit (SU) 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 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 sometimes distinguishes authorization due to a program being in either supervisor state or system key from authorization through APF. 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 can not use these keywords.

When MVS 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 the discussion under "Restricting Load Module Access" for exceptions to this rule.)

You should recognize the distinction between an authorized program and an authorized user. To use restricted functions in MVS, 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 must use existing data set security facilities to place the program in a library protected by RACF or by a password. If the program is an authorized program, the library must be an authorized library.

Authorized Libraries

APF-authorized programs must reside in authorized libraries. The authorized libraries in MVS are:

- SYS1.LINKLIB
- SYS1.SVCLIB
- SYS1.LPALIB (only during an IPL; see note 2 under "Notes")
- Installation authorized libraries

To allow an installation to authorize libraries, MVS 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.

MVS also considers as authorized any library contained in the LNKLSTxx members of SYS1.PARMLIB (which denotes libraries to be concatenated to SYS1.LINKLIB) when a program in the library is located by a search through the link list. However, if these libraries are accessed through either JOBLIB or STEPLIB DD statements, MVS does not consider them authorized unless the installation has also placed the names of the libraries in IEAAPFxx. As long as a load module 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 on SYS1.PARMLIB members IEAAPFxx and LNKLSTxx, see the **Initialization and Tuning Guide**.)

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.

Using APF

APF allows an MVS 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

MVS 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 IEAAP00 member of SYS1.PARMLIB. (See the description of this member in the *Initialization and Tuning Guide*.)

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 306 abend results. The authorized program is responsible

for recovering from the 306 abend in a way that does not allow the accessed module to execute.

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 8, 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          DSNAME=SYS1.LINKLIB,DISP=OLD
//SYSLIN        DD          DSNAME=SYS1.LOADSET,DISP=(OLD,PASS),
//              DD          UNIT=SYSDA
//              DD          *
SETCODE          AC(1)
NAME             MOD1(R)
/*
```

Figure 8. 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 9 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 9. Authorization Rules

The rules shown in Figure 9 are also true when the ATTACH macro instruction is used unless the RSAPF keyword is specified. A routine that specifies RSAPF=YES and is running in supervisor state or system key can attach programs residing outside APF-authorized libraries if:

- the caller is not running APF-authorized
- the caller is attaching a subtask in problem state
- the attached task's TCB key is 8 or more

The newly attached subtask does not run APF-authorized.

However, if the subtask comes from an APF-authorized library and is link edited with the APF-authorized attribute, then the step 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

Installation 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.
- Ensures 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 mixup 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 exists. 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.

Changing System Status

The MODESET macro instruction alters selective fields of the program status word (PSW). The standard form of MODESET may be coded in two separate ways: one form generates an SVC and the other form generates inline code.

Generating an SVC

This form of MODESET, executable 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 turned on or off. When bit 15 is one, the processor is in the problem state. 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, executable only in supervisor state, 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 may be coded to indicate the key to be set in the current PSW.

The following keys may be set:

- 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.

Interprocessor Communications (IPC)

Interprocessor communications is a function in a multiprocessing system that provides communication between processors sharing the same control program. Those executing functions which require a processor or program action on one or more processors will 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.

Service Classes

The SIGP instruction provides twelve distinct hardware functions to support two classes of IPC services - *direct* and *remote*:

Direct class - These services are defined for those control program functions that require the modification or sensing of the physical state of one of the configured processors. Ten of the twelve SIGP hardware functions are defined as IPC direct services, and are accessible 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.

Initial Program Reset: The specified processor performs initial program reset. The execution of the reset does not affect other processors and does not affect channels not configured to the processor being reset. The reset operation is not necessarily completed during the execution of the SIGP instruction.

Program Reset: The specified processor performs program reset. The execution of the reset does not affect other processors and does not affect channels not configured to the processor being reset. The reset operation is not necessarily completed 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.

Initial Microprogram Load: The specified processor performs initial program reset and then initiates the initial-microprogram-load function. The latter function is the same as that which is performed as part of manual initial micro program loading. The operation is not necessarily completed during the execution of the SIGP instruction.

Initial Processor 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 configured to the specified processor, to be reset. The reset operation is not necessarily completed during the execution of the SIGP instruction.

Processor 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.

Remote class - These services are defined for those control program functions which require the execution of a software function on one of the configured processors. The two remaining SIGP functions are defined as remote services:

External Call: An "external call" 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. Only one external-call condition can be kept pending in a processor at a time. The external-call function is accessible via the RPSGNL macro instruction.

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. The emergency-signal function is accessible via the RISGNL macro instruction.

Status Conditions

Eight status conditions are defined whereby the issuing and specified processors can indicate their response to the designated hardware function. The status conditions are contained in register 0 and are:

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.

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 completed. The condition may be due to 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. The condition may 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.

Not Ready: This condition exists when the specified processor uses reloadable control storage to perform a function and the required microprogram is not loaded. The not-ready condition may be indicated in response to all functions except IMPL.

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.

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, since 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.

For more details on the SIGP instruction, see *System Programming Library: MVS Diagnostic Techniques*.

Event Completion

Cross Memory POST

The POST macro instruction is used to signify the completion of an event by one routine to another routine. Usually, the completion of the event is posted in the user's address space.

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 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 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 which 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 may or may not be posted, depending on existing conditions. See *OS/VS2 System Logic Library* for details.) Since SPOST invokes the PURGEDQ SVC, see the description of PURGEDQ for the restrictions on its use.

Bypassing the POST Routine

The problem programmer may bypass the POST routine whenever the corresponding WAIT has not yet been issued. To do this, he must issue a TEST UNDER MASK (TM) instruction to determine if the wait bit in the ECB is on. If the wait bit is on, the normal POST routine must be executed. If the wait bit is not on, a COMPARE AND SWAP (CS) instruction is issued. The compare operand should reflect the ECB content with the wait bit off, the swap operand should have the post bit on and contain the desired post completion code. If the wait bit has been set on after the ECB has been tested with the TM instruction, the CS will fail

(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.

Figure 10 demonstrates an example of this use of POST.

	L	RX, ECB	Get contents of ECB.
	L	RY, =X'40000000'	Completion bit and code to be compared and swapped
*			
CSLOOP	TM	ECB, X'80'	Wait bit on?
	BO	DOPOST	If yes, then execute post.
	N	RX, =X'7FFFFFFF'	Compare operand has wait bit off.
	CS	RX, RY, ECB	Compare and swap completion bit and code.
*			
	BE	POSTDONE	Branch if CS is successful.
DOPOST	POST	ECB	
POSTDONE	EQU	*	

Figure 10. Bypassing the POST Routine

Waiting for Event Completion

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 is accomplished by specifying **BRANCH=YES** on the EVENTS macro instruction. If this parameter is used, the branch entry routine performs all normal WAIT processing and ECB initialization. **BRANCH=YES** may be specified in conjunction with either **WAIT=YES**, **WAIT=NO**, or **ECB=**.

- If **WAIT=YES** is specified, control will later be returned to the dispatcher, even though there may be ECBs posted to the EVENTS table. EVENTS frees the local lock. Before the EVENTS macro instruction with the **WAIT=YES** option is specified, the caller is responsible for establishing 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 **WAIT=NO** is not specified, control is later returned to the caller. 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. The 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".

The process that establishes a POST exit is responsible for deleting that exit prior to 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 prior to the initialization of the extended ECB. This sequence avoids the possibility of an initialized extended ECB being posted prior to the initialization of the ECB extension.

The ECB extension is two words long, begins on a word boundary, and can be from any subpool. 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 is of the format shown in Figure 11.




Figure 11. ECB Extension (ECBE)

where:

- VALUE - is one byte containing a value from 1-255. A value of 1 indicates that the POST exit function is being requested. All other function codes are reserved.
- POST DATA - When VALUE is 1 (that is, contains X'01') this field contains the address of the exit routine to be given control when the POST occurs.

The extended ECB must conform to current requirements for ECBs and should be initialized to the format shown in Figure 12 via compare and swap while holding the local lock.

If the compare and swap fails and if the ECB is pre-posted, the user should perform the appropriate POST exit functions in order to replace those ordinarily performed by the already concluded POST processing.




Figure 12. Extended ECB

ECB initialization must use both compare and swap and the local lock to ensure proper serialization so that posts are not lost. The reason for this is that some system routines initialize ECBs with compare and swap while other system routines use the local lock for serialization. Therefore, both serialization techniques must be used in interfaces with general system services.

The ECB must be initialized only after it is eligible for posting.

POST Interface with Exit Routines

Prior to giving the exit control, POST checks to ensure that the user's exit routine address identified in the ECB extension denotes a valid POST exit routine. Even though POST thereby makes sure that a valid system exit receives control, the exit routine must ensure that an unauthorized routine has not counterfeited the extended ECB/ECB extension pair.

The user's exit routine receives control from POST locally locked in supervisor state, key zero. The routine must not release the local lock and should be able to process in both SRB and TCB mode. The register contents at entry to the user exit routine are:

Register	Contents
0	Address of the currently posted ECB
1	Address of related ECB extension
14	Return address
15	Exit routine entry point address
2-13	Unpredictable

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, key zero.

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 may issue POST only via the POST entry point, IEA0PT03. Details of the interface are in "Branch Entry to the POST Macro Instruction".

Because of the save area recursion within POST, a POST exit routine may not post another extended ECB unless it does so by specifying a cross address space post. Any attempt to activate another POST exit prior to the completion of the current exit results in a 702 abend. If you must post another extended ECB from a POST exit routine, you should either have your routine issue a cross address space 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 13. Once initialized, the subsystem dispatcher waits on a list of ECBs. Each list entry identifies an ECB in an EQT.

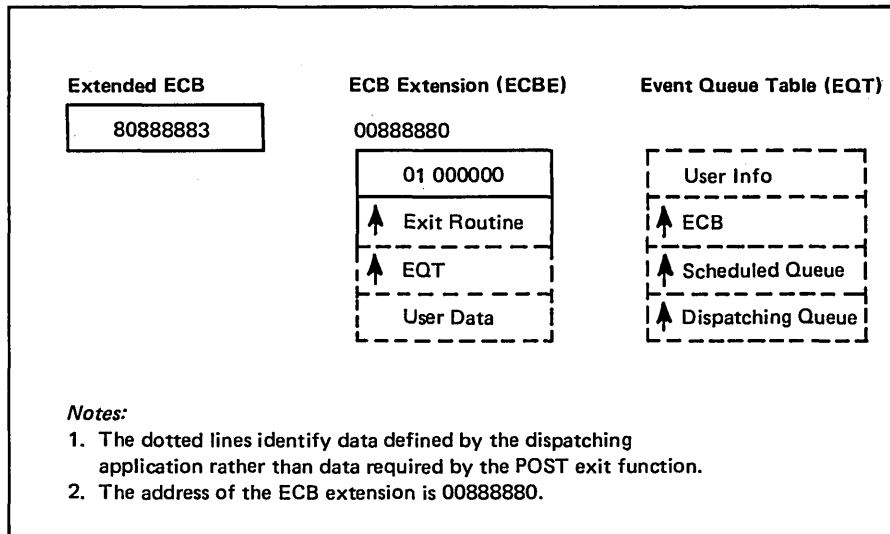


Figure 13. Data Areas Post Exit Example

As soon as any system routine posts an extended ECB, the subsystem exit routine identified in the ECBE receives control. 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 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 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 Macro Instruction

Branch entry to the POST macro instruction provides all the normal ECB and RB POST processing. In general, the caller must hold the local lock and be in supervisor state, key zero. Upon completion of the POST process, control returns to the caller in supervisor state, key zero with the local lock.

Figures 14, 15, and 16 respectively illustrate the POST function and the branch entry points through which those functions can be performed, the input parameters to POST, and the output parameters from POST.

Functions		Entry Points			
	IEAOPT01	IEAOPT02	IEAOPT03*	IEAOPT0E	
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 IEAOPT01. 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 IEAOPT0E.
 **The local lock does not need to be held for a cross address space POST at this entry point.

Figure 14. POST Function and Branch Entry Points

Registers	IEAOPT01 (CVTOPT01)	IEAOPT02 (CVTOPT02)	IEAOPT03 (CVTOPT03)	IEAOPT0E (CVTOPT0E)
0	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 ⁴	ECB Address
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
<p>¹ If cross address space post, optionally contains the storage protection key of the ECB. 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 byte of register is zero; if cross address space POST, set high-order byte of register to X'80'. 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.</p>				

Figure 15. POST Branch Entry Input

Entry Points	Registers Saved and Restored
IEAOPT01 ¹	0-9, 12, ² 13, ² 14
IEAOPT02	0-9, 12-14
IEAOPT03	0-14
IEAOPT0E	2-14
<p>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.</p>	

Figure 16. POST Branch Entry Output

Branch Entry to the WAIT Macro Instruction

Branch entry to the WAIT macro instruction 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 the local lock and be in key zero, supervisor state. While holding the local lock and before branching to WAIT, the caller must re-establish the PSW and register return environment in its RB and TCB. When WAIT is invoked, only the local lock should be held by the caller. WAIT performs the following functions:

- Store 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).
- Release the local lock.
- Return 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 results in the setting of 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).

Suspension and Resumption of Request Block Tasks

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.

Using the SUSPEND Macro Instruction

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 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 uses. When used with the SUSPEND macro, however, the wait count field is known as the suspend count field, even though the function it performs for both macros is the same.

The SUSPEND macro does not have an immediate effect on the issuer as the WAIT macro does. Instead, the effect is delayed, depending on the type of suspension requested by the macro user. 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 event completion, the RB that issues a SUSPEND RB=CURRENT should exit to the dispatcher (for example, issue a CALLDISP macro with the BRANCH=YES option). Using the SUSPEND macro this way minimizes potential performance problems because the RB in this case must either be disabled or must hold the local lock. Keeping suspension time minimal also minimizes other potential problems the program may experience by limiting the time in which the RB is unable to cause any synchronous interrupts (for example, 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.

The following examples show typical SUSPEND macro instruction sequences:

Example 1:

SUSPEND RB=PREVIOUS

1. Type 2 SVC routine receives control.
2. The SVC suspends the macro issuer's task.
3. Event completion (for example, an I/O event) occurs.
4. The macro issuer's task resumes (exit from the SVC routine).

Example 2:

SUSPEND RB=CURRENT

1. Macro issuer acquires the local lock.
2. The macro suspends processing of the current RB.
3. Event completion occurs: stimulus for resumption initiated.
4. Macro issuer releases local lock.
5. Macro issuer issues CALLDISP BRANCH=YES.
6. Normal processing resumes.

Here are some other considerations for using the SUSPEND macro instruction:

- SUSPEND may be issued only by a routine executing under protection key 0.
- SUSPEND requires that the CVT mapping macro be included.
- When the issuer requests (or assumes 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 may issue SUSPEND, and then only for the current task. SUSPEND may not 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-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 prior to 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. 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. Since the issuer must also co-ordinate the SUSPEND - RESUME sequence, the event completion must not occur until after the SUSPEND RB=CURRENT macro takes effect. The issuing function should be either a Type 1 or a Type 6 SVC because the SUSPEND function sets the RBSCF (suspend count field) to one and decreases the ASCBTCBS count. Holding the local lock or disabling will prevent the normal dispatcher processing. The CALLDISP function may be used by Type 2, 3, or 4 SVCs to exit.
- The SUSPEND-RESUME sequence must not be intermixed with the WAIT-POST sequence on a single RB because both sequences use the same count field for control of the functions. Since 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 may have only one SUSPEND outstanding at a time. There may be no subsequent SUSPEND macros issued until a RESUME occurs for the outstanding SUSPEND macro.

Using the RESUME Macro Instruction

The RESUME macro instruction provides an efficient means for indicating the completion of an event. The RESUME macro specifies the TCB and RB that were suspended. Control returns to the issuer who specifies RETURN=Y. Issuers in SRB mode can use the option RETURN=N to have control passed to the resumed task.

Note that the RESUME macro instruction:

- Can be issued only by a routine executing in supervisor state under protection key 0.
- Can be issued only in the address space in which the task being resumed resides.
- Can not be issued from a disabled interrupt exit (DIE) because the address space in which the DIE routine will execute is unpredictable.
- Can not be issued by a disabled routine.
- Must be able to acquire the local lock unconditionally under the following circumstances:
 - The TCB (task) being resumed is active on another processor.
 - The stage 3 exit effector is in the process of enqueueing an interrupt request block onto the task control block.
 - RCT QUIESCE is preparing to swap out the address space specified by PSAAOLD.

Therefore, any user of RESUME who executes disabled or who holds a lock higher than the local lock must ensure that the macro does not require the local lock.

- The RETURN=N option can be issued only by SRBs. Because this option constitutes an exit from the SRB, the caller can not hold any locks when issuing RESUME.
- RESUME requires that the CVT mapping macro be included.

Using the TCTL Macro Instruction

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 chain, but only if it passes all the dispatchability tests normally made by the dispatcher.

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 macro or a RESUME RETURN=N, the routine will abnormally terminate.
- 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, can not hold any locks.

Using the BRANCH=YES Option of the CALLDISP Macro Instruction

The CALLDISP macro instruction includes the BRANCH=YES option to allow 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.

Here are some considerations for using the BRANCH=YES option on the CALLDISP macro instruction:

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

CVT
IHAPSA
IKJTCB
IHARB

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 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 system integrity. The related topic of security in regard to the physical environment of a computing system is discussed in **Data Security and Data Processing**, G320-1370 through G320-1375.

Restricted functions (those that can be requested only by authorized programs) are documented separately from non-restricted functions in **OS/VS2 System Programming Library: Data Management**, GC26-3830.

These publications describe the restricted functions and the means of authorizing appropriate programs to use those functions. Macro instructions that require the issuing program to be authorized are listed in the contents and described in the text of these books. Review the contents of these books to determine what functions require authorization.

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 control program employs to maintain system integrity.

Elimination of Potential Integrity Exposures

MVS 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 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 fetch 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.

The elimination of this problem requires that system-key routines 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 convention that system-key routines obtain the protection key of the user before accessing the user-specified

area of storage. For example, MVS 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 may have many data sets allocated). Inadequate validity checking in this situation can create an integrity exposure, since an unauthorized 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 (since 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, 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 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 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 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 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, etc., 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 prior to 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, 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 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 maintains such items in system storage, or in a separate address space in the case of some APF-authorized programs.

It may 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 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.

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, etc.) 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.

Installation personnel should understand the distinction between an authorized program and an authorized user. A program that utilizes a restricted function must be authorized. Any user, however, can submit a job to execute an authorized program. Any program that must be restricted to an individual or a class of users must be restricted by means of existing data set security facilities. The program must be placed on a separate password-protected library other than LINKLIB, SVCLIB, or LPALIB. IEH utilities might be placed on a separate password-protected library to limit their use to system programmers.

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 present data security measure currently being used.

This facility 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. 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, GC28-0722**.

RACDEF Macro Instruction

This macro instruction can be used to define, modify, and delete resource (for example, tape volumes and DASD data set) profiles for RACF.

The resource manager (a data management program) responsible for establishing and maintaining the resources issues the RACDEF macro instruction to define or delete their profiles to RACF.

RACINIT Macro Instruction

This 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 card. 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 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.

RACHECK Macro Instruction

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.

System authorization checking is performed by RACF when a resource manager (for example, data management OPEN) which 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 users' 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.

RACLIST Macro Instruction

This 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.

Interruption, Recovery/Termination, and Dumping Services

The supervisor offers many services to assist in the detection and processing of abnormal conditions during system execution. Certain types of abnormal conditions are detected by the hardware and cause program interruptions to occur (for example, an attempt to execute an instruction with an invalid operation code). Other abnormal conditions are detected by the software (for example, an attempt to open a data set that is not defined to the system causes an ABEND to be issued by the Open routine).

The supervisor provides facilities for writing exit routines to handle specific types of interruptions and abnormal conditions. The supervisor initiates the recovery/termination process 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.

The services discussed in this chapter include:

- Program interruption processing (SPIE macro instruction)
- Recovery/termination (CALLRTM macro instruction)
- Functional recovery routines (SETFRR macro instruction)
- Task recovery (STAE and ESTAE macro instructions)
- Installation-written clean-up routines
- Virtual storage dumping (SDUMP macro instruction and CHNGDUMP command)

SPIE Processing

The SPIE macro instruction provides a problem program with a means of specifying an error exit routine in response to one or more program error interruptions. SPIE and its related services are discussed in detail in *OS/VS2 Supervisor Services and Macro Instructions*.

For the problem programmer, interruptions 1-15 may be specified in the SPIE macro instruction. For the installation-authorized system programmer, interruption 17 may also be specified. Interruption 17 designates page faults and may be specified so that a user-written SPIE routine gets control on a page fault before any supervisor routine. The user-provided SPIE routine only gets control in problem program state and in the key of the TCB(TCBPKF) on a page fault for the program issuing the SPIE 17. The SPIE 17 routine covers page faults at the task level and any RBs executing under the task for which the SPIE was issued. Since the SPIE 17 user may not be in supervisor state, the routine generally can not resolve the page fault in a manner compatible with the operation of the system. Hence, the SPIE 17 routine usually resolves the page fault by invoking PGLOAD/PGFIX functions of the paging supervisor.

The SPIE 17 routine gets control in problem program state and in the key of the TCB(TCBPKF). If the program is in supervisor state at the time the SPIE 17 is issued, it is be abnormally terminated with a 30E abend completion code. If the program is in supervisor state and takes a page fault while the SPIE 17 routine is active, the SPIE 17 exit routine does *not* get control. Supervisor routines resolve the page fault and continue program processing without abending the program.

During SPIE processing, the Program Check First Level Interrupt Handler (FLIH) passes control directly to the SPIE routine after some set-up processing via a LPSW. The program check FLIH sets a recursion indicator to cover any PIE/PICA references during the setup processing done to handle a page fault incurred in the problem program. If a page fault occurs on a reference to the PIE/PICA, the paging supervisor is given control to handle the original page fault. Processing continues in the problem program once the page fault is resolved. The

SPIE exit routine does not receive control at all since the FLIH is not able to obtain the information needed by the SPIE routine as input parameters.

If interruption 17 is used, the programmer must page fix the PIE, PICA, and SPIE exit programs and data areas. The SPIE exit routine must be aware that page faults can occur after issuing the SPIE macro instruction for interruption 17 and prior to fixing the required control blocks. If the page fault occurs at this time, the program FLIH tries to pass control to the SPIE 17 routine after the setup processing mentioned. If the PIE/PICA can be validly referenced, control is passed to the SPIE 17 routine. If the SPIE 17 routine page faults, the page fault is resolved by the paging supervisor if the routine is not running disabled. A disabled page fault causes an OC4 abend. Once the page fault is resolved, normal processing continues in the SPIE 17 routine.

It is important to note that the SPIE 17 routine may not get control on every page fault due to the recursion logic mentioned above.

Recovery/Termination

The recovery/termination manager (RTM) monitors the flow of control of software recovery processing and supplies the services of normal and abnormal task and address space termination. Its purpose is to select the appropriate recovery or termination process according to the status of the system and the request of its invokers.

The RTM may be called to perform its recovery and termination services on behalf of the caller or to direct its services to another routine. It is invoked by two macros: the ABEND macro and the CALLRTM macro.

The recovery/termination process is invoked for the following events:

- Unanticipated program checks - Except those protected by SPIE routines.
- Machine checks.
- Invalid issuance of an SVC while locked, disabled, or in SRB mode.
- I/O error on page - in request.
- Restart Key - An operator-initiated recovery action, requested by pressing the console RESTART key.
- Request by an authorized caller to terminate a task or a memory.
- ABEND macro - This may be system-issued or user-issued.

When one of these events occurs, the RTM is given control. If a recovery exit has been specified (via SETFRR, STAE, or ESTAE) it is invoked to recover or clean up for the process in control. Should this recovery routine be unable to recover from the incident (request termination or fail itself) the previously established recovery exit is invoked. This process is called percolation. In the event that all recovery routines are unable to recover, the process is terminated.

The recovery routines are given control in LIFO (last in, first out) order. If all recovery routines established via SETFRR percolate, the related task, if one exists, is abended. Then the STAE/ESTAE routines which were created by the task are invoked.

Invoking the Recovery/Termination Manager

CALLRTM

A routine should use CALLRTM to direct the recovery/termination services to a task or routine other than itself or its ancestors (callers). CALLRTM may be issued only by key 0 privileged routines. Control is returned to the issuer of the macro if the TYPE=ABTERM or TYPE=MEMTERM options are coded.

The following locking, workarea, and special considerations should be noted when using CALLRTM.

TYPE=ABTERM

If TYPE=ABTERM is specified in the CALLRTM macro instruction, the RTM processing is directed towards another task. In this situation, the following locking and save area requirements should be considered:

- 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 pass a save area via register 13. If dump options are supplied, they must be contained in fixed pages.
- 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 a save area is not required.
- If the ASID parameter is specified, the ABTERM function is scheduled as a service request block (SRB). Although there is no specific lock requirement, the caller must pass the address of an 18-word work area in register 13. In this way, the ASID parameter allows processing across address spaces and allows processing in the current address space for routines that cannot acquire the proper locks (possibly because of hierarchy conflicts).

TYPE=MEMTERM

If TYPE=MEMTERM is specified in the CALLRTM macro instruction, the RTM processing is directed towards an address space. In this situation, the following information should be considered:

- If the ASID parameter is specified as 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: This 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.

Since the MEMTERM process 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 is 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 FRR, when it determines that uncorrectable translation errors are occurring in the address space.
- RTM, when it determines that task recovery and termination cannot take place in the current address space.
- Region control task, when it has determined that the address space may become permanently deadlocked--that is, unusable--or that the status of the address space is unpredictable due to 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.
- ASM recovery routine, when it determines that uncorrectable translation errors are occurring while ASM is using the control register of another address space to update the LSQA of the address space.
- 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 bypassing of the task resource manager processing.

ABEND

The ABEND macro instruction should be used by any routine, including supervisor state, locked, disabled or SRB routines, to request the services of the RTM to be directed to itself (cause entry into its recovery routine) or to its callers. The issuer of ABEND should remove its own recovery exit if it wishes its caller to receive the services of the RTM. Control is never returned to the issuer of the macro (except by using the STA/ESTA/FRR retry mechanism). See OS/VS2 Supervisor Services and Macro Instructions for a description of the ABEND macro.

Types of Recovery Routines

Functional Recovery Routines (FRRs)

FRRs are recovery routines established to protect locked, disabled, or SRB mode routines. An FRR is identified to the RTM by coding the SETFRR macro instruction. When a functional recovery routine is invoked it runs in the state of the system (enabled or disabled) and with the locks that were held at the time of the error, or as modified by previous FRRs.

Task Recovery Routines (STAE/STAI ESTAE/ESTAI)

Task recovery is accomplished through STAE/STAI or ESTAE/ESTAI routines. Issuance of the STAE or ESTAE macro instruction, or the ATTACH macro instruction with the STAI or ESTAI parameter, allows the user to intercept a scheduled ABEND. Control is given to a user-specified exit routine in which the user may diagnose the cause of the ABEND, and perform pre-termination processing, or specify a retry address if he wishes to prevent the termination.

Establishing Recovery Routines

Functional Recovery Routines

The SETFRR macro instruction provides control program functions with the ability to define their recovery in the FRR LIFO stack which is used during system recovery management. The LIFO stack is maintained by the recovery termination manager and contains the addresses of the FRRs established to protect a single path through supervisor control and SRB code.

The issuer of SETFRR must be key 0 because the stack is maintained in protected storage. Furthermore, the SETFRR issuer must be disabled, locked, or in SRB mode (and therefore supervisor state) to maintain FRR stack integrity. The FRR stacks are serialized by disablement or ownership of a global spin lock. For owners of suspend locks (CMS or LOCAL) and SRB functions, the stack is saved and restored as part of the paths operating environment by the supervisor control functions on interruptions and redispach.

If RTM is invoked, the last FRR established is given control and executes in the system mode (locks held, disablement, SRB mode) at the time of the error.

The FRR must indicate to the RTM the action to be performed. This is done via settings in the system diagnostic work area (SDWA), which is used for communication between FRRs and the RTM. If the FRR requests percolation (no retry), the previously established FRR is given control. Each FRR is given control in LIFO order until retry is requested or the stack is exhausted.

When FRR processing is exhausted, the RTM determines if a task should be terminated (current or related task if in SRB mode). The RTM then sets the task up for ABEND and task recovery takes place if task recovery exits exist, otherwise, the task is abnormally terminated.

An FRR environment is canceled when a routine issues SETFRR with the delete option, when that FRR requests percolation, or when the system becomes enabled in task mode and an interruption and redispach occurs. The FRR entry should be deleted before the function returns to its caller, otherwise, the FRR may get control for its caller's error.

All SETFRR users must include the DSECTs for the FRR stack (via the IHAFRRS mapping macro instruction) and the PSA (via the IHAPSA mapping macro instruction) prior to using the SETFRR macro instruction. In addition, all disabled, locked, and SRB routines which define recovery must be key 0 supervisor state when using the SETFRR macro instruction. It is necessary to copy IHAPSA from AMODGEN into MACLIB.

Note: The size of all FRR stacks satisfies the recovery needs of the control program. If additional FRRs are placed on the stack, thereby causing the size to be exceeded, the routine issuing the SETFRR macro instruction is abnormally terminated. Any user-written routines outside the control program may add one, and only one, FRR to the stack; if more than one is added, abnormal termination may occur. This applies to all of the recovery stacks, including the normal stack. The normal FRR stack is used by control program routines which are invoked on behalf of the user.

Task Recovery Routines

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 which 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 exit routine receives control. The STAE exit 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 exits receive control (return code = 16).

Programs running under a single TCB may 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 which 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.

ESTAE, FESTAE, and ESTAI Exit Routines: The ESTAE and FESTAE macro instructions, like the STAE macro instruction, cause a task and RB-related recovery environment to be created. Use of the ESTAI option on the ATTACH macro instruction also identifies a recovery routine to the control program, but the routine is associated with the subtask created via ATTACH. Furthermore, ESTAI recovery routines are propagated to all lower-level subtasks of the subtask created with the ATTACH that specified the ESTAI parameter. (See Figure 17.)

If a task is scheduled for abnormal termination, the recovery routine specified by the most recently issued ESTAE macro instruction is given control and executes under a program request block created by the SYNCH service routine. On return, the exit routine may indicate whether a retry routine should be scheduled or whether termination should continue. If it requests that termination continue, the next ESTAE routine for the task receives control. If all ESTAE routines request that termination continue, or if none exist, the ESTAI routines, if any, receive control.

Before the initial recovery routine receives control, the purge and asynchronous processing requests specified when the exit was created are performed by the control program. The I/O processing requested is performed only for the first exit routine selected. Subsequent routines receive an indication of the I/O processing previously performed, but no additional I/O processing is performed. The asynchronous processing request, however, is performed for each routine.

Each ESTAE exit established by a task is eligible to receive control. If an error occurs, the most recently created ESTAE is entered. If it requests that termination continue, or it fails itself, the next ESTAE exit, if any, is entered. This continues until an ESTAE exit requests retry or all exits for the task are exhausted.

Both STAE and ESTAE exits can exist for the same task. However, only one STAE exit receives control. STAI and ESTAI exits receive control after all ESTAE exits and one STAE exit, if any, have been processed.

An ESTAE environment is canceled when the RB which creates it goes away (unless it issues an XCTL and specified the XCTL=YES parameter on the ESTAE macro instruction), when the ESTAE macro instruction is issued with the CANCEL option, the exit routine fails, or the exit routine requests that termination continue.

Figure 17 demonstrates the queuing structure of the ESTAE routines and the propagation of ESTAI to subtasks.

The fast ESTAE (FESTAE) macro instruction allows an SVC to establish an ESTAE environment with minimal overhead and no locking requirements. The ESTAE exit activated by FESTAE receives control in the same sequence and under the same conditions as though created by the ESTAE macro instruction.

To use the fast ESTAE capability, SVCs must code the FESTAE macro instruction as described in Part II of this publication. Restrictions governing the use of FESTAE are as follows:

- Only an SVC executing under an SVRB in protection key 0 may use FESTAE.
- FESTAE may be used only once to create an ESTAE exit, which means that:
 - Any SVC needing to change its exit address must use branch entry ESTAE services.
 - Any SVC needing more than a single ESTAE exit must use SVC 60 or branch entry to get additional exits.

FESTAE provides a 24 byte parameter area that the user may optionally employ. The name of the parameter area is RBFEPARM, and it is in the SVRB. The ESTAE exit receives this parameter area when an error occurs. Hence, the user of FESTAE can clear (zero) and initialize the parameter area with appropriate information (for example, tracking data) that may be useful to the ESTAE exit. Clearing the parameter area before use is necessary to ensure no spurious data remains in it from previous processing.

FESTAE users must also include the following DSECTS for the FESTAE macro expansion:

IHARB
IKJTCB
IHAPSA
IHASCB

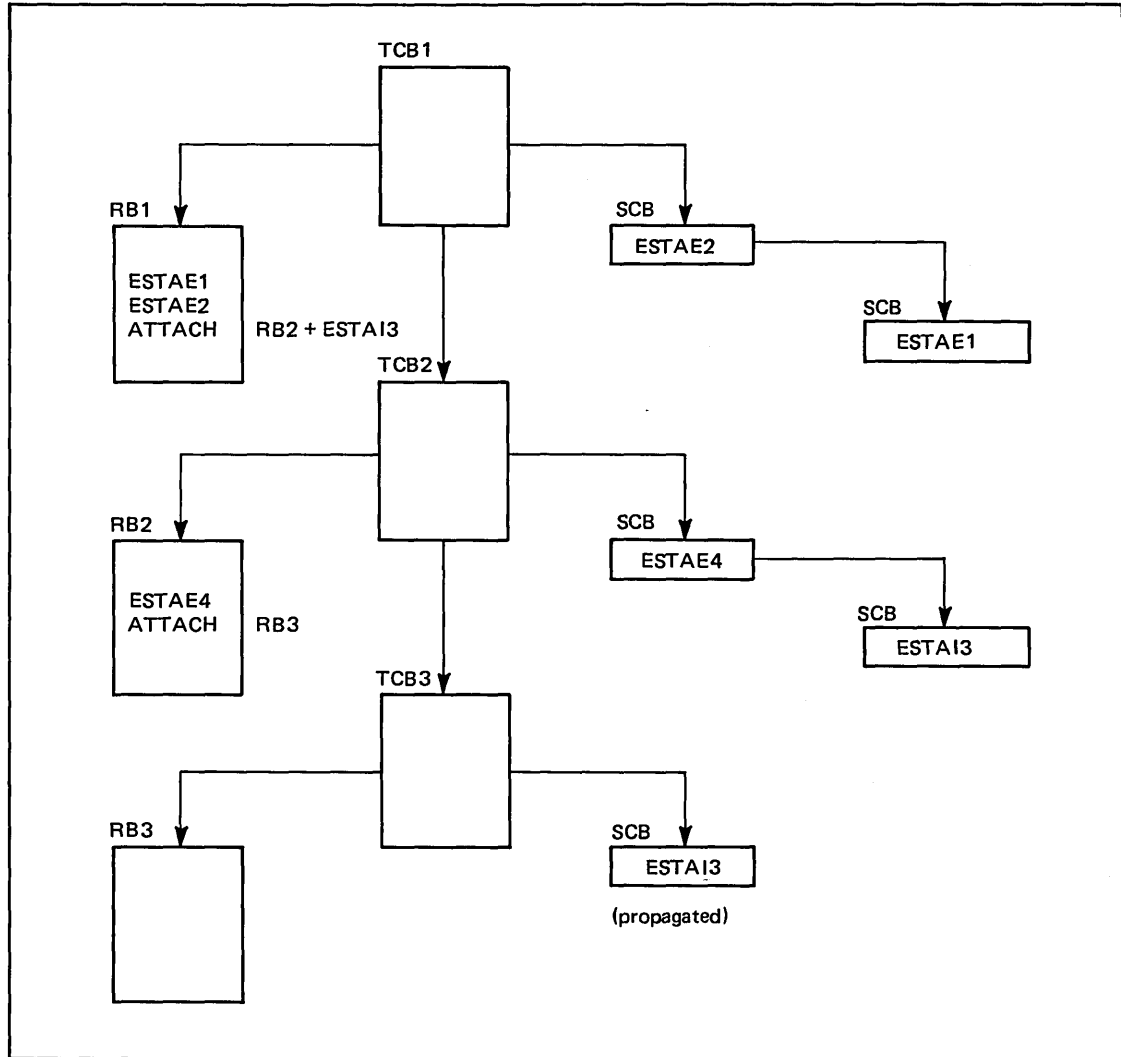


Figure 17. ESTAE Environment

RTM/Recovery Routine Interface

Interface to Functional Recovery Routines

Prior to giving control to FRRs the RTM locates and initializes a work area that contains information about the error. This work area is called the system diagnostic work area (SDWA) and is 512 bytes long. The first word of the SDWA contains the address of the six-word parameter area returned by the SETFRR macro instruction when the PARMAD keyword is specified on the SETFRR macro instruction.

FRRs represented on the system recovery stack receive control via the LPSW instruction. The most recent FRR address on the stack is merged with the saved error PSW so that the FRR gets control in the system state representing the processor status at the time of error. Note: The PSW is always enabled for machine checks and DAT.

Upon entry to the FRR, parameter registers are as follows:

Register	0	Address of a 200-byte FRR work area.
Register	1	Address of the SDWA.
Register	14	Return address.
Register	15	Address of the FRR.

Note that register 13 is not part of the FRR interface. Any register may be used without saving it, but caution should be used to maintain the return address supplied in register 14.

The locks held and disablement are the same as at the time of error, except for percolated-to FRRs in which case, lock status may change if previous exits requested that locks be freed.

The System Diagnostic Work Area (SDWA) contains information pertaining to the error (that is, registers and PSW). The SDWAFMID field contains an indication of the memory in which an error occurred if recovery is being initiated in another memory. If this field is not zero, no reference may be made to private (local) area.

An FRR should use the lock freeing capabilities of the RTM (via SETRP) to free locks obtained by the mainline, if that is the action desired. In any event, an FRR must not free the last global lock causing enablement, or the local lock. The RTM must be used in these cases. If the RTM is used to free all locks required, the above checks can be avoided.

RTM freeing of locks is only honored on percolation. Freeing or obtaining locks for or in retry situations must be done by the retry routine.

During FRR processing, the FRR should use the SETRP macro instruction to inform RTM of the action it desires to be performed when the FRR completes. This macro instruction initializes the system diagnostic work area (SDWA) with the desired options.

The SETRP macro instruction is described in *Supervisor Services and Macro Instructions*, with the exception of several restricted parameters which are described in this publication.

Interface to Task Recovery Routines

Interface to a STAE/STAI Exit: Prior to going to a STAE/STAI recovery routine, the control program attempts to obtain and initialize a work area that contains information about the error. The SDWA work area is 512 bytes long. 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	code 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 exit will be given control. No other STAE exits 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 exit.

For the following situations, STAE/STAI exits 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.
- 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.

Interface to an ESTAE/ESTAI Exit: Before going to an ESTAE/ESTAI recovery routine, RTM attempts to obtain and initialize a work area that contains information about the error. This work area is the system diagnostic work area (SDWA) and is 512 bytes long. The first word of the SDWA contains the address of the parameter list specified on the ESTAE macro instruction or the ESTAI parameter of the ATTACH macro instruction.

Upon entry to the ESTAE exit 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 no I/O processing was performed.
Register	1	address of the SDWA.
Register	13	save area address.
Register	14	return address.
Register	15	entry point address.

If no SDWA was available:

Register	0	code 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	entry point address.

During ESTAE/ESTAI processing, ESTAE/ESTAI should use the SETRP macro instruction to inform RTM of the action it wants done when ESTAE/ESTAI completes. This macro instruction initializes the system diagnostic work area (SDWA) with the desired options.

If RTM could not provide a work area, a register save area is not provided either. If no SDWA is available, register 14 must be saved and used as the return register to RTM.

When the ESTAE or ESTAI routine has completed its processing, it should return to RTM via the contents of register 14. Register 15 should contain one of the following return codes if an SDWA was not obtained:

Return Code	Action
0	Continue the termination. Any ESTAE exits established prior to this one receive control.
4	A retry routine is to be scheduled, and its address is placed in register 0.
16	Valid only for ESTAI. Termination should be continued; no further ESTAI processing should be performed.

When an ESTAE routine requests retry, the RB queue is terminated up to, but not including, the RB of the program that issued the ESTAE macro instruction. This is done by pointing the RB old PSW to an SVC 3 instruction. In addition, open DCBs that can be associated with the purged RBs are closed and queued I/O requests associated with these DCBs being closed are deleted from the I/O restore chain.

The RB queue 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 be canceled by this RB purge. Examples of these effects are as follows:

- Subtasks created by a program to be purged
- Resources allocated by the ENQ macro instruction
- DCBs that exist in dynamically acquired virtual storage

If there are quiesced restorable input/output operations, they can be restored in the ESTAE 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.

RTM/Retry Routine Interface

FRR-Requested Retry Routine

If an FRR requests that a retry routine be given control by specifying a return code of 4 on SETRP, the following interface is established:

- Registers 0 - 14 are the registers at the time of error with the exception of registers that the FRR requested be changed.
- Locks held and disablement are the same as on exit from the FRR.
- Protect key is 0.
- Register 15 contains the retry routine address.
- The SWDA and associated 200-byte work area are not available to the retry routine.

Task Recovery Retry Routines

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 user's 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 which 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:

- Subtask 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 OS/VS2 System Programming Library: 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 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 512-byte 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 the PIRL; or 0 if I/O is not restorable.

ESTAE/ESTAI Retry Routines: If the ESTAE/ESTAI routine is specified, the following actions will be performed by the control program prior to scheduling it:

- A dump is provided if requested.
- FREEMAIN of SDWA is performed, if requested.
- Registers are updated with user-supplied values, if requested.
- The RB queue is purged up to the level of the retrying RB.

Outstanding WTORs are not purged prior to scheduling an ESTAE routine.

Retry routines run at the RB associated with the requestor of the ESTAE exit routine causing the retry. RBs on the RB queue are always purged to the level of RB associated with the ESTAE exit prior to the scheduling of the retry routine.

ESTAE retry routines get control in the key in which the ESTAE macro instruction was issued. Retry routines should not reissue ESTAE to maintain the same exit. They may, however, issue ESTAE to add or change exits.

Like the exit routine, the retry routine must be in storage when the exit routine determines that retry is to be attempted. If not already resident within the program, the retry routine may be brought into storage via the LOAD macro instruction by either the user's program or exit routine.

If an SDWA was obtained, the user has a choice of interfaces to his retry address. The user can set (in the SDWA) the registers he wishes to have and request that they be passed to the retry address by coding RETREGS=YES on the SETRP macro instruction. This alternative is most often used when retrying into mainline processing.

If no SDWA was obtained or if RETREGS=NO was specified on SETRP, only parameter registers are passed to the retry address. This alternative is more often used if a special retry routine is to get control.

The parameter registers are as follows:

If no SDWA was obtained:

Register	0	12
Register	1	address of the user parameter list established via ESTAE or ATTACH with ESTAI.
Register	2	pointer to the PIRL if I/O was quiesced and is restorable; otherwise, 0.
Register	14	address of supervisor-assisted exit linkage.
Register	15	entry point address of retry routine.

If an SDWA was obtained and the exit did not request register update, or freeing of SDWA:

Register	0	0
Register	1	address of SDWA.
Register	2	unpredictable.
Register	14	address of supervisor-assisted exit linkage.
Register	15	entry point address of retry routine.

If an SDWA was obtained and the exit did not request register update, but did request freeing of SDWA:

Register	0	20
Register	1	pointer to the user parameter list established via ESTAE or ATTACH with ESTAI.
Register	2	pointer to the PIRL if I/O was quiesced and is restorable; otherwise, 0.
Register	14	address of supervisor-assisted exit linkage.
Register	15	entry point address of retry routine.

Recovery Routine Guidelines

This section is intended to assist in the writing of recovery routines. The actions a recovery routine should take are highly dependent on the function being recovered, therefore, these guidelines are general and intended to serve as suggestions.

The first consideration is that the recovery routine be beneficial. In general, if a function acquires resources which may be requested by another function, or is 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 within a subpool which is not task-related (for example, subpool

231). Another case when a recovery routine should be established is when data areas, queues, data sets, etc. which are used by more than one function are manipulated. The recovery routine in this case should maintain integrity in case of failure.

Recovery routines may also be used to:

- Intercept errors and perform clean-up processing.
- Intercept expected program checks and perform desired action.
- Isolate an error to a particular section of processing and continue further processing if possible.
- Intercept its own abends and provide tailored dumps.

The second consideration is what type of recovery routine be established. If the function holds a lock, is physically disabled, or is an SRB, an FRR can be used to intercept errors. If the function is running under a task and holds a lock during some portion of its processing, an ESTAE could be used to catch errors in its processing, but the lock is freed before the ESTAE is given control. Losing the locked status can be tolerated when the lock is used only to read a queue to prevent another from changing it. Also, if the function is running as an enabled, unlocked, SRB, an ESTAE associated with its related task could be used to catch errors in the SRB. If an FRR is not required, but a recovery routine is necessary, an ESTAE routine should be used.

If the function attaches any subtasks, then it can also provide recovery for the subtask. Specify the ESTAI/STAI parameter on the ATTACH macro instruction to provide recovery.

If the function obtains any SQA or CSA for an address space, and the address space is terminated by the operator (for example, through use of the FORCE operator command), then the resource clean-up routine must free the areas obtained. This is necessary because the recovery routine is not invoked when the address space is terminated. If the address space is not terminated, and an abend occurs, then the recovery routine is invoked. If the recovery routine does not retry, then it must free the SQA and CSA.

FRRs

If it is decided that recovery should be via an FRR, the following information should be reviewed:

- Syntax of the SETFRR macro which is documented in Part II of this manual.
- Guidelines for using SETFRR which are documented in this publication under "Establishing Functional Recovery Routines."
- The interface to functional recovery routines which is described earlier in this section.
- Syntax of the SETRP (SET return parameters) macro instruction which is described in *OS/VS2 Supervisor Services and Macro Instructions* with a description of the restricted parameters in Part II of this publication.
- Interface to FRR requested retry routines which is described earlier in this section.
- The contents of the SDWA (mapped by the IHASDWA mapping macro instruction). All error/recovery information available to an FRR is contained in this work area, and the commentary in this data area serves as the documentation. This data area is described in the *OS/VS2 System Programming Library: Debugging Handbook*.

Task Recovery

If it is decided that recovery via FRR is not necessary, ESTAE recovery should be used. Discussions earlier in this section concerning STAE routines are primarily documented as they are supported for OS/VS2 Release 1 compatibility.

Before designing an ESTAE routine, the following information should be reviewed:

- The use of the ESTAE macro instruction which is described in **OS/VS2 Supervisor Services and Macro Instructions** with a description of restricted parameters in Part II of this publication.
- The uses of the FESTA macro instruction, described earlier in this section.
- Rules concerning establishing recovery routines (ESTAE/ESTAI exit routines) earlier in this section.
- The interface to ESTAE/ESTAI exits which is documented in this section under "Interface to Functional Recovery Routines."
- Syntax of the SETRP macro instruction which is described in **OS/VS2 Supervisor Services and Macro Instructions** with a description of restricted parameters in Part II of this publication.
- Interface to ESTAE/ESTAI retry routines described earlier in this section.
- The contents of the SDWA (mapped by the IHASDWA mapping macro instruction). All error/recovery information available to an ESTAE/ESTAI exit is contained in this work area and the commentary in this data area serves as the documentation. This data area is described in **OS/VS2 System Programming Library: Debugging Handbook**.

For any recovery routine that is being written, the following information should be reviewed:

- When a task recovery routine receives control it should first examine the code in register 0 to see if a SDWA was provided. If an SDWA was not provided (register 0=12) a save area is not pointed to by register 13 and the registers should not be saved. (Affects task recovery routines--ESTAE/ESTAI.)
- An SDWA is always provided to FRRs. (Affects functional recovery routines.)
- A recovery routine should not assume the registers in the SDWA are its own. Many reasons preclude this such as errors in called routines which have no recovery, errors in an asynchronous routine (for example SRB and IRB). The safest method to assure a successful retry is to save volatile information in the parameter list passed to the recovery routine and use those registers, addresses, and so forth, for retry. For example, the issuer of the macro instruction can save the base register and data register for the function in the parameter list. This enables the recovery routine to reference the areas that belong to the function. (Affects both task recovery routines--ESTAE/ESTAI, and functional recovery routines.)
- The default dump data set for SETRP is the data set that the user of the function 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, then either the data sets specified on the DD statement should be secure or the SDUMP macro instruction should be used to take the dump. SDUMP can include a DCB for a secure data set, or the SYS1.DUMP data set can be the default.
- If an FRR requests a dump via SETRP, the following should be considered:
 1. No dump is taken by the supervisor if retry is performed before the error is percolated to task recovery.
 2. A dump is taken by the supervisor if all FRRs percolate and no subsequent recovery routine suppresses the dump.
 3. A dump is taken by the supervisor at the task recovery level (after the system is enabled). If volatile information is required, an SDUMP macro instruction should be issued instead with the volatile information moved into the 4K SQA buffer described later in this section. Volatile information can also be obtained by specifying SUMLIST=ADDRLIST, BRANCH=YES, and SDATA=SUMDUMP in the SDUMP macro instruction. Then a summary dump will be taken while the system is disabled. The areas indicated by the SUMLIST parameter are included in the summary dump.

(These considerations affect functional recovery routines.)

- Dump options specified on SETRP are accumulated in the SDWA. During percolation these options are merged with any dump options specified on an ABEND macro instruction, a CALLRTM macro instruction, or specified by other recovery routines. Also, the CHNGDUMP operator command can add to or override dump options. The supervisor takes one dump as described by the dump options. If the recovery routine requests a retry, then the sequence of events is SETRP, take the dump, and then retry. If the recovery routine does not request a retry, then the sequence of events is SETRP, percolate, and then take the dump. (Affects both task recovery routines--ESTAE/ESTAI, and functional recovery routines.)
- If the abend was caused by a system routine, the recovery routine should issue an SDUMP macro instruction instead of taking the default dump for SETRP. If SDUMP is used, then the following should be considered:

1. SDUMP should specify a title that summarizes the problem and function. 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:

```
0C4 ABEND IN OPEN, ERRMOD=IFGORROA,
RECVRMOD=IGG020FC, JOBN=C49JAC1A, STEPN=GO,
SDWAVRA=8417F0
```

Note: The Print Dump Service will insert the first 62 characters of the title on its output pages.

2. SDUMP should usually specify all SDATA parameters. If a system area is not required to analyze the dump, then use the appropriate SDATA parameters to omit it from the dump.

(Affects task recovery routines -- ESTAE/ESTAI, and functional recovery routines.)

- Valuable function-related information can be lost by the time a dump or LOGREC entry is created. For example, information in control blocks used for communication with other address spaces can be changed before a dump is taken. "Footprint" areas should be used to save this information. The available "footprint" areas include the function's work areas, the LOGREC and dump variable recording area in SDWAVRA, the SDUMP 4K SQA buffer, and any areas the SUMLIST parameter specifies in a SDUMP macro instruction. Document any use of "footprint" areas. This documentation is very useful when analyzing the dump or LOGREC entries.

(Affects task recovery routines -- ESTAE/ESTAI, and functional recovery routines.)

- If LOGREC recording is active (the default for FRRs or specified on SETRP), then the name of the module that failed and the name of the recovery module should be saved in the SDWA (SDWAMODN, SDWAC SCT, and SDWAREXN).

(Affects task recovery routines -- ESTAE/ESTAI, and functional recovery routines.)

- 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 a dump (indicated by the SDWACTS and SDWAMABD bits being on) then another dump is not necessary. When several recovery routines are written for the same component, then 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, then the higher recovery routine should issue one.

(Affects task recovery routines -- ESTAE/ESTAI, and functional recovery routines.)

- Recovery routines that cleanup critical areas, such as queues of control blocks, should issue an ESTAE or SETFRR macro instruction. Then if an error occurs in the original recovery routine, control is passed to the recovery routine that was specified on the ESTAE or SETFRR. This routine can do the necessary cleanup before control is passed to the next routine (above the failing routine) in the recovery hierarchy. For example, if a system service invoked through a macro instruction in a recovery routine abends, then control is passed to the recovery routine specified on the ESTAE or SETFRR. This routine cannot retry. After it completes any necessary cleanup, it returns control to the supervisor. The supervisor passes control to the next routine (above the failing routine) in the recovery hierarchy. The next recovery routine can retry, because the error occurred in a recovery routine that was lower in the hierarchy.
- The following is additional information about some SDWA fields:

SDWAPERC - indicates this recovery is being percolated to; however, this does not indicate if a task recovery routine is being percolated to from an FRR.

SDWAFMID - is zero if recovery is taking place in the address space which suffered the error.

SDWACLUP - indicates retry is not permitted; resources should be freed in exit.

(These affect both task recovery routines--ESTAE/ESTAI, and functional recovery routines.)
- An ESTAE exit may request, via a SETRP parameter, that the control program free the SDWA instead of freeing it itself in a retry routine. (Affects task recovery routines--ESTAE/ESTAI.)
- An ESTAE exit may specify, via a SETRP parameter, what the contents of its registers should be on entry to the retry routine.

(Affects task recovery routines--ESTAE/ESTAI.)
- An ESTAE exit remains in effect when its retry routine receives control. It need not reissue the ESTAE to reestablish itself. (Affects task recovery routines--ESTAE/ESTAI.)

When using ESTAE/ESTAI routines the following should be considered:

- ESTAE or ATTACH may be issued by the routine.
- If an ESTAE/ESTAI exit requests termination or fails ESTAE/ESTAI percolation and the accumulation of dump options occurs.
 - The asynchronous exit indicator is reset according to the new exit's request.
 - I/O options for the new exit are ignored.
 - A new SDWA is initialized.
 - The new exit is scheduled.
- If all recovery exits (STAE/STAI and ESTAE/ESTAI) fail or indicate termination, the task is terminated.
- If a non-jobstep task issues ABEND with the STEP option, exits are entered for the non-jobstep task. If retry is not requested, the jobstep is terminated with the ABEND code, and only TERM exits are entered.
- ESTAE exits receive control with the same status (supervisor or problem program state) that existed at the time the program issued the ESTAE macro instruction to queue the exit. ESTAE exits created by a program running under any control program protection key (keys 0-7) receives control in key 0; otherwise ESTAE exits receive control with the same protection key as the program that established the exit.
- ESTAI exits receive control in the key of the TCB of the task that created them.

- In the following cases:
 - forced logoff
 - job step timer expiration
 - wait time limit for job step exceeded
 - ABEND condition because of DETACH of an incomplete subtask
 - ESTAI was issued by a subtask and the attaching task abnormally terminates

The following actions occur:

- ESTAE exit routines are scheduled if TERM=YES was specified as a parameter when ESTAE was issued.
- All such routines which may exist get control in LIFO order.
- Any ESTAI exit previously suppressed via return code 16, or any exit previously entered which specified return code 0, is not entered again during TERM processing.
- Retry indications on return are ignored.
- If the TERM option is used on the ESTAE macro instruction issued by an ESTAE exit, it is ignored.

Although the ESTAE routines should issue SETRP to allow the system to free the SDWA, the freeing could also be accomplished by the retry routine. In this case, it is important to note that the ESTAE recovery routine created under any control program protection key receives an SDWA in key 0 storage. Therefore, if the retry routine is executing under a key other than key 0, it must issue MODESET to become key 0 before issuing the FREEMAIN.

Clean-Up Routines

Task and 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. It is the responsibility of the resource managers invoked to establish clean-up routines to “clean up” the queues and control blocks associated with the resources.

The responsibilities of the clean-up routines 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 clean-up routine 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.

Support for Installation-Written Clean-Up Routines

In order to support installation-written clean-up routines, a CSECT is provided into which an installation can assemble the names of subsystem clean-up modules. These modules are given control at the beginning of both the task and address space termination processes to do any special clean-up processing required by the subsystems. (The processing described above is performed by the IBM system routines.) After the CSECT is assembled, it is used to replace the existing CSECT IEAVTRML in load module IGC0001C in SYS1.LPALIB. The installation-written modules must be placed in SYS1.LINKLIB (or a library concatenated to

SYS1.LINKLIB via a LNKSTxx member of PARMLIB) or SYS1.LPALIB. If the routine is not present in one of these libraries, the system will not IPL.

Initially, the CSECT consists of four 12-byte entries of all zeros. Each of the first three 12-byte entries can contain a module name in the first 8 bytes; the last 4 bytes of each entry are reserved and should contain zeroes. The last entry must consist of all zeroes.

A typical entry for the CSECT may appear as follows:

```
DC CL8 'MODULENM'  
DC XL4 '00'
```

Programming Considerations

All clean-up routines of the resource manager use a standard interface, available through the IHARMPL mapping macro. On entry to the clean-up routines, the register contents are as follows:

Register	Contents
1	pointer to a 4-byte field that contains the address of the interface block
13	pointer to a standard save area
14	return address
15	entry point of clean-up routine
0,2-12	unpredictable

Registers 0-14 must be saved and restored by the clean-up routine; register 15 is used to pass a return code back to termination. A return code of 0 indicates a successful clean-up, and a return code of 4 indicates an unsuccessful clean-up.

The clean-up routines receive control on all task and address space terminations prior to any of the control program resource manager, and receive control in key 0, supervisor state, with no locks held. Each clean-up routine must acquire and release any locks it may need to do its processing.

For task termination, the clean-up routine executes under the TCB of the terminating task, and executes in the address space of the terminating task.

For address space termination, the clean-up routine executes under a task in the address space of the master scheduler. The clean-up routine will be able to examine the ASCB for the address space, queues, and other control blocks which reside in the common area; nothing in the private area for the terminating address space is accessible.

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 whenever 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 *Operators Library: System Commands* and *System Programming Library: 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 *Operator's Library: System Commands*, or *System Programming Library: TSO*. The descriptions in those publications also explain the various operands used in the following examples. For more information on designing effective SLIP traps, see *System Programming Library: Diagnostic Techniques*.

SLIP Command Examples

These examples suggest ways to use the SLIP command to set and modify SLIP traps.

Example 1: Obtaining a Dump with SQA Control Blocks

A control section within load module CHECKLOG experiences a program check and 0C4 abend at a location displaced 3E44 into the load module whenever user QCHECK6 logs on. As a result of this error, the system supplies only an abend dump. CHECKLOG, however, is a global routine with many of its control blocks in the system queue area (SQA). To get a dump that includes SQA control blocks, the system programmer therefore issues the following SLIP command:

```
SLIP SET, ID=CLOG, ACTION=SVCD, ERRYP=PROG, MODE=(TCB, PP, PKEY, EVERY),  
COMP=0C4, JOBNAME=QCHECK6, LPAMOD=(CHECKLOG, 3E48),  
SDATA=(SQA, RGN, TRT, SUM), END
```

ID	identifies this SLIP trap as "CLOG".
ACTION	specifies that an SVC dump is to be taken when the trap matches.
ERRYP	specifies a program check (PROG) as the error condition that must exist for the action specified by ACTION to occur.
MODE	specifies TCB, problem program, and problem program key as the mode in which the system must be.

COMP	specifies the applicable completion code (the 0C4 abend).
JOBNAME	identifies "QCHECK6" as the user whose logon produces the abend.
LPAMOD	names the program (CHECKLOG) and identifies the offset (3E48) indicated by the PSW at the time of the error.
SDATA	specifies the system information to be included in the SVC dump.
END	indicates the end of this SLIP command.

Note: Although the 0C4 abend results from the four-byte instruction at displacement 3E44 in the load module, the PSW at the time of the error indicates 3E48, the next sequential instruction. It is this latter displacement that the SLIP command must reference.

Example 2: Obtaining a Dump with Queue Elements and Control Blocks

The system programmer suspects an IBM error in the DEQ SVC routine because whenever program DVTRTN executes, it abnormally terminates even though its parameter list is correct. The resulting abend dump does not include queue control blocks and queue elements. To get a dump that includes them, the system programmer issues the following SLIP command:

```
SLIP SET, ID=QELS, ACTION=SVCD, COMP=X30, ERRRTYP=ABEND, JSPGM=DVTRTN,
SDATA=( SQA, RGN, TRT, SUM ), END
```

ID	identifies this SLIP trap as "QELS".
ACTION	specifies that an SVC dump is to be taken when the trap matches.
COMP	specifies the applicable system completion codes.
ERRRTYP	specifies an abend as the error condition that must exist for the action specified by ACTION to occur.
JSPGM	identifies "DVTRTN" as the job step program that must be executing for this error interception.
SDATA	specifies the system information to be included in the SVC dump.
END	indicates the end of this SLIP command.

Example 3: Replacing an SVC Dump with a Standalone Dump

The system programmer suspects that a system routine's FRR makes unauthorized use of system real storage space ("low core") whenever a machine check occurs during its processing. Even though the intent of the FRR is to perform recovery cleanup and to free the CMS lock, the use of low core necessitates a re-IPL of the system. To get a meaningful dump (the eventual SVC dump is of little relevance to this problem), the system programmer issues the following SLIP command:

```
SLIP SET, ACTION=WAIT, ERRRTYP=MACH, MODE=( GLOCS, RECV, EVERY ),
ADDRESS=( 2E048, 2E25C ), END
```

ACTION	indicates that the desired error action is for the system to be placed in a wait state. The system can be restarted by pressing the RESTART key.
ERRRTYP	specifies a machine (MACH) check as the error condition that must exist for the action specified in ACTION to occur.
MODE	specifies that the system mode at error interception is to be the holding of a global suspend lock (GLOCS), and a recovery routine (RECV) is in control.

ADDRESS specifies the range of addresses in which the error is to occur as 2E048 - 2E25C.

END indicates the end of this SLIP command.

Example 4: Setting SLIP Traps for an Application Program

An installation has just installed an application program named **BANKER**. Because the application is new, the installation expects a reasonable period of debugging. To intercept the expected errors, the installation system programmer adds the following SLIP command to parmlib member **COMMNDxx**:

```
SLIP SET,PVTMOD=BANKER,END
```

PVTMOD identifies "BANKER" as the private area load module that must be in control when the error occurs; **END** denotes the end of this SLIP command. This command produces an SVC dump for any errors in **BANKER**.

Assume, however, that **BANKER** uses **ESTAE** to check invalid parameter input. This causes frequent program checks and then subsequent retries to indicate invalid parameters to the user. To eliminate the dumps resulting from invalid parameter input, the system programmer enters the following additional command:

```
SLIP SET,ACTION=IGNORE,PVTMOD=(BANKER,0,4C),ERRTYP=PROG,END
```

This command causes program checks that occur within offsets in the range 0 - 4C in **BANKER** (the part that checks parameter input) to be ignored, while still permitting abends in the rest of **BANKER** to cause SVC dumps. **ERRTYP** identifies program check interruptions as the kind of errors to be intercepted, and **END** indicates the end of this SLIP command.

Example 5: Setting a SLIP Trap Using the Instruction Fetch PER Event

A system programmer suspects that a particular routine is abending either because invalid parameters are being passed or because there is an error in the program. In order to obtain information about the problem, the system programmer issues the following SLIP command:

```
SLIP SET,IF,ENABLE,ACTION=SVCD,RANGE=CD3100,END
```

The routine still abends but the SVC dump taken upon entry to the abending routine proves whether or not the parameters are valid. The parameters used in the SLIP trap are as follows:

IF specifies that an instruction fetch **PER** trap is set.

ENABLE specifies that the SLIP trap is enabled or eligible for checking.

ACTION specifies that a SVC dump is to be taken when the trap matches.

RANGE specifies the virtual storage location (CD3100) to be monitored.

END indicates the end of this SLIP command.

Because the default **MATCHLIM** and **PRCNTLIM** operands are taken, the trap is automatically disabled after the SVC dump is scheduled or if **PER** checking requires more than 10% of the system.

Example 6: Setting a SLIP Trap Using the Storage Alternation PER Event

Location **CD3010** contains an address and is normally modified by many programs. Occasionally it is being set to zero thus, causing system problems. In order to obtain information about the error, the system programmer sets the following SLIP trap:

```
SLIP SET,SA,ENABLE,ACTION=SVCD,RANGE=(CD3100,CD3013)  
DATA=(CD3010,EQ,00000000),END
```

When locations CD3010 through CD3013 are modified to zero, an SVC dump is scheduled. The parameters used in the SLIP trap are:

SA specifies that a storage alternation PER trap is set.
ENABLE specifies that the trap is eligible for checking.
ACTION specifies that a SVC dump is to be taken when the trap matches.
RANGE specifies a four byte storage range to be monitored for storage alteration.
DATA specifies that the four bytes of storage beginning at location CD3010 must be zero.
END indicates the end of this SLIP command.

Because the default MATCHLIM and PRCNTLIM operands are taken, the trap is automatically disabled after the SVC dump is scheduled or if PER checking requires more than 10% of the system.

Example 7: Setting a SLIP Trap Using the Storage Alteration PER Event and the ACTION = IGNORE Option

A system programmer suspects that location CD3010 (in LPA), which is normally modified by JES (ASID=2), is being modified by some other program. The system programmer issues the following three commands:

```
SLIP SET,SA,DISABLE,ACTION=SVCD,RANGE=CD3010,END  
SLIP SET,SA,DISABLE,ACTION=IGNORE,ASID=( 2 ),END  
SLIP MOD,ENABLE,ALL
```

Because SLIP traps default to enabled when entered, the first and second commands use the DISABLE keyword to allow a second trap to be created before the first trap is enabled. Because SLIP traps are processed in a last-in-first-out (LIFO) order, the second trap is checked first for a match on any storage alteration PER event. Any modifications to location CD3010 by JES (ASID=2) are ignored; any other modifications cause a SVC dump to be taken. The parameters used in the SLIP traps are:

SA specifies that a storage alternation PER trap is set.
DISABLE specifies that the SLIP trap is not enabled when entered in the system. A SLIP MOD command enables the trap.
ACTION specifies what action is to take place when the trap matches.
RANGE specifies the virtual storage location (CD3010) to be monitored.
ASID specifies the address space that must be in control.
END indicates the end of this SLIP command.

Because the default MATCHLIM and PRCNTLIM operands are taken, the first trap is automatically disabled after the SVC dump is scheduled or if PER checking requires more than 10% of the system. Because the action specified is IGNORE, the second trap is not automatically disabled after the SVC dump is taken.

Example 8: Setting a SLIP Trap Using the Successful Branch PER Event

A system programmer wants to collect trace records of the path taken in module MOD01 starting at X'108' through X'4FC' during execution of JOBX. In order to collect trace records, GTF must be active with the trace option for SLIP (TRACE=SLIP) in effect. The system programmer then issues the following SLIP command:

```
SLIP SET,SB,ENABLE,LPAMOD=(MOD01,108,4FC),JOBNAME=JOBX,  
ACTION=TRDUMP,MATCHLIM=20,END
```

Each time the trap is matched, a GTF trace record is written. After 20 matches (MATCHLIM=20), the trap is disabled and a SVC dump is scheduled. The parameters used in the SLIP command are as follows:

SB	specifies that a successful branch PER trap is set.
ENABLE	specifies that the trap is eligible for checking.
LPAMOD	specifies the virtual storage range to be monitored for successful branches (starting at offset X'108' from the beginning of the module and ending at offset X'4FC').
JOBNAME	specifies that the monitoring takes place in the address space in which JOBX executes.
ACTION	specifies that trace records for each event are collected when a match occurs. A SVC dump is scheduled when the trap is disabled.
MATCHLIM	specifies that the trap is to be disabled after matching 20 times.
END	indicates the end of this SLIP command.

Because the default is taken for the PRCNTLIM operand, the trap is automatically disabled if PER checking requires more than 10% of the system.

Dumping Virtual Storage

Any user can issue the SNAP and ABEND macro instructions to request dumps. (These macro instructions are described in *OS/VS2 Supervisor Services and Macro Instructions*.) In addition, the system programmer can also use the SDUMP macro instruction to provide dumps of virtual storage, and the CHNGDUMP command to influence the contents of the dump.

Using the SDUMP Macro Instruction

The SDUMP macro instruction can be used by system routines to provide fast unformatted dumps of virtual storage. SDUMP invokes SVC DUMP to provide the services. Only one SVC DUMP may be taken in the system at any one time.

Issuers of SVC Dump with entry by SVC must be authorized via APF or have a control program key. SVC Dump can be initiated by a branch entry (BRANCH=YES specified on SDUMP) for callers who cannot issue an SVC. The branch entry caller must be key 0, supervisor state, and must be in SRB mode, or own a lock, or be disabled (with a supervisor bit on in the PSASUPER field of the Prefixed Save Area). The branch entry interface uses standard linkage conventions. Branch entry callers must include the CVT mapping macro instruction with the PREFIX=YES parameter.

The SVC Dump routine can schedule a dump in the address space specified by the ASID parameter or in the address spaces specified by the ASIDLST parameter of SDUMP. SDUMP issuers using the ASID or ASIDLST parameters must ensure that the areas referenced by the SDUMP parameter list are addressable from the address spaces in which the SVC Dump routine will execute. The LIST parameter should be used on the SDUMP macro instruction instead of the STORAGE parameter by programs specifying ASID or ASIDLST and not executing in a common area of storage. Also, the storage list addressed by the LIST parameter must be in a common area of storage for these programs.

SQA Buffer

A 4K buffer is reserved in the system queue area for the callers of SVC DUMP. A user can reserve the buffer 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 buffer is pointed to from the CVTSDBF field of the CVT. 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, it must be assumed that a dump is in progress and the caller must continue processing as if a dump could not be taken. If the bit is not set, the caller must set the bit before filling the buffer and calling SVC DUMP.

The first word in the buffer is the actual virtual storage address of the data. The next halfword is the length of the data. A copy of the data follows this 6-byte descriptor field. More 6-byte fields and data may then be specified in this buffer. If the entire buffer is not filled, the last data area must be followed by a 6-byte zero descriptor field to indicate the end of meaningful data.

An address field of X'FFFFFFFF' indicates that the information that follows is in key-length-data format, mapped by the IHAVRA mapping macro. The length field gives the total length of the formatted information.

Using the CHNGDUMP Command

The CHNGDUMP operator command can be used to alter the system dump options for SYSABEND, SYSMDUMP, SYSUDUMP, and SDUMP. The options are changed by updating the option lists in the RTCT (Recovery/Termination Control Table). The system dump options can either add to or override the options supplied with the ABEND, CALLRTM, SDUMP, SETRP, and SNAP macro instructions. For a complete description of the CHNGDUMP command refer to *Operator's Library: OS/VS2 MVS System Commands*.

In order to tailor the dump, the dump option merging may take place in several successive stages. Suppose an error occurred at a low program level, and the recovery exit at that level specified certain dump options on a DUMPOPT parameter. Suppose also that there are two other recovery exits between the lowest level exit and the recovery exit that actually precipitates the dump. One of these two exits also specifies certain dump options on a DUMPOPT parameter.

As the recovery effort percolates up toward the top level recovery exit involved, all the dump options from all the exits specifying them are merged with the parmlib options to produce the combination of options for the dump the system actually takes.

Getting More Than One SYSMDUMP

MVS 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 OS/VS2 MVS JCL 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 cataloged, the system writes an EOF 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 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 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 "Writing a WTO/WTOR Exit Routine" and "Programming Conventions for WTO/WTOR Routines" in the "Miscellaneous Services" section later in Part I.)

SNAP/ABEND User Exits

A user exit from SNAP/ABEND DUMP allows an installation to gather and format information to be included in a SNAP/ABEND dump to a data set described by a SYSABEND or a SYSUDUMP DD statement. The exits receive control automatically during the control block formatting phase of every SNAP and ABEND dump for which the CB option was requested. The exit provides an area in which the installation's routine can build a print line and the address of an IBM-supplied print routine to which the installation routine can pass the line to be printed.

The user's formatting routines must conform to or establish the following conditions:

1. The installation must link-edit the exit routine into either SYS1.LPALIB or a LINKLST library. Only one exit routine can be link-edited into SYS1.LPALIB with the load module name of IGC0905A. The IBM supplied exit routine, IEAVADUS, can be replaced with an installation written routine. IGC0905A can be link-edited to include the new IEAVADUS. The SNAP/ABEND dump routine CSECT IEAVAD08 of load module IGC0005A loads IGC0905A, and then branches and links to it.

If a list of exits is used instead, then one or more exit routines can be link-edited into SYS1.LPALIB or into a LINKLST library with any load module name. The installation must also add the load module's name to the IEAVADFM list of user exits. IEAVADFM is a CSECT in load module IGC0005A, and each entry is eight characters, padded to the right with blanks. If an entry is to be ignored, then the installation must change it to eight blanks. The end of the list is indicated by four bytes of hexadecimal zeros. The IBM-supplied version of IEAVADFM has nine words of hexadecimal zeros.

2. Entry Specifications:

- a. The exit receives control in protection key zero, supervisor state, with no locks held.
- b. The following register contents are available to the user's routine on entry:
 1. Register 1 - points to a parameter list (see below)
 2. Register 13 - points to a standard 18 word save area
 3. Register 14 - address to which the formatting routine should return control
 4. Register 15 - formatting routine entry point address
 5. Registers 2-12 are irrelevant
3. The parameter list pointed to by general register 1 contains the following information that is useful for user formatting routines:

OFFSET IN HEX	FIELD DESCRIPTION
+ 0	Fullword address of the TCB for task being displayed.
+ 4	Halfword address space identifier.
+ 6	One byte number denoting the subpool containing the save area.
+ 7	One byte flag field.
+ 8	Fullword address of the 120 character output buffer for print lines.
+ C	Fullword address of the print routine (IEAVAD81).
+10	Fullword address of the CVT.
+14	Address of memory access service routine.
+18	Address of format service routine.
+1C	Four fullwords that can be used by the user's formatting routine.
+2C	Two fullwords for output from the format service routine.

Note: This parameter list is mapped by the IHAABDPL system mapping macro. The parameter list includes all the fields of the Print Dump Service aid's parameter list so user formatting routines can be invoked by either SNAP/ABEND or Print Dump. The memory access service routine is a no-operation routine for SNAP/ABEND. (It always returns a successful return code and the same address that is was supplied.)

4. Operation:

The user's formatting routine should build one print line at a time in the buffer provided, and should use BALR to branch to the IBM-supplied print routine, which in turn prints the line on the dump data set. Offsets are recommended for all formatted control blocks that are longer than one output line. (One line generally formats 20 hexadecimal characters.) The print routine saves registers, prints the line, blanks the buffer, restores the registers, and returns control to the user's routine via register 14. No registers are necessary as input to this routine.

Because it works through the IBM-supplied print routine, the formatting routine has no direct access to the carriage controls. In order to cause a skipped line in the dump output, it is therefore necessary to pass a blank buffer to the print routine. Similarly, the print routine also handles page ejects. The user exit routine can use format patterns to format data in the output buffer, by interfacing with the IBM-supplied format service routine. The service routine can also convert data to printable hexadecimal. This service routine is the same routine as is provided by the PRDMP Service Aid, and the details of its interface are given in the Appendix C of the OS/VS2 MVS System Programming Library: Service Aids.

When processing is complete, the exit routine should restore SNAP's registers and return via a branch to the address provided in register 14 on entry. Register 15 should contain a

zero return code. A return code of 12 suppresses the remainder of the abend dump. If the same exit routine is executed under the AMDPRDMP Service Aid, you should be aware that the print dump does not suppress the remainder of the dump for a return code of 12. The return should be made in protection key 0, supervisor state, with no locks held (the same state as when entered). For an example of a formatting routine see module IEAVTFMT.

5. The routines must be reentrant.

Notes:

- a. The user's routines must not free either the entry parameter list or the print buffer in order to avoid an abnormal termination later in the SNAP/ABEND routine.
- b. Recovery - SNAP is covered by an ESTAE exit. Each user formatting routine should set up its own recovery to handle any abends encountered during the formatting process. This routine should either recover and continue formatting, or recover and return to SNAP with a zero return code. A non-zero return code is interpreted as a GETMAIN failure, resulting in the following message in the dump data set:

```
USER/PP CONTROL BLOCKS UNAVAILABLE
```

The dump is truncated due to lack of storage. The recovery routine should not continue formatting if an x37 abend occurs, since no space remains in the dump data set. Before the recovery routine returns to SNAP it should free all the storage that it has obtained.

If the formatting routine does not establish recovery, or if the recovery exit specifies continue-with-termination after an abend, SNAP terminates this control block formatter entirely and continues with the next portion of the dump, if any.

- c. Interface to the print routine (IEAVAD81):

Entry: via BALR 14,15 for each line to be written.

State: protection key 0, supervisor state, no locks held.

Registers on entry: register 1 - points to the parameter list
register 13 - points to the save area
register 14 - points to the return address
register 15 - points to the IEAVAD81 routine.

Registers saved and restored: 14 through 12.

Return: via register 14.

- d. To remove a formatting routine from the system, link-edit a copy of module IEFBR14 into SYS1.LPALIB with the name IGC0905A, or remove the entry from the IEAVADFM list of user exits by replacing the entry with eight blanks.

Correlating Diagnostic Material

The RTM correlates diagnostic material by providing error identifiers on related pieces of material. The error identifier (ID) appears in system message IEA911A and is part of the SVC dump header information. The RTM also adds it to a SYS1.LOGREC record.

The IEA911A message text is:

```
COMPLETE          SYS1.DUMPnn
PARTIAL           DUMP ON
FOR              ASID (xx,xx,....)
[ERRORID=SEQyyyyy CPUzz ASIDaaaa TIMEhh.mm.ss.t]
```

The meanings for the lowercase alphabetic symbols in the message text are as follows:

nn	qualifies the direct access data set name
ddd	indicates the device address containing a tape volume
(xx,xx,...)	indicates the ASID(s) of the address space(s) included in the SVC Dump
yyyyy	represents the sequence number portion of the error ID
zz	represents the CPU identifier portion of the error ID
aaaa	represents the address space identifier portion of the error ID
hh.mm.ss.t	represents the time (hours, minutes, seconds, and tenths of seconds) portion of the error ID.

The AMDPRDMP service aid puts the error ID from the RTM on the title page of SVC dumps in the following format:

```
ERRORID FOR THIS DUMP = SEQyyyyy CPUzz ASIDaaaa TIMEhh.mm.ss.t
```

The explanations of the SEQ, CPU, ASID, and TIME elements of this character sequence are the same as for the IEA911A message text.

If there is no error ID available for a given dump, AMDPRDMP prints the following message where it would ordinarily print the error ID, as explained above:

```
NO ERRORID ASSOCIATED WITH THIS DUMP
```

An illustration of how these messages look on the dump's title page appears in OS/VS2 System Programming Library: Service Aids, GC28-0674.

When an error ID is available, the RTM also adds it to the SYS1.LOGREC record, either for hard machine checks or for software errors. The format of the record's error ID is:

```
ERRORID = SEQyyyyy CPUzz ASIDaaaa TIMEhh.mm.ss.t
```

The fields are as previously explained for message IEA911A. When there is no error ID to associate with the LOGREC record, the record contains the following message in place of the error ID:

```
NO ERRORID ASSOCIATED WITH THIS RECORD
```

Illustrations of the error ID in the SYS1.LOGREC record appear in OS/VS2 System Programming Library: SYS1.LOGREC Error Recording, GC28-0677.

The RTM supplies these error IDs and their respective time stamps under the following circumstances:

- It generates an entire new error ID when:
 - The RTM receives control as a result of a machine check.
 - RTM1¹ receives control in SLIH mode² and:
 - 1) the entry is non-recursive
 - 2) no processing is in progress on behalf of RTM2.
 - There is a direct entry to RTM2.
- It supplies only a new time stamp when:
 - RTM1 receives control in SLIH mode and the error ID is recursive³.
 - RTM1 receives control in SLIH mode and RTM2 processing is in progress.
 - RTM2 receives control recursively for abnormal termination.
 - RTM2 receives control from RTM1 as a result of a machine check or software error in which RTM1 has performed some processing due to the mode of the system at the time of the error (non-task mode).

Notes:

1. RTM1 and RTM2 are the two major parts of the RTM. In general, RTM1 processes SLIH requests and CALLRTM macro calls, while RTM2 is an SVC 13 processor (hence, is task-oriented).
2. SLIH mode is an entry to the RTM from a FLIH for such errors as program checks, machine checks, SVC errors, etc.
3. Recursive indicates a re-entry into the RTM while processing a previous error.

For those cases in which the RTM supplies a new time stamp, the sequence number correlates the original error with subsequent errors, and the new time stamp distinguishes unique errors within the recovery path. For example, if an error occurs during the recovery of a previous error, the RTM1 records an ID for the new error on SYS1.LOGREC with the same sequence number as the original error, but with a different time stamp.

Not all errors detected in the RTM cause recording to SYS1.LOGREC and/or SVC dumps. Because of this, the sequence numbers in the error ID may not follow exact sequential order. The RTM skips a number if, for example, it receives control for an apparent error that is retried by an FRR or ESTAE exit routine with no SYS1.LOGREC record or SVC dump. This recovery technique is used by some system service routines to check the validity of parameters. A number is skipped if the recovery routine does not request recording.

Establishing a Timer Disabled Interrupt Exit

Timer supervision provides a function called SETDIE 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 SETDIE function is available only to programs executing in supervisor state with protection key zero. SETDIE allows users to initiate a real time interval by branch entry to the user-written DIE routine. When the time interval expires, the user's DIE routine gains control as an extension of the timer second level interrupt handler. 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.

Using SETDIE

The caller of SETDIE can be executing in either task control block (TCB) or service request block (SRB) mode, but must be in protection key zero and supervisor state. The entry point to SETDIE 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 SETDIE must provide the following input environment:

- 1) Register 1 must contain the address of a user-supplied task queue element (TQE) whose fields are available from the IHATQE mapping macro. This user TQE must:
 - be a contiguous block of 128 bytes
 - reside in SQA
 - include the following field initialization:

TQE AID -- Zero or a valid ASID. This is important in case of an address space failure (see "Obtaining and Freeing the TQE").

TQE VAL -- the desired real time interval (a 64 bit unsigned binary number with bit 51=1 microsecond).

TQE EXIT -- address of the user's DIE.

- 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 registers 0 and 11-13 occurs upon return from SETDIE. 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: SETDIE obtains the dispatcher lock if the caller does not already hold it. After completing its processing, SETDIE releases the lock if the caller did not previously hold it.

Recovery Considerations

SETDIE does not establish its own recovery routine. Any system program calling SETDIE should have its own FRR or ESTAE routine. A program check occurs in SETDIE if the caller is not both in protection key zero and in supervisor state.

The DIE routine executing out of the timer SLIH gains control under timer supervision's FRR. When it gets control, this FRR tries to repair any damage to the system's real time queue and then relinquishes control -- "percolates" -to the next higher level in the recovery hierarchy. The DIE itself can optionally establish its own FRR, which should also terminate by percolation to let the timer supervision FRR gain control. If the DIE routine causes a program check and has no FRR, the external FLIH abends the current task. The current task might or might not be the one related to the DIE. The task abends with a 3FC system code. IMS, CICS, and most other system tasks cannot handle this situation in their ESTAE exits. The same thing can occur if there is an FRR, and the FRR percolates instead of retrying after the program check.

DIE Characteristics

Entry to the DIE routine is in supervisor state, with protection 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 SETDIE.
- Register 2-12 are as they were upon entry to the SETDIE routine (or as changed by a previous DIE entry -- see "DIE Execution").
- Register 14 contains the return address.
- Register 15 contains the entry point of the DIE routine.

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 protection key zero, and disabled. All locks obtained by the routine must be released, but no registers need be saved.

DIE Execution: must be like the execution of an interrupt handler because it executes as an extension of the timer SLIH. Specifically, DIE executes under the following restrictions:

- 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.
- DIE cannot reference any private storage areas.
- DIE must execute disabled. Hence it cannot cause a page fault.
- DIE cannot request the local lock or the CMS lock because these are suspend locks and may therefore already be in use. Furthermore, the DIE routine cannot assume whether or not these locks are held upon entry.
- 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. The TQE fields must have remained unchanged except for the following fields:

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".

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.

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 SETDIE function is similar to the TQE ENQUEUE function, the routines differ in the following respects:

- Although TQE ENQUEUE expects an already established and completed TQE as input, SETDIE 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 SETDIE, it must be set to the desired interval. SETDIE then converts it to the proper clock comparator value.
- TQE ENQUEUE assumes that the clocks are functioning correctly. SETDIE must use the clocks directly and therefore verifies (rather than assumes) that the clocks are functioning correctly. SETDIE is therefore capable of advantageously using alternate clocks in a multiprocessing environment in which one or more clocks have failed.

Task Queue Element Control

The major aspects of controlling the task queue element (TQE) associated with the user's DIE routine are:

- Obtaining and freeing the TQE
- Serializing the usage of each TQE
- Time-of-day clock failure
- Interval cancellation

Descriptions of each of these aspects follow.

Obtaining and Freeing the TQE: is the responsibility of the user of the SETDIE function because the TQE resides in SQA. Thus, the user 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 the user for a failing address space only if the TQE is enqueued on the real time queue and has field TQEAID set to the ASID of the failing address space.

Before freeing the TQE, however, the user must ensure that it is not currently on the real time queue. The user can accomplish this in several ways:

- 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) The TQE must not have been altered (other than in the fields previously described) by the user.

2) The interface for the TQE DEQUEUE routine is described in the section "Interval Cancellation".

Serializing the Usage of Each TQE: is also the user's responsibility. This included the execution of the SETDIE, TQE ENQUEUE, and TQE DEQUEUE routines for a given TQE because these routines update the supplied TQE. The user must 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 protection 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.

Virtual Storage Management

GETMAIN and FREEMAIN respectively allocate and free one or more areas of virtual storage. The KEY parameter allows a user executing in protection key zero to specify the protection key for storage he requests, and subsequently, the CHANGKEY macro instruction allows the user to alter that protection key. Most of the functions of GETMAIN and FREEMAIN are available to all users, but the entire function of CHANGKEY and some of the GETMAIN and FREEMAIN functions are available only to programs executing in supervisor state under protection key zero.

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, as explained below. Although the branch entries require the user to do more work, they are significantly more efficient than the SVC entries and do save some system overhead. There is no SVC entry to CHANGKEY, however. Therefore the user must invoke its function via a branch entry.

The BRANCH Parameter

Branch entry to the GETMAIN or FREEMAIN macro instructions is accomplished by specifying BRANCH=YES on the macro instructions. If the BRANCH 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 prior to entry. (*Note:* If the BRANCH parameter is not used, it is still necessary for the current branch entry user of the macro instruction to alter his code to include the preloading of the ASCB address in register 7, and to hold the local lock.)

An additional branch entry point (GLBRANCH) 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). It is necessary, however, to hold the SALLOC lock before entering the routine. Although the TCB address and ASCB address are not required for this entry, register 4 must be loaded with the address of the global save area pointed to by the CVT; this is done by the macro expansion.

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

Branch entry to the CHANGKEY macro instruction is also accomplished by specifying BRANCH=YES on the macro instruction. However, the only entry requirement is that the caller must hold the local lock prior to entry.

The KEY Parameter

Since branch entry users are required to be executing in 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 new subpools - 227, 228, 229, 230, 231, and 241. These subpools allow both global and local storage to be obtained in the requestor's storage protection 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.

Real Storage Management

The real storage manager (RSM) administers the use of real storage and directs the movement of virtual pages between auxiliary storage and real storage in page size (4096 bytes) blocks. It makes all addressable virtual storage in each address space appear as real storage. Only virtual pages necessary for program execution are kept in real storage. The remainder reside on auxiliary storage. RSM employs the auxiliary storage manager (ASM) of the data manager 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 I/O 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 page frames upon request from a pool of available frames, thereby associating virtual addresses with real storage addresses. Frames are repossessed upon termination of use, 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 specified as a system page that must be resident in real storage. 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 paging services provided include the following:

- PGFIX -- Fix virtual storage contents.
- PGFREE -- Free virtual storage contents.
- PGLOAD -- Load virtual storage areas into real storage.
- PGOUT -- Page out virtual storage areas from real storage.
- PGRLSE -- Release virtual storage contents.

The PGFIX and PGFREE functions are available only to authorized system functions and users. The PGLOAD, PGOUT, and PGRLSE are not restricted and are available to all users; these functions are described in *OS/VS2 Supervisor Services and Macro Instructions*.

Fixing/Freeing Virtual Storage Contents

The PGFIX and PGFREE macro instructions provide complementary functions. The PGFIX macro instruction makes specified storage areas resident in real storage and ineligible for page-out as long as the virtual address space of the requesting TCB remains in real storage. The PGFREE macro instruction makes specified storage areas, which were previously fixed via the PGFIX macro instruction, eligible for page-out. Real frames fixed by PGFIX are not considered pageable until the same number of PGFREE and PGFIX requests have been issued for any virtual area.

Page fixing ties up valuable real storage and is usually detrimental to system performance unless utilization of the resources is extremely high.

In the PGFIX function, you have the option of specifying the relative real time duration anticipated for the fix. If you specify LONG=Y, 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 may be required to avoid an assignment of a frame to the V=R area or an area that might be varied offline. If you specify LONG=N, the time duration of the fix is assumed to be relatively short. A long-term PGFIX is assumed if this option is not specified.

In both the PGFIX and PGFREE 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, you specify RELEASE=Y; otherwise, you specify RELEASE=N. If you specify PGFIX with RELEASE=Y, the PGRLSE function is performed before the PGFIX function.

If you specify PGFREE with RELEASE=Y, the PGFREE 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.

Note: PGFIX does not prevent pages from being paged out when an entire virtual address space is swapped out of real storage. Consequently, the user of PGFIX cannot assume a constant real address mapping for fixed virtual areas in most cases.

Completion Considerations

Under normal circumstances, you can reverse the effect of a PGFIX via a PGFREE when the need for a PGFIX ceases. However, a PGFIX request sometimes completes asynchronously if it requires a page-in operation. In such cases, it may be necessary to explicitly purge PGFIX operations.

For this reason, the PGFIX 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. The ECB is posted when all requested pages are fixed in real storage.

Explicit purging of a PGFIX is carried out in one of two ways:

- If the PGFIX is known to be complete, the PGFIX is reversed through the usual PGFREE function.
- If there is any possibility that the PGFIX has not been posted as complete, the PGFREE should be issued with an ECB address supplied. This ECB parameter identifies the event control block that was supplied as an input parameter with the PGFIX being purged. Note that for the purpose of canceling a PGFIX request that has not yet completed, the ECB must uniquely identify the PGFIX request. Consequently, to provide for explicit purging, you must ensure that the ECB address for any incomplete PGFIX can be located in a purge situation, and that the ECB has not been reused at the time the PGFIX is to be canceled.

The PGFREE function always completes immediately and requires no ECB address except for purging considerations.

Virtual Subarea List (VSL)

The virtual subarea list provides the basic input to the page services functions: PGFIX, PGFREE, PGLOAD, PGRLSE, and PGOUT. The list consists of one or more doubleword entries; each entry is on a fullword boundary and describes an area of virtual storage. The list must be non-pageable (for example, in SQA or LSQA) and contained in the address space to be processed.

Note: If the list specifies a PGFREE of the page in which the list resides, the fix count for the page must be greater than one.

Each parameter list entry has the format shown in Figure 18.

Byte	0	1	2	3	4	5	6	7
	FLAGS	START	ADDRESS		FLAGS	END	ADDRESS	+ 1

Figure 18. Virtual Subarea List Entries

Byte 0 Flags:

Bit 0	(1...)	This bit indicates that bytes 1-3 are a chain pointer to the next VSL entry to be processed; bytes 4-7 are ignored, but the checking of this bit is subject to the setting of byte 4, bit 1. This feature allows several parameter lists to be chained as a single logical parameter list.
Bit 1	(.1..)	PGFIX is to be performed; reserved, set by macro instruction.
Bit 2	(..1.)	PGFREE is to be performed; reserved, set by macro instruction.
Bit 3	(...1)	PGLOAD is to be performed; reserved, set by macro instruction.
Bit 4	(.... 1...)	PGRLSE is to be performed; reserved, set by macro instruction.
Bit 5	(.... .1..)	Reserved.
Bit 6	(.... ..1.)	Long-term PGFIX is to be performed; reserved, set by macro instruction.
Bit 7	(.... ...1)	Reserved.

Start Address:

The virtual address of the origin of the virtual area to be processed.

Byte 4 Flags:

Bit 0	(1...)	This flag indicates the last entry of the list. It is set in the last doubleword entry in the list.
Bit 1	(.1..)	When this flag is set, the entry in which it is set is ignored. This bit takes precedence over byte 0, bit 0.
Bit 2	(..1.)	Reserved.
Bit 3	(...1)	This flag indicates that a return code of 4 was issued from a page service function other than PGRLSE.
Bit 4	(.... 1...)	Reserved.
Bit 5	(.... .1..)	PGOUT is to be performed; reserved, set by macro instruction.
Bit 6	(.... ..1.)	KEEPREAL option of PGOUT is to be performed; reserved, set by macro instruction.
Bit 7	(.... ...1)	Reserved.

End Address + 1:

The virtual address of the byte immediately following the end of the virtual area.

Reconfiguration Using Vary Storage

Vary storage permits users of multiprocessing IBM System/370 Models 158 and 168 to more easily specify reconfigurable storage units for use with the other CPU as a uniprocessor. To identify the storage units that the operating system attempts to preserve for reconfiguration, define the reconfigurable storage unit parameter (RSU=nn) in the IEASYSxx member, where nn is the number of storage units you want to take offline for the reconfiguration.

Note: If power may be turned off for one processor, all of the high address storage units, except the highest unit, must be assigned to that processor. For example, in a system with eight storage units (0-7), units 3-6 should be attached to the processor that may have power turned off.

The RSU=nn parameter defines storage units which contain pages that can be moved in the event of an off-line vary. The system attempts to preserve the frames in these storage units. These frames are known as the non-preferred area. Frames used for long term resident pages are clustered in a preferred area. These long-term pages include SQA, non-swappable LSQA, common area fixed pages, and non-swappable private fixed pages.

LSQA and private area fixed pages are called non-swappable if they belong to a non-swappable address space. An address space is non-swappable if:

- The program name is in the Program Properties Table with the appropriate flags set. For further details see "Assigning Special Program Properties" in *OS/VS2 System Programming Library: Job Management, GC28-0627*.
- ADDRSPC=REAL is specified on the JOB or EXEC JCL statements.
- The address space issues the SYSEVENT TRANSWAP. The use of SYSEVENT DONTSWAP also makes an address space non-swappable; however, the address space is considered to be non-swappable for only a short duration, and associated LSQA and private area pages are not necessarily put into preferred storage.

Figure 19 indicates that when the system is IPLed, the nucleus and V=R areas are assigned to the low-end of processor storage and the SQA is assigned to the high-end of processor storage. The operating system then defines the low and high end processor storage units as preferred areas (in this case 0, 1, and 7). Next, the non-preferred area is defined as requested by the RSU=nn specification and coinciding with the next available high-end processor storage unit (6). Offline processor storage units (5 and 2) are spanned at this time, but omitted. If the RSU=nn parameter value exceeds the available processor storage units, then all units are defined as non-preferred except those previously assigned as preferred at IPL for the nucleus, V=R and SQA.

During IPL, and after the non-preferred area has been defined, the operating system defines the remainder of the processor storage units as the preferred area, for long-term resident pages. In Figure 19, RSU=2 was specified, making processor storage units 6 and 4 non-preferred. Any processor storage units varied online after IPL are designated as non-preferred. The remaining processor storage unit (3) is defined preferred.

Processor Storage Units	Data Mapping	
7	Preferred (SQA at IPL)	②
6	Non-preferred	⑤
Offline 5	Unassigned	
4	Non-preferred	⑥
3	Preferred	⑦
Offline 2	Unassigned	
1	Preferred	④
0	V=R	③
	Nucleus	①

Will be non-preferred if units are brought online

Sequence of IPL assignments

Figure 19. IPL Designation of Processor Storage Units

While jobs are processing, short-term pages are assigned to any available processor storage frame in either non-preferred or preferred areas. Long-term pages are assigned only to processor storage frames in the preferred area. A condition may arise where a long-term page requires processor storage space and there are no preferred area frames available. If this situation occurs, one of the following happens:

- If a short-term page is using a frame in preferred storage, that short-term page is removed and the new long-term page is assigned to the vacated space. This is called immediate steal.
- If all of the frames in the preferred area are being used for long-term pages, the operating system looks for a frame in the non-preferred area. It will select a non $V=R$ frame if possible. When the system finds such a frame, if it is non $V=R$, the entire processor storage unit is converted from a non-preferred area to a preferred area. This is called dynamic expansion.

Whenever a non-preferred area is converted into a preferred area, the system operator cannot expect to place as many processor storage units offline as were originally designated by the $RSU=nn$ parameter. Therefore, the first time this conversion takes place, message IEA988I is issued to the system console notifying the operator that the preferred area has expanded and that reconfigurability may be impaired. The operator can then determine the processor storage units that are still being preserved for reconfiguration by using the DISPLAY MATRIX (D M) command.

Multiprocessing Configuration Considerations

On a 168 MP, and 158 MP Model 3 the processors and the storage (in two but not necessarily equal portions) are each separately powered. On the 158 MP Model 1, however, each processor and its half of the storage are powered by the same source. If a processor is powered down, its half of storage is also powered down. Therefore, in a 158 MP Model 1 system, if an installation intends to vary one processor (CPU 0) offline for maintenance, the user should assign both high and low address ranges to the other processor (CPU 1). This allows CPU 0 to be powered down without requiring an IPL.

Notes:

- The highest (preferred storage area) and lowest address ranges cannot be varied offline.
- When assigning high and low address ranges to the powered up processor, ensure (1) that the highest address range is large enough to include all the SQA and preferred storage area, and (2) that the lowest range is large enough to contain all of the nucleus.
- Since the 158 MP Model 3 uses the Alternate Power Down feature to supply power independently to both processors, and storage, its associated console and channels can be varied offline and powered on and off without an IPL. Its storage can remain online to the powered up processor.

If a 158 MP or 168 MP is configured as two MVS uniprocessors, both systems should be configured at the highest and lowest addresses for an MP system, to allow reconfiguration to an MP system without IPL. This is done automatically by specifying STORAGE=highest address in the CTRLPROG macro instruction during system generation. For example, in a 4-megabyte system, each uniprocessor should be configured at the following addresses:

0-512
512-1024
3072-3584
3584-4096

To do this, specify STORAGE=4,096,000. This leaves a 2-megabyte "hole" for reconfigured storage; to reconfigure to an MP system, you can fill this hole with the following address ranges by first adjusting the hardware configuration panel and then entering the proper VARY commands:

1024-1536
1536-2048
2048-2560
2560-3072

Note: The 158 MP Model 1 and 158 MP Model 3 are compatible and can be interconnected in an MP configuration. However, when these models are interconnected the Alternate Power Down and Asymmetric storage features on the 158 MP Model 3 are not functional.

If there is more than one console, one must be specified as the primary console. One reason is that the communications task must know which to use during IPL. Do not use the same address for both consoles. In general, no two devices should have the same address. If identical addresses are used, failure of one console is considered failure of both; with 3066 consoles, errors on one can result in alteration of the display on the other.

In a multiprocessing environment, the concurrent execution of tasks makes it difficult to predict the contents of dynamic fields in all system control blocks, (that is, fields whose contents change during processing from IPL time). Thus, all read-only and write-only references in user programs to these dynamic fields must be evaluated for their impact on the user's program and/or the system control program, and should be changed or eliminated accordingly. Otherwise, the effect (such as program abend or system failure) is dependent upon the user program's recovery scheme and/or system control program's use of that changed field.

Additional services are provided by the supervisor which do not appropriately fit into the previous chapters. These services, discussed in this chapter, include:

- Operator messages (WTO and WTOR macro instructions)
- User-written message routing exit routines
- Service management facilities (SCHEDULE and PURGEDQ macro instructions)
- Stage 1 exit effector (CIRB macro instruction)
- User-written SVC routines
- Missing interruption handler
- Power warning feature support

Writing Operator Messages

The WTO and WTOR macro instructions allow you to write messages to the operator. The WTOR macro instruction also allows you to request a reply from the operator. A complete description of the use of these two macro instructions is found in *OS/VS2 Supervisor Services and Macro Instructions*.

Routing the Message

The WTO and WTOR macro instructions have two special parameters, MSGTYP and MCSFLAG. The MSGTYP parameter specifies how the message is to be routed; the MCSFLAG parameter specifies that the macro expansion is to set bits in the MCSFLAG field as indicated by each name coded. Only programmers familiar with MCS should use these parameters, since using them improperly could impede the entire message routing scheme.

If MSGTYP=Y is specified, the message type specifies that two bytes are to be reserved in the WTO or WTOR macro expansion so that flags can be set to describe what MSGTYP functions are desired. If Y is specified, two bytes of zeros are to be included in the macro expansion at displacement WTO (or WTOR + 8) + 12 + the total length of the message text, descriptor code, and routing code fields. If MSGTYP=N is specified, or if the MSGTYP parameter is omitted, the two bytes are not needed and the message is routed as specified in the ROUTCDE parameter.

The bit definitions for MSGTYP=Y are:

Bit 0	MONITOR JOB NAMES
Bit 1	MONITOR STATUS
Bit 2-4	Reserved
Bit 5	MONITOR SESS
Bit 6-15	Reserved

When MSGTYP=Y is specified, the issuer of the WTO or WTOR macro instruction that contains the MSGTYP information must set the appropriate message identifier bit in the MSGTYP field of the macro expansion. The MCSFLAG field in the macro expansion has been set to zero, indicating that the MSGTYP field is to be used for the message routing criteria. When the message type is identified by the system, the message is routed to all consoles and TSO terminals in operator modes that had requested that particular type of information. If there are no consoles or terminals requesting that particular type of information, the 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.

Writing a Multiple-Line Message

The WTO macro instruction is used to write a multiple-line message to one or more operator consoles. System programs (supervisor state, protection key 0-7, or APF-authorized) may 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 use more than one WTO.

When using more than one request, the first WTO supplies the first 255 lines of the message. Other 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 put this identifier in the leftmost three bytes of register zero.

Message Routing Exit Routines

This topic provides detailed information on how to write user exit routines that modify the routing and descriptor codes of WTO or WTOR messages for the VS2 operating system. Information is provided on inserting this exit routine into the resident portion of the control program. In addition, a description of the characteristics and configuration of MCS is supplied.

Characteristics of MCS

The multiple console support (MCS) facility routes messages to different functional areas according to the type of information that the message contains. In MCS, a functional area is defined as one or more operator's consoles that are doing the same type of work. Some examples of functional areas are: (1) the tape pool area, (2) the disk pool area, and (3) the unit record pool area. Each WTO and WTOR macro instruction is assigned one or more routing codes that are used to determine the destination of the message. There are fifteen possible routing codes. When the message is ready, the routing codes assigned to the message are compared to the routing codes assigned to each console. If any of the routing codes match, the message is sent to that console.

If the standard routing codes provided on application and system messages do not cover special situations at an installation, the routing codes can be modified by coding a user exit routine. The exit routine receives control before the routing of messages so users can examine the message text and modify the message's routing and descriptor codes. The system uses the modified routing codes to route the message. Descriptor codes explained later in this chapter provide a mechanism for message presentation and deletion.

Automatic console switching occurs when permanent hardware errors are detected. Command-initiated console switching is provided to permit restructuring of the system console configuration and the hard copy log by system operators. Consoles can be moved into or out of functional areas at any time during system operation.

A hard copy log records messages, operator and system commands, and operator and system responses to commands. The hard copy log can be a console device or it can be the system log (SYSLOG). The number and type of messages recorded on the log is optional. The installation can record a selected group of messages, or it can record all messages. If commands are recorded, the system automatically records command responses.

Whenever possible, the hardcopy function should be delegated to an output-only device (such as a printer) or to the system log.

Programming Conventions for WTO/WTOR Routines

The programming conventions for the WTO/WTOR exit routines are summarized below:

- Exit routine is part of the resident control program. The program should be loaded on a page boundary.
- Exit routine is any size.
- Exit routine may allow interruptions. The routine receives control with no locks held; it should return control with no locks held.
- Exit routine is reenterable and serially reusable. Macro instructions whose expansions store information into an online parameter list should not be used.
- IEFCVXIT is name of routine.
- Registers must be saved at entry and restored prior to returning.
- Exit routine may issue WAIT, XCTL, WTO, or WTOR macro instructions.

Note: The WAIT macro instruction should not be used when the exit routine is entered under the console communications task. Doing so permanently terminates console communications.

- Exit routine is part of WTO SVC. If the exit routine terminates abnormally, the WTO request is terminated.
- Exit from the routine is via the RETURN macro instruction. An exit may occur when there is no currently owned region. GETMAIN should not be issued for subpools that represent space within a region (0 through 127, 240 or 250 through 252). Because the exit routine executes as a part of the control program, subpools such as 229 or 230 can be used.
- Format of text and codes is:

Message text (128 characters padded with blanks).

Routing codes (4 bytes). Descriptor codes (4 bytes).

In the routing code field, a bit setting of "1" indicates that the WTO or WTOR was assigned that particular routing code. Bit assignments and their meanings are:

Bit	Assignment	Meaning
Byte 0		
Bit 0	Routing code 1	Master Console Action
Bit 1	Routing code 2	Master Console Information
Bit 2	Routing code 3	Tape Pool
Bit 3	Routing code 4	Direct Access Pool
Bit 4	Routing code 5	Tape Library
Bit 5	Routing code 6	Disk Library
Bit 6	Routing code 7	Unit Record Pool
Bit 7	Routing code 8	Teleprocessing Control
Byte 1		
Bit 0	Routing code 9	System Security
Bit 1	Routing code 10	System Error/Maintenance
Bit 2	Routing code 11	Programmer Information
Bit 3	Routing code 12	Emulators
Bit 4	Routing code 13	Available for Customer Use
Bit 5	Routing code 14	Available for Customer Use
Bit 6	Routing code 15	Available for Customer Use
Bit 7	Routing code 16	Reserved
Byte 2		Reserved
Byte 3		Reserved

In the descriptor code field, a bit setting of "1" indicates that the WTO or WTOR was assigned that particular descriptor code. Bit assignments and their meanings are:

Bit	Assignment	Meaning
Byte 0		
Bit 0	Descriptor code 1	System failure
Bit 1	Descriptor code 2	Immediate action required
Bit 2	Descriptor code 3	Eventual action required
Bit 3	Descriptor code 4	System status
Bit 4	Descriptor code 5	Immediate command response
Bit 5	Descriptor code 6	Job status
Bit 6	Descriptor code 7	Application program/processor
Bit 7	Descriptor code 8	Out-of-line message
Byte 1		
Bit 0	Descriptor code 9	DISPLAY or TRACK command response
Bit 1	Descriptor code 10	Dynamic status displays
Bit 2	Descriptor code 11	Critical eventual action required
	12 through 16	Reserved
Byte 2		
		Reserved
Byte 3		
		Reserved

Messages Not Using Routing Codes

There are certain messages that the exit routine does not see. These are messages that have the MSGTYP parameter in the WTO or WTOR macro instruction coded with the JOB NAMES, STATUS, or Y parameter, multiple-line WTOs (including status displays), and messages that are being returned to the requesting console, such as a response to a DISPLAY A command. Routing of these messages depends on criteria other than the routing codes; therefore, the system bypasses the exit routine.

Writing a WTO/WTOR Exit Routine

To modify the standard routing codes and descriptor codes, a WTO/WTOR exit routine must be written. This routine becomes part of the control program. When a message's routing code field is used by the operating system to route the message, the routine receives control prior to the routing of the message. When the routine receives control, register 1 contains a pointer to a word that points to the first word of the message text. The message text field is 128 bytes long, followed by a four-byte routing code field and a four-byte descriptor code field. The exit routine can examine but not modify the message text.

A message is sent only to those locations specified in the modified routing codes. All messages with modified routing codes are sent to the hard copy log when the log is included in the operating system. When the log is not included, the exit routine must not suppress messages that contain a routing code of 1, 2, 3, 4, 7, 8, or 10 because messages with these codes are necessary for system maintenance. Message suppression is turning off all routing codes of a message by setting the routing code field to zero, thus causing the message to be discarded. WTO messages can be suppressed. If a WTOR message is suppressed, it is sent to the master console by the operating system.

Adding a WTO/WTOR Exit Routine to the Control Program

The WTO/WTOR exit routine is standard. If the user does not specify one, the IBM-supplied module IEECVXIT is included.

To enter the exit routine into the control program before system generation, the linkage editor should be used to replace the dummy WTO/WTOR exit routine IEECVXIT in SYS1.AOSC5 with the WTO/WTOR exit routine. The linkage editor should be instructed to load IEECVXIT on a page boundary.

To enter the exit routine into the control program after system generation, the linkage editor should be used to replace the dummy WTO/WTOR exit routine (CSECT IEBCVXIT in load module IGC0003E) in the SYS1.LPALIB with the user-written WTO/WTOR exit routine.

Service Management

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 would be required by using existing task management services. The facilities provided by service management are those services required to:

- 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 first two facilities are accomplished via the SCHEDULE macro instruction; the third facility is accomplished via the PURGEDQ macro instruction.

The main features of this support are:

- A control block, called a service request block (SRB), represents a service request. This block is similar to a TCB but requires less information to be specified. It identifies a unit of work to the dispatcher.
- The SCHEDULE macro service 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 old task structure. In addition, SRBs may be scheduled to different address spaces either at a priority equal to the specified address space or at an independent priority higher than any address space.
- The service management facilities provide an increase in parallel processing on multiple processors. The programs benefiting the 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 almost all 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 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 the one that took the interruption.
- Process the interruption enabled except where specific serialization through locks is used.
- To switch from the random address space where the interruption was taken to the address space of the user which originally requested the I/O. This latter capability provides the interruption handler routine with addressability to the user's control blocks necessary to complete the interruption processing.

Scheduling Service Requests

The introduction of a service request into the queue of work is accomplished via the SCHEDULE macro instruction. To use this macro, you must provide the following information:

- The address of a previously obtained and formatted service request block (SRB) that is to represent the request until it is actually dispatched. The contents of the SRB supplied define the attributes of the routine to be given control. Once the service routine is given control, the SRB is no longer needed by the system and may be released.
- The priority of the request relative to other requests in the system. The service may be scheduled at either local or global priorities.

The SCHEDULE macro instruction does not obtain storage for the SRB. It simply causes the indicated SRB to be queued at the appropriate priority. When the request subsequently becomes the highest priority work in the system, the dispatcher dispatches it in the address space specified by the SRB.

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 *OS/VS2 System Programming Library: Debugging Handbook*, section 6, "Data Areas". The information required within the SRB at the time of SCHEDULE is 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 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.

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 may be dispatched. (The nth bit set "on" indicates that the SRB may be dispatched on the processor with physical address n.)

SRBPRIOR - specifies a code representing the priority level to which the SRB is being scheduled. The codes are assigned as follows:

0	SYSTEM
4	NONQ

SRBRMTR - contains the address of a resource manager termination routine. This routine is responsible for cleaning up SRBs that have been scheduled but not yet dispatched. The address also serves as a function identifier to PURGEDQ when purging SRBs from the queues. If 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, then this field must contain a valid non-zero address of a routine to clean up that SRB. The routine must be commonly-addressable by all address spaces.

SRBPTCB - contains the address of a TCB associated with the asynchronous routine, which serves two purposes:

- 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, a zero value should be specified.

SRBASID - contains the ASID of an address space associated with the asynchronous routine. If a non-zero 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, a zero value may be specified.

Priorities

Services may 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 will be 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 usage (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 in the process of being quiesced. System level service requests are not to be 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, 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 may have been suspended due to lock requests or page faults. Since page fault processing and rescheduling of the suspended SRBs makes use of SRBs themselves, 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:

- All routines are entered in supervisor state, enabled, and unlocked. The routines cannot enter problem program mode and must establish a recovery environment. Upon entry to the 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 may not issue SVCs. (However, ABEND may be issued.)
- The routines are non-preemptable -- that is, although the routines are run enabled and may be interrupted by an asynchronous interruption, they will not lose control to higher priority tasks or SRBs until they give up control voluntarily. However, service routines may lose control due to synchronous events which cause suspension of the program in control -- for example, 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 may 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 may request any lock through the locking interface (such as SETLOCK).
- The routines must return control to the address supplied in register 14, and must return control in supervisor state with no locks held. All cleanup must be performed prior to exiting.

Purging Service Requests

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, 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.

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

PURGEDQ Parameters

The inputs to the PURGEDQ macro instruction are specified by the ASID, ASIDTCB, and RMTR parameters.

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 caller of PURGEDQ can purge the SRBs associated with a specific task by coding the ASIDTCB parameter. The ASIDTCB parameter specifies the address of a doubleword associated with 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 the acceptable values for the ASIDTCB parameter, and the meaning of the values:

Bytes 0 - 7 zero	All SRBs are to be purged.
Bytes 0 - 1 reserved Bytes 2 - 3 nonzero ASID Bytes 4 - 7 zero	All SRBs associated with the specified address space (SRBPASID field) are to be purged.
Bytes 0 - 1 reserved Bytes 2 - 3 nonzero ASID Bytes 4 - 7 nonzero TCB address	All SRBs associated with the specified address space (SRBPASID field) and the specified task (SRBPTCB field) are to be purged.

All other values are unacceptable and produce unpredictable results.

The RMTR parameter specifies the address of the resource manager termination routine.

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	SRB 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 key 0 and no locks held.

The RMTR must not leave supervisor state or issue ESTAE with branch entry. The RMTR must return control enabled, in supervisor state with key zero and no locks held and via a BR14. 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 desires 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.

Asynchronous Exit Routines

The following facility is available whereby an authorized user can request an asynchronous exit routine to execute on behalf of a specific task. Prior to 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 is invoked by the authorized user to create and initialize an Interrupt Request Block (IRB) that identifies to the system the user's asynchronous exit routine. Interface to the Stage 1 Exit Effector is through the Create Interrupt Request Block (CIRB) macro instruction. A detailed description of the CIRB macro instruction is found in Part II of this manual.

The Stage 2 Exit Effector routine is invoked by the authorized user to schedule his exit routine for execution. Input to the Stage 2 Exit Effector is an Interrupt Queue Element (IQE) that has been 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 relating 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. Execution of the exit routine occurs only once for each invocation of the Stage 2 Exit Effector by the caller for that exit routine. Reuse of previously defined control blocks is possible however. Thus, one interface to the Stage 1 Exit Effector can suffice for two or more invocations of the exit routine.

The Stage 3 Exit Effector is invoked by the MVS dispatcher to queue the IRB for the exit routine 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

Interface to the Stage 1 Exit Effector is provided via the CIRB macro and is normally by a Type 1 SVC interface. Also, a branch entry interface is available by specifying BRANCH=YES as one of the macro options. Specific instructions for using the CIRB macro instruction are in Part II. 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, is contiguous in storage to the interruption request block 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 as shown in Figure 20.

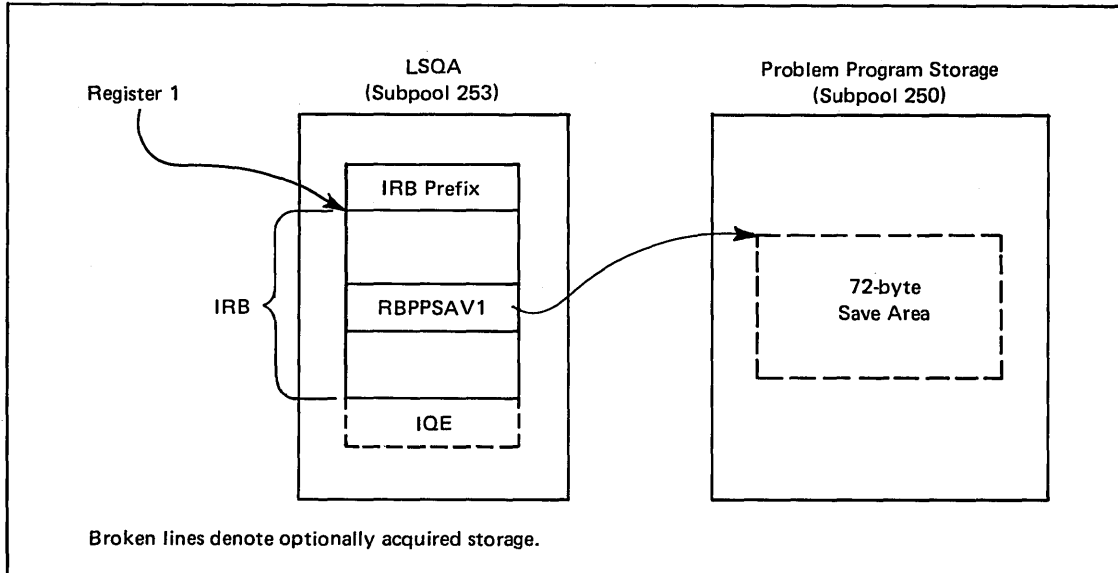


Figure 20. Asynchronous Exit Data Area Configuration

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: RQEs are only for use by IOS.
- 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.
 - Mode: Supervisor state if MODE=SUPR on the CIRB macro instruction; problem program state if MODE=PP on CIRB.
- RBSIZE -- size of the IRB (including the size of the IQE if the CIRB specification included the WKAREA parameter).
- RBNEXAV -- address of the IQE if WKAREA was specified.
- RBPPSAV1 -- address of the problem program save area if SVAREA was specified.

Stage 2 Scheduling

Initialization of the IQE is the responsibility of the user and is necessary to define to the system the task the exit routine is to execute under. The fields to be initialized are:

- IOEPARAM -- address of the parameter list (if wanted) 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 form of the IQE address.
2-13	Irrelevant.
14	Return address.
15	Irrelevant.

Note: Upon return, registers 0, and 2-14 are unchanged and register 1 contains a true (non-complemented) form of the IQE address.

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 to the MVS 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 may 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 initialization of the IQE fields is correct.

Execution and Termination Characteristics

The following asynchronous exit routine characteristics 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 the user repeating 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.

Writing SVC Routines

User-written SVC routines become part of the control program, so you must follow the same programming conventions used by SVC routines supplied with VS2. Five types of SVC routines are supplied with VS2, and the programming conventions for each type differ. The general characteristics of the five types are described in the following text, and the programming conventions for all types are shown in tabular form.

Characteristics of SVC Routines

All SVC routines receive control in key zero and supervisor state. You should keep the following characteristics in mind when deciding what type of SVC routine to write:

- Location of the routine -- Your SVC routine can be either in storage at all times as part of the resident control program, or in the fixed or pageable link pack area. Types 1, 2, and 6 SVC routines are part of the resident control program, and types 3 and 4 are in the link pack areas. No transient areas are provided in MVS.
- Size of the routine -- SVC routines are not limited in size, but should be kept under one page if disabled global locks are obtained.
- Design of the routine -- All SVC routines must be reenterable. If you wish to aid system facilities in recovering from machine malfunctions, your SVC routines must be refreshable.
- Authorization -- At nucleus initialization, all SVC routines that are to be loaded into the fixed or pageable LPA must be contained in SYS1.SVCLIB, SYS1.LPALIB, or SYS1.LINKLIB. After the initial load of type 4 SVCs, subsequent loads may be contained in any authorized library.
- Serialization -- In MVS, locking has generally replaced processor disablement as the technique for serializing multiple processor functions. If you write SVC routines which must serialize with other parts of the control program, you must use the same locking conventions as the control program. If you write two or more SVC routines which must serialize with each other, you can use either the locking facilities or the ENQ/DEQ services.

SVC routines can receive control with one or more locks held. During system generation, you must define which locks are to be acquired for your SVC routines. (For more information on locking, see the discussion under Resource Control.)

SVC routines are normally entered enabled. However, an SVC routine will be entered disabled if it was specified that a disabled spin lock is to be acquired for the routine. (See reference code 2 on the following page.)

Programming Conventions for SVC Routines

The programming conventions for the five types of SVC routines are summarized in Figure 21. Details about many of the conventions are in the reference notes that follow. The notes are referred to by the numbers in the last column of the figure. 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 the first byte of the routine or load module, and must be on a doubleword boundary				Same	
Number of routine	Numbers assigned to your SVC routines should be in descending order from 255 through 200				Same	
Name of routine	IGCnnn	IGCnnn	IGC00nnn	IGCcssnn	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				Same	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 issue XCTL macro instruction	No	Yes	Yes	Yes	No	8
May pass control to what other types of SVC routines	None	Any	Any	Any	None	9
Type of linkage with other SVC routines	Not Applicable	Issue supervisor call (SVC) instruction			Not Applicable	10
Exit from SVC routine	Branch using return register 14				T6 EXIT or BR 14	11
Method of abnormal termination	ABEND	ABEND			ABEND	
Recovery	FRR	ESTAE or FRR			FRR	12

Figure 21. 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	The following conventions on locking requirements apply: <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 note 9).
0		Type 6 may not request any locks.
4	all	You must use the following conventions when naming SVC routines: <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. • Type 3 must be named IGC00nnn; nnn is the signed decimal number of the SVC routine. • Type 4 must be named IGCssnnn; nnn is the signed decimal number of the SVC routine, and ss is the number of the load module minus one. For example, ss is 01 for the second load module of the routine.
5	all	Before your SVC routine receives control, the contents of all registers are saved. <p>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	The SVRB is no longer extendable, but is a fixed size of 224 bytes. When a type 2, 3, or 4 SVC routine receives control, register 5 contains the address of the SVRB within 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 recovery termination manager when a ESTAE is issued within an SVC routine.
7	2,3,4,6	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 or 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	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.
9	all	No SVC routines except ABEND may be called if locks are held. ABEND may be called at any time.
10	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.
11	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.
12	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 terminate 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

You insert SVC routines into the control program during the system generation. Before your SVC routine can be inserted, the routine must be a member of a cataloged partitioned data set named SYS1.name, where the name is a name of your choice. The data set must be an authorized library (see discussion on APF).

The following text describes the information you must supply during system generation. You will find a description of the macro instructions required during system generation in the **OS/VS2 System Programming Library: System Generation Reference** publication.

Specifying SVC Routines

You use the SVCTABLE macro instruction to specify the SVC number, the type of SVC routine, and the locks that are required for your SVC.

Inserting SVC Routines During the System Generation Process

To insert an SVC routine during system generation, the name of the partitioned data set and the names of the members to be included must be specified in the DATASET macro instruction:

- For each type 1 or 2 SVC, a member (containing one or more SVC routines) should be specified in the DATASET macro instruction that is being used to define SYS1.NUCLEUS.
- For each type 3 or 4 SVC, a member (containing only one SVC routine) should be specified in the DATASET macro instruction that is being used to define SYS1.LPALIB.

Type 5 SVC Facility

The Type 5 SVC facility permits a user to indicate (at system generation time) that one or more user-written SVCs are added to the system after the completion of system generation. This facility is only used in cases where the SVC routines for the user-written SVC are not available at the time of the system generation. This facility further permits the user to specify the authorization characteristic and lock requirements as input parameters to the system generation macro instruction SVCTABLE.

SVC Table Entries

Each entry in the SVC table has the following format:

Offset	Length	Name	Description
0	4	SVCEP	SVC Entry Point Address
4	2	SVCATTP1	Attributes
	SVCTP1	Type 1 SVC
	1...	SVCTP2	Type 2 SVC
	11..	SVCTP34	Type 3 or 4 SVC
	..1.	SVCTP6	Type 6 SVC
 1..	SVCAPF	APF Authorized 1-Authorized
1.	SVCESR	SVC is a part of the ESR
1.	SVCNP	Non-premptive SVC
6	2	SVCLOCKS	Lock Attributes
	1...	SVCLL	Local Lock Needed
	.1.	SVCCMS	CMS Lock Needed
	..1.	SVCSRMR	SRM Lock Needed
	...1	SVCALLOC	SALLOC Lock Needed
 1..	SVCDISP	Dispatcher Lock Needed

If a Type 5 SVC is specified at system generation time, the following entry is generated:

IGCRETRN	Entry Point Address (of IGCRETRN)
Type 1	
Authorization	Default is unauthorized, unless the user specifies authorized.
Locks	Default is no locks; user may specify any lock or combination of locks needed by his routine.

Note: IGCRETRN is an existing routine which zeroes register 15 and returns via register 14.

Subsequent to system generation, the user is expected to supply the SVC routine and to update the entry in the SVC table. Assuming that the SVC attributes generated at system generation are correct, the user need only supply the address of the entry point to the SVC routine. If the attributes are not correct for the SVC as generated by system generation, then AMASPZAP can be used to update or replace all attributes stored in the SVC table entry.

Note: AMASPZAP is not practical for use in updating the entry point address of the SVC routine. For nucleus resident Type 1, Type 2, or Type 6 SVC routines, the user should dynamically update the entry point address in the SVC table after each IPL by using a routine which executes prior to any need for that SVC. For Type 3 or Type 4 SVC routines named according to SVC naming conventions and link-edited into SYS1.LPALIB, a NIP RIM initializes the entry point address in the SVC table for those SVCs.

Type 6 SVC Facility

The Type 6 SVC performs functions similar to the Type 1 SVC, but, by cutting the instruction path lengths for receiving and releasing control nearly in half, it offers performance advantages over the Type 1 SVC. It does not always require the local lock, as noted later.

The Type 6 SVC also provides a more efficient way for a program to change from TCB mode to SRB mode processing. When using a Type 1 SVC for this purpose, a program must itself schedule an SRB, the SRB must then go through queuing and dequeuing operations and eventually be dispatched. Using the Type 6 SVC, normally results in immediate scheduling and dispatching of the SRB.

The Type 6 SVC executes under the control of the first level interrupt handler (FLIH) and therefore has similar function and limitations. The Type 6 SVC has three exit options:

- Return to the caller directly
- Return to the dispatcher
- Dispatch an SRB (service request block)

All these returns must be to the SVC FLIH through the T6EXIT macro instruction. For specific information about how to code this macro, refer to Part II of this publication in the section "T6EXIT -- Type 6 Exit". The following overall considerations for using this macro also apply:

- The Type 6 SVC exit via the T6EXIT macro must be different from the exit used by the other types of SVCs. To ensure this difference, the caller should issue the macro or use the exit address in register 14 upon entry to the SVC when exiting. Using the T6EXIT macro 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 addr.	N/A
BR 14	As on entry	Returned to caller	Returned to caller	Returned to caller

- Plans for modifying existing (Type 1,2,3,4) SVCs to Type 6 SVCs, should include these factors:
 - The system neither acquires nor releases any locks for Type 6 SVCs. Currently written SVCs have locks acquired in the SVC FLIH according to the SVCTABLE options.
 - Because a Type 6 SVC executes disabled, it has exclusive use of the processor. Type 6 SVCs should be short enough to minimize their adverse effect on performance.
 - Because Type 6 SVCs have control of the processor (causing the SVRB save areas to be unavailable), PSA and LCCA areas are available to take the place of SVRB save areas.
 - When a Type 6 SVC is executing, no other task-related activity can occur concurrently. To indicate this situation, the TCBACTIV flag is on.
 - A Type 6 SVC should establish recovery via the SETFRR macro instruction.
- If a Type 6 SVC uses the RETURN=SRB exit option on the T6EXIT macro instruction, register 1 points to an SRB input format (see Figure 22).

SRB

Common Name: Service Request Block

Macro ID: IHASRB

Created by: Control Program Routines

Size: 44

Pointed by: SPL or SERVICE MANAGER QUEUES

Function: The I/O Supervisor uses the SRB to dispatch I/O processing for a request. It identifies the address space in which processing is to be done. It is also used as input to the SCHEDULE macro when scheduling a routine for asynchronous EXECUTION.

Offsets	Length	Name	Description
0 (0)	4	SRB	
0 (0)	4	SRBID	EBCDIC ACRONYM FOR SRB
4 (4)	4	SRBFLNK	FORWARD CHAIN FIELD
8 (8)	4	SRBASCB	PTR TO ASCB OF ADDRESS SPACE SRB IS TO BE DISPATCHED TO
12 (C)	8	SRBFLC	SRB AREA MOVED TO LOW CORE
12 (C)	2	SRBCPAFF	CPU AFFINITY MASK
14 (E)	2	SRBPASID	PURGEDQ ASID IDENTIFIER
16 (10)	4	SRBPTCB	PURGEDQ TCB IDENTIFIER
20 (14)	4	SRBEP	ENTRY POINT OF ROUTINE
24 (18)	4	SRBRMTR	ADDRESS OF RESOURCE MGR RTN
28 (1C)	4	SRBPARAM	USER PARAMETER
32 (20)	4	SRBSAVE	SAVE AREA POINTER
36 (24)	1	SRBPKF	PROTECT KEY INDICATION
37 (25)	1	SRBPRIOR	PRIORITY LEVEL INDICATION
....	SRBPSYS	SYSTEM PRIORITY LEVEL
....	.1..	SRBPNONQ	NON QUIESCEABLE PRIORITY
38 (26)	2		RESERVED
40 (28)	4		RESERVED

Note: The SVC FLIH dispatches an SRB only for the current address space. If the ASCB is unequal to the current address space ASCB an abend results.

Figure 22. SRB Input Format

Non-Preemptable SVCs

It is possible at system generation to denote SVCs so that the system subsequently considers them non-preemptable for I/O operations. After system generation, the SVC first level interrupt handler (FLIH) recognizes any SVC thus denoted through an indicator in the SVC table and then initializes the non-preemptable state for it. To handle this state, the I/O FLIH returns control directly to the non-preemptable SVC task at any time when there would normally be an I/O interrupt that would preempt it.

The following three restrictions apply to non-preemptable SVCs:

- The non-preemptable state should be of short duration -the life of the SVC should be the maximum.
- 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 causes the issuing SVC to lose its non-preemptable quality.
- If a non-preemptable SVC issues a STAX DEFER=NO macro instruction, the SVC remains non-preemptable until entry to the exit, RTM, or ABEND.

Subsystem SVC Screening

Subsystem SVC screening allows a system routine to define those SVCs that a specific task may 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 may validly request that SVC function. If so, the SVC FLIH processes the request as a normal SVC request. If not, the SVC FLIH gives control to a subroutine supplied by the subsystem routine that activated SVC screening for that task, instead of to the requested SVC.

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 in the main body of the TCB and a one-word field in the TCB extension. This field contains the address of the subsystem screen table, which provides the interface between the SVC FLIH and the subsystem subroutine. These fields are:

- TCBSVCS -screen flag bit, which indicates SVC screening in effect for this task when on (one).
- TCBSVCA2 -- address of the subsystem screen table.

The subsystem subroutine that receives control as a result of an invalid SVC request executes 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. This area must come from either LSQA (subpool 253-255) or SQA (subpool 245) or must be fixed to prevent a page fault. The subsystem screen table contains two areas as follows:

1) SSTSVCN -- 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
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.

Missing Interruption Handler

The missing interruption handler (MIH) checks whether expected MIH conditions occur within a specified period of time. The MIH conditions are pending device ends, channel ends, DDR swaps, or MOUNT commands. If an expected MIH condition does not occur, the operator is informed so that steps can be taken to correct the situation before system performance is affected.

MIH detects a pending MIH condition in the UCB for the device. After a time interval elapses, MIH checks the UCB and, if the MIH condition is still pending, takes the following actions:

- For all MIH conditions, MIH issues a message informing the operator of the pending condition and the necessary action to correct it.
- For all missing channel ends or device ends from critical devices, such as paging and SYSRES devices, MIH informs the I/O supervisor (IOS) so that IOS can restart the device.
- MIH records all missing I/O interruptions on SYS1.LOGREC for further analysis.

MIH provides a primary and a secondary level of processing. The primary level applies to all devices; all devices are checked for MIH conditions at an installation-specified primary time interval. The secondary level applies to all critical devices: paging and SYSRES devices defined at IPL and additional installation-defined devices. These devices are listed in the secondary look-up table. Because a MIH condition on one of these critical devices can impact system performance, the secondary time level must be smaller than the primary time interval.

The module IGFINTVL, which resides in the nucleus, contains the primary and secondary time intervals, the installation-defined UCB entries, the address of the MIH usage area containing the secondary look-up table, and the exit index table. The offsets and default values are as follows:

Offsets	Length	Default Value	Description
0 (0)	8 bytes	00030000	Primary time interval
8 (8)	8 bytes	00001500	Secondary time interval
16 (10)	48 bytes	N/A	Installation-defined device address table
64 (40)	16 bytes	N/A	MIH usage area
80 (50)	16 bytes	N/A	Exit index table

The MIH usage area contains a pointer to the secondary look-up table. The secondary look-up table contains the UCBs for all paging and SYSRES devices defined at IPL. MIH adds the entries from the installation-defined device address table into the secondary look-up table. The entries are moved in the order entered by the installation until the installation-defined table is exhausted or the secondary look-up table is full. MIH converts all device addresses to its equivalent UCB address. A zero on invalid device address is ignored. After an IOGEN the installation should verify the device addresses in the device address table.

Each exit index table contains the addresses for the first and second level exits and the two exit constants. The last 4-bytes are reserved for MIH. The first entry is used by the Mass Storage System (MSS) and is discussed later in this publication. MSS devices must not be listed in the installation-defined device table.

Choosing a Time Interval

An installation can use the IBM-supplied time intervals or specify its own intervals. The AMASPZAP program can be used to modify the values; an example appears in Figure 23. When modifying the values, consider the following points:

- The time intervals are expressed as EBCDIC in the form HHmmsshh (hours, minutes, seconds, hundredths). MIH ignores the hundredths. If the time interval is changed to 0 or a value other than an EBCDIC 0-9 (F0-F9) is specified, the default value is used.
- To ensure that critical devices are checked for pending conditions on a timely basis, MIH requires that the secondary time interval be shorter than the primary time interval. Otherwise, MIH sets the primary interval equal to the secondary time interval. The primary time must also be a multiple of the secondary value. If not, MIH makes the primary value a multiple of the secondary value. When the remainder from the division is greater than half the secondary time interval, MIH adjusts the primary time value upwards.

An example of Modifying the CSECT IGFINTVL

To modify ICFINTVL, use the AMASPZAP program. The modifications do not take affect until the next IPL of the system. The following modifications are made in Figure 23:

- The primary interval is changed from 3 minutes to 3 minutes and 30 seconds.
- The secondary interval is changed from 15 seconds to 10 seconds
- Four device addresses (0120, 0275, 03FB, and 0410) are entered into the installation-defined device address area so that they are included in the secondary look-up table.
- The MSS constants are changed from 4 to 5.

```
//MODIFY   JOB
//STEP     EXEC      PGM=AMASPZAP
//SYSPRINT DD        SYSOUT=A
//SYSLIB   DD        DSNAME=SYS1.NUCLEUS,DISP=OLD
//SYSIN    DD        *
NAME       IEANUC01  IGFINTVL
VERIFY    0000      F0F0F0F3F0F0F0F0  PRIMARY INTERVAL
VERIFY    0008      F0F0F0F0F1F5F0F0  SECONDARY INTERVAL
VERIFY    0010      0000000000000000  ENTER DEVICE ADDRESSES
VERIFY    0058      F0F4F0F4          MSS NVALUE YVALUE
REP       0000      F0F0F0F3F3F0F0F0  PRIMARY INTERVAL
REP       0008      F0F0F0F0F1F0F0F0  SECONDARY INTERVAL
REP       0010      0120027503F80410  DEVICE ADDRESSES
REP       0058      F0F5F0F5          MSS NVALUE YVALUE
/*
```

Figure 23. Changing the Missing Interrupt Handler Time Interval

Mass Storage System Missing Interrupts

Mass Storage System devices tend to require more time to respond with interrupts than do other devices. Therefore, Mass Storage System devices use time intervals that are multiples of the basic time interval. The time required by the 3851 to respond with Channel End or Device End varies depending on the I/O request and on the amount of queued stage/destage work. This is similarly true for 3330V (3330 virtual unit) Attention Interrupts that are associated with cylinder faults and for 3330V pack change interrupts that are associated with MOUNT requests. Therefore, to minimize premature determination of these missing interrupts, a value N is added to the exit index table of the IGFINTVL csect of MIH.

The normal default MIH time interval is 3 minutes, but the Mass Storage System exit routine (ICBTMSS0) uses the value N to allow checking only every Nth MIH time period. This value N is a ratio of the number of normal time periods per effective Mass Storage System time period. For example, with a basic time interval of 3 minutes and a ratio value of N=4, the effective Mass Storage System time interval would be 4x3=12 minutes.

Channel End and Device End from 3330Vs can be delayed when a cylinder fault occurs. This happens when a heavy workload causes queuing of messages in the Staging Adapter and Mass Storage Control communication buffer. The length of such delays, however, does not tend to be as long as delays that are caused by queued stage/destage work. Therefore, to minimize premature determination of missing 3330V Channel End and Device End interrupts, a ratio value Y is used in the way the ratio value N is used for all other Mass Storage System interrupt checking. This value Y is also added to the exit index table of the IGFINTVL csect of MIH.

The actual frequency of checking for missing Mass Storage System device interrupts is proportional to whatever basic time interval the user has chosen. The best checking frequency for a given user is dependent on his individual system workload characteristics and under-commitment or over-commitment of Mass Storage Facility space and staging space. A default ratio value of 4 has been chosen for both N and Y and, like the basic time interval, the ratio values can be modified by the user through the use of the SPZAP Service Aid as indicated in Figure 23. Modification of either the basic time interval or a ratio value affects the frequency of discovery of apparent overdue Mass Storage System interrupts and corresponding frequency of operator notification through message IGF991E.

The N value and Y value ratio fields must each contain an EBCDIC value from X'F0F1' to X'F9F9'. The N value and Y value defaults to X'F0F4'.

Intercepting Hot I/O Interrupts

Hot I/O occurs when a hardware malfunction causes repeated I/O interrupts. If left undetected, hot I/O interrupts can cause the system to loop or system queue area (SQA) to be exhausted. The I/O supervisor (IOS) attempts to detect the hot I/O so that the problem can be resolved before the system requires a re-IPL.

Repeated I/O interrupts indicate a hot I/O problem. When the number of repeated interrupts exceeds any one of four installation-defined threshold values (hot device, hot channel, hot control unit, or time out), the system assumes that there is a hot I/O condition. The system issues an operator message or enters a restartable wait state and prompts the operator to take corrective action.

The timeout threshold detects hot I/O in situations where interrupts might be quite far apart. The threshold value must be set low enough to detect the interrupts. The timeout threshold value is incremented each time a repeated interrupt occurs after a specified time interval.

The module IECVHIDT, which resides in the nucleus, contains the four threshold values and the time interval value that is used with the timeout threshold. The offsets and IBM-supplied default values are as follows:

OFFSET	LENGTH	DEFAULT VALUE	ELEMENT DESCRIPTION
0(0)	2	X'64'	Hot channel threshold
2(2)	2	X'64'	Hot device threshold
4(4)	2	X'08'	Timeout threshold
6(6)	2	X'04'	Timeout interval in seconds
8(8)	2	X'800'	Hot control unit threshold

The possible range of each threshold value is from 0 to X'7FFF'. A threshold value of 0 specifies that hot I/O is not to be detected for that element, while too low a value might result in false detection.

The values can be replaced using AMASPZAP. The following example modifies the control unit threshold to eliminate hot I/O detection and increases the threshold for hot channel interrupts.

```
//MODIFY   JOB      MSGLEVEL=(1,1)
//STEP     EXEC     PGM=AMASPZAP
//SYSPRINT DD      SYSOUT=A
//SYSLIB   DD      DSN=SYS1.NUCLEUS,DISP=OLD
//SYSIN    DD      *
           NAME     IEANUC01  IECVHIDT
           VERIFY   0000  0064
           VERIFY   0008  0800
           REP      0000  0090
           REP      0008  0000
/*
```

If IOS detects hot I/O and later investigation determines that no hardware malfunction occurred, increase the appropriate threshold value. If hot I/O causes the exhaustion of SQA before detection and recovery, decrease the appropriate threshold value.

Adding Code to the Power Warning Feature Support

The power warning feature support, along with its supporting hardware, prevents the loss of information in real storage at the occurrence of a utility power disturbance. The supporting hardware must include an uninterruptable power supply to provide alternate power and equipment to signal the power warning feature support routines when a disturbance occurs.

The power warning feature support, after receiving the signal of a power disturbance, and determining the significance of the disturbance, can transfer the contents of real storage to disk. After utility power is restored, the customer can use the power warning feature support "restore" routine to refresh the contents of real storage from disk.

Adding code to the machine check handler appendage -- You can add code that executes after receipt of the signal that indicates that a sustained power disturbance has occurred. Your code can then permit transfer of real storage to a warn data set or have control returned to the supervisor.

Adding code to the master scheduler initialization module -- You can add code that executes when a warn data set contains an image of real storage and the system operator chooses FORM during system IPL. Your code executes just before the warn data sets are erased and reformatted.

Note: For details on the warn data set, see *OS/VS2 System Programming Library: System Generation Reference*, GC26-3792.

Note: Since adding code to the power warning feature support requires considerable programming skill, before attempting any addition you should carefully examine the complexities involved.

Writing Code for the Machine Check Handler Appendage

You can insert code that executes when the power warning feature support is entered due to a power warning interrupt that would normally cause the transfer of real storage to the warn data set. After your code executes you can either cause control to be transferred to the dump routine or have control returned to the machine check handler for the system to continue processing.

As shown in Figure 24, the machine check handler appendage consists of three parts: 1) the warning appendage, 2) your code, and 3) the dump routine.

The warning appendage routine, after a power interruption, receives control from the system's machine check handler. This warning appendage monitors the power interruption during the time delay you specified in the CTRLPROG macro instruction at system generation. If a machine check occurs during the time delay, the remainder of the time delay is canceled.

Normally your code receives control at the end of the time delay, assuming the power interruption is still present. If the utility power returns before the expiration of the time delay, control returns to the supervisor, via the machine check handler.

The dump routine transfers the contents of real storage to a warn data set. After execution of the dump routine, the system enters a wait state.

To include code:

1. Use ICFBDE50 as csect the name of your code.
2. Assemble csect ICFBDE50.
3. Link ICFBF00 with the REPLACE option, replacing csect ICFBDE50 with your csect ICFBDE50 into DISTLIB AOSCE.
4. Link IEANUC01 from SYS1.NUCLEUS, include ICFBDF00 from DISTLIB AOSCE and replace in SYS1.NUCLEUS.
5. Update the SMPADS to reflect this change for regression messages.

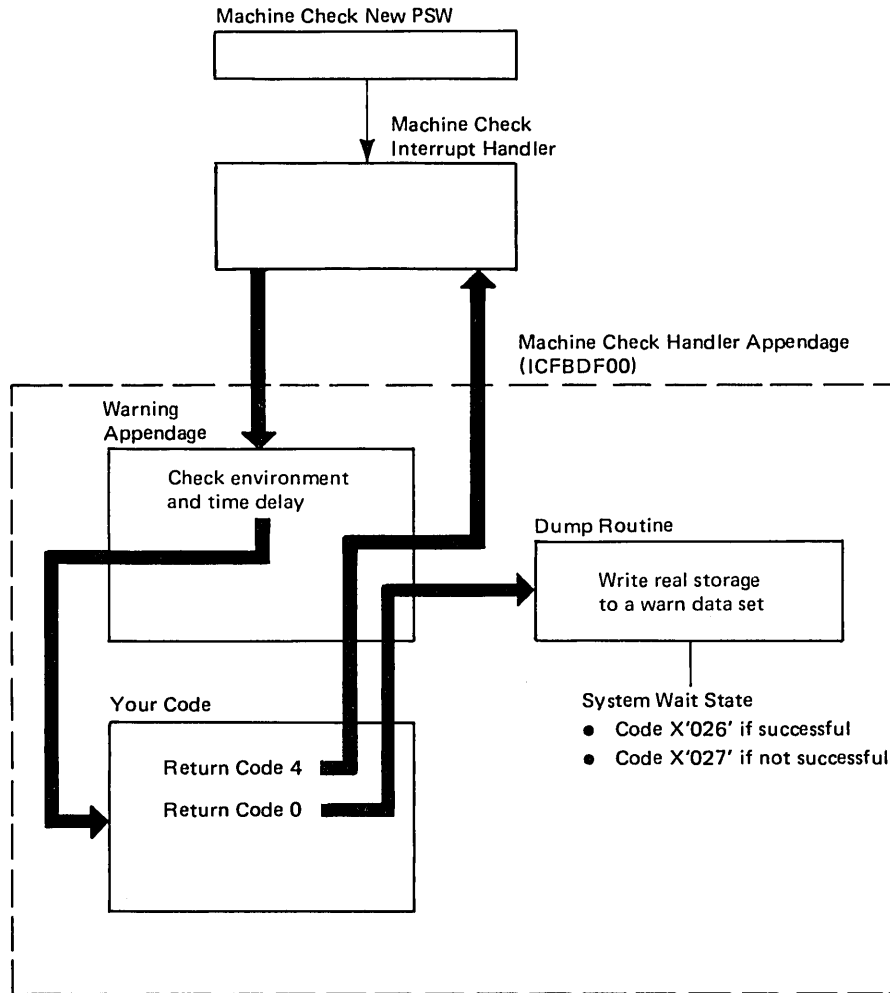


Figure 24. Logical Placement of Your Code in the Machine Check Handler Appendage

Coding Considerations

At entry to your code, Register 9 addresses the PWF Communication area, Register 14 contains the return address, and Register 15 contains the address of your first instruction. Your code must save and restore all the general registers except register 15, prior to the exit from your code. You must also restore all control registers and all real storage outside of your inserted code. Register 15 must be set to 0 if you wish real storage to be transferred to a warn data set, or be set to 4 if you wish to resume system operations with all power warning interrupts disabled.

Your code is entered:

- Key zero.
- Disabled for all interrupts.
- In supervisor state.

Your code must not:

- Use supervisor services.
- Contain Address Constants A-type or V-type since your code may be relocated.

Note: With a multiprocessing system, your program is executed by only one processor.

Writing Code for the Master Scheduler Initialization Routine

You can write a routine that executes when there is a real storage image on a warn data set, just before warn data sets are erased and reformatted.

As shown in Figure 25, during IPL, if the warn data set contains information from real storage, the system operator can choose to respond either REST or FORM. REST causes real storage to be refreshed with the contents of the warn data set. FORM causes a transfer of control to your code; after your code returns control, the warn data sets are erased and reformatted by the power warning feature support. If you have not replaced csect ICFBIE50 FORM immediately causes the warn data set to be erased and reformatted.

To include code:

1. Use ICFBIE50 as the csect name of your code.
2. Assemble csect ICFBIE50.
3. Replace csect ICFBIE50 with your csect by performing the following link edit steps:
 - a. Include your replacement csect for ICFBIE50.
 - b. Include ICFBIF00 from DISTLIB AOSCE and re-link to AOSCE.
 - c. Include ICFBIF00 (containing your csect) from DISTLIB AOSCE.
 - d. Include IEEMB860 (the master scheduler initialization module) from SYS1.LINKLIB and replace in SYS1.LINKLIB.
4. Update the SMPADS to reflect this change for regression messages.

Master Scheduler Initialization Routine

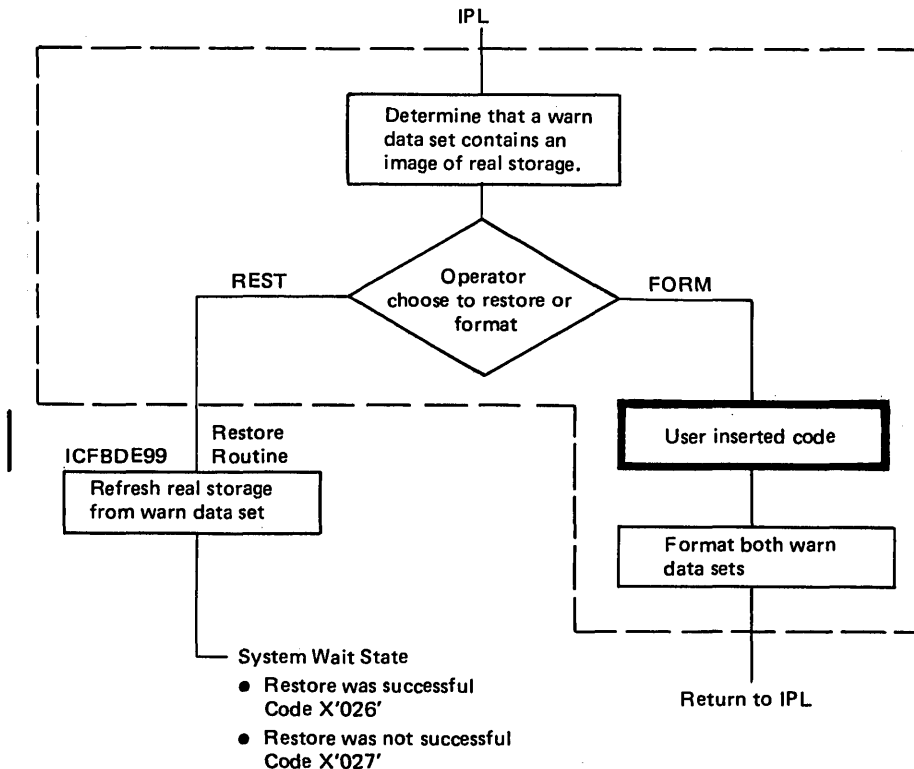


Figure 25. Logical Placement of Your Code in Master Scheduler Initialization Module

Coding Considerations:

At entry to your code: Register 13 addresses the register save area, Register 14 contains the return address, and Register 15 contains the address of your first instruction. Register 1 points to a word that contains the address of the TIOT for the warn data set.

If your code attempts to read information on the warn data set, it should first reference the control record to get vital information about the data set. The control record is the first record on logical cylinder 0 on each warn data set. This record indicates if the warn data set contains information from real storage and indicates its format. The format of the control record is shown in Figure 26.

When your code gets control, the environment is as follows:

- Your code is pageable.
- Your code is transient.
- Your code cannot obtain permanent storage within the region.
- Consoles are available for your code to write to operator.
- Job Scheduler and SYSIN/SYSOUT services are not available.

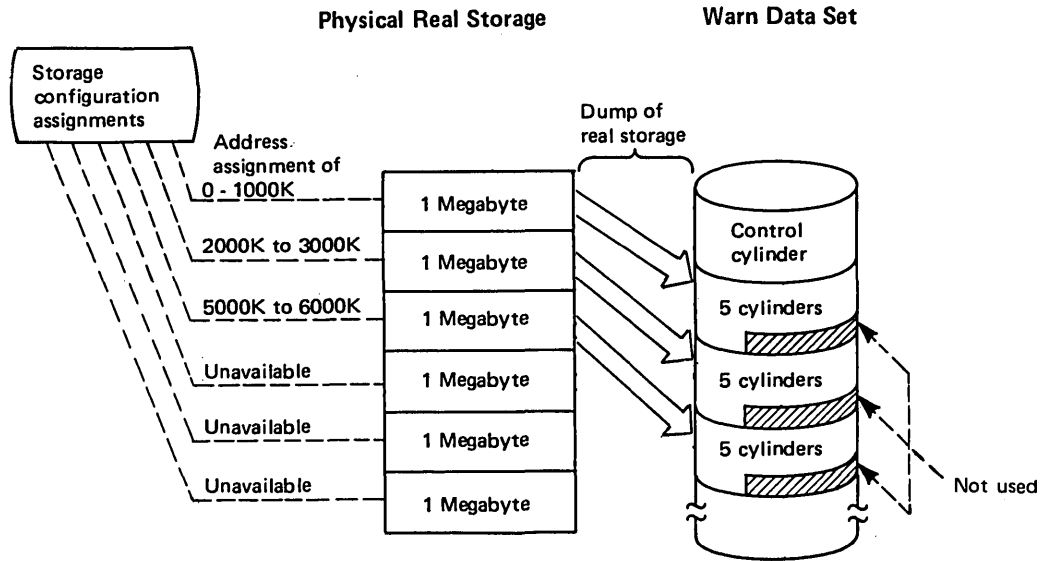
If your system is an M158 or M168 multiprocessor, you must consider that the first address after the last unused address starts a new cylinder on the warn data set. As shown in Figure 27, addresses 0, 2000K, and 5000K start new cylinders. An unused address refers to a location in real storage that is not available because of settings of the console switches.

Control Track Record

This record is located on cylinder 0 on each warn data set.

Offset	Size/Bits Length	Name	Description
0 (0)	4	ICFCTID	This is the identifier of the first word of the control track record.
4 (4)	128	ICFCTCF	It always contains CNTL. This area contains 128 byte-indicators. One indicator can be allocated to each of 128 cylinders. (128 cylinders will accommodate 16 megabytes of real storage.) Each byte indicator is structured as follows: 00.. Indicates no data on this cylinder. 01.. Real storage has been transferred to this cylinder and this cylinder contains no defective tracks. 10.. Real storage has been transferred to this cylinder, and this cylinder contains a defective track. ...xx xxxx Contains the track number of either the defective or spare track. (The spare track is the last track on the cylinder.)
132(84)	1	ICFCTFLA	Status flags. 00.. 00.. The data set is formatted, but contains no data. 10.. 00.. This data set contains a successful transfer from real storage. 00.. 10.. This data set contains a partial transfer from real storage. 10.. 01.. This data set contains a successful transfer from real storage, but at least one track was found defective. ...xx ...xx Reserved. Set to zero.
133(85)	3	—	Reserved.
136(88)	4	ICFCTTS	Track size. Number of bytes in each track.
140(8C)	4	ICFCTAWA	The real storage address of the PWF Communication area in real storage.
144(90)	128	ICFCTB11	This area contains eight 16-byte fields. Each 16-byte field represents a contiguous area of real storage on this data set.
144(90)	16	ICFCTB11	Information concerning 1st contiguous area of real storage.
	4	ICFCTB11	Contains the real storage address of the 1st byte represented by this field. (In this case, this byte contains 0's.)
	4	ICFCTB12	Contains the cylinder and track on this data set where this contiguous real storage begins.
	4	ICFCTB13	Contains the real storage address of the last byte represented by this field.
	4	ICFCTB14	Contains the cylinder and track on this data set where this contiguous real storage ends.
160(AO)	16	ICFCTB21-24	Information concerning 2nd contiguous area of real storage.
176(BO)	16	ICFCTB31-34	Information concerning 3rd contiguous area of real storage.
192(CO)	16	ICFCTB41-44	Information concerning 4th contiguous area of real storage.
208(DO)	16	ICFCTB51-54	Information concerning 5th contiguous area of real storage.
224(EO)	16	ICFCTB61-64	Information concerning 6th contiguous area of real storage.
240(FO)	16	ICFCTB71-74	Information concerning 7th contiguous area of real storage.
256(100)	16	ICFCTB81-84	Information concerning 8th contiguous area of real storage.
272(110)	8	ICFCTST	Contains a true reading (binary) of the time-of-day clock when system was last IPLed: or if after a dump, the time at which processing of last PLD began.
280(118)	8	ICFCTED	Contains a true reading (binary) of time-of-day clock at the end of the last real storage transfer to the warn data set.
288(120)	4	ICFCTPC	Total number of tracks of each cylinder.
292(124)	4	ICFCTRDA	The address of this device. This field is set just before entering the restore routine.

Figure 26. Control Track Record



Note: Addresses 1000K to 2000K, and 3000K to 5000K are unused.
The first address after a unused address starts a new cylinder.

Figure 27. Storage Assignments on MP Systems

Limiting User Region Size - IEALIMIT

An installation can enforce a region-size limit by writing an exit routine that is invoked once per step when the initiator is establishing region size. If an installation-written exit routine does not exist, an IBM-supplied routine receives control.

The installation-written exit routine, which replaces the IBM-supplied routine, must be named IEALIMIT and must be link-edited into the nucleus. The routine must observe standard linkage conventions.

Upon entry to the IEALIMIT routine, the register contents are as follows:

Register 0	number of bytes requested by the application program for its region (specified implicitly through the REGION parameter or implicitly through the default JCL value)
Register 1	same as register 0
Register 13	address of standard save area
Register 14	return address
Register 15	entry point address for the IEALIMIT routine

Upon exit from the IEALIMIT routine, the register contents must be as follows:

Register 0	number of bytes to be used as the region parameter value
Register 1	number of bytes to be used as the limit on all types of requests from subpools 0-127, 251 and 252
Registers 2-13	restored
Registers 14-15	irrelevant

If the input register 1 is non-zero when the IBM-supplied IEALIMIT receives control, then IEALIMIT adds 64K to the contents of register 1 and returns. Register 0 remains unchanged. The register 1 value is used to limit the total allocation of storage from subpools 0-127, 251, and 252.

If the IBM-supplied IEALIMIT routine receives control and the input register 1 contains a zero, then IEALIMIT returns a zero in register 1 and no limit is assigned (to a job, a started program, or a TSO user). No limit is set only when the REGION parameter is not specified and the default value is zero.

When no limit is set, sufficient space within a region may be obtained via repeated small GETMAINS or via a single large GETMAIN, such that no space remains in the private area for use by the system. This situation is likely to occur when a variable GETMAIN is issued which specifies such a large maximum value that most or all of the space remaining in the private area is allocated to the requestor. Therefore, it is strongly recommended that a region size be specified on the JOB or EXEC statement, or that the default region size for the job class not be zero.

After the IEALIMIT routine determines the appropriate limit, it must pass back to IEAVPRT0, via register 1, a numeric value that represents the imposed limit in bytes. As noted above, a zero returned in register 1 indicates that a limit is not imposed. The IEALIMIT routine should pass back, in register 0, a value that is less than that in register 1. Both register 0 and 1 should be rounded to a multiple of 4K. IEAVPRT0 stores register 0 as the REGION parameter value and register 1 as the IEALIMIT value for future reference (that is, for use in processing subsequent GETMAINS as described below).

The REGION parameter value (register 0) should be less than the IEALIMIT value (register 1) to protect against programs that issue variable-length GETMAINS with very large maximums, and then do not immediately FREEMAIN part of that space or FREEMAIN such a small amount that a subsequent GETMAIN (possibly issued by a system service) causes the job to fail. As an example, suppose that a program issues a variable-length GETMAIN with a maximum of 2^{24-1} bytes. If the REGION parameter value is not less than the IEALIMIT value, all the space in the region up to the IEALIMIT value is allocated, and any subsequent GETMAIN that cannot be satisfied from free space in an already existing subpool causes the job to fail. If however, the REGION parameter value is made less than the IEALIMIT value, only space up to the REGION parameter value is allocated for the GETMAIN. Thus, an amount of space equal to the IEALIMIT value minus the REGION parameter value remains for subsequent GETMAINS.

The REGION parameter value specifies the maximum amount of storage that can be allocated to a job by any single variable-length GETMAIN request. The IEALIMIT value specifies the maximum total storage that can be allocated to a job by any combination of GETMAINS. The relationship between the REGION parameter value and the IEALIMIT value and their effect upon both fixed-length and variable-length GETMAINS is shown in Figure 28.

Type of GETMAIN	Results
Fixed-length:	
IEALIMIT value minus currently alloc space \geq request	Satisfied
IEALIMIT value minus currently alloc space $<$ request	Rejected
Variable-length: (REGION parm $<$ IEALIMIT or REGION parm $>$ IEALIMIT and IEALIMIT=0)	
REGION parm minus currently alloc space \geq max	Maximum is allocated.
min \leq REGION parm minus currently alloc space \leq max	Unallocated is allocated.
REGION parm minus currently alloc space \leq min	Minimum is allocated as long as IEALIMIT is not exceeded (in which case the request fails unless IEALIMIT=0).
Variable-length: (REGION parm \geq IEALIMIT and IEALIMIT \neq 0)	
IEALIMIT minus currently alloc space \geq max	Maximum is allocated.
min \leq IEALIMIT minus currently alloc space \leq max	Unallocated amount is allocated.
IEALIMIT minus currently alloc space \leq min	Minimum is allocated as long as IEALIMIT is not exceeded (in which case the request fails).

Figure 28. The Effects of IEALIMIT and REGION Values on Various GETMAINS

For example, assume that application program A has the following characteristics:

IEALIMIT value	150K
REGION-parameter value	100K
Currently allocated space	80K

Program A issues the following variable-length GETMAIN requests in the order indicated:

1. Request 5K-10K: 10K is allocated; currently allocated space is now 90K. Because the amount currently allocated (80K) does not exceed the REGION-parameter value (100K) and because the amount unallocated (20K - relative to the REGION-parameter value) is greater than the maximum amount requested (10K), the maximum is allocated.
2. Request 5K-100K: 10K is allocated; currently allocated space is now 100K. Because the amount unallocated (10K - relative to the REGION-parameter value) is between the minimum and maximum, the amount unallocated is allocated.
3. Request 40K-100K: 40K is allocated; currently allocated space is now 140K. Because the amount unallocated (0K - relative to the REGION-parameter value) is less than the minimum amount requested (40K), the minimum amount is allocated.
4. Request 15K-50K: the GETMAIN request fails. The amount unallocated (0K - relative to the REGION-parameter value) is less than the minimum amount requested (15K). If the minimum amount were allocated, the currently allocated amount would become 155K, which exceeds the IEALIMIT value (150K). Therefore, the request fails.

Part II: Reference — Macro Instructions

You can communicate service requests to the control program using a set of macro instructions provided by IBM. Most of the macro instructions have no restrictions on use by applications programmers, however, use of some of the macro instructions should be restricted to systems programmers and installation-approved personnel.

This section describes those Supervisor macro instructions that should be installation-controlled. Some macro instructions should be totally restricted in use; these are described fully in this book. Other macro instructions are not restricted in use, but do contain some parameters that should be restricted; in these cases, only the parameters that should be restricted are fully described in this book. In all cases, however, the format of the complete macro instruction is described.

Figure 29 lists all macro instructions described in this book, and indicates which ones are fully described and which ones are partially described.

Macro Instruction	Fully Described	Partially Described
CALLDISP	X	
CALLRTM	X	
CHANGKEY	X	
CIRB	X	
DSGNL	X	
EXTRACT	X	
FESTAE	X	
MODESET	X	
NIL	X	
OIL	X	
PGFIX	X	
PGFREE	X	
PURGEDQ	X	
QEDIT	X	
RACDEF	X	
RACINIT	X	
RACLIST	X	
RESERVE	X	
RESUME	X	
RISGNL	X	
RPSGNL	X	
SCHEDULE	X	
SDUMP	X	
SETFRR	X	
SETLOCK	X	
SPOST	X	
STAE	X	
SUSPEND	X	
SYNCH	X	
TCTL	X	
TESTAUTH	X	
T6EXIT	X	
ATTACH		X
DEQ		X
ENQ		X
ESTAE		X
EVENTS		X
FREEMAIN		X
GETMAIN		X
POST		X
RACHECK		X
SETRP		X
SPIE		X
STATUS		X
WTO		X
WTOR		X

Figure 29. Macro Instruction Coverage

The macro instructions are available only when programming in the assembler language, and are processed by the assembler program using macro definitions supplied by IBM and placed in the macro library when the system was generated. The processing of the macro instruction by the assembler program results in a macro expansion, generally consisting of data and executable instructions in the form of assembler language statements. The data fields are the parameters to be passed to the requested control program routine. The executable instructions generally consist of a branch around the data, instructions to load registers, and either a branch instruction or a supervisor call (SVC) to give control to the proper program. The exact macro expansion appears as part of the assembler output listing.

Macro Instruction Forms

When written in the standard form, some of the macro instructions result in instructions that store into an inline parameter list. The option of storing into an out-of-line parameter list is provided to allow the use of these macro instructions in a reenterable program. You can request this option through the use of list and execute forms. When list and execute forms exist for a macro instruction, their descriptions follow the description of the standard form.

Use the list form of the macro instruction to provide a parameter list to be passed either to the control program or to a problem program, depending on the macro instruction. The expansion of the list form contains no executable instructions; therefore registers cannot be used in the list form.

Use the execute form of the macro instruction in conjunction with one or two parameter lists established using the list form. The expansion of the execute form provides the executable instructions required to modify the parameter lists and to pass control to the required program.

The descriptions of the following macro instructions assume that the standard begin, end, and continue columns are used -- for example, column 1 is assumed as the begin column. To change the begin, end, and continue columns, code the ICTL instruction to establish the coding format you wish to use. If you do not use ICTL, the assembler recognizes the standard columns. To code the ICTL instruction, see *OS/VS - DOS/VS - VM/370 Assembler Language*.

Coding the Macro Instructions

The table appearing near the beginning of each macro instruction indicates how the macro instruction is to be coded. The table does not attempt to explain the meanings of the parameters; the parameters are explained following the table.

Figure 30 presents a sample macro instruction, TEST, and summarizes all the coding information that is available for it. The table is divided into three columns.

(A)	(B)	(C)
<i>name</i>		<i>name</i> : symbol. Begin <i>name</i> in column 1.
	b	One or more blanks must precede TEST.
(A1) →	TEST	
	b	One or more blanks must follow TEST.
<hr/>		
	MATH	
(A2) →	HIST	
	GEOG	
	,DATA= <i>data addr</i>	<i>data addr</i> : RX-type address, or register (2) - (12).
(B1) →	,LNG= <i>data length</i>	<i>data length</i> : symbol or decimal digit, with a maximum value of 256.
	,FMT=HEX	Default: FMT=HEX
(B2) →	,FMT=DEC	
	,FMT=BIN	
	,PASS= <i>value</i>	<i>value</i> : symbol, decimal digit, or register (1) or (2) - (12). Default: PASS=65

Figure 30. Sample Macro Instruction

- The first column, (A), contains those parameters that are required for that macro instruction. If a single line appears in that column, (A1), the parameter on that line is required and must be coded. If two or more lines appear together, (A2), the parameter appearing on one and only one of the lines must be coded.
- The second column, (B), contains those parameters that are optional for that macro instruction. If a single line appears in that column, (B1), the parameter on that line is optional. If two or more lines appear together, (B2), although the entire parameter is optional, if you elect to make an entry, one and only one of the lines may be coded.
- The third column, (C), provides additional information for coding the macro instruction. When substitution of a variable is required, the following classifications should be understood:

symbol: any symbol valid in the assembler language. That is, an alphabetic character followed by 0-7 alphameric characters, with no special characters and no blanks.

decimal digit: any decimal digit up to the value indicated in the parameter description. If both symbol and decimal digit are indicated, an absolute expression is also allowed.

register (2) - (12): one of general registers 2 through 12, specified within parentheses, previously loaded with the right-adjusted value or address indicated in the parameter description. The unused high-order bits must be set to zero. The register may be designated symbolically or with an absolute expression.

register (0): general register 0, previously loaded as indicated under register (2) - (12) above. Designate the register as (0) only.

register (1): general register 1, previously loaded as indicated under register (2) - (12) above. Designate the register as (1) only.

RX-type address: any address that is valid in an RX-type instruction (for example, LA).

A-type address: any address that may be written in an A-type address constant.

default: a value that is used in default of a specified value, and that is assumed if the parameter is not coded.

Use the parameters to specify the services and options to be performed, and write them according to the following general rules:

- If the selected parameter is written in all capital letters (for example, STEP, DUMP, or RET=USE), code the parameter exactly as shown.
- If the selected parameter is written in italics (for example, *value* or *comp code*), substitute the indicated value, address, or name.
- If the selected parameter is a combination of capital letters and italics separated by an equal sign (for example, EP=*entry point*), code the capital letters and equal sign as shown, and then make the indicated substitution for the italics.
- Read the table from top to bottom.
- Code commas and parentheses exactly as shown.
- Positional parameters (parameters without equal signs) appear first and must be coded in the order shown. Keyword parameters (parameters with equal signs) may be coded in any order.
- If a parameter is selected, read column 3 before proceeding to the next parameter. Column 3 will often contain notes pertaining to restrictions on coding the parameter.

Continuation Lines

You can continue the parameter field of a macro instruction on one or more additional lines according to the following rules:

1. Enter a continuation character (not blank, and not part of the parameter coding) in column 72 of the line.
2. Continue the parameter field on the next line, starting in column 16. All columns to the left of column 16 must be blank.

You can code the parameter field being continued in one of two ways. Code the parameter field through column 71, with no blanks, and continue in column 16 of the next line; or truncate the parameter field by a comma, where a comma normally falls, with at least one blank before column 71, and then continue in column 16 of the next line. Figure 31 shows an example of each method.

NAME1	OP 1	OPERAND1, OPERAND2, OPERAND3, OPERAND4, OPERAND5, OPERAND6	X
			THIS IS ONE WAY
NAME2	OP 2	OPERAND1, OPERAND2, OPERAND3,	X
		OPERAND4	WAY
			X

Figure 31. Continuation Coding

ATTACH — Create a New Task

The ATTACH macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the JSTCB, SM, SVAREA, KEY, DISP, JSCB, TID, NSHSPV, NSHSPL, and RSAPF parameters. These parameters are restricted in use and should only be used with tasks in supervisor mode and/or having a system protection key. If they are used with other tasks, the default values are used.

The syntax of the complete ATTACH macro instruction is shown as follows. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

The standard form of the ATTACH macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede ATTACH.
ATTACH	
b	One or more blanks must follow ATTACH.
EP= <i>entry name</i>	<i>entry name</i> : symbol.
EPLOC= <i>entry name addr</i>	<i>entry name addr</i> : A-type address, or register (2) - (12).
DE= <i>list entry addr</i>	<i>list entry addr</i> : A-type address, or register (2) - (12).
,DCB= <i>dcb addr</i>	<i>dcb addr</i> : A-type address, or register (2) - (12).
,LPMOD= <i>limit prior nmb</i>	<i>limit prior nmb</i> : symbol, decimal digit, or register (2) - (12).
,DPMOD= <i>disp prior nmb</i>	<i>disp prior nmb</i> : symbol, decimal digit, or register (2) - (12).
,PARAM=(<i>addr</i>)	<i>addr</i> : A-type address, or register (2) - (12).
,PARAM=(<i>addr</i>),VL=1	Note: <i>addr</i> is one or more addresses, separated by commas. For example, PARAM=(<i>addr,addr,addr</i>)
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (2) - (12).
,ETXR= <i>exit rtn addr</i>	<i>exit rtn addr</i> : A-type address, or register (2) - (12).
,GSPV= <i>subpool nmb</i>	<i>subpool nmb</i> : symbol, decimal digit, or register (2) - (12).
,GSPL= <i>subpool list addr</i>	<i>subpool list addr</i> : A-type address, or register (2) - (12).
,SHSPV= <i>subpool nmb</i>	<i>subpool nmb</i> : symbol, decimal digit, or register (2) - (12).
,SHSPL= <i>subpool list addr</i>	<i>subpool list addr</i> : A-type address, or register (2) - (12).
,SZERO=YES	Default: SZERO=YES
,SZERO=NO	
,TASKLIB= <i>dcb addr</i>	<i>dcb addr</i> : A-type address, or register (2) - (12).
,STAI=(<i>exit addr</i>)	<i>exit addr</i> : A-type address, or register (2) - (12).
,STAI=(<i>exit addr,parm addr</i>)	<i>parm addr</i> : A-type address, or register (2) - (12).
,ESTAI=(<i>exit addr</i>)	
,ESTAI=(<i>exit addr,parm addr</i>)	Note: PURGE may be specified only if STAI or ESTAI is specified.
,PURGE=QUIESCE	Default for STAI: PURGE=QUIESCE
,PURGE=NONE	Default for ESTAI: PURGE=NONE
,PURGE=HALT	
,ASYNCH=NO	Note: ASYNCH may be coded only if STAI or ESTAI is specified.
,ASYNCH=YES	Default for STAI: ASYNCH=NO
	Default for ESTAI: ASYNCH=YES
,TERM=NO	Note: TERM may be specified only if ESTAI is specified.
,TERM=YES	Default: TERM=NO
,JSTCB=NO	Default: JSTCB=NO
,JSTCB=YES	
,SM=PROB	Default: SM=PROB
,SM=SUPV	
,SVAREA=YES	Default: SVAREA=YES
,SVAREA=NO	
,KEY=PROP	Default: KEY=PROP
,KEY=ZERO	
,DISP=YES	Default: DISP=YES
,DISP=NO	
,JSCB= <i>jscb addr</i>	<i>jscb addr</i> : A-type address, or register (2) - (12).
,TID= <i>task id</i>	<i>task id</i> : decimal digits 0-255, or register (2) - (12).
	Default: TID=0
,NSHSPV= <i>subpool nmb</i>	<i>subpool nmb</i> : symbol, decimal digit, or register (2) - (12).
,NSHSPL= <i>subpool list addr</i>	<i>subpool list addr</i> : A-type address, or register (2) - (12).
,RSAPF=NO	Default: RSAPF=NO
,RSAPF=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.)

,JSTCB = NO

,JSTCB = YES

specifies whether the attached task is a new job step (YES) or a task in the present job step (NO). If YES is specified, the address of the TCB of the newly created task is placed in the TCBJSTCB field of the TCB; if NO is specified, the TCBJSTCB field of the task using ATTACH is propagated to the new task.

,SM = PROB

,SM = SUPV

specifies that the system is to run in problem program mode (PROB) or in supervisor mode (SUPV) when executing the attached task.

,SVAREA = YES

,SVAREA = NO

specifies whether a save area is needed for the attaching task. If YES is specified, the ATTACH routine obtains a 72-byte save area. If both attaching and attached task share subpool zero, the save area is obtained there; otherwise, it is obtained from a new 4K-byte block.

,KEY = PROP

,KEY = ZERO

specifies whether the protection key of the newly created task should be zero (ZERO) or copied from the TCBPKF field of the TCB for the task using ATTACH (PROP).

,DISP = YES

,DISP = NO

specifies whether the subtask is to be dispatchable (YES) or nondispatchable (NO). (Note: If DISP=NO is specified, the attaching task must use the STATUS macro instruction to reset the TCBANDSP nondispatchability bit to 0, before the ATTACH processing can be completed for the new task.)

,JSCB = *jscb addr*

specifies the address of the job step control block. If specified, the JSCB is used for the new task. Otherwise, the JSCB of the attaching task is also used for the new task.

,TID = *task id*

specifies the task identifier to be placed in the TCBTID field of the attached task.

,NSHSPV = *subpool nmb*

,NSHSPL = *subpool list addr*

specifies the virtual storage subpool number 236 or 237, or the address of a list of virtual storage subpool numbers 236 and 237. The subpools specified will not be shared with the subtask.

If NSHSPL is specified, the first byte of the list contains the number of bytes remaining in the list; each of the following bytes contains a virtual storage subpool number.

,RSAPF = YES

specifies that the attached subtask may come from an unauthorized library. If, however, it comes from an APF-authorized library and is link-edited with the APF-authorized attribute, then the step begins execution with APF authorization.

The default (RSAPF=NO) applies unless the following conditions are met:

- The caller is running in supervisor state, system key (0-7), or both.
- The caller is running non-APF authorized.
- The attached task is ATTACHED in the problem program state and with a non-system key.

,RELATED = *value*

specifies information used to self-document macro instructions by "relating" functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
ATTCH1 ATTACH      EP=MYJOB,ECB=MYECB,RELATED=( DETCH1 ,
                  'CREATE SUBTASK' )
DETCH1 DETACH      ( 1 ),RELATED=( ATTCH1 , 'DETACH SUBTASK' )
```

When control is returned, register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	Successful completion.
04	ATTACH was issued in a STAE exit; processing not completed.
08	Insufficient storage available for control block for STAI/ESTAI request; processing not completed.
0C	Invalid exit routine address or invalid parameter list address specified with STAI parameter; processing not completed.
14	Authorized task specifying JSTCB=YES was not itself a job step task; processing not completed.
18	Attempt to create a new subtask would result in both job step tasks and non-job step tasks being subtasks of current task; processing not completed.

Note: For any return code other than 00, register 1 is set to zero upon return.

Note: The program manager processing for ATTACH is performed under the new subtask after control has been returned to the originating task. Therefore, it is possible for the originating task to obtain return code 00, and still not have the subtask successfully created (for example, if the entry name could not be found by the program manager). In such cases, the new subtask is abnormally terminated.

Example 1

Operation: Attach program SYSPROGM, which will run with protection key 0 and in supervisor mode. Subpool 0 is not to be shared, and the new task is not to have a savearea provided.

```
ATTACH EP=SYSPROGM,KEY=ZERO,SM=SUPV,SZERO=NO,SVAREA=NO
```

Example 2

Operation: Attach as a new job step the program name addressed in register 7. The new task is to run in problem program mode, a savearea is to be provided, a job step control block is provided, subpool 0 is not to be shared, a task library DCB is provided, and the new task is to be nondispatchable.

```
ATTACH EPLOC=( 7 ),SM=PROB,JSTCB=YES,SVAREA=YES,SZERO=NO,
      JSCB=( 5 ),DISP=NO,TASKLIB=( 8 )
```

CALLDISP — Force Dispatcher Entry

The CALLDISP macro instruction expands into an SVC that results in the caller's status being saved in the current TCB/RB, then the dispatcher is entered. The dispatcher then searches for the highest priority ready work to dispatch. When this task is redispached, control is returned to the next sequential instruction.

The CALLDISP macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
‡	One or more blanks must precede CALLDISP.
CALLDISP	
‡	One or more blanks must follow CALLDISP.

BRANCH=NO	<i>Default</i> : NO
BRANCH=YES	

The parameters are explained as follows:

BRANCH=NO

BRANCH=YES

specifies whether the branch entry (BRANCH=YES) or the SVC entry to the dispatcher is to be used. The default, BRANCH=NO, results in the use of the SVC entry.

Example 1

Operation: Pass control to a higher priority ready task.

```
CALLDISP
```

CALLRTM — Call Recovery/Termination Manager

The CALLRTM macro instruction is used to direct the services of the recovery/termination manager to a task or address space other than itself or its callers. The recovery/termination manager selects the appropriate recovery or termination process according to the status of the system and the requests of its invokers.

CALLRTM may be used only by key zero supervisor state routines. After execution of the macro instruction, control usually returns to the caller. If, however, the current task is abnormally terminated (ABTERM), control does not return to the caller. If the current address space is terminated (MEMTERM), control might or might not return to the caller before the MEMTERM takes effect.

The CALLRTM macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
	One or more blanks must precede CALLRTM.
CALLRTM	
	One or more blanks must follow CALLRTM.

TYPE=ABTERM	
TYPE=MEMTERM	
,COMPCOD= <i>comp code</i>	<i>comp code</i> : symbol, decimal digit, or register (2) - (12).
,ASID= <i>asid</i>	<i>asid</i> : decimal digits 0 - 32,765 or register (2) - (15).
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : 0, or register (2) - (12). Note: This parameter may only be specified with TYPE=ABTERM.
,DUMP=YES	Default: DUMP=YES
,DUMP=NO	Note: This parameter may only be specified with TYPE=ABTERM.
,STEP=NO	Default: STEP=NO
,STEP=YES	Note: This parameter may only be specified with TYPE=ABTERM.
,DUMPOPT= <i>parm list addr</i>	<i>parm list addr</i> : register (3) - (15).

The parameters are explained as follows:

TYPE = ABTERM
TYPE = MEMTERM

specifies that the services of the recovery/termination manager are being directed towards another task (ABTERM) or that an address space is to be terminated (MEMTERM). For MEMTERM, all recovery processing in the address space is circumvented.

Note: If MEMTERM is specified, the caller must pass the address of an 18-word save area in register 13.

,COMPCOD = comp code

specifies the system completion code associated with the abnormal termination. This parameter can be specified as a hexadecimal code (x'80A'), a decimal code (2058), or a register containing a hexadecimal code; in all cases, the result is hexadecimal.

,ASID = asid

specifies the address space identifier of the address space to be terminated (for MEMTERM) or the address space identifier of the address space containing the TCB of the task to be terminated (for ABTERM).

,TCB = tcb addr

specifies the TCB address of the task to be terminated.

,DUMP = YES

,DUMP = NO

specifies that a dump is (YES) or is not (NO) to be taken. If the DUMPOPT parameter is not also specified, the contents of the dump are defined by the //SYSABEND, //SYSMDUMP, or //SYSUDUMP DD statement and the system or user defined defaults.

,STEP = NO

,STEP = YES

specifies that the entire job step is (YES) or is not (NO) to be abnormally terminated.

,DUMPOPT = parm list addr

specifies the address of a parameter list valid for the SNAP macro instruction. The parameter list is used to produce a tailored dump, and may be created using the list form of the SNAP macro instruction, or a compatible list may be created. The system dump options specified by the CHNGDUMP operator command can add to or override this parameter list. All recovery routines entered for the failure can also add to the list of dump options. The TCB, DCB, and STRHDR options available on SNAP are ignored if they appear in the parameter list; the TCB used is for the task that received ABEND, the DCB used is provided by the ABDUMP routine. If a //SYSABEND, //SYSMDUMP, or //SYSUDUMP DD statement is not provided, the DUMPOPT parameter is ignored.

Example 1

Operation: Terminate the current address space with a completion code of 123.

```
CALLRTM      TYPE=MEMTERM,COMPCOD=123,ASID=0
```

Example 2

Operation: Schedule ABTERM of the TCB whose address is specified in register 8. The ABTERM of this TCB takes place in the address space identified by the ASID specified in register 5, and has a completion code of 123.

```
CALLRTM      TYPE=ABTERM,COMPCOD=123,ASID=( 5 ),TCB=( 8 )
```


CHANGKEY — Change Virtual Storage Protection Key

The CHANGKEY macro instruction changes the protection key and fetch protection status of one or more pages of virtual storage. The CHANGKEY function is available only for use by system components that execute in supervisor state under protection key zero.

The CHANGKEY function is valid for virtual storage obtained by GETMAIN in page multiples from problem program subpools. The requesting program must hold the local lock, but should not hold the SALLOC because execution of the CHANGKEY macro instruction results in the release of SALLOC.

The CHANGKEY macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
<i>b</i>	One or more blanks must precede CHANGKEY.
CHANGKEY	
<i>b</i>	One or more blanks must follow CHANGKEY.

R,BA= <i>page addr</i> ,EA= <i>page addr</i>	<i>page addr</i> : A-type address or register (1)-(13). Note: The R-type macro expansion alters the contents of register 2. EA should not be specified as (1).
L,LISTAD: <i>list addr</i>	<i>list addr</i> : A-type address or register (1)-(13).
,KEY= <i>stor key</i>	<i>stor key</i> : Decimal digit 0-15 or register (0) or register (3)-(13).
,BRANCH=YES	Required.

The explanation of the parameters follows:

R,BA = *page addr*,EA = *page addr*

L,LISTAD = *list addr*

specifies the type of CHANGKEY request:

R indicates a request to change the key of a single area of virtual storage.

L indicates a request to change the key of one or more areas of virtual storage.

BA specifies the address of the first byte of the first page of the virtual storage area whose key is to be changed.

EA specifies the address of the first byte of the last page of the virtual storage area whose key is to be changed.

Note: $BA \leq EA$

LISTAD specifies the address of the first double-word of a variable length parameter list in fixed storage. The first word of each element is defined as BA above and the second word of each element as EA above. The last element in the parameter list is indicated by the high order bit in the second word when the bit is one.

,KEY = stor key

specifies the new storage key and fetch protection status for the virtual storage areas specified. If the *stor key* specification is a decimal digit, then the supervisor assumes the user wants fetch protection. If the user does not want fetch protection, he should specify the protection key he wants in bits 24-27 of a register and leave bit 28 at zero to indicate that he doesn't want fetch protection.

,BRANCH = YES

The only entry available into the CHANGKEY service routine is branch entry.

Note: The requestor must have addressability to the CVT.

Upon completion of the CHANGKEY macro instruction, register 15 contains a zero return code. Register 0 contains the original protection key and fetch protection flag of the first page in the virtual storage area specified in bits 24-28.

Example 1

Operation: Change the storage key and ensure fetch protection of a single page of virtual storage addressed by register 5.

```
CHANGKEY R,BA=(REG5),EA=(REG5),KEY=8,BRANCH=YES
```

Example 2

Operation: Change the storage key and ensure fetch protection of two non-contiguous pages of virtual storage addressed by PAGE1 and PAGE2 respectively.

```
CHANGKEY L,LISTAD=PLIST,KEY=10,BRANCH=YES
```

```
•  
•  
•
```

PLIST	DC 2A(PAGE1)	FIRST ELEMENT IN LIST
	DC A(PAGE2)	BA PART OF SECOND ELEMENT
	DC AL1(X'80')	INDICATES LAST ELEMENT IN LIST
	DC AL3(PAGE2)	EA PART OF SECOND ELEMENT

CIRB — Create Interruption Request Block

The CIRB macro instruction is included in SYS1.MACLIB and must be included in the system at system generation time if the macro instruction is to be used. The issuing of this macro instruction causes a supervisor routine (called the exit effector routine) to create an interruption request block (IRB). In addition, other parameters of this macro instruction may specify the building of a register save area and/or a work area to contain interruption queue elements, which are used by supervisor routines in scheduling the execution of user exit routines.

Note: The IRB address is returned in register 1.

The CIRB macro instruction is written as follows:

<i>name</i>	<i>name:</i> symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede CIRB.
CIRB	
␣	One or more blanks must follow CIRB.

<i>EP=entry point addr</i>	<i>entry point addr:</i> RX-type address, or register (0) or (2) - (12).
,KEY=PP	Default: KEY=PP
,KEY=SUPR	
,MODE=PP	Default: MODE=PP
,MODE=SUPR	
,SVAREA=NO	Default: SVAREA=NO
,SVAREA=YES	
,RETIQE=YES	Default: RETIQE=YES
,RETIQE=NO	
,STAB=(DYN)	
,WKAREA= <i>workarea size</i>	<i>workarea size:</i> Decimal digit, or register (2) - (12). Default: zero
,BRANCH=NO	Default: BRANCH=NO
,BRANCH=YES	
,RETRN=NO	Default: RETRN=NO
,RETRN=YES	Note: This parameter has meaning only if RETIQE=NO is specified above.

The parameters are explained as follows:

EP = entry point addr

specifies the address of the entry point of the user's asynchronous exit routine.

,KEY = PP

,KEY = SUPR

specifies whether the asynchronous exit routine operates with a key of zero (SUPR) or with a key obtained from the TCB of the task issuing the CIRB macro instruction (PP).

,MODE = PP

,MODE = SUPR

specifies whether the asynchronous exit routine executes in problem program (PP) or supervisor (SUPR) mode.

,SVAREA = NO

,SVAREA = YES

specifies whether to obtain a 72-byte register save area from the virtual storage assigned to the problem program. If a save area is requested, CIRB places the save area address in the IRB. The address of this area is passed to the user routine via register 13.

,RETIQE = YES

,RETIQE = NO

specifies whether the associated queue elements are request queue elements (YES) or interruption queue elements (NO).

,STAB = (DYN)

specifies that the IRB (including the work area) is freed by EXIT.

Branch Entry Interface

For **BRANCH=YES**, the branch entry interface is as follows:

- The caller must be in supervisor state, key zero, and own the local lock and no locks above the SALLOC in the locking hierarchy.
- The caller must pass a TCB address in register 4 to be used by GETMAIN when allocating space for the IRB and for the problem program save area. Also, if a problem key is specified in the **KEY=** parameter of the CIRB, the TCBPKF field of that TCB is used.
- Upon return, register 1 contains the address of the created IRB, registers 0, and 2-14 are unchanged, and register 15 is unpredictable.
- Control is returned in supervisor state, key zero, with the same locks held as on entry.

Note: If the **STAB** parameter is omitted from the CIRB macro instruction, the IRB will remain available for later use by the task issuing the macro.

,WKAREA = workarea size

specifies the size, in doublewords, of the work area to be included in the IRB. The area may be used to build IQE's. The maximum size is 255 double words.

,BRANCH = NO

,BRANCH = YES

specifies that branch linkage (YES) or SVC linkage (NO) to CIRB will be provided.

,RETRN = NO

,RETRN = YES

specifies that the IQE is (YES) or is not (NO) returned to the available queue when the asynchronous exit terminates.

Example 1

Operation: Create an IRB to be used in scheduling an asynchronous exit. The exit is scheduled via the IQE interface to Stage 2 Exit Effector, and receives control in the supervisor state. The IRB is freed when it terminates. The exit receives control at the IQERTN label.

```
CIRB    EP=IQERTN,MODE=SUPR,RETIQE=NO,STAB=(DYN),BRANCH=NO
```

Example 2

Operation: Create an IRB to be used in scheduling an asynchronous exit. The RQE interface to Stage 2 Exit Effector is used to schedule the routine. The exit gets control at the RQETEST label.

```
CIRB    EP=RQETEST,KEY=SUPR,MODE=SUPR,STAB=(DYN),BRANCH=NO
```

DEQ — Release a Serially Reusable Resource

The DEQ macro instruction is described in the *OS/VS2 Supervisor Services and Macro Instructions*, with the exception of the RMC, GENERIC, TCB, and UCB parameters. These parameters are restricted in use and should only be used with tasks that are authorized. The UCB parameter is used to release a device that was reserved with a RESERVE macro instruction.

The syntax of the complete DEQ macro instruction appears below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in *OS/VS2 Supervisor Services and Macro Instructions*.

The standard form of the DEQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede DEQ.
DEQ	
␣	One or more blanks must follow DEQ.

(
<i>qname addr</i>	<i>qname addr</i> : A-type address, or register (2) - (12).
, <i>rname addr</i>	<i>rname addr</i> : A-type address, or register (2) - (12).
,	<i>rname length</i> : symbol, decimal digit, or register (2) - (12).
, <i>rname length</i>	Note: <i>rname length</i> must coded if a register is specified for <i>rname addr</i> .
,	Default: STEP
,STEP	
,SYSTEM	
,SYSTEMS	
, <i>var1234</i>	<i>var1234</i> : The preceding 4 parameters may be repeated up to 65,535 times.
)	
,RET=HAVE	
,RET=NONE	
,RMC=NONE	Default: RMC=NONE
,RMC=STEP	Default: GENERIC=NO
,GENERIC=NO	Note: If GENERIC=YES is specified, you must also specify RET=HAVE above.
,GENERIC=YES	
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : A-type address, or register (2) - (12).
	Note: TCB cannot be specified with RMC above.
,UCB= <i>ucb addr</i>	<i>ucb addr</i> : RX-type address, or register (2) - (12).
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained as follows. (The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instructions*.)

,RMC = NONE
 ,RMC = STEP
 ,GENERIC = NO
 ,GENERIC = YES

specifies optional parameters available to the system programmer:

RMC specifies that the reset-must-complete function is not to be used (NONE) or that the requesting task is to release the resources and terminate the must complete function (STEP). The NONE or STEP subparameter must agree with the subparameter specified in the SMC parameter of the corresponding ENQ macro instruction.

GENERIC specifies whether or not (YES or NO) all queue elements for the task under the specified major name are dequeued, regardless of whether they have control of the resource.

,TCB = *tcb addr*

specifies a register that points to a TCB or specifies the address of a fullword on a fullword boundary that points to a TCB on whose behalf the DEQ is to be done. The caller (not the directed task) is abnormally terminated if the RET parameter is omitted and an attempt is made to DEQ a resource not requested or not owned by the directed task.

,UCB = *ucb addr*

specifies the address of a fullword that contains the address of a UCB for a reserved device that is now being released.

Return codes are provided by the control program only if RET=HAVE is designated. If all of the return codes for the resources named in DEQ are 0, register 15 contains 0. If any of the return codes are not 0, register 15 contains the address of a virtual storage area containing the return codes as shown in Figure 32.




Figure 32. Return Code Area Used by DEQ

The return codes are placed in the parameter list resulting from the macro expansion in the same sequence as the resource names in the DEQ macro instruction. The return codes are shown below.

Hexadecimal Code	Meaning
0	The resource has been released.
4	The resource has been requested for the task, but the task has not been assigned control. The task is not removed from the wait condition. (This return code could result if DEQ is issued within an exit routine which was given control because of an interruption.)
8	Control of the resource has not been requested by the active task, or the resource has already been released.

Example 1

Operation: Unconditionally release control of the resource in Example 1 of ENQ, and reset the "must-complete" state.

```
DEQ ( MAJOR1 , MINOR1 , 8 , STEP ) , RMC=STEP
```

Example 2

Operation: Conditionally release control of the resource in Example 2 of ENQ.

```
DEQ ( MAJOR2 , MINOR2 , 4 , SYSTEM ) , TCB=( R2 ) , RET=HAVE
```

Example 3

Operation: Unconditionally release control of the resource (device) in Example 1 of RESERVE.

```
DEQ ( MAJOR3 , MINOR3 , , SYSTEMS ) , UCB=( R3 )
```


DEQ (List Form)

The list form of the DEQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede DEQ.
DEQ	
␣	One or more blanks must follow DEQ.

(
<i>qname addr</i>	<i>qname addr</i> : A-type address.
,	<i>rname addr</i> : A-type address.
<i>rname addr</i>	
,	<i>rname length</i> : symbol or decimal digit.
<i>rname length</i>	
,	Default: STEP
,STEP	
,SYSTEM	
,SYSTEMS	
, <i>var1234</i>	<i>var1234</i> : The preceding 4 parameters may be repeated up to 65,535 times.
,RET=HAVE	Default: RET=NONE
,RET=NONE	
,RMC=NONE	Default: RMC=NONE
,RMC=STEP	Default: GENERIC=NO
,GENERIC=NO	Note: If GENERIC=YES is specified, you must also specify RET=HAVE above.
,GENERIC=YES	
,TCB=0	Note: TCB cannot be specified with RMC above, and must be specified on the list form if used on the execute form.
,UCB= <i>ucb addr</i>	<i>ucb addr</i> : A-type address.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=L	

The parameters restricted in use are explained under the standard form of the DEQ macro instruction. The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instructions**.

DEQ (Execute Form)

The execute form of the DEQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede DEQ.
DEQ	
␣	One or more blanks must follow DEQ.

(Note: (and) are the beginning and end of a parameter list. The entire list is optional. If nothing in the list is desired, then (,), and all parameters between (and) should not be specified. If something in the list is desired, then (,), and all parameters in the list should be specified as indicated at the left.
<i>qname addr</i>	<i>qname addr</i> : RX-type address, or register (2) - (12).
,	
<i>rname addr</i>	<i>rname addr</i> : RX-type address, or register (2) - (12).
,	
<i>rname length</i>	<i>rname length</i> : symbol, decimal digit, or register (2) - (12).
,	
,STEP	
,SYSTEM	
,SYSTEMS	
, <i>var1234</i>	<i>var1234</i> : The preceding 4 parameters may be repeated up to 65,535 times.
)	Note: See note opposite (above).
,RET=HAVE	
,RET=NONE	
,RMC=NONE	Note: If GENERIC=YES is specified, you must also specify RET=HAVE above.
,RMC=STEP	
,GENERIC=NO	
,GENERIC=YES	
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : RX-type address, or register (2) - (12). Note: TCB cannot be specified with RMC above, and must be specified on the execute form if used on the list form.
,UCB= <i>ucb addr</i>	<i>ucb addr</i> : RX-type address, or register (2) - (12).
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E , <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters restricted in use are explained under the standard form of the DEQ macro instruction. The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instructions**.

	Code	Action
SENSE	01	State of specified processor is to be sensed
START	04	Start function
STOP	05	Stop function
RESTART	06	Restart function
IPR	07	Initial program reset function
PR	08	Program reset function
SSS	09	Stop and store status function
IMPL	0A	Initial microprogram load function
ICPUR	0B	Initial processor reset function
CPUR	0C	processor reset function

Note: Codes 0A, 0B, and 0C are only valid on a Model 168.

,CPU = PCCA addr

specifies the address of the physical configuration communication area (PCCA) of the CPU on which the function is to be executed.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal

Code	Meaning
00	Function successfully initiated, but not necessarily completed.
04	Function not completed because access path to the addressed processor was busy or the addressed processor was in a state where it could not accept and respond to the function code.
08	Function unsuccessful. Status returned in register 0.
12	Specified CPU is either not installed, not configured into the system, or powered off.
16	CPU is a uniprocessor and does not have signal processor sending and receiving capabilities.

With a return code of 8, register 0 contains:

Bits	Meaning
0	Equipment check
1-23	Reserved
24	External call pending
25	Stopped
26	Operator intervening
27	Check stop
28	Not ready
29	Reserved
30	Invalid function
31	Receiver check

Example 1

Operation: The state of the processor whose PCCA is located at PCCA is requested. If the processor is executing or is in a wait state, register 15 contains a 0 return code; otherwise, register 0 contains a return code of 8 with status indicators.

DSGNL SENSE, CPU=PCCA

Example 2

Operation: The processor whose PCCA address is in register 1 will be placed in the STOP state.

DSGNL STOP, CPU=(1)

ENQ — Request Control of a Serially Reusable Resource

The ENQ macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the SMC, ECB, and TCB parameters. These parameters are restricted in use and should only be used with tasks that are authorized.

The syntax of the complete ENQ macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

The standard form of the ENQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede ENQ.
ENQ	One or more blanks must follow ENQ.
␣	

(
<i>qname addr</i>	<i>qname addr</i> : A-type address, or register (2) - (12).
, <i>rname addr</i>	<i>rname addr</i> : A-type address, or register (2) - (12).
,	Default: E
,E	
,S	
,	<i>rname length</i> : symbol, decimal digit, or register (2) - (12).
, <i>rname length</i>	Default: assembled length of <i>rname</i>
	Note: <i>rname length</i> must be coded if a register is specified for <i>rname addr</i>.
,	Default: STEP
,STEP	
,SYSTEM	
,SYSTEMS	
, <i>var12345</i>	<i>var12345</i> : The preceding 5 parameters may be repeated up to 65,535 times.
)	
,RET=CHNG	Default: RET=NONE
,RET=HAVE	
,RET=TEST	
,RET=USE	
,RET=NONE	
,SMC=NONE	<i>ecb addr</i> : A-type address, or register (2) - (12).
,SMC=STEP	<i>tcb addr</i> : A-type address, or register (2) - (12).
,ECB= <i>ecb addr</i>	Default: SMC=NONE
,TCB= <i>tcb addr</i>	Note: ECB cannot be specified with RET above. ECB and TCB can be specified together. If TCB and ECB are specified, then RET cannot be specified. If TCB is specified but not ECB, then RET=CHNG, TEST or USE must be specified above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.)

,SMC = NONE
 ,SMC = STEP
 ,ECB = *ecb addr*
 ,TCB = *tcb addr*

specifies optional parameters available to the system programmer:

SMC specifies that the set-must-complete function is not to be used (NONE) or that it is to place other tasks for the step nondispatchable until the requesting task has completed its operations on the resource (STEP).

ECB specifies the address of an ECB, and conditionally requests all of the resources named in the macro instruction. If the return code for one or more requested resources is 4 and the request is not nullified by a corresponding DEQ, the ECB is posted when all the requested resources (specifically, those that initially received a return code of 4) are assigned to the requesting task.

TCB specifies a register that points to a TCB or specifies the address of a fullword on a fullword boundary that points to a TCB on whose behalf the ENQ is to be done.

Return codes are provided by the control program only if you specify RET=TEST, RET=USE, RET=CHNG, RET=HAVE, or ECB=; otherwise return of the task to the active condition indicates that control of the resource has been assigned to the task. If all return codes for the resources named in the ENQ macro instruction are 0, register 15 contains 0. If any of the return codes are not 0, register 15 contains the address of a storage area containing the return codes, as shown in Figure 33.




Figure 33. Return Code Area Used by ENQ

The return codes are placed in the parameter list resulting from the macro expansion in the same sequence as the resource names in the ENQ macro instruction. The return codes are shown below.

Hexadecimal Code	Meaning
0	For RET=TEST, the resource was immediately available. For RET=USE, RET=HAVE, or ECB=, control of the resource has been assigned to the active task.
4	For RET=CHNG, the status of the resource has been changed to exclusive. For RET=TEST or RET=USE, the resource is not immediately available. For RET=CHNG, the status cannot be changed to exclusive. For ECB=, the ECB will be posted when available.
8	For RET=TEST, RET=USE, RET=HAVE, or ECB=, a previous request for control of the same resource has been made for the same task. Task has control of resource. For RET=CHNG, the resource has not been queued.
20	If bit 3 is on - shared control of resource; if bit 3 is off - exclusive control. A previous request for control of the same resource has been made for the same task. Task does not have control of resource.

Example 1

Operation: Unconditionally request exclusive control of a serially reusable resource that is known only within the address space (STEP), and place other tasks for the step nondispatchable until the requesting task has completed its operations on the resource.

```
ENQ      ( MAJOR1 , MINOR1 , E , 8 , STEP ) , SMC=STEP
```

Example 2

Operation: Conditionally request control of a serially reusable resource in behalf of another task. The resource is known by more than one address space, and is only wanted if immediately available.

```
ENQ      ( MAJOR2 , MINOR2 , S , 4 , SYSTEM ) , TCB=( R2 ) , RET=USE
```


ENQ (List Form)

The list form of the ENQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede ENQ.
ENQ	
␣	One or more blanks must follow ENQ.

(
<i>qname addr</i>	<i>qname addr</i> : A-type address.
,	
<i>rname addr</i>	<i>rname addr</i> : A-type address.
,	
E	Default: E
,	
S	
,	
<i>rname length</i>	<i>rname length</i> : symbol or decimal digit. Default: assembled length of <i>rname</i>
,	
STEP	Default: STEP
,	
SYSTEM	
,	
SYSTEMS	
,	
<i>var12345</i>	<i>var12345</i> : The preceding 5 parameters may be repeated up to 65,535 times.
)	
,	
RET=CHNG	Default; RET=NONE
,	
RET=HAVE	
,	
RET=TEST	
,	
RET=USE	
,	
RET=NONE	
,	
SMC=NONE	<i>ecb addr</i> : A-type address.
,	Default: SMC=NONE
SMC=STEP	
,	Note: ECB cannot be specified with RET above.
ECB= <i>ecb addr</i>	Note: TCB or ECB must be specified on the list form if it is used on the execute form. ECB and TCB can be specified together. If TCB and ECB are specified, then RET cannot be specified. If TCB is specified but not ECB, then RET=CHNG, TEST or USE must be specified above.
,	
TCB=0	
,	
RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,	
MF=L	

The parameters restricted in use are explained under the standard form of the ENQ macro instruction. The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instructions*.

ENQ (Execute Form)

The execute form of the ENQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede ENQ.
ENQ	
␣	One or more blanks must follow ENQ.
(Note: (and) are the beginning and end of a parameter list. The entire list is optional. If nothing in the list is desired then (,) and all parameters between (and) should not be specified. If something in the list is desired, then (,) and all parameters in the list should be specified as indicated at the left.
<i>qname addr</i>	<i>qname addr</i> : RX-type address, or register (2) - (12).
,	
<i>rname addr</i>	<i>rname addr</i> : RX-type address, or register (2) - (12).
,	
,E	
,S	
,	
<i>rname length</i>	<i>rname length</i> : symbol, decimal digit, or register (2) - (12).
,	
,STEP	
,SYSTEM	
,SYSTEMS	
, <i>var12345</i>	<i>var12345</i> : The preceding 5 parameters may be repeated up to 65,535 times.
)	Note: See note opposite (above.
,RET=CHNG	
,RET=HAVE	
,RET=TEST	
,RET=USE	
,RET=NONE	
,SMC=NONE	<i>ecb addr</i> : A-type address, or register (2) - (12).
,SMC=STEP	<i>tcg addr</i> : A-type address, or register (2) - (12).
,ECB= <i>ecb addr</i>	Note: ECB cannot be specified with RET above.
,TCB= <i>tcg addr</i>	Note: ECB and TCB can be specified together. If TCB and ECB are specified, then RET cannot be specified. If TCB is specified but not ECB, then RET=CHNG, TEST or USE must be specified above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters restricted in use are explained under the standard form of the ENQ macro instruction. The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instructions**.

Note: If ECB (or TCB) is specified in the execute form, ECB (or TCB=0) must be specified in the list form.

ESTAE — Extended STAE

The ESTAE macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the RECORD, BRANCH, and SVEAREA parameters.

The syntax of the complete ESTAE macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

The standard form of the ESTAE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede ESTAE.
ESTAE	
␣	One or more blanks must follow ESTAE.

<i>exit addr</i>	<i>exit addr</i> : A-type address, or register (2) - (12).
0	
,CT	Default: CT
,OV	
,PARAM= <i>list addr</i>	<i>list addr</i> : A-type address, or register (2) - (12).
,XCTL=NO	Default: XCTL=NO
,XCTL=YES	
,PURGE=NONE	Default: PURGE=NONE
,PURGE=QUIESCE	
,PURGE=HALT	
,ASYNCH=YES	Default: ASYNCH=YES
,ASYNCH=NO	
,TERM=NO	Default: TERM=NO
,TERM=YES	
,BRANCH=NO	Default: BRANCH=NO
,BRANCH=YES, SVEAREA= <i>save area</i>	<i>save addr</i> : A-type address, or register (2) - (12) or (13).
KEY=SAVE	<i>storage key</i> : any numeral in the range 0-15
KEY= <i>storage key</i>	
,RECORD=NO	Default: RECORD=NO
,RECORD=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters reserved for supervisor use are explained as follows. (The explanation of the other parameters is as explained in **OS/VS2 Supervisor Services and Macro Instructions**.)

,PURGE=QUIESCE

specifies that all outstanding requests for I/O operations are to be saved when the first ESTAE exit is taken (QUIESCE). The user's retry routine can restore the outstanding I/O requests.

,BRANCH=NO

,BRANCH=YES,SVEAREA=*save addr*

specifies that an SVC 60 entry to the ESTAE service routine is to be performed (NO) or that a branch entry is to be performed (YES). The save area is a 72-byte area used to save the general registers. If the caller is not in key zero, the KEY parameter must be specified.

KEY=SAVE

KEY=*storage key*

specifies that supervisor state users who are not in key zero may use the branch entry interface to the ESTAE service processor.

If the user specifies KEY=SAVE, the system saves the current PSW protection key in register 2 and issues a set protection key instruction (SPKA) to change to protection key zero. When the ESTAE service processor returns control, the original PSW key is restored from register 2. Because of the above, the user should save register 2 before the macro expansion and restore it afterwards. Specifying KEY=SAVE causes register 2 to lose its contents during the macro expansion.

On the other hand, if the user knows the current PSW protection key, he may specify it directly in the form KEY=(0-15) to eliminate saving and restoring the original protection key. This procedure eliminates an IPK instruction and prevents the use of register 2 in the macro expansion.

,RECORD=NO

,RECORD=YES

specifies that the SDWA workarea is not written to SYS1.LOGREC (NO) or that the entire SDWA workarea, both fixed and variable, is written to SYS1.LOGREC (YES).

Control returns to the instruction following the ESTAE macro instruction. When control returns, register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	Successful completion of ESTAE request.
04	ESTAE OV was specified with a valid exit address, but the current exit is either nonexistent, not owned by the user's RB, or is not an ESTAE exit.
0C	Cancel or an exit address equal to zero was specified, and either there are no exits for this TCB, the most recent exit is not owned by the caller, or the most recent exit is not an ESTAE exit.
10	An unexpected error was encountered while processing this request.
14	ESTAE was unable to obtain storage for an SCB.

Example 1

Operation: Take the ESTAE exit specified by register 4, allow asynchronous exit processing, do not allow special error processing, do not branch enter SVC 60, and default to CT (create) and PURGE=NONE.

ESTAE (4), ASYNCH=YES, TERM=NO, BRANCH=NO

ESTAE (List Form)

The list form of the ESTAE macro instruction is used to construct a remote control program parameter list.

The list form of the ESTAE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede ESTAE.
ESTAE	
␣	One or more blanks must follow ESTAE.

<i>exit addr</i>	<i>exit addr</i> : A-type address.
,PARAM= <i>list addr</i>	<i>list addr</i> : A-type address.
,PURGE=NONE	Default: PURGE=NONE
,PURGE=QUIESCE	
,PURGE=HALT	
,ASYNCH=YES	Default: ASYNCH=YES
,ASYNCH=NO	
,TERM=NO	Default: TERM=NO
,TERM=YES	
,RECORD=NO	Default: RECORD=NO
,RECORD=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=L	

The parameters restricted in use are explained under the standard form of the ESTAE macro instruction. The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instructions*.

ESTAE (Execute Form)

A remote control program parameter list is used in, and can be modified by, the execute form of the ESTAE macro instruction. The control program parameter list can be generated by the list form of the ESTAE macro instruction. If the user desires to dynamically change the contents of the remote ESTAE parameter list, he may do so by coding a new exit address and/or a new parameter list address. If exit address or PARM= is coded, only the associated field in the remote ESTAE parameter list is changed. The other field remains as it was before the current ESTAE request was made.

The execute form of the ESTAE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede ESTAE.
ESTAE	
␣	One or more blanks must follow ESTAE.

<i>exit addr</i>	<i>exit addr</i> : RX-type address, or register (2) - (12).
0	
,CT	<i>list addr</i> : RX-type address, or register (2) - (12).
,OV	
,PARAM= <i>list addr</i>	
,XTCL=NO	
,XCTL=YES	
,PURGE=NONE	
,PURGE=QUIESCE	
,PURGE=HALT	
,ASYNCH=YES	
,ASYNCH=NO	
,TERM=NO	
,TERM=YES	
,BRANCH=NO	<i>save addr</i> : RX-type address, or register (2) - (12) or (13).
,BRANCH=YES,	
SVEAREA= <i>save addr</i>	
,RECORD=NO	
,RECORD=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters restricted in use are explained under the standard form of the ESTAE macro instruction. The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instructions**.

Example 1

Operation: Take the ESTAE exit labeled ADDR, allow synchronous exit processing, halt I/O, allow special error processing, branch enter SVC 60, provide 72-byte save area at SADDR, and execute the execute form of the macro instruction. EXEC is the label of the ESTAE parameter list built by the list form of the macro instruction.

```
ESTAE  ADDR, ASYNCH=YES, PURGE=HALT, TERM=YES, BRANCH=YES,  
       SVEAREA=SADDR, MF=( E, EXEC )
```


EVENTS — Wait for Events

The EVENTS macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the BRANCH parameter. This parameter is restricted in use and should only be used by programs executing in supervisor state, under protection key zero, and owning the local lock.

The syntax of the complete EVENTS macro instruction is shown below. However, only the explanation of the restricted parameter is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

The EVENTS macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede EVENTS.
EVENTS	
b	One or more blanks must follow EVENTS.
ENTRIES= <i>n</i> mb	<i>n</i> mb: decimal digits 1-32767.
ENTRIES=DEL, TABLE= <i>tab addr</i>	<i>tab addr</i> : symbol, RX-type address, or register (2) -(12).
TABLE= <i>tab addr</i>	Note: If the ENTRIES parameter is specified as indicated in the first two formats, no other parameters may be specified.
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : symbol, RX-type address, or register (2) - (12).
,LAST= <i>last addr</i>	<i>last addr</i> : symbol, RX-type address, or register (2) - (12).
	Note: If LAST is specified, WAIT must also be specified.
,WAIT=YES	
,WAIT=NO	
,BRANCH=NO	Default: BRANCH=NO
,BRANCH=YES	

The parameters restricted in use are explained as follows. (The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instructions**.)

,BRANCH=NO
,BRANCH=YES

specifies that an SVC entry (BRANCH=NO) or a branch entry (BRANCH=YES) is to be performed.

FRS is the address of the save area used by the control program to save the floating point registers 0, 2, 4, and 6 when the task is not active.

AETX is the address of the end of task exit routine specified in the ETXR parameter of the ATTACH macro instruction used to create the task.

PRI is the current limit (third byte) and dispatching (fourth byte) priorities of the task. The two high-order bytes are set to zero.

CMC is the task completion code. If the task is not complete, the field is set to zero.

TIOT is the address of the task input/output table.

COMM is the address of the command scheduler communications list. The list consists of a pointer to the communications event control block and a pointer to the command input buffer.

TSO is the address of a byte in which a high bit of 1 indicates a TSO address space, and a high bit of 0 indicates a non-TSO address space.

PSB is the address of the protected storage control block, which is extracted from the job step control block.

TJID is the address space identifier (ASID) for a TSO address space, and zero for a non-TSO address space.

ASID is the address space identifier.

Example 1

Operation: Provide information from all the fields of the indicated TCB except ASID. WHERE is the label of the answer area, ADDRESS is the label of a fullword that contains the address of the subtask TCB for which information is to be extracted.

```
EXTRACT WHERE, ADDRESS, FIELDS=( ALL, TSO, COMM, PSB, TJID )
```

Example 2

Operation: Provide information from the current TCB, as above.

```
EXTRACT WHERE, 'S', FIELDS=( ALL, TSO, COMM, PSB, TJID )
```

EXTRACT (List Form)

The list form of the EXTRACT macro instruction is used to construct a remote control program parameter list.

The list form of the EXTRACT macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede EXTRACT.
EXTRACT	
␣	One or more blanks must follow EXTRACT.

<i>answer addr</i>	<i>answer addr</i> : A-type address.
,'S'	<i>tcb addr</i> : A-type address.
<i>,tcb addr</i>	<i>Default</i> : 'S'
<i>,FIELDS=(tcb info)</i>	<i>tcb info</i> : any combination of the following, separated by commas:
	ALL PRI TSO
	GRS CMC PSB
	FRS TIOT TJID
	AETX COMM ASID
<i>,MF=L</i>	

The parameters are explained under the standard form of the EXTRACT macro instruction, with the following exception:

,MF=L
specifies the list form of the EXTRACT macro instruction.

EXTRACT (Execute Form)

The execute form of the EXTRACT macro instruction uses, and can modify, a remote control program parameter list. If the FIELDS parameter restricted in use is coded in the execute form, any TCB information specified in a previous FIELDS parameter is canceled and must be respecified if required for this execution of the macro instruction.

The execute form of the EXTRACT macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede EXTRACT.
EXTRACT	
␣	One or more blanks must follow EXTRACT.

<i>answer addr</i>	<i>answer addr</i> : RX-type address, or register (2) - (12).
,S'	<i>tcb addr</i> : RX-type address, or register (2) - (12).
, <i>tcb addr</i>	
,FIELDS=(<i>tcb info</i>)	<i>tcb info</i> : any combination of the following, separated by commas:
	ALL PRI TSO
	GRS CMC PSB
	FRS TIOT TJID
	AETX COMM ASID
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the EXTRACT macro instruction, with the following exception:

,MF = (E, *ctrl addr*)

specifies the execute form of the EXTRACT macro instruction using a remote control program parameter list.

FESTAE — Fast Extended STAE

The FESTAE macro instruction allows an SVC to establish an ESTAE environment with minimal overhead and no locking requirements. The ESTAE exit activated by FESTAE receives control in the same sequence and under the same conditions as though created by the ESTAE macro instruction.

The FESTAE macro expansion has no external linkage. The macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede FESTAE.
FESTAE	
b	One or more blanks must follow FESTAE.

0,WRKREG=<i>work reg addr</i> EXITADR=<i>exit addr</i>,	<i>work reg addr</i> : A-type address or register (1)-(14)
RBADDR=<i>svrb addr</i>	<i>exit addr</i> : A-type address or register (1)-(14) <i>svrb addr</i> : A-type address or register (1)-(14)
,TCBADDR=<i>tcb addr</i>	<i>tcb addr</i> : A-type address or register (1)-(14)
,PARAM=<i>list addr</i>	<i>list addr</i> : A-type address or register (1)-(14)
,XCTL=NO ,XCTL=YES	Default: XCTL=NO
,PURGE=NONE ,PURGE=HALT ,PURGE=QUIESCE	Default: PURGE=NONE
,ASYNCH=YES ,ASYNCH=NO	Default: ASYNCH=YES
,TERM=NO ,TERM=YES	Default: TERM=NO
,RECORD=NO ,RECORD=YES	Default: RECORD=NO
,ERRET=<i>label</i>:	<i>label</i> : any valid assembler name

The parameters are explained as follows:

0,WRKREG

specifies that the most recent ESTAE exit is canceled. A work register must be specified for use by the FESTAE macro expansion.

EXITADR

specifies a register that contains the address of an ESTAE exit routine to be entered if the task terminates abnormally. This register is used as a work register that is not reloaded. If 0 (the cancel option) is specified, the user must be certain that the ESTAE is the current exit. If the user cannot determine this, he should use the branch entry service instead.

RBADDR

specifies a register that contains the address of the current SVRB prefix (RBPRFX). RBADDR must be specified if EXITADR has also been specified. The specified register is not altered.

TCBADDR

specifies the register containing the current TCB address. This register is not altered, and its specification results in the generation of more efficient code.

PARAM

specifies the register containing the address of a user-defined parameter list that has data to be used by the ESTAE exit routine. The exit routine receives this address upon its scheduling for execution. The use of this parameter list is optional, but the user should zero out any spurious data it may contain whether or not he intends to use it. If the user does not select the PARAM option, the exit routine receives instead the 24 byte parameter area in the SVRB. In this case, the user must locate this SVRB parameter area and initialize it with appropriate data.

ERRET

specifies a label within the CSECT issuing the FESTAE for which addressability has been established. The FESTAE macro branches to this label if it is returning a code other than zero. This option saves the user the instructions necessary to check the return code. If the user does not specify the ERRET option, control returns instead to the instruction immediately following the FESTAE macro instruction. The return code is in register 15.

All the other FESTAE parameters have the same meaning as their ESTAE counterparts, as described in *OS/VS2 Supervisor Services and Macro Instructions*.

Upon conclusion of FESTAE processing, control resumes at the instruction following the FESTAE macro instruction. Register 15 then contains one of the following return codes:

Hexadecimal Code	Meaning
00	Successful completion of the FESTAE request.
08	A previous create has been issued with BRANCH=YES or FESTAE for this SVRB; the request has been ignored.
0C	Cancel has been specified under one of the following conditions: <ol style="list-style-type: none"> 1) There are no exits for this TCB. 2) The most recent exit is not owned by the caller. 3) The most recent exit was not created by FESTAE.

Example 1:

Operation: Take the ESTAE exit specified by register 2, allow asynchronous exit processing, do not allow special error processing, default to PURGE=NONE, and pass the parameter list pointed to by register 7 to the ESTAE exit routine.

```
FESTAE  EXITADR=( REG2 ), RBADDR=( REG3 ), TCBADDR=( REG6 ), x
        PARAM=( REG7 ), ASYNCH=YES, TERM=NO
```

FREEMAIN — Free Virtual Storage

The FREEMAIN macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the **BRANCH** and **KEY** parameters. These parameters are restricted in use and should only be used by programs executing in supervisor state, under protection key zero.

Note: This macro instruction requires that the CVT mapping macro be assembled as a DSECT into the caller's program.

The syntax of the complete FREEMAIN macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

The standard form of the FREEMAIN macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede FREEMAIN.
FREEMAIN	
b	One or more blanks must follow FREEMAIN.
LC,LA= <i>length addr</i>	<i>length addr</i> : A-type address, or register (2) - (12).
LU,LA= <i>length addr</i>	<i>length value</i> : symbol, decimal digit, or register (2) - (12). If R is specified, register (0) may also be specified.
L,LA= <i>length addr</i>	<i>subpool nmb</i> : symbol, decimal digit 0-127, or register (0) or (2)
VC	- (12). If R,SP=(0) is specified, the high order byte of register 0 must contain the subpool number and the low order 3 bytes must contain the length value.
VU	Note: For a subpool FREEMAIN, if the formats RC,SP= <i>subpool nmb</i> or RU,SP= <i>subpool nmb</i> or R,SP= <i>subpool nmb</i> are specified, no other parameters may be specified.
V	
EC,LV= <i>length value</i>	
EU,LV= <i>length value</i>	
E,LV= <i>length value</i>	
RC,LV= <i>length value</i>	
RC,SP= <i>subpool nmb</i>	
RU,LV= <i>length value</i>	
RU,SP= <i>subpool nmb</i>	
R,LV= <i>length value</i>	
R,SP= <i>subpool nmb</i>	
,A= <i>addr</i>	<i>addr</i> : A-type address, or register (2) - (12).
,SP= <i>subpool nmb</i>	<i>subpool nmb</i> : symbol, decimal digit 0-127, or register (0) or (2) - (12). If R,SP=(0) is specified, the high order byte of register 0 must contain the subpool number and the low order 3 bytes must contain the length value.
,BRANCH=YES	Note: BRANCH=(YES,GLOBAL) may only be specified with RC or RU above. Also, the macro expansion uses register 4 for the address of the global save area pointed to by the CVT. The previous contents of register 4 is overridden. BRANCH requires inclusion of the CVT mapping macro.
,BRANCH=(YES,GLOBAL)	
,KEY= <i>nmb</i>	<i>nmb</i> : decimal digits 0-15, or register (2) - (12). Note: This parameter may be specified only with BRANCH and RC or RU above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instruction.)

,BRANCH = YES

,BRANCH = (YES,GLOBAL)

specifies that a branch entry is to be used instead of an SVC entry. If (YES,GLOBAL) is specified, the GLBRANCH entry point to service global storage requests without the need for the local memory lock will be used; the SALLOC lock must be held.

If BRANCH=YES is specified, the caller must pre-load register 4 with the TCB address, pre-load register 7 with the ASCB address, and hold the local address space lock prior to entry. Register 7 will not contain the ASCB address when control is returned to the caller. Register 3 will also be destroyed if RC or RU was specified.

If `BRANCH=(YES,GLOBAL)` is specified, registers 4 and 7 need not contain the TCB and ASCB addresses; and registers 3 and 4 are changed when control is returned to the caller. Additionally, the SP parameter may only designate subpools 227, 228, 231, 239, 241, or 245.

The FREEMAIN macro instruction, with `BRANCH=(YES,GLOBAL)` specified, requires that the IHAWSAVT mapping macro be assembled as a DSECT into the caller's program.

Note: Subpool freemains can be issued only for the following subpools: 0-128, 229, 230, 233, 236, 237, and 250-253. Any attempt to issue a subpool freemain for any other subpool causes a 478 abend.

`,KEY=key nmbr`

specifies the key (in bits 24-27 of the register) in which the requested storage was obtained. This parameter applies to subpools 227, 228, 229, 230, 231, and 241, and allows both global and local storage to be freed in the requestor's storage protection key. When control returns, register 15 may contain the following return code:

Hexadecimal Code	Meaning
8	Part of area being freed is still fixed.

The parameters restricted in use are explained under the standard form of the FREEMAIN macro instruction. The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instructions*.

Example 1

Operation: Free 400 bytes of storage addressed by register 1 via a branch entry. If the storage is successfully freed, register 15 contains 0; otherwise, register 15 contains a nonzero value.

```
FREEMAIN EC, LV=400, A=(1), BRANCH=YES
```

Example 2

Operation: Free 48 bytes of the storage (addressed by register 5) in subpool 231. Register 3 has been preset to contain the storage key of the storage to be released. If the request is unsuccessful, the caller is abnormally terminated.

```
FREEMAIN RU, LV=48, A=(5), SP=231, KEY=(3), BRANCH=(YES, GLOBAL)
```

FREEMAIN (List Form)

The list form of the FREEMAIN macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede FREEMAIN.
FREEMAIN	
b	One or more blanks must follow FREEMAIN.

LC	
LU	
L	
VC	
VU	
V	
EC	
EU	
E	
,LA= <i>length addr</i>	<i>length addr</i> : A-type address.
,LV= <i>length value</i>	<i>length value</i> : symbol or decimal digit. Note: LA may only be specified with LC, LU, or L above. Note: LV may only be specified with EC, EU, or E above.
,A= <i>addr</i>	<i>addr</i> : A-type address.
,SP= <i>subpool nmb</i>	<i>subpool nmb</i> : symbol or decimal digit 0 - 127.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=L	

The parameters restricted in use are explained under the standard form of the FREEMAIN macro instruction. The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.

GETMAIN — Allocate Virtual Storage

The GETMAIN macro instruction is described in the OS/VS2 Supervisor Services and Macro Instructions, with the exception of the BRANCH and KEY parameters. These parameters are restricted in use and should only be used by programs executing in supervisor state, under protection key zero.

Note: This macro instruction requires that the CVT mapping macro instruction be assembled as a DSECT into the caller's program.

The syntax of the complete GETMAIN macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in OS/VS2 Supervisor Services and Macro Instructions.

The standard form of the GETMAIN macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede GETMAIN.
GETMAIN	
␣	One or more blanks must follow GETMAIN.
LC,LA= <i>length addr</i> ,A= <i>addr</i>	<i>length addr</i> : A-type address, or register (2) - (12).
LU,LA= <i>length addr</i> ,A= <i>addr</i>	<i>length value</i> : symbol, decimal digit, or register (2) - (12). If R is specified, register (0) may also be specified.
VC,LA= <i>length addr</i> ,A= <i>addr</i>	<i>addr</i> : A-type address, or register (2) - (12).
VU,LA= <i>length addr</i> ,A= <i>addr</i>	
EC,LV= <i>length value</i> ,A= <i>addr</i>	
EU,LV= <i>length value</i> ,A= <i>addr</i>	
RC,LV= <i>length value</i>	
RU,LV= <i>length value</i>	
R,LV= <i>length value</i>	
,SP= <i>subpool nmb</i> r	<i>subpool nmb</i> r: symbol, decimal digit 0 - 127, or register (0) or (2) - (12). Note: If R,LV=(0) is specified above, SP may not be specified.
,BNDRY=DBLWD	Default: BNDRY=DBLWD
,BNDRY=PAGE	Note: This parameter may not be specified with R above.
,BRANCH=YES	Note: BRANCH=(YES,GLOBAL) may only be specified with RC or RU above. Also, the macro expansion uses register 4 for the address of the global save area pointed to by the CVT. The previous contents of register 4 is overridden. BRANCH requires inclusion of the CVT mapping macro. The macro expansion also uses register 3.
,BRANCH=(YES,GLOBAL)	
,KEY= <i>key number</i>	<i>key nmb</i> r: decimal digits 0-15, or register (2) - (12). Note: This parameter may be specified only with BRANCH and RC or RU above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.)

,BRANCH= YES

,BRANCH= (YES,GLOBAL)

specifies that a branch entry is to be used instead of an SVC entry. If (YES,GLOBAL) is specified, the GLBRANCH entry point to service global storage requests without the need for the local memory lock is used; the SALLOC lock must be held.

If BRANCH=YES is specified, the caller must pre-load register 4 with the TCB address, pre-load register 7 with the ASCB address, and hold the local address space lock prior to entry. Register 7 does not contain the ASCB address when control is returned to the caller. Register 3 will also be destroyed if RC or RU was specified.

If BRANCH=(YES,GLOBAL) is specified, registers 4 and 7 need not contain the TCB and ASCB addresses; and register 3 and 4 are changed when control returns to the caller. Additionally, the SP parameter may only designate subpools 227, 228, 231, 239, 241, or 245.

The parameter LV=(0) for BRANCH=(YES,GLOBAL) must contain only the length value. The subpool id is specified with the SP parameter.

The GETMAIN macro instruction, with this parameter specified, requires that the IHAWSAVT mapping macro be assembled as a DSECT into the caller's program.

,KEY=key nmb

specifies the key (in bits 24-27 of the register) in which the requested storage is to be obtained. This parameter applies to subpools 227, 228, 229, 230, 231, and 241, and allows both global and local storage to be obtained in the requester's storage protection key. When control is returned, register 15 may contain the following return code:

Hexadecimal Code	Meaning
8	On request for SQA or LSQA, no real storage page is available.

Example 1

Operation: Obtain 248 bytes of storage from the user's region via a branch entry. If the routine is in supervisor state, subpool 252 is used; otherwise, subpool 0 is used. If the storage cannot be obtained, the caller is abnormally terminated.

```
GETMAIN EU, LV=248, A=AREAADDR, BRANCH=YES
```

Example 2

Operation: Obtain one page of storage from the common service area, and cause the acquired storage to be initialized with a storage key of 9. A return code of 0 (if successful) or 4 (if unsuccessful) is returned.

```
GETMAIN RC, LV=4096, SP=231, BRANCH=( YES, GLOBAL ), BNDRY=PAGE, KEY=9
```

GETMAIN (List Form)

The list form of the GETMAIN macro instruction is written as follows:

<i>name</i>	<i>name</i> : Begin <i>name</i> in column 1.
␣	One or more blanks must precede GETMAIN.
GETMAIN	
␣	One or more blanks must follow GETMAIN.

LC	
LU	
VC	
VU	
EC	
EU	
<i>,LA=length addr</i>	<i>length addr</i> : A-type address.
<i>,LV=length value</i>	<i>length value</i> : symbol or decimal digit.
	Note : LA may be specified with EC or EU above.
	Note : LV may not be specified with LC, LU, VC or VU above.
<i>,A=addr</i>	<i>addr</i> : A-type address.
<i>,SP=subpool nmb</i>	<i>subpool nmb</i> : symbol or decimal digit 0-127.
<i>,BNDRY=DBLWD</i>	Default : BNDRY=DBLWD
<i>,BNDRY=PAGE</i>	
<i>,RELATED=value</i>	<i>value</i> : any valid macro keyword specification.
<i>,MF=L</i>	

The parameters restricted in use are explained under the standard form of the GETMAIN macro instruction. The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instructions*.

GETMAIN (Execute Form)

The execute form of the GETMAIN macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede GETMAIN.
GETMAIN	
␣	One or more blanks must follow GETMAIN.

LC	
LU	
VC	
VU	
EC	
EU	
<i>,LA=length addr</i>	<i>length addr</i> : RX-type address or register (2) - (12).
<i>,LV=length value</i>	<i>length value</i> : symbol, decimal digit, or register (2) - (12). Note : LA may not be specified with EC or EU above. Note : LV may not be specified with LC, LU, VC, or VU above.
<i>,A=addr</i>	<i>addr</i> : RX-type address, or register (2) - (12).
<i>,SP=subpool nmb</i>	<i>subpool nmb</i> : symbol, decimal digit 0-127, or register (0) or (2) - (12).
<i>,BNDRY=DBLWD</i>	Default : BNDRY=DBLWD
<i>,BNDRY=PAGE</i>	
<i>,BRANCH=YES</i>	
<i>,RELATED=value</i>	<i>value</i> : any valid macro keyword specification.
<i>,MF=(E ,ctrl prog)</i>	<i>ctrl prog</i> : RX-type address, or register (1) or (2) - (12).

The parameters restricted in use are explained under the standard form of the GETMAIN macro instruction. The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.

MODESET — Change System Status

The MODESET macro instruction is used to change system status by altering the PSW key or mode indicator. It causes a supervisor routine (IEAVMODE) to alter the RB old program status word (RBOPSW) so that the desired PSW is loaded when MODESET returns to the caller. MODESET also generates inline code that saves and/or changes the protection key in the current PSW.

The standard form of the MODESET macro instruction has two forms: the form that generates an SVC and the form that generates inline code. The form that generates inline code uses the SPKA instruction (see *IBM System/370 Principles of Operation*), and is executable only in supervisor state. The form that generates an SVC is executable by users in supervisor state, under protection key 0-7, or APF-authorized.

The standard form of the MODESET macro instruction that generates inline code is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede MODESET.
MODESET	
␣	One or more blanks must follow MODESET.

EXTKEY= <i>key</i>	<i>key</i> : one of the following:												
KEYADDR= <i>key addr</i>	<table style="width: 100%; border: none;"> <tr> <td>SCHED</td> <td>SRM</td> <td>ZERO</td> </tr> <tr> <td>JES</td> <td>SUPR</td> <td>TCB</td> </tr> <tr> <td>RSM</td> <td>DATAMGT</td> <td>RBT1</td> </tr> <tr> <td>VSM</td> <td>TCAM</td> <td>RBT234</td> </tr> </table>	SCHED	SRM	ZERO	JES	SUPR	TCB	RSM	DATAMGT	RBT1	VSM	TCAM	RBT234
SCHED	SRM	ZERO											
JES	SUPR	TCB											
RSM	DATAMGT	RBT1											
VSM	TCAM	RBT234											
	<i>key addr</i> : A-type address or register (2).												
	Note: The WORKREG parameter is required if the following are specified:												
	<table style="width: 100%; border: none;"> <tr> <td>EXTKEY=TCB</td> <td>EXTKEY=RBT234</td> </tr> <tr> <td>EXTKEY=RBT1</td> <td>KEYADDR=A-type address</td> </tr> </table>	EXTKEY=TCB	EXTKEY=RBT234	EXTKEY=RBT1	KEYADDR=A-type address								
EXTKEY=TCB	EXTKEY=RBT234												
EXTKEY=RBT1	KEYADDR=A-type address												
	Note: The WORKREG parameter should be register 1-15 if one of these four parameters is specified because WORKREG is used as a base register on the SPKA instruction. WORKREG=0 sets the PSW key to zero.												
,SAVEKEY= <i>old key addr</i>	<i>old key addr</i> : A-type address or register (2).												
	Note: If KEYADDR=(2) is specified above, then SAVEKEY=(2) cannot be specified.												
	Note: The WORKREG parameter is required if SAVEKEY=A-type address is specified.												
,WORKREG= <i>reg</i>	<i>reg</i> : decimal digits 0 - 15.												
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.												

The parameters are explained as follows:

EXTKEY = *key*

KEYADDR = *key addr*

specifies the key to be set in the current PSW or the address of the key.

SCHED — Scheduler key.

JES — Job entry subsystem key.

RSM — Real storage management key.

VSM — Virtual storage management key.

SRM — System resource management key.

SUPR — Supervisor key.

DATAMGT — Data management key.

TCAM — Telecommunications access method key.

ZERO — Key of zero is to be set.

TCB — Key is to be obtained from TCB field TCBPKF.

RBT1 — Key is to be obtained from the RBOPSW field of the active RB of type 1 SVC routine issuing MODESET.

RBT234 — Key is to be obtained from the RBOPSW field of the RB preceding SVRB of type 2, 3, or 4 SVC routine issuing MODESET.

KEYADDR specifies a location 1 byte in length which contains the key in bit positions 0-3.

If register (2) is specified, the key is contained in bit positions 24-27 (bits 28-31 are ignored). This parameter permits a previously saved key to be restored. If TCB, RBT1 or RBT234 is specified as the key address, the TCB mapping macro IKJATCB is required.

The user is expected to establish addressability to the TCB with a USING statement.

,SAVEKEY = *old key addr*

specifies a location 1 byte in length where the current PSW key is to be saved, in bit positions 0-3. If register (2) is specified, the key is left in register 2.

,WORKREG = *reg*

specifies the register into which the contents of register 2 are to be saved while performing the SAVEKEY function, or the working register to be used by the EXTKEY or KEYADDR function. If WORKREG=2 is specified, no register saving takes place.

,RELATED = *value*

specifies information used to self-document macro instructions by 'relating' functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R,LV=4096,RELATED=(FREE1,'GET STORAGE')
FREE1   FREEMAIN   R,LV=4096,A=(1),RELATED=(GET1,'FREE STORAGE')
```

Note: This form of the macro instruction does not generate any return codes.

The standard form of the MODESET macro instruction that generates an SVC is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede MODESET.
MODESET	
b	One or more blanks must follow MODESET.

KEY=ZERO	Note: KEY is required only if no other parameter is specified.
KEY=NZERO	
,MODE=PROB	Note: MODE is required only if no other parameter is specified.
,MODE=SUP	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained as follows:

KEY = ZERO

KEY = NZERO

specifies that the PSW key (bits 8-11) is to be either set to zero (ZERO) or set to the value in the caller's TCB (NZERO).

,MODE = PROB

,MODE = SUP

specifies that the PSW mode indicator (bit 15) is to be either turned on (PROB) or turned off (SUP).

Note: This form of the macro instruction does not generate any return codes.

Example 1

Operation: Save the current PSW key, and change the key to that of the scheduler.

```
MODESET EXTKEY=SCHED,SAVEKEY=KEYSAVE,WORKREG=1
```

Example 2

Operation: Change to supervisor mode and key zero.

```
MODESET KEY=ZERO,MODE=SUP
```

MODESET (List Form)

The list form of the MODESET macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede MODESET..
MODESET	
␣	One or more blanks must follow MODESET.

RELATED= <i>value</i> ,	<i>value</i> : any valid macro keyword specification.
MF=L	

The parameters are explained under the standard form of the MODESET macro instruction, with the following exception:

MF = L

specifies the list form of the MODESET macro instruction.

MODESET (Execute Form)

The execute form of the MODESET macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede MODESET.
MODESET	
␣	One or more blanks must follow MODESET.

RELATED= <i>value</i> ,	<i>value</i> : any valid macro keyword specification.
MF=(E , <i>list addr</i>)	<i>list addr</i> : RX-type address, or register (1).

The parameters are explained under the standard form of the MODESET macro instruction, with the following exception:

MF=(E ,*list addr*)

specifies the execute form of the MODESET macro instruction, using a parameter list address.

NIL — Provide a Lock Via an AND IMMEDIATE (NI) Instruction

The NIL macro instruction is used to provide a lock on a byte of storage on which an AND IMMEDIATE (NI) instruction is to be executed. The byte of storage exists in a multiprocessing environment and the possibility exists that the byte may be changed by another CPU at the same time. Storage modification during NIL processing is accomplished in the same manner as that used by the COMPARE AND SWAP (CS) instruction.

For details on the AND IMMEDIATE and COMPARE AND SWAP instructions, see **IBM System/370 Principles of Operation**.

The NIL macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede NIL.
NIL	
␣	One or more blanks must follow NIL.

<i>byte addr</i>	<i>byte addr</i> : A-type address.
<i>,mask</i>	<i>mask</i> : symbol or self defining term.
<i>,REF=stor addr</i>	<i>stor addr</i> : A-type address.
<i>,WREGS=(reg1,reg2,reg3)</i>	<i>reg1</i> : symbol, or decimal digits 0-16.
<i>,WREGS=(reg1,reg2)</i>	<i>reg2</i> : symbol, or decimal digits 1-16.
<i>,WREGS=(reg1,,reg3)</i>	<i>reg3</i> : symbol, or decimal digits 0-16.
<i>,WREGS=(,reg2,reg3)</i>	Default for <i>reg1</i> : 0
<i>,WREGS=(reg1)</i>	Default for <i>reg2</i> : 1
<i>,WREGS=(,reg2)</i>	Default for <i>reg3</i> : 2
<i>,WREGS=(,reg3)</i>	

The parameters are explained as follows:

byte addr

specifies the address of the byte to which the AND function is to be applied.

,mask

specifies the value to be ANDed to the byte at the address specified above.

,REF=stor addr

specifies the address of a storage location on a fullword boundary. This address provides the means by which the COMPARE AND SWAP instruction may be executed. The address must be less than or equal to the byte address specified above, and the difference between the addresses must be less than 4096. The two addresses must be addressable via the same base register.

,WREGS=(reg1,reg2,reg3)

,WREGS=(reg1,reg2)

,WREGS=(reg1,,reg3)

,WREGS=(,reg2,reg3)

,WREGS=(reg1)

,WREGS=(,reg2)

,WREGS=(,reg3)

specifies the work registers to be used to perform the COMPARE AND SWAP instruction.

reg1 is used to contain the "old" byte; *reg2* is used to contain the "updated" byte; and *reg3* is used to contain the mask.

Example 1

Operation: Provide a lock on the byte of storage specified by the address STRESTAT. UCBOB is the address used to reference the byte, and FSRTECGS is the mask used.

NIL SRTESTAT,FSRTECGS,WREGS=(15,4,5),REF=UCBOB

OIL — Provide a Lock Via an OR IMMEDIATE (OI) Instruction

The OIL macro instruction is used to provide a lock on a byte of storage on which an OR IMMEDIATE (OI) instruction is to be executed. The byte of storage exists in a multiprocessing environment and the possibility exists that the byte may be changed by another CPU at the same time. Storage modification during OIL processing is accomplished in the same manner as that used by the COMPARE AND SWAP (CS) instruction.

For details on the OR IMMEDIATE and COMPARE AND SWAP instructions, see **IBM System/370 Principles of Operation**.

The OIL macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede OIL.
OIL	
␣	One or more blanks must follow OIL.

<i>byte addr</i>	<i>byte addr</i> : A-type address.
<i>,mask</i>	<i>mask</i> : symbol or self defining term.
<i>,REF=stor addr</i>	<i>stor addr</i> : A-type address.
<i>,WREGS=(reg1,reg2,reg3)</i>	<i>reg1</i> : symbol, or decimal digits 0-16.
<i>,WREGS=(reg1,reg2)</i>	<i>reg2</i> : symbol, or decimal digits 0-16.
<i>,WREGS=(reg1,,reg3)</i>	<i>reg3</i> : symbol, or decimal digits 0-16.
<i>,WREGS=(,reg2,reg3)</i>	Default for <i>reg1</i> : 0
<i>,WREGS=(reg1)</i>	Default for <i>reg2</i> : 1
<i>,WREGS=(,reg2)</i>	Default for <i>reg3</i> : 2
<i>,WREGS=(,,reg3)</i>	

The parameters are explained as follows:

byte addr

specifies the address of the byte to which the OR function is to be applied.

,mask

specifies the value to be ORed to the byte at the address specified above.

,REF=stor addr

specifies the address of a storage location on a fullword boundary. This address provides the means by which the COMPARE AND SWAP instruction may be executed. The address must be less than or equal to the byte address specified above, and the difference between the addresses must be less than 4096. The two addresses must be addressable via the same base register.

,WREGS=(reg1,reg2,reg3)

,WREGS=reg1,reg2)

,WREGS=(reg1,,reg3)

,WREGS=(,reg2,reg3)

,WREGS=(reg1)

,WREGS=(,reg2)

,WREGS=(,,reg3)

specifies the work registers to be used to perform the COMPARE AND SWAP instruction.

reg1 is used to contain the "old" byte; *reg2* is used to contain the "updated" byte; and *reg3* is used to contain the mask.

Example 1

Operation: Provide a lock on the byte of storage specified by the address SRTESTAT. UCBOB is the address used to reference the byte, and SRTECHGS is the mask used.

```
OIL SRTESTAT,SRTECHGS,WREGS=(15,4,5),REF=UCBOB
```


,EA=*end addr*

specifies the end address + 1 of the virtual area to be fixed.

,LONG=Y

,LONG=N

specifies that the relative real time duration anticipated for the fix is long (Y) or short (N).

,RELEASE=N

,RELEASE=Y

specifies that the contents of the virtual area is to remain intact (N) or be released (Y).

,RELATED=*value*

specifies information used to self-document macro instructions by "relating" functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R,LV=4096,RELATED=(FREE1,'GET STORAGE')
FREE1   FREEMAIN  R,LV=4096,A=(1),RELATED=(GET1,'FREE STORAGE')
```

Upon completion, register 15 contains one of the following return codes:

Hexadecimal

Code	Meaning
00	Operation completed normally; ECB posted complete.
04	Operation abnormally terminated with a 171 abend. Operation incomplete because of invalid address in virtual subarea list entry; ECB posted complete. See OS/VS Message Library: VS2 System Codes , GC38-1016 for a complete description of the register contents after a 171 abend.
08	Operation proceeding; ECB will be posted when all requested pages are fixed in real storage.
10	Operation abnormally terminated with a 171 abend. Virtual subarea list entry or ECB address invalid; no ECB is posted. See OS/VS Message Library: VS2 System Codes , GC38-1016 for a complete description of the register contents after a 171 abend.

The ECB is unchanged if the request was initiated but not complete (return code 8), or if an ABEND was issued with return code 10. Otherwise, the ECB is posted complete with code:

- 0 - operation completed successfully.
- 4 - operation incomplete because of invalid address in VSL entry.

If the return code issued is 8, the ECB is posted asynchronously when paging I/O has completed, with code:

- 0 - operation completed successfully.
- 4 - operation incomplete because of paging error; requesting TCB will be abnormally terminated.

The ECB code is posted in the low-order 3 bytes of the ECB, and is right-justified.

Example 1

Operation: Fix a single byte of virtual storage addressed by register 3. Note that the full 4096-byte page containing the specified byte is actually be fixed. The storage is long fixed.

```
PGFIX R,A=(R3),ECB=(R5)
```

Example 2

Operation: Fix virtual storage without using a virtual subarea list. Storage is long fixed.

```
PGFIX R,A=(R3),EA=(R4),ECB=ECB1
```

Example 3

Operation: Fix, but not long-fix, virtual storage, and ensure that the pages fully included in the address range are forfeited before fixing the area specified by registers 3 and 4.

```
PGFIX R,A=(R3),EA=(R4),ECB=(R5),LONG=N,RELEASE=Y
```

PGFREE — Free Virtual Storage Contents

The PGFREE macro instruction makes virtual storage areas that fixed via the PGFIX macro instruction eligible for page-out. The PGFREE function is available only to authorized system functions and users.

Note that a fixed page is not considered pageable until the number of PGFREEs issued for the page is equal to the number of PGFIXes previously issued for that page. That is, a page is not automatically removed from real storage as the result of issuing a PGFREE macro instruction.

The standard form of the PGFREE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
 	One or more blanks must precede PGFREE.
PGFREE	
 	One or more blanks must follow PGFREE.

L	
,LA=list addr	<i>list addr</i> : A-type address, or register (1) or (2) - (12).
R	
,A=start addr	<i>start addr</i> : A-type address, or register (1) or (2) - (12).
,ECB=ecb addr	<i>ecb addr</i> : A-type address, or register (0) or (2) - (12).
,EA=end addr	<i>end addr</i> : A-type address, or register (2) - (12) or (15). Default : <i>start addr</i> + 1
,RELEASE=N	Default : RELEASE=N
,RELEASE=Y	Note : RELEASE=Y may only be specified with EA above.
,RELATED=value	<i>value</i> : any valid macro keyword specification.

The parameters are explained as follows:

- L**
specifies that a parameter list is being supplied with this request.
- ,LA = list addr**
specifies the address of the first entry of a virtual subarea list.
- R**
specifies that no parameter list is being supplied with this request.
- ,A = start addr**
specifies the start address of the virtual area to be freed.
- ,ECB = ecb addr**
specifies the address of the ECB that was used in a prior PGFIX request. This parameter is used if there is any possibility that the ECB for the previously issued PGFIX was not posted complete.
- ,EA = end addr**
specifies the end address + 1 of the virtual area to be freed.

,RELEASE = N

,RELEASE = Y

specifies that the contents of the virtual area is to remain intact (N) or be released (Y).

,RELATED = *value*

specifies information used to self-document macro instructions by "relating" functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R, LV=4096, RELATED=( FREE1, 'GET STORAGE' )
FREE1   FREEMAIN   R, LV=4096, A=( 1 ), RELATED=( GET1, 'FREE STORAGE' )
```

When control is returned, register 15 contains one of the following return codes:

Hexadecimal

Code	Meaning
00	Operation completed normally.
04	Operation abnormally terminated. Operation incomplete because of invalid address in virtual subarea list entry.
10	Operation abnormally terminated. Virtual subarea list entry or ECB address invalid.

Example 1

Operation: Free the storage in Example 1 of standard-form PGFIX.

```
PGFREE R, A=( R3 )
```

Example 2

Operation: Free the storage in Example 2 of standard-form PGFIX.

```
PGFREE R, A=( R3 ), EA=( R4 )
```

Example 3

Operation: Free the storage in Example 3 of standard-form PGFIX, and forfeit the pages full included in the address range.

```
PGFREE R, A=( R3 ), EA=( R4 ), ECB=( R5 ), RELEASE=Y
```

POST — Signal Event Completion

The POST macro instruction is described in the OS/VS2 Supervisor Services and Macro Instructions, with the exception of the ASCB and ERRET parameters. These parameters are restricted in use and should only be used with tasks in supervisor state, APF-authorized, or with protection key 0-7.

The syntax of the complete POST macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in OS/VS2 Supervisor Services and Macro Instructions.

The standard form of the POST macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede POST.
POST	
␣	One or more blanks must follow POST.
<i>ecb addr</i>	<i>ecb addr</i> : RX-type address, or register (2) - (12).
<i>,comp code</i>	<i>comp code</i> : symbol, decimal or hexadecimal digit, or register (0) or (2) - (12). Range of values: 0 - 2 ³⁰ -1 Default: 0
<i>,ASCB=addr,ERRET=err addr</i>	<i>addr</i> : RX-type address, or register (2) - (12).
<i>,ASCB=addr,ERRET=err addr,</i>	<i>err addr</i> : RX-type address, or register (2) - (12).
<i>ECBKEY=key</i>	<i>addr</i> : RX-type address, or register (2) - (12). <i>err addr</i> : RX-type address, or register (2) - (12). <i>key</i> : symbol, decimal or hexadecimal digit, or register (2) - (12). Range of values: 0 - 15 (decimal) Default: none. Note: If the register form is specified, bits 24-27 of the register must contain the key.
<i>,RELATED=value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.)

,ASCB=addr,ERRET=err addr

specifies the address of the ASCB of the address space containing the ECB being posted, and the address of the routine to be given control when an error condition resulting from a POST failure is detected.

,ASCB=addr,ERRET=err addr, ECBKEY=key

specifies the address of the ASCB containing the ECB being posted, the address of the routine to be given control when an error condition resulting from a POST failure is detected, and the storage protection key of the ECB to be posted. If the ECB does not identify a current wait condition against it, the ECB is checked against the key before it is updated with the post completion code. Otherwise, the ECB is checked against the protection key of the waiting task.

Example 1

Operation: Post an event control block whose address is ECB, where the address space containing the ECB has an ASCB specified by register 5, and where ERRRTN is the routine to be given control on error conditions.

```
POST    ECB,ASCB=(REG5),ERRET=ERRRTN
```

Example 2

Operation: Post the ECB from example 1 with a hexadecimal completion code of 3FF.

```
POST    ECB,X'3FF',ASCB=(REG5),ERRET=ERRRTN
```


POST (List Form)

The list form of the POST macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
	One or more blanks must precede POST.
POST	
	One or more blanks must follow POST.

<i>ecb addr</i>	<i>ecb addr</i> : A-type address.
,ASCB= <i>addr</i> ,ERRET= <i>err addr</i>	<i>addr</i> : A-type address. <i>err addr</i> : A-type address.
,ASCB= <i>addr</i> ,ERRET= <i>err addr</i> , ECBKEY=YES	<i>addr</i> : A-type address <i>err addr</i> : A-type address
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=L	

The parameters are explained under the standard form of the POST macro instruction, with the following exceptions:

- ,MF=L
specifies the list form of the POST macro instruction.
- ,ASCB=*addr*,ERRET=*err addr*,ECBKEY=YES
specifies that the storage protection key of the ECB is defined in the execute form of the POST macro instruction.

POST (Execute Form)

The execute form of the POST macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede POST.
POST	
b	One or more blanks must follow POST.

<i>ecb addr</i>	<i>ecb addr</i> : RX-type address, or register (2) - (12).
<i>,comp code</i>	<i>comp code</i> : symbol, decimal or hexadecimal digit, or register (0) or (2) - (12). Range of values: 0 - 2 ³⁰ -1
<i>,ASCB=addr,ERRET=err addr</i>	<i>addr</i> : RX-type address, or register (2) - (12). <i>err addr</i> : RX-type address, or register (2) - (12).
<i>,ASCB=addr,ERRET=err addr,ECBKEY=key</i>	<i>addr</i> : RX-type address, or register (2) - (12). <i>err addr</i> : RX-type address, or register (2) - (12). <i>key</i> : symbol, decimal or hexadecimal digit, or register (2) - (12). Range of values: 0 - 15 (decimal). Default: none. Note: If the register form is specified, bits 24-27 of the register must contain the key.
<i>,RELATED=value</i>	<i>value</i> : any valid macro keyword specification.
<i>,MF=(E ,prob addr)</i>	<i>prob addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the POST macro instruction, with the following exception:

,MF = (E ,prob addr)
specifies the execute form of the POST macro instruction using a remote control program parameter list.

PURGEDQ (Execute Form)

The execute form of the PURGEDQ macro instruction uses a remote control program parameter list. The parameter list is constructed using the list form of PURGEDQ.

The execute form of the PURGEDQ macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede PURGEDQ.
PURGEDQ	
␣	One or more blanks must follow PURGEDQ.

RMTR= <i>RMTR addr</i>	<i>RMTR addr</i> : RX-type address, or register (2) - (12).
,ASID= <i>ASID addr</i>	<i>ASID addr</i> : RX-type address, or register (2) - (12).
,ASIDTCB= <i>TCB addr</i>	<i>TCB addr</i> : RX-type address, or register (2) - (12).
,MF=(<i>E, ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the PURGEDQ macro instruction, with the following exception:

,MF = (*E, ctrl addr*)

specifies the execute form of the PURGEDQ macro instruction, using a remote control program parameter list.

Example 1

Operation: All SRBs scheduled into the address space designated by register 6 are to be purged by the RMTR routine IEAVRSPG. Register 1 is a pointer to the parameter list, and register 7 indicates that all SRBs are to be purged.

```
PURGEDQ MF=(E,R1),ASID=(R6),ASIDTCB=(R7),RMTR=IEAVRSPG
```

QEDIT — Command Input Buffer Manipulation

The QEDIT macro instruction generates the required entry parameters and processes the command input buffer for the following uses:

- Dechaining and freeing of a command input buffer (CIB) from the CIB chain for a task.
- Setting a limit for the number of CIBs that may be simultaneously chained for a task.

The QEDIT macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede QEDIT.
QEDIT	
␣	One or more blanks must follow QEDIT.

ORIGIN= <i>CIB addr ptr</i>	<i>CIB addr ptr</i> : RX-type address, or register (2) - (12).
,BLOCK= <i>CIB addr</i>	<i>CIB addr</i> : RX-type address, or register (2) - (12).
,CIBCTR= <i>CIB nmb</i>	<i>CIB nmb</i> : decimal digit, with a maximum value of 255.

The parameters are explained as follows:

ORIGIN = *CIB addr ptr*

specifies the address of the pointer to the first CIB chain for the task. This address is obtained using the EXTRACT macro instruction. If ORIGIN is the only parameter specified, the caller must be executing under system key 0-7; in this case, the entire CIB chain is freed.

,BLOCK = *CIB addr*

specifies the address of the CIB to be freed from the CIB chain for a task.

,CIBCTR = *CIB nmb*

specifies the limit for the number of CIBs to be chained at any time for a task.

Note: When using any address returned from the EXTRACT macro instruction as input to the QEDIT macro instruction, the user must use IEZCOM to establish addressability based on the address returned by EXTRACT.

Example 1

Operation: Free the entire CIB chain, where register 8 contains the address of the pointer to the CIB chain.

```
QEDIT  ORGIN=( 8 )
```

Example 2

Operation: Free the CIB whose address is in register 5 from the CIB chain. Register 8 contains the address of the pointer to the CIB chain.

```
QEDIT  ORIGIN=( 8 ), BLOCK=( 5 )
```

RACDEF — Define a Resource to RACF

The RACDEF macro instruction is used to define, modify, or delete resource profiles for RACF. RACF uses the profiles to perform RACHECK authorization checking.

The RACDEF caller must be authorized (APF-authorized, in system key 0-7, or in supervisor state).

A RACF user can change or add the RACDEF parameters, OWNER, LEVEL, UACC, or AUDIT by means of the RACDEF pre-processing routine. This routine is described in OS/VS2 MVS Resource Access Control Facility (RACF): Installation Reference Manual.

The standard form of the RACDEF macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACDEF.
RACDEF	
b	One or more blanks must follow RACDEF.
ENTITY= <i>resource name addr</i>	<i>resource name addr</i> A-type address, or register (2) - (12).
,VOLSER= <i>vol addr</i>	<i>vol addr</i> : A-type address, or register (2) - (12). Note : VOLSER is required only for CLASS='DATASET' and DSTYPE not equal to M.
,TYPE=DEFINE	
,TYPE=DEFINE,NEWNAME = <i>new dsn addr</i>	<i>new dsn addr</i> : A-type address, or register (2) - (12).
,TYPE=ADDVOL,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : A-type address, or register (2) - (12).
,TYPE=CHGVOL,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : A-type address, or register (2) - (12).
,TYPE=DELETE	Default : TYPE=DEFINE
,DSTYPE=N	Default : DSTYPE=N
,DSTYPE=V	
,DSTYPE=M	
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address, or register (2) - (12). Default : zero.
,CLASS=' <i>classname</i> '	<i>classname</i> : DASDVOL, DATASET, or TAPEVOL
,CLASS= <i>class name addr</i>	<i>class name addr</i> : A-type address, or register (2) - (12). Default : CLASS='DATASET'
,MENTITY= <i>entity addr</i>	<i>entity addr</i> : A-type address, or register (2) - (12). Default : zero.
,MVOLSER= <i>volser addr</i>	<i>volser addr</i> : A-type address, or register (2) - (12). Default : zero.
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address, or register (2) - (12). Default : zero.
,UNIT= <i>unit addr</i>	<i>unit addr</i> : A-type address, or register (2) - (12).
,SPECIAL=YES	Default : SPECIAL=NO
,SPECIAL=NO	
,OWNER= <i>owner name addr</i>	<i>owner name addr</i> : A-type address, or register (2) - (12).
,LEVEL= <i>number</i>	Default : zero.
,LEVEL= <i>reg</i>	<i>reg</i> : register (2) - (12).
,UACC=ALTER	
,UACC=CONTROL	
,UACC=UPDATE	
,UACC=READ	
,UACC=NONE	
,UACC= <i>reg</i>	<i>reg</i> : register (2) - (12).
,DATA= <i>data addr</i>	<i>data addr</i> : A-type address, or register (2) - (12).
,AUDIT=ALL	
,AUDIT=SUCCESS	
,AUDIT=FAILURES	
AUDIT=NONE	
,AUDIT= <i>reg</i>	<i>reg</i> : register (2) - (12).

The parameters are explained as follows:

ENTITY = *resource name addr*

specifies the address of the name of the resource that is to be defined to, modified, or deleted from RACF. The resource name is a 44-byte DASD data set name for CLASS='DATASET' or a 6-byte volume serial name for CLASS='DASDVOL' or CLASS='TAPEVOL'. The name must be left justified in the field and padded with blanks. The length of all other resource names is determined by the class descriptor table.

,VOLSER = *vol addr*

specifies the address of the volume serial number:

- For TYPE=ADDVOL, of the new volume to be added to the definition of the data set.
- For TYPE=ADDVOL and CLASS='TAPEVOL', of the new volume being added to the tape volume set identified by ENTITY.
- For TYPE=DEFINE and CLASS='DATASET', of the catalog (for a VSAM data set), or of the volume on which the data set resides (for a non-VSAM data set).

The volume serial number is optional if DSTYPE=M is specified.

The field pointed to by the specified address contains the volume serial number (padded to the right with blanks, if necessary, to make six characters).

,TYPE = DEFINE

,TYPE = DEFINE,NEWNAME = *new dsn addr*

,TYPE = ADDVOL,OLDVOL = *old vol addr*

,TYPE = CHGVOL,OLDVOL = *old vol addr*

,TYPE = DELETE

specifies the type of action to be taken:

- TYPE=DEFINE - The definition of the resource is added to the RACF data set, and the current user is established as the owner of the defined entity. If NEWNAME is specified, the address points to a 44-byte field containing the new name for the data set that is to be renamed. NEWNAME is only valid with CLASS='DATASET'.
- TYPE=ADDVOL - The new volume is added to the definition of the specified resource. For the DATASET class, the OLDVOL address specifies a previous volume of a multivolume data set. For the TAPEVOL class, the ENTITY address specifies a previous volume of a tape volume set.
- TYPE=CHGVOL - The volume serial number in the definition of the specified resource is changed from the old volume serial number identified in OLDVOL to the new volume serial number identified in the VOLSER parameter.
- TYPE=DELETE - The definition of the resource is removed from the RACF data set. (If a multivolume data set or a tape volume set, only the specified volume is removed from the definition.)

,DSTYPE = N

,DSTYPE = V

,DSTYPE = M

specifies whether the data set associated with the request is VSAM (V), non-VSAM (N), or a model profile (M). DSTYPE should only be specified for CLASS='DATASET'.

,INSTLN = *parm list addr*

specifies the address of an area that is to contain parameter information meaningful to the RACDEF installation exit. This information is passed to the installation exit when it is given control from the RACDEF routine.

The INSTLN parameter can be used by an application program acting in the capacity of resource manager that needs to pass information to the RACDEF installation exit.

,CLASS = '*classname*'

,CLASS = *class name addr*

specifies that a profile is to be defined, modified, or deleted for a data set or tape volume.

If an address is specified, the address must point to a one byte length field followed by the class name (for example, DATASET or TAPEVOL).

,MENTITY = *entity addr*

specifies the address of the name of the resource whose profile is to be used as a model in defining the ENTITY profile. MENTITY is valid for the DATASET class only, and is valid only if TYPE=DEFINE is specified and NEWNAME= is not specified. The name is contained in a 44 byte field pointed to by the specified address. The name is left justified in the field and padded with blanks.

,MVOLSER = *volser addr*

specifies the address of the volume serial number of the volume on which the MENTITY data set resides (for a non-VSAM data set), or of the catalog (for a VSAM data set). The field pointed to by the specified address contains the volume serial number, padded to the right with blanks, if necessary, to make six characters.

,ACEE = *acee addr*

specifies the address of the accessor control environment element (ACEE) to be used during RACDEF processing. If no ACEE is specified, or if a register that contains zeros is specified, RACF uses the main ACEE. The main ACEE is pointed to by the ASXBSENV field(x'C8') in the address space extension block.

,UNIT = *unit addr*

specifies the address of a field containing unit information. UNIT is valid only if TYPE=CHGVOL or TYPE=DEFINE is specified. If a unit address is specified, the unit information in the data set profile is replaced by the unit information pointed to by this unit address. The unit address must point to a field containing a one-byte length field (whose value can range from 4 to 8) followed by the actual unit information. If the value in the length field is 4, the unit information is assumed to contain a copy of the information in the UCBTYP field of the UCB; otherwise, the unit information is assumed to be in the generic form (for example, 3330-1).

,SPECIAL = YES

,SPECIAL = NO

specifies whether or not a RACDEF operation is to be completed if the requestor has the SPECIAL authority. SPECIAL=YES is required for a TSO user who defines a data set (either for another user or in a group that does not have the necessary authority) to RACF using ADDSD.

,OWNER = *owner name addr*

specifies the address of a field containing the profile owner's name. OWNER is valid only if TYPE=DEFINE is specified. The owner's name must be a valid (RACF-defined) userid. The address must point to an 8-byte field containing the owner's name, left-justified and padded with blanks.

Note: RACF does not check the validity of the userid if it has been added or modified by the RACDEF pre-processing exit.

,LEVEL = *number*

,LEVEL = *reg*

specifies a level value for the profile. LEVEL is valid only if TYPE=DEFINE is specified. The level number must be a valid decimal number in the range 00 to 99. If a register is specified, its low-order byte must contain the binary representation of the number.

Note: RACF does not check the validity of this number if it has been added or modified by the RACDEF pre-processing exit.

,UACC = ALTER

,UACC = CONTROL

,UACC = UPDATE

,UACC = READ

,UACC = NONE

,UACC = *reg*

specifies a universal access authority for the profile. UACC is valid only if TYPE=DEFINE is specified. UACC must contain a valid universal access authority (ALTER, CONTROL, UPDATE, READ, or NONE). If a register is specified, the low-order byte must contain one of the following valid universal access authorities:

X'80' - ALTER
X'40' - CONTROL
X'20' - UPDATE
X'10' - READ
X'01' - NONE

Note: RACF does not check the validity of the universal access authority if it has been added or modified by the RACDEF pre-processing exit.

,DATA=*data addr*

specifies the address of a field that contains up to 255 characters of installation-defined data to be placed in the profile. DATA is valid only if TYPE=DEFINE is specified. The data address must point to a field containing a one-byte length field (whose value can range from 0 to 255) followed by the actual installation-defined data.

,AUDIT = ALL

,AUDIT = SUCCESS

,AUDIT = FAILURES

,AUDIT = NONE

,AUDIT = *reg*

specifies the types of access that are to be logged to the SMF data set. AUDIT is valid only if TYPE=DEFINE is specified. AUDIT must contain a valid audit value (ALL, SUCCESS, FAILURES, or NONE). If a register is specified, its low-order byte must contain one of the following valid audit types:

X'80' - ALL

X'40' - SUCCESS

X'20' - FAILURES

X'10' - NONE

Note: RACF does not check the validity of the audit type if it has been added or modified by the RACDEF pre-processing exit.

When control is returned, register 15 contains one of these return codes:

Hexadecimal

Code	Meaning
00	RACDEF has completed successfully.
04	For DEFINE, the resource name was previously defined. For NEWNAME, the new resource name was previously defined. For DELETE, the resource name was not previously defined.
08	For DEFINE or DELETE, RACDEF was failed by the installation exit, the data set name was not qualified by a RACF-defined userid or group name, or the user was not authorized to create group data sets. For TYPE=DEFINE, the user is not a RACF defined user. For ADDVOL, the old volume was not defined.
0C	For NEWNAME, the old data set name was not defined; or if the generation data group (GDG) modelling function is active, an attempt was made to rename a GDG name to a name that requires the creation of a new profile.
10	For DEFINE with MENTITY, the model resource was not defined.

September 30, 1981

RACDEF (List Form)

The list form of the RACDEF macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede RACDEF.
RACDEF	
␣	One or more blanks must follow RACDEF.

ENTITY= <i>resource name addr</i>	<i>resource name addr</i> : A-type address. Note : ENTITY must be specified on either the list or the execute form of the macro.
,VOLSER= <i>vol addr</i>	<i>vol addr</i> : A-type address. Note : VOLSER is required only for CLASS='DATASET' and DSTYPE not equal to M.
,TYPE=DEFINE	
,TYPE=DEFINE,NEWNAME= <i>new dsn addr</i>	<i>new dsn addr</i> : A-type address.
,TYPE=ADDVOL,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : A-type address.
,TYPE=CHGVOL,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : A-type address.
,TYPE=DELETE	Default : TYPE=DEFINE
,DSTYPE=N	Default : DSTYPE=N
,DSTYPE=V	
,DSTYPE=M	
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address. Default : zero.
,CLASS='classname'	<i>classname</i> : DASDVOL, DATASET, or TAPEVOL
,CLASS= <i>class name addr</i>	<i>class name addr</i> : A-type address. Default : CLASS='DATASET'
,MENTITY= <i>entity addr</i>	<i>entity addr</i> : A-type address. Default : zero.
,MVOLSER= <i>volser addr</i>	<i>volser addr</i> : A-type address. Default : zero.
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address Default : zero.
,UNIT= <i>unit addr</i>	<i>unit addr</i> : A-type address. Default : SPECIAL=NO
,SPECIAL=YES	
,SPECIAL=NO	
,OWNER= <i>owner name addr</i>	<i>owner name addr</i> : A-type address. Default : zero.
,LEVEL= <i>number</i>	<i>level number addr</i> : A-type address.
,LEVEL= <i>level number addr</i>	
,UACC=ALTER	
,UACC=CONTROL	
,UACC=UPDATE	
,UACC=READ	
,UACC=NONE	
,UACC= <i>uacc addr</i>	<i>uacc addr</i> : A-type address.
,DATA= <i>data addr</i>	<i>data addr</i> : A-type address.
,AUDIT=ALL	
,AUDIT=SUCCESS	
,AUDIT=FAILURES	
,AUDIT=NONE	
,AUDIT= <i>audit addr</i>	<i>audit addr</i> : A-type address.
,MF=L	

The parameters are explained under the standard form of the RACDEF macro instruction, with the following exception:

,MF=L

specifies the list form of the RACDEF macro instruction.

RACDEF (Execute Form)

The execute form of the RACDEF macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACDEF.
RACDEF	
b	One or more blanks must follow RACDEF.
ENTITY= <i>resource name addr</i>	<i>resource name addr</i> : RX-type address, or register (2) - (12). Note : ENTITY must be specified on either the list or the execute form of the macro.
,VOLSER= <i>vol addr</i>	<i>vol addr</i> : RX-type address, or register (2) - (12). Note : VOLSER is required only for CLASS='DATASET' and DSTYPE not equal to M.
,TYPE=DEFINE	
,TYPE=DEFINE,NEWNAME <i>new dsn addr</i>	<i>new dsn addr</i> : RX-type address, or register (2) - (12).
,TYPE=ADDVOL,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : RX-type address, or register (2) - (12).
,TYPE=CHGVOL,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : RX-type address, or register (2) - (12).
,TYPE=DELETE	
,DSTYPE=N	
,DSTYPE=V	
,DSTYPE=M	
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : RX-type address, or register (2) - (12).
,CLASS= <i>class name addr</i>	<i>class name addr</i> : RX-type address, or register (2) - (12).
,MENTITY= <i>entity addr</i>	<i>entity addr</i> : RX-type address, or register (2) - (12).
,MVOLSER= <i>volser addr</i>	<i>volser addr</i> : RX-type address, or register (2) - (12).
,ACEE= <i>acee addr</i>	<i>acee addr</i> : RX-type address, or register (2) - (12).
,UNIT= <i>unit addr</i>	<i>unit addr</i> : RX-type address, or register (2) - (12).
,SPECIAL=YES	
,SPECIAL=NO	
,OWNER= <i>owner name addr</i>	<i>owner name addr</i> : RX-type address, or register (2) - (12).
,LEVEL= <i>number</i>	
,LEVEL= <i>reg</i>	<i>reg</i> : register (2) - (12).
,UACC=ALTER	
,UACC=CONTROL	
,UACC=UPDATE	
,UACC=READ	
,UACC=NONE	
,UACC= <i>reg</i>	<i>reg</i> : register (2) - (12).
,DATA= <i>data addr</i>	<i>data addr</i> : RX-type address or register (2) - (12).
,AUDIT=ALL	
,AUDIT=SUCCESS	
,AUDIT=FAILURES	
,AUDIT=NONE	
,AUDIT= <i>reg</i>	<i>reg</i> : register (2) - (12).
,MF=(<i>E,ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the RACDEF macro instruction, with the following exception:

,MF = (E,ctrl addr)

specifies the execute form of the RACDEF macro instruction using a remote control program parameter list.

September 30, 1981

RACHECK — Check RACF Authorization

The RACHECK macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the PROFILE, CSA, ACEE, and LOG parameters. These parameters are restricted in use and must only be used by programs that are authorized (APF-authorized, in system key 0-7, or in supervisor state).

The syntax of the complete RACHECK macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

The standard form of the RACHECK macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACHECK.
RACHECK	
b	One or more blanks must follow RACHECK.
PROFILE= <i>profile addr</i>	<i>profile addr</i> : A-type address, or register (2) - (12).
ENTITY=(<i>resource name addr</i>)	<i>resource name addr</i> : A-type address, or register (2) - (12).
ENTITY=(<i>resource name addr</i> ,CSA)	
,VOLSER= <i>vol addr</i>	<i>vol addr</i> : A-type address, or register (2) - (12). Note : VOLSER is required only for CLASS='DATASET' and DSTYPE not equal to M.
,CLASS=' <i>classname</i> '	<i>classname</i> : DATASET, DASDVOL, or TAPEVOL
,CLASS= <i>classname addr</i>	<i>classname addr</i> : A-type address, or register (2) - (12). Default : CLASS='DATASET'
,ATTR=READ	<i>reg</i> : register (2) - (12). Default : ATTR=READ
,ATTR=UPDATE	
,ATTR=CONTROL	
,ATTR=ALTER	
,ATTR= <i>reg</i>	
,DSTYPE=N	Default : DSTYPE=N
,DSTYPE=V	
,DSTYPE=M	
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address, or register (2) - (12). Default : zero.
,LOG=ASIS	Default : LOG=ASIS
,LOG=NOFAIL	
,LOG=NONE	
,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : A-type address, or register (2) - (12). Default : zero.
,APPL=' <i>applname</i> '	<i>applname addr</i> : A-type address, or register (2) - (12).
,APPL= <i>applname addr</i>	
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address, or register (2) - (12).
,OWNER= <i>userid addr</i>	<i>userid addr</i> : A-type address, or register (2) - (12).

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.)

PROFILE=*profile addr*

ENTITY=(*resource name addr, CSA*)

PROFILE=*profile addr* specifies that RACF authorization checking is to be performed for the resource whose profile is pointed to be the specified address.

ENTITY=(*resource name addr, CSA*) specifies that RACF authorization checking is to be performed for the indicated resource, and that a copy of the profile is to be maintained in main storage. The storage acquired for the profile is obtained from the common storage area (CSA), and must be freed by the issuer of RACHECK when the profile is no longer needed. If CSA is specified (and the return code produced by the RACHECK macro instruction is 00 or 08), the address of the profile is returned in register 1.

By establishing and maintaining a resource profile, the resource manager can reduce the I/O required to perform RACF authorization checks on highly-accessed resources.

,LOG = ASIS

,LOG = NOFAIL

,LOG = NONE

specifies the types of access attempts to be recorded on the SMF data set:

ASIS - Attempts to be recorded are as specified on the ADDSD or ALTDSD command that was issued for the data set or the RDEFINE or RALTER command for the tape or DASD volume.

NOFAIL - If the authorization check fails, the attempt is not recorded. If the authorization check succeeds, the attempt is recorded as in ASIS.

NONE - The attempt is not to be recorded.

,ACEE=*acee addr*

specifies the address of the accessor control environment element (ACEE) to be used for checking authorization during RACHECK processing. If no ACEE is specified, RACF uses the main ACEE for the address space. The main ACEE is pointed to by the ASXBSENV field(X'C8') of the address space extension block.

RACHECK (List Form)

The list form of the RACHECK macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACHECK.
RACHECK	
b	One or more blanks must follow RACHECK.
PROFILE= <i>profile addr</i>	<i>profile addr</i> : A-type address.
ENTITY=(<i>resource name addr</i>)	<i>resource name addr</i> : A-type address.
ENTITY=(<i>resource name addr</i> , CSA)	Note: PROFILE or ENTITY is required on either the list or the execute form of the macro.
,VOLSER= <i>vol addr</i>	<i>vol addr</i> : A-type address. Note: VOLSER is required only for CLASS='DATASET' and DSTYPE not equal to M.
,CLASS=' <i>classname</i> '	<i>classname</i> : DATASET, DASDVOL, or TAPEVOL
,CLASS= <i>classname addr</i>	<i>classname addr</i> : A-type address. Default: CLASS='DATASET'
,ATTR=READ	Default: ATTR=READ
,ATTR=UPDATE	
,ATTR=CONTROL	
,ATTR=ALTER	
,DSTYPE=N	Default: DSTYPE=N
,DSTYPE=V	
,DSTYPE=M	
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address. Default: zero.
,LOG=ASIS	Default: LOG=ASIS
,LOG=NOFAIL	
,LOG=NONE	
,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : A-type address. Default: zero.
,APPL=' <i>applname</i> '	
,APPL= <i>applname addr</i>	<i>applname addr</i> : A-type address. Default: zero.
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address. Default: zero.
,OWNER= <i>userid addr</i>	<i>userid addr</i> : A-type address.
,MF=L	

The parameters restricted in use are explained under the standard form of the RACHECK macro instruction. The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.

RACHECK (Execute Form)

The execute form of the RACHECK macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACHECK.
RACHECK	
b	One or more blanks must follow RACHECK.

PROFILE= <i>profile addr</i>	<i>profile addr</i> : RX-type address, or register (2) - (12).
ENTITY=(<i>resource name addr</i>)	<i>resource name addr</i> : RX-type address, or register (2) - (12).
ENTITY=(<i>resource name addr</i> ,CSA)	Note: PROFILE or ENTITY is required on either the list or the execute form of the macro.
,VOLSER= <i>vol addr</i>	<i>vol addr</i> : RX-type address, or register (2) - (12). Note: VOLSER is required only for CLASS='DATASET' and DSTYPE not equal to M.
,CLASS= <i>classname addr</i>	<i>classname addr</i> : RX-type address, or register (2) - (12).
,ATTR=READ	
,ATTR=UPDATE	
,ATTR=CONTROL	
,ATTR=ALTER	
,ATTR= <i>reg</i>	<i>reg</i> : register (2) - (12).
,DSTYPE=N	
,DSTYPE=V	
,DSTYPE=M	
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : RX-type address, or register (2) - (12).
,LOG=ASIS	
,LOG=NOFAIL	
,LOG=NONE	
,OLDVOL= <i>old vol addr</i>	<i>old vol addr</i> : RX-type address. Default: zero.
,APPL= <i>applname addr</i>	<i>applname addr</i> : RX-type address or registers (2) - (12).
,ACEE= <i>acee addr</i>	<i>acee addr</i> : RX-type address, or registers (2) - (12).
,OWNER= <i>userid addr</i>	<i>userid addr</i> : RX-type address, or registers (2) - (12).
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters restricted in use are explained under the standard form of the RACHECK macro instruction. The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.

RACINIT — Identify a RACF-Defined User

The RACINIT macro instruction is used to provide Resource Access Control Facility (RACF) user identification and verification. The macro instruction identifies a user and verifies that the user is defined to RACF and has supplied a valid password and/or operator ID card.

To issue the RACINIT macro instruction with the NEWPASS keyword, the calling module must be authorized (APF-authorized, in system key 0-7, or in supervisor state). To issue the RACINIT macro without the NEWPASS keyword, the calling module must either be authorized (APF-authorized, in system key 0-7, or in supervisor state) or in the RACF-authorized caller table and fetched from an authorized library.

The standard form of the RACINIT macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede RACINIT.
RACINIT	
␣	One or more blanks must follow RACINIT.
USERID= <i>userid addr</i>	<i>userid addr</i> : A-type address, or register (2) - (12). Default: zero.
,PASSWRD= <i>password addr</i>	<i>password addr</i> : A-type address, or register (2) - (12). Default: zero.
,START= <i>procname addr</i>	<i>procname addr</i> : A-type address, or register (2) - (12). Default: zero.
,NEWPASS= <i>new password addr</i>	<i>new password addr</i> : A-type address, or register (2) -(12). Default: zero.
,GROUP= <i>group addr</i>	<i>group addr</i> : A-type address, or register (2) - (12). Default: zero.
,PGMNAME= <i>programmer name addr</i>	<i>programmer name addr</i> : A-type address, or register (2) - (12). Default: zero.
,ACTINFO= <i>account addr</i>	<i>account addr</i> : A-type address, or register (2) - (12). Default: zero.
,OIDCARD= <i>oid addr</i>	<i>oid addr</i> : A-type address, or register (2) - (12). Default: zero.
,TERMID= <i>terminal addr</i>	<i>terminal addr</i> : A-type address, or register (2) - (12). Default: zero.
,JOBNAME= <i>jobname addr</i>	<i>jobname addr</i> : A-type address, or register (2) - (12). Default: zero.
,ENVIR=CREATE ,ENVIR=CHANGE ,ENVIR=DELETE	Default: ENVIR=CREATE Note: ENVIR=CHANGE may not be specified with USERID=, PASSWRD=, START=, NEWPASS=, ACTINFO=, PGMNAME=, OIDCARD=, or TERMID= above. Note: ENVIR=DELETE may not be specified with USERID=, PASSWRD=, START=, NEWPASS=, GROUP=, ACTINFO=, PGMNAME=, OIDCARD=, or TERMID= above.
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address, or register (2) - (12). Default: zero.
,APPL= <i>applname</i> ,APPL= <i>applname addr</i>	<i>applname addr</i> : A-type address, or register (2) - (12).
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address, or register (2) -(12).
,SUBPOOL= <i>subpool number</i>	<i>subpool number</i> : decimal digit 0-255.
,SMC=YES ,SMC=NO	Default: SMC=YES.
,PASSCHK=YES ,PASSCHK=NO	Default: PASSCHK=YES

The parameters are explained as follows:

USERID = *userid addr*

specifies the user identification of the user who has entered the system. The address points to a one-byte length field, followed by the userid.

,PASSWORD = *password addr*

specifies the currently defined password of the user who has entered the system. The address points to a one-byte length field, followed by the password.

,START = *procname addr*

specifies the PROC name of a started task. The address points to an eight-byte area containing the PROC name (left-justified and padded with blanks, if necessary).

The START parameter is not used by RACINIT authorization checking, but it is passed to the installation exit.

,NEWPASS = *new password addr*

specifies the password which is to replace the user's currently defined password. The address points to a one-byte length field, followed by the password.

,GROUP = *group addr*

specifies the group specified by the user who has entered the system. The address points to a one-byte length field, followed by the group name.

,PGMNAME = *programmer name addr*

specifies the address of the name of the user who has entered the system. This twenty byte area is passed to the RACINIT installation exit; it is not used by the RACINIT routine.

,ACTINFO = *account addr*

specifies the address of a field containing accounting information. This 144 byte area is passed to the RACINIT installation exit; it is not used by the RACINIT routine. The accounting field, if supplied, should have the following format:

- First byte of field contains the number (binary) of accounting fields.
- Following bytes contain accounting fields, where each entry for an accounting field contains the length of the field (one byte) followed by the field.

,OIDCARD = *oid addr*

specifies the address of the currently defined operator ID card of the user who has entered the system. The address points to a one byte length field followed by the operator ID card.

,TERMINAL = *terminal addr*

specifies the address of the identifier for the terminal through which the user is accessing the system. The address points to an eight byte area containing the terminal identifier. The area must reside in a non-task-related storage subpool.

,JOBNAME = *jobname addr*

specifies the address of the JOB name of a background job. The address points to an eight byte area containing the JOB name (left justified and padded with blanks, if necessary).

The JOBNAME parameter is not used by RACINIT authorization checking, but it is passed to the installation exit.

,ENVIR = CREATE

,ENVIR = CHANGE

,ENVIR = DELETE

specifies the action to be performed by the user initialization component:

CREATE - The user should be verified and an access environment element created.

CHANGE - The access environment pointed to by the ASXB should be modified according to other parameters specified on RACINIT.

DELETE - The access environment element pointed to by the ASXB should be deleted.

This parameter should only be used if a previous RACINIT has completed successfully.

,INSTLN = *parm list addr*

specifies the address of an area containing parameter information meaningful to the RACINIT installation exit. This area is passed to the installation exit when the exit is given control from the RACINIT routine.

The INSTLN parameter can be used by an installation having a user verification or job initiation application, and wanting to pass information from one installation module to the RACINIT installation exit.

,APPL = '*applname*'

,APPL = *applname addr*

specifies the name of the application issuing the RACINIT. If an address is specified, the address must point to an 8-byte application name, left justified and padded with blanks, if necessary.

,ACEE = *acee addr*

specifies the address of the ACEE.

For ENVIR=CHANGE or ENVIR=DELETE:

specifies the address of a fullword that contains the address of the accessor control environment element (ACEE) to be changed or deleted. If ACEE= is not specified, RACF uses the main ACEE for the address space. The main ACEE is pointed to by the ASXBSENV field(x'C8') in the address space extension block.

For ENVIR=CREATE:

specifies the address of a full word into which the RACINIT function will place the address of the ACEE created. If an ACEE is not specified, a pointer to the ACEE is placed in the ASXBSENV field(x'C8') of the address space extension block.

,SUBPOOL = *subpool number*

specifies the storage subpool from which the ACEE and related storage are obtained.

,SMC = YES

,SMC = NO

specifies the use of the step-must-complete function of RACINIT processing. SMC=YES specifies that RACINIT processing should continue to place other tasks for the step non-dispatchable. SMC=NO specifies that the step-must-complete function is not used.

Note: SMC=NO should not be used if DADSM ALLOCATE/SCRATCH functions execute simultaneously in the same address space as the RACINIT function.

,PASSCHK = YES

,PASSCHK = NO

specifies that the user's password is to be verified. PASSCHK=YES specifies that RACINIT verifies the user's password. PASSCHK=NO specifies that the user's password is not verified.

When control is returned, register 15 contains one of these return codes:

Hexadecimal Code	Meaning
00	RACINIT has completed successfully.
04	The user profile is not defined to RACF.
08	The password is not authorized.
0C	The password has expired.
10	The new password is invalid.
14	The user is not defined to the group.
18	RACINIT was failed by the installation exit.
1C	The user access has been revoked.
20	RACF is not active.
24	The user's access to the specified group has been revoked.
28	Operator ID card is required but not supplied.
2C	Operator ID card is invalid for specified user.
30	The user is not authorized to use the terminal.

RACINIT (List Form)

The list form of the RACINIT macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede RACINIT.
RACINIT	
␣	One or more blanks must follow RACINIT.

USERID= <i>userid addr</i>	<i>userid addr</i> : A-type address. Default: zero.
,PASSWRD= <i>password addr</i>	<i>password addr</i> : A-type address. Default: zero.
,START= <i>procname addr</i>	<i>procname addr</i> : A-type address. Default: zero.
,NEWPASS= <i>new password addr</i>	<i>new password addr</i> : A-type address. Default: zero.
,GROUP= <i>group addr</i>	<i>group addr</i> : A-type address. Default: zero.
,PGMNAME= <i>programmer name addr</i>	<i>programmer name addr</i> : A-type address. Default: zero.
,ACTINFO= <i>account addr</i>	<i>account addr</i> : A-type address. Default: zero.
,OIDCARD= <i>oid addr</i>	<i>oid addr</i> : A-type address. Default: zero.
,TERMINAL= <i>terminal addr</i>	<i>terminal addr</i> : A-type address. Default: zero.
,JOBNAME= <i>jobname addr</i>	<i>jobname addr</i> : A-type address. Default: zero.
,ENVIR=CREATE ,ENVIR=CHANGE ,ENVIR=DELETE	Default: ENVIR=CREATE Note: ENVIR=CHANGE may not be specified with USERID=, PASSWRD=, START=, NEWPASS=, ACTINFO=, PGMNAME=, OIDCARD=, or TERMINAL= above. Note: ENVIR=DELETE may not be specified with USERID=, PASSWRD=, START=, NEWPASS=, GROUP=, ACTINFO=, PGMNAME=, OIDCARD=, or TERMINAL= above.
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address. Default: zero.
,APPL= <i>applname</i> ,APPL= <i>applname addr</i>	<i>applname addr</i> : A-type address. Default: zero.
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address. Default: zero.
,SUBPOOL= <i>subpool number</i>	<i>subpool number</i> : decimal digit 0-255.
,SMC=YES ,SMC=NO	Default: SMC=YES.
,PASSCHK=YES ,PASSCHK=NO	Default: PASSCHK=YES
,MF=L	

The parameters are explained under the standard form of the RACINIT macro instruction, with the following exception:

MF=L

specifies the list form of the RACINIT macro instruction.

RACINIT (Execute Form)

The execute form of the RACINIT macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede RACINIT.
RACINIT	
␣	One or more blanks must follow RACINIT.
USERID= <i>userid addr</i>	<i>userid addr</i> : RX-type address, or register (2) - (12).
,PASSWRD= <i>password addr</i>	<i>password addr</i> : RX-type address, or register (2) - (12).
,START= <i>procname addr</i>	<i>procname addr</i> : RX-type address, or register (2) - (12).
,NEWPASS= <i>new password addr</i>	<i>new password addr</i> : RX-type address, or register (2) - (12).
,GROUP= <i>group addr</i>	<i>group addr</i> : RX-type address, or register (2) - (12).
,PGMNAME= <i>programmer name addr</i>	<i>programmer name addr</i> : RX-type address, or register (2) - (12).
,ACTINFO= <i>account addr</i>	<i>account addr</i> : RX-type address, or register (2) - (12).
,OIDCARD= <i>oid addr</i>	<i>oid addr</i> : RX-type address, or register (2) - (12).
,TERMID= <i>terminal addr</i>	<i>terminal addr</i> : RX-type address, or register (2) - (12).
,JOBNAME= <i>jobname addr</i>	<i>jobname addr</i> : RX-type address, or register (2) - (12).
,ENVIR=CREATE ,ENVIR=CHANGE ,ENVIR=DELETE	Note: ENVIR=CHANGE may not be specified with USERID=, PASSWRD=, START=, NEWPASS=, ACTINFO=, PGMNAME=, OIDCARD=, or TERMID= above. Note: ENVIR=DELETE may not be specified with USERID=, PASSWRD=, START=, NEWPASS=, GROUP=, ACTINFO=, PGMNAME=, OIDCARD=, or TERMID= above.
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : RX-type address, or register (2) - (12).
,APPL= <i>applname addr</i>	<i>applname addr</i> : RX-type address, or register (2) - (12).
,ACEE= <i>acee addr</i>	<i>acee addr</i> : RX-type address, or register (2) - (12).
,SUBPOOL= <i>subpool number</i>	<i>subpool number</i> : decimal digit 0-255.
,SMC=YES ,SMC=NO	
,PASSCHK=YES ,PASSCHK=NO	
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the RACINIT macro instruction, with the following exception:

,MF=(E,*ctrl addr*)
specifies the execute form of the RACINIT macro instruction using a remote control program parameter list.

RACLIST — Build In-Storage Profiles

RACLIST is used to build in-storage profiles for RACF defined resources. RACLIST processes only those resources described by class descriptors. The primary advantage of using the RACLIST macro is to use the resource grouping function and to improve resource authorization checking performance.

The module calling the RACLIST macro instruction must either be authorized (APF-authorized, in system key 0-7, or in supervisor state) or in the RACF-authorized caller table and fetched from an authorized library.

The standard form of the RACLIST macro is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACLIST.
RACLIST	
b	One or more blanks must follow RACLIST.

CLASS='classname'	
CLASS=classname addr	classname addr: A-type address or register (2) - (12).
,LIST=list addr	list addr: A-type address or register (2) - (12). Default: zero.
,ACEE=acee addr	acee addr: A-type address or register (2) - (12). Default: zero.
,INSTLN=parm list addr	parm list addr: A-type address or register (2) - (12).
,APPL='applname'	
,APPL=applname addr	applname addr: A-type address or register (2) - (12). Default: zero.
,SUBPOOL=(sub#1,sub#2)	sub#,sub#2: decimal digit 0-255.
,ENVIR=CREATE	
,ENVIR=DELETE	Default: ENVIR=CREATE.
,OWNER=YES	
,OWNER=NO	Default: OWNER=NO.

The parameters are explained below:

CLASS='classname'

CLASS=classname addr

specifies that RACLIST is to build an in-storage profile for the resources of the specified class. If an address is specified, the address must point to an 8-byte field containing the class name, left justified and padded with blanks, if necessary. The class name must be defined by a class descriptor; if not, the class is not considered to be defined.

,LIST = *addr*

specifies the address of a list of resource names for which RACLIST is to build the in-storage profiles. The list consists of a 2-byte field containing the number of the names in the list, followed by one or more variable length names. Each name consists of a 1-byte-length field, followed by the name. A zero in the 2-byte field causes the operand to be omitted. If LIST= is omitted, in-storage profiles are built for all the profiles defined to RACF in the given class as well as each member for a resource grouping associated with the specified class.

Note: This operand can be specified only with ENVIR=CREATE. If ENVIR=DELETE is specified, the RACLIST macro instruction issues a return code of 18.

,ACEE = *acee addr*

specifies the address of the accessor control environment element (ACEE). The ACEE points to the in-storage profiles. If an ACEE is not specified, RACF uses the main ACEE to obtain the list of the in-storage profiles. The main ACEE is pointed to by the ASXBSENV field(x'08') of the address space extension block. If an ACEE is not specified and there is no main ACEE, the in-storage profiles are not constructed.

,INSTLN = *parm list addr*

specifies the address of an area that contains parameter information for the RACLIST installation exit. The address is passed to the installation exit when the exit is given control by the RACLIST routine. The INSTLN parameter can be used by an application or an installation program to pass information to the RACLIST installation exit.

,APPL = '*applname*'**,APPL = *applname addr***

specifies the name of the application requesting the authorization checking. This information is not used for the authorization checking process but is made available to the installation exit(s). If an address is specified, it should point to an 8-byte area containing the application name, left justified and padded with blanks, if necessary.

,SUBPOOL = (*sub#1,sub#2*)

specifies the subpool numbers of the storage into which the components of the in-storage profiles are to be built. Sub#1 represents the subpool of the profile index. Sub#2 represents the subpool of the profile proper. If the subpools are not specified they default to subpool 255. Registers can be used to specify sub#1 and sub#2.

,ENVIR = CREATE**,ENVIR = DELETE**

specifies the action to be performed by the RACLIST macro.

CREATE - In-storage profiles for the specified class are to be built. The RACLIST function issues a return code of 18, if an in-storage list currently exists for the specified class.

DELETE - The in-storage profiles for the specified class are to be freed. If class is not specified, the in-storage profiles for all classes are freed.

Note: It is the responsibility of the user issuing the RACLIST macro to assure that no multi-tasking that results in the issuing of a RACHECK, FRACHECK, RACINIT, or RACLIST macro instruction occurs at the same time that the RACLIST occurs.

,OWNER = YES

,OWNER = NO

specifies that the resource owner is to be placed in the profile access list with the ALTER authority. If the OWNER= operand is omitted, the default is NO.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal

Code	Meaning
00	RACLIST function completed successfully.
04	Unable to perform the requested function. Register 0 contains additional codes as follows: <ul style="list-style-type: none">0 - Unable to establish an ESTAE environment.1 - The function code (the third byte of the parameter list) does not represent a valid function. X'01' represents the RACF manager; X'02' represents the RACLIST macro.
08	The specified class is not defined to RACF.
0C	An error was encountered during RACLIST processing.
10	RACF and/or the resource class is not active.
14	RACLIST installation exit error occurred.
18	Parameter list error.
1C	RACF CVT does not exist (RACF is not installed) or an insufficient level of RACF is installed.

Note: If the resource class specified by the CLASS= operand is inactive, RACLIST does not build the in-storage profiles and a code of 0C is returned. If the resource group class is not active, RACLIST builds an in-storage profile but only from the individual resource profiles; resource group profiles are ignored.

RACLIST (List Form)

The list form of the RACLIST macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACLIST.
RACLIST	
b	One or more blanks must follow RACLIST.

CLASS= <i>'classname'</i>	
CLASS= <i>classname addr</i>	<i>classname addr</i> : A-type address. Default: zero.
,LIST= <i>list addr</i>	<i>list addr</i> : A-type address. Default: zero.
,ACEE= <i>acee addr</i>	<i>acee addr</i> : A-type address. Default: Zero.
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address. Default: zero.
,APPL= <i>'applname'</i>	
,APPL= <i>applname addr</i>	<i>applname addr</i> : A-type address. Default: zero.
,SUBPOOL=(<i>sub#1;sub#2</i>)	<i>sub#1,sub#2</i> : decimal digit 0-255. Default: 255.
,ENVIR=CREATE	
,ENVIR=DELETE	Default: ENVIR=CREATE.
,OWNER=YES	
,OWNER=NO	Default: OWNER=NO.
,MF=L	

The parameters are explained under the standard form of the RACLIST macro instruction with the following exception:

MF=L

specifies the list form of the RACLIST macro instruction.

RACLIST (Execute Form)

The execute form of the RACLIST macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RACLIST.
RACLIST	
b	One or more blanks must follow RACLIST.

CLASS= <i>classname addr</i>	<i>classname addr</i> : RX-type address or register (2) - (12).
,LIST= <i>list addr</i>	<i>list addr</i> : RX-type address or register (2) - (12).
,ACEE= <i>acee addr</i>	<i>acee addr</i> : RX-type address or register (2) - (12).
,INSTLN= <i>parm list addr</i>	<i>parm list addr</i> : A-type address or register (2) - (12).
,APPL= <i>applname addr</i>	<i>applname addr</i> : RX-type address or register (2) - (12).
,SUBPOOL=(<i>sub#1,sub#2</i>)	<i>sub#1,sub#2</i> : decimal digit 0-255.
,ENVIR=CREATE	
,ENVIR=DELETE	
,OWNER=YES	
,OWNER=NO	
,MF=(<i>E,ctrl addr</i>)	<i>ctrl addr</i> : RX-type address or register (2) - (12).

The parameters are explained under the standard form of the RACLIST macro instruction with the following exception:

,MF=(*E,ctrl addr*)
specifies the execute form of the RACLIST macro instruction using a remote control program parameter list.

RESERVE — Reserve a Device (Shared DASD)

The RESERVE macro instruction is used to reserve a device for use by a particular system; it must be issued by each task needing device reservation. The RESERVE macro instruction protects the issuing task from interference by other tasks in the system and locks out the other processor. When the reserving program no longer needs the reserved device, it should issue a DEQ macro instruction, specifying the UCB parameter, to release the device. If a task issues two RESERVE instructions for the same resource without an intervening DEQ, an abnormal termination results unless the second RESERVE specifies the keyword parameter RET= or ECB=. (If a 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.) If a DEQ is not issued for a particular device, termination routines release devices reserved by a terminating task.

To use the shared DASD option in higher level languages, an assembler language subroutine should be written to issue the RESERVE macro instruction. The following information should be passed to this routine: *ddname*, *qnameaddress*, *rnameaddress*, *rnamelength*, and RET parameter.

The standard form of the RESERVE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
<i>b</i>	One or more blanks must precede RESERVE.
RESERVE	
<i>b</i>	One or more blanks must follow RESERVE.

(
<i>qname addr</i>	<i>qname addr</i> : A-type address, or register (2) - (12).
<i>,rname addr</i>	<i>rname addr</i> : A-type address, or register (2) - (12).
,	Default: E
,E	
,S	
,	
<i>,rname length</i>	<i>rname length</i> : symbol, decimal digit, or register (2) - (12).
,SYSTEMS	
)	
,RET=TEST	
,RET=USE	
,RET=HAVE	
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (2) - (12).
,UCB= <i>ucb addr</i>	<i>ucb addr</i> : A-type address, or register (2) - (12).
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained as follows:

- (specifies the beginning of the resource description.
- qname addr* specifies the address in virtual storage of an 8-character name. The name should not start with SYS, so that it will not conflict with system names. Every task issuing RESERVE against the same resource must use the same *qname* and *rname* to represent the resource.
- ,rname addr* specifies the address in virtual storage of the name used in conjunction with *qname* to represent a single resource. The name can be qualified, and must be from 1 to 255 bytes long.
- ,
- ,E
- ,S specifies whether the request is for exclusive (E) or shared (S) control of the resource. If the resource is modified while under control of the task, the request must be for exclusive control; if the resource is not modified, the request should be for shared control.
- ,
- ,rname length* specifies the length of the *rname* described above. If this parameter is omitted, the assembled length of the *rname* is used. You can specify a value between 1 to 255 to override the assembled length, or you may specify a value of 0. If 0 is specified, the length of the *rname* must be contained in the first byte at the *rname addr* specified above.
- ,SYSTEMS specifies that the resource is shared between systems.
-) specifies the end of the resource description.
- ,RET = TEST
- ,RET = USE
- ,RET = HAVE specifies a conditional request for all the resources named above.
 - RET = TEST the availability of the resources is to be tested, but control of the resources is not requested.
 - RET = USE control of the resources is to be assigned to the active task only if the resources are immediately available.
 - RET = HAVE control of the resources is requested only if a request has not been made previously for the same task.
- ,ECB = *ecb addr* specifies the address of an ECB, and conditionally requests the resource named in the macro instruction. If the return code for one or more requested resources is 4 and the request is not nullified by a corresponding DEQ, the ECB is posted when all the requested resources (specifically, those that initially received a return code of 4) are assigned to the requesting task.
- ,UCB = *ucb addr* specifies the address of a fullword that contains the address of the UCB for the device to be reserved. The UCB must be allocated to the job step before RESERVE is issued unless the issuer is in supervisor state, system key, or APF-authorized.

,RELATED = value

specifies information used to self-document macro instructions by "relating" functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R,LV=4096,RELATED=( FREE1, 'GET STORAGE' )
FREE1   FREEMAIN   R,LV=4096,A=( 1 ),RELATED=( GET1, 'FREE STORAGE' )
```

Return codes are provided by the control program only if you specify RET=TEST, RET=USE, RET=HAVE, or ECB=; otherwise, return of the task to the active condition indicates that control of the resource has been assigned to the task. If return code for the resource named in the RESERVE macro instruction is 0, register 15 contains 0. If the return code is not 0, register 15 contains the address of a storage area containing the return codes, as shown in Figure 34.




Figure 34. Return Code Area Used by RESERVE

The return code is placed in the parameter list resulting from the macro expansion. The return codes are shown below.

Hexadecimal Code	Meaning
0	For RET=TEST, the resource was immediately available. For RET=USE, RET=HAVE, or ECB=, control of the resource has been assigned to the active task.
4	For RET=TEST or RET=USE, the resource is not immediately available. For ECB=, the ECB will be posted when available.
8	A previous request for control of the same resource has been made for the same task. Task has control of resource. If bit 3 is on — shared control of resource; if bit 3 is off — exclusive control.
20	A previous request for control of the same resource has been made for the same task. Task does not have control of resource.

Example 1

Operation: Unconditionally reserve exclusive control of a device. The length of the rname is allowed to default.

```
RESERVE ( MAJOR3 , MINOR3 , E , , SYSTEMS ) , UCB=( R3 )
```

RESERVE (List Form)

The list form of the RESERVE macro is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede RESERVE.
RESERVE	
␣	One or more blanks must follow RESERVE.

(
<i>qname addr</i>	<i>qname addr</i> : A-type address.
,	<i>rname addr</i> : A-type address.
<i>rname addr</i>	
,	
,E	
,S	
,	<i>mame length</i> : symbol or decimal digit.
<i>rname length</i>	Note: <i>rname length</i> must be coded if a register is specified for <i>mame addr</i> above.
,	
,SYSTEMS	
)	
,RET=TEST	
,RET=USE	
,RET=HAVE	
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address.
,UCB=0	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=L	

The parameters are explained under the standard form of the RESERVE macro instruction, with the following exception:

,MF=L
specifies the list form of the RESERVE macro instruction.

RESERVE (Execute Form)

The execute form of the RESERVE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RESERVE.
RESERVE	
b	One or more blanks must follow RESERVE.

(Note: (and) are the beginning and end of a parameter list. The entire list is optional. If nothing in the list is desired, the (,), and all parameters between (and) should not be specified. If something in the list is desired, then (,), and all parameters in the list should be specified as indicated at the left.
<i>qname addr</i>	<i>qname addr</i> : RX-type address, or register (2) - (12).
,	
<i>rname addr</i>	<i>rname addr</i> : RX-type address, or register (2) - (12).
,	
,E	
,S	
,	
<i>rname length</i>	<i>rname length</i> : symbol, decimal digit, or register (2) - (12).
,	
,SYSTEMS	
)	
,RET=TEST	
,RET=USE	
,RET=HAVE	
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : RX-type address, or register (2) - (12).
,UCB= <i>ucb addr</i>	<i>ucb addr</i> : RX-type address, or register (2) - (12).
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the RESERVE macro instruction, with the following exception:

,MF=(E, *ctrl addr*)
specifies the execute form of the RESERVE macro instruction using a remote control program parameter list.

RESUME — Resume Execution of a Suspended Request Block Task

The RESUME macro instruction causes suspended RBs to resume execution.

The RESUME macro instruction is coded as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
<i>b</i>	One or more blanks must precede RESUME.
RESUME	
<i>b</i>	One or more blanks must follow RESUME.

TCB=(4)	Default: TCB address contents of register (4).
TCB= <i>tcbaddr</i>	<i>tcbaddr</i> : A-type address or registers (2) - (12).
RB=(5)	Default: RB address contents of register (5).
RB= <i>rbaddr</i>	<i>rbaddr</i> : A-type address or registers (2) - (12).
RETURN=Y	Default: Y
RETURN=N	

The parameters are explained as follows:

TCB=(4)

TCB=*tcbaddr*

specifies the TCB address of the task to be resumed. Register (4) is the default; it is assumed to contain the TCB address.

RB=(5)

RB=*rbaddr*

specifies the address of the RB to be resumed. Register (5) is the default; it is assumed to contain the address of the RB to be resumed. The specification of the RB operand determines which RB will have its suspend count decremented (which RB will be made ready for resumption of execution).

RETURN=Y

RETURN=N

specifies whether control is to return to the caller (RETURN=Y). If the caller is an SRB and the specification is RETURN=N, either the TCTL function receives control or an abend occurs.

The RESUME macro instruction uses registers as follows:

Register	Use
0-1	Work registers
2-3	Unchanged
4	TCB address
5	RB address
6-10	Unchanged
11-13	Work registers
14	Return point
15	EPA of the RESUME routine, or return code, if RETURN=Y was specified.

Example 1

Operation: Resume execution of the task specified in the address labeled CURRTCB. Use the request block address in register 5. Pass control back to the task (the issuer is currently in SRB mode, and this step terminates SRB mode processing).

```
RESUME      TCB=CURRTCB, RB=( 5 ), RETURN=N
```

RISGNL — Issue Remote Immediate Signal

The RISGNL macro instruction uses the emergency signal (ES) function of the signal processor (SIGP) instruction to invoke the execution of a specified software program on one of the processors in a tightly coupled multiprocessing system. The program may be requested to execute in parallel or serially with the function requesting the program.

Ten of the twelve SIGP hardware functions are defined as direct services and are accessible via the DSGNL macro instruction. The other SIGP function is accessible via the RPSGNL macro instruction.

The RISGNL macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede RISGNL.
RISGNL	
␣	One or more blanks must follow RISGNL.

PARALLEL	
SERIAL	
,CPU= <i>PCCA addr</i>	<i>PCCA addr</i> : RX-type address, or register (1).
,EP= <i>entry name addr</i>	<i>entry name addr</i> : RX-type address, or register (12).
,PARM= <i>parm addr</i>	<i>parm addr</i> : RX-type address, or register (11).

The parameters are explained as follows:

PARALLEL
SERIAL

specifies that control is to be returned to the caller when the specified receiving routine has been given control (PARALLEL) or has completed execution (SERIAL) on the designated processor.

,CPU=*PCCA addr*

specifies the address of the physical configuration communication area (PCCA) of the processor on which the function is to be performed.

,EP=*entry name addr*

specifies the address of the entry name of the receiving routine to be executed on the specified processor.

,PARM=*parm addr*

specifies the address of a user-defined fullword parameter to be passed to the receiving routine.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	Specified receiving routine has been given control or has completed execution, as requested.
04	Function not initiated because addressed processor not online.
08	Function unsuccessful. Emergency signal could not be generated on processor. Status returned in register 0.
12	Function unsuccessful. Specified processor is either not installed, not configured into system, or powered off.
16	Processor is a uniprocessor and does not have signal processor sending and receiving capabilities.
20	Processor alive bit was turned off during the remote immediate window spin routine.

With a return code of 8, register 0 contains:

Bits	Meaning
0	Equipment check
1-24	Reserved
25	Stopped
26	Operator intervening
27	Check stop
28	Not ready
29	Reserved
30	Invalid function
31	Receiver check

Example 1

Operation: The routine whose address is in register 12 is to be given control on the processor whose PCCA address is in register 1. The routine will execute in parallel with the caller who invoked RISGNL.

```
RISGNL  PARALLEL,CPU=( 1 ),EP=( 12 )
```

Example 2

Operation: The routine whose address is in register 12 is to be given control on the processor whose PCCA address is in register 1. The routine will complete before the caller of RISGNL receives control again. Register 11 will contain the address of a parameter to be passed.

```
RISGNL  SERIAL,CPU=( 1 ),EP=( 12 ),PARM=( 11 )
```

RPSGNL — Issue Remote Pendable Signal

The RPSGNL macro instruction uses the external call (EC) function of the signal processor (SIGP) instruction to invoke the execution of one of six software programs on one of the processors in a tightly coupled multiprocessing system.

Ten of the twelve SIGP hardware functions are defined as direct services and are accessible via the DSGNL macro instruction. The other SIGP function is accessible via the RISGNL macro instruction.

The RPSGNL macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede RPSGNL.
RPSGNL	
b	One or more blanks must follow RPSGNL.

SWITCH
SIO
RQCHECK
GTFCRM
MODE
MFITCH

,CPU=PCCA addr *PCCA addr*: RX-type address, or register (1).

The parameters are explained as follows:

SWITCH
SIO
RQCHECK
GTFCRM
MODE
MFITCH

specifies the action to be performed:

SWITCH Memory/task switch function

SIO IOS start I/O function

RQCHECK Timer supervision TOE check function, to ensure that TOE in real time queue is being timed.

GTFCRM GTF function, to modify monitor call control registers

MODE RMS function, to modify RMS-oriented control registers

MFITCH MF1 function, to issue TCH instructions on processor to which channels are attached.

,CPU=PCCA addr

specifies the address of the physical configuration communication area (PCCA) of the processor on which the function is to be executed.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	Specified processor is online and has been notified that the specified service is to be executed.
04	Function not initiated because addressed processor not online.
08	Function unsuccessful. External call signal could not be generated on processor. Status returned in register 0.
12	Specified processor is either not installed, not configured into system, or powered off.
16	Processor is a uniprocessor and does not have signal processor sending and receiving capabilities.

With a return code of 8, register 0 contains:

Bits	Meaning
0	Equipment check
1-25	Reserved
26	Operator intervening
27	Check stop
28	Not ready
29-30	Reserved
31	Receiver check

Example 1

Operation: The service routine of memory switch is to be given control on the processor whose PCCA address is in register 1.

```
RPSGNL SWITCH,CPU=(1)
```

Example 2

Operation: The IOS start I/O routine is to be given control on the processor whose PCCA address is in register 1.

```
RPSGNL SIO,CPU=(1)
```

SCHEDULE — Schedule System Services for Asynchronous Execution

The SCHEDULE macro instruction schedules system services for asynchronous execution. These services may be scheduled for execution in any address space and may be scheduled at either global or local priorities.

Services scheduled at a global priority have a priority that is greater than, and independent of, any address space priority. Services scheduled at a local priority have the priority of the specific address space they execute in, but still have a priority greater than that of any task within the address space. To use SCHEDULE you must be in supervisor state, key zero.

Note: This macro instruction requires that the SRB and CVT mapping macros be assembled as a DSECT into the caller's program.

The SCHEDULE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede SCHEDULE.
SCHEDULE	
␣	One or more blanks must follow SCHEDULE.

SRB= <i>SRB addr</i>	<i>SRB addr</i> : RX-type address, or register (1) or (2) - (12).
,SCOPE=LOCAL	Default: SCOPE=LOCAL.
,SCOPE=GLOBAL	

The parameters are explained as follows:

SRB=*SRB addr*

specifies the address of the service request block (SRB).

,SCOPE=LOCAL

,SCOPE=GLOBAL

specifies whether the service is to be scheduled at a local or global priority.

Example 1

Operation: Schedule an SRB at a global priority.

```
SCHEDULE    SRB=( 1 ),SCOPE=GLOBAL
```

Example 2

Operation: Schedule an SRB at a local priority.

```
SCHEDULE    SRB=( 1 ),SCOPE=LOCAL
```

SDUMP — Dump Virtual Storage

The **SDUMP** macro instruction provides a dumping capability for the system routines. It invokes **SVC Dump** to provide a fast unformatted dump of virtual storage to a data set. It is intended to be used by system routines that suffer errors.

SVC Dump is available only to authorized programs. Issuers of **SDUMP** with entry by **SVC** must be authorized via **APF** or have a control program key. Branch entry callers must be key zero, supervisor state, and must be in **SRB** mode, or own a lock, or be disabled (with supervisor bit on).

The service of initiating an **SVC Dump** in any address space is provided for callers who need to dump address spaces other than the one in which they are running. A branch entry to this service is also provided for callers who wish a dump of their own or another address space but cannot issue an **SVC**. The system operator can take **SVC Dumps** by issuing the **DUMP** command. For more information see *Operator's Library: OS/VS2 MVS System Commands*.

The standard form of the **SDUMP** macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must follow SDUMP.
SDUMP	
b	One or more blanks must follow SDUMP.

HDR='dump title'	<i>dump title</i> : from 1 to 100 characters.
HDRAD=dump title addr	<i>dump title addr</i> : A-type address, or register (2) - (12).
,DCB=dcb addr	<i>dcb addr</i> : A-type address, or register (2) - (12).
,ASID=ASID addr	<i>ASID addr</i> : A-type address, or register (2) - (12).
,ASIDLST=list addr	<i>list addr</i> : A-type address, RX-type address, or register (2) - (12).
,ECB=ecb addr	<i>ecb addr</i> : A-type address, or register (2) - (12). Note : If ECB is specified, ASID or ASIDLST must also be specified.
,SDATA=(data code)	<i>data code</i> : any combination of the following, separated by commas: PSA CSA NUC SWA LSQA ALLPSA or NOALLPSA (abbreviated NOALL) RGN SQA or NOSQA LPA SUMDUMP (abbreviated SUM) or TRT NOSUMDUMP (abbreviated NOSUM) Note : Executing the SDUMP macro results in the ALLPSA, SQA, and SUMDUMP storage areas being dumped unless excluded by the NOALLPSA, NOSQA, or NOSUMDUMP parameter.
,STORAGE=(strt addr,end addr)	<i>strt addr</i> : A-type address, or register (2) - (12).
,LIST=list addr	<i>end addr</i> : A-type address, or register (2) - (12). <i>list addr</i> : A-type address, or register (2) - (12). Note : One or more pairs of addresses may be specified, separated by commas. For example: ,STORAGE=(strt addr,end addr,strt addr,end addr)
,BUFFER=NO	Default : BUFFER=NO
,BUFFER=YES	
,QUIESCE=YES	Default : QUIESCE=YES
,QUIESCE=NO	
,BRANCH=NO	Default : BRANCH=NO
,BRANCH=YES	Note : If BRANCH=YES is specified, ASID or ASIDLST must also be specified.
,SUMLIST=list addr	<i>list addr</i> : A-type address, RX-type address, or register (2) - (12).

The parameters are explained as follows:

HDR = *'dump title'*

HDRAD = *dump title addr*

specifies the title or address of the title to be used for the dump. If HDR is specified, the title must be 1-100 characters enclosed in apostrophes, although the apostrophes do not appear in the actual title. If HDRAD is specified, the first byte at the indicated address specifies the length of the title in bytes.

Note: The Print Dump Service Aid inserts the first 62 characters of the title on its output pages.

,DCB = *dcb addr*

specifies the address of a previously opened data control block for the data set that is to contain the dump. If this parameter is omitted, one of the SYS1.DUMP data sets is used. The data control block must be addressable from all the address spaces in which the SVC Dump routine executes. The control blocks built by OPEN must also be addressable from the address spaces. The DCB must support EXCP.

The DCB must reference device type supported by SVC dump. Eligible device types are unlabeled 9-track 2400-series tape devices (or tape devices compatible with the 2400-series) and any direct access devices supported by the system that have a track size of at least 4104 bytes. (4104 bytes equals 1 SVC dump output record.) The following are examples of supported direct access devices:

- 2314
- 2319
- 3330
- 3330-1
- 2305-1
- 2305-2
- 3340/3344
- 3350

,ASID = *ASID addr*

,ASIDLST = *list addr*

specifies the address of a halfword or a list of halfwords containing the hexadecimal address space identifier of an address space to be dumped. If register notation is used, the low order halfword of the register contains the address space identifier of the address space to be dumped. If both parameters are omitted, the current address space will be dumped. If 0 is specified for the address space identifier, a dump is scheduled in the current address space.

No private area storage will be included in the dump for the specified address space(s) if either of the following events occurred:

- No SDATA parameters were specified that apply to the private area of the requested address space(s).
- The CHNGDUMP operator command was used to set an overriding parameter in the system dump options list that limits SVC Dumps to areas outside of the private area.

The ASID list can contain a maximum of 15 address space identifiers. The high order bit of the halfword containing the last identifier of the list must be set to 1, and all other high order bits must be set to 0.

,ECB = *ecb addr*

specifies the address of a fullword containing the address of an event control block that is posted on completion of a scheduled dump. If this parameter is omitted, the caller is not notified of the completion of the scheduled dump. The fullword and the event control block must be addressable from all the address spaces in which the SVC Dump routine executes.

,SDATA=(data code)

specifies the system control program information to be dumped:

SQA — The system queue area(subpools 227, 228, 239, and 245).

ALLPSA — All of the prefixed storage areas in the system.

PSA — The prefixed storage area for the current processor.

NUC — The nucleus.

LSQA — The local system queue area for each address space being dumped (subpools 229, 230, 233-235, and 253-255).

RGN — The allocated pages in the private area of each address space being dumped. This includes the LSQA and the SWA.

LPA — The active link pack area modules and SVCs for each address space being dumped.

TRT — The GTF trace buffers if GTF tracing is active, or the supervisor trace table if it is not active. If a dump occurs in a GTF address space, no attempt is made to include trace information. If master trace is active, the master trace table is also included in the dump.

CSA — The common service area subpools (subpools 231 and 241).

SWA — The scheduler work area subpools for each address space being dumped (subpools 236 and 237).

NOALLPSA(NOALL) — Only the prefixed storage area for the current processor are dumped.

NOSQA — The system queue area is not dumped.

NOSUMDUMP(NOSUM) — A summary dump is not included in the SVC Dump.

SUMDUMP(SUM) — A summary dump is written to the same dump data set as the other portions of the SVC Dump.

If the **BRANCH=YES** parameter is specified, the following system areas are included in the summary dump output:

- Any storage areas specified with the **SUMLIST** parameter.
- The physical configuration communication area (**PCCA**), the logical configuration communication area (**LCCA**), and the prefixed storage area (**PSA**) for each functioning processor.
- The supervisor trace table, if tracing is requested and GTF tracing is not active. Note that GTF trace records are not dumped.
- The interrupt handler save area (**IHSA**) and 2K of storage before and after every valid, unique address in the registers that are saved in the **IHSA**.
- Any system diagnostic work area (**SDWA**) associated with the failure of the system routine, and 2K of storage before and after every valid, unique address in the registers that are saved in the **SDWA** at the time of the error.
- The global, processor, and local work/save area vector tables (**WASVTG**, **WSAVTC**, **WSAVTL**) and the work/save areas pointed to by the addresses in these vector tables.
- 2K of storage before and after the instruction counter values of the External Old PSW, Program Check Old PSW, I/O Old PSW, and Restart Old PSW saved in the **PSA** of the current processor.
- The functional recovery routine (**FRR**) stack for the current processor.

If the **BRANCH=NO** parameter is in effect, the following system areas are included in the summary dump output:

- Any storage areas specified with the **SUMLIST** parameter.
- The supervisor trace table, if tracing is requested and GTF tracing is not active. Note that GTF trace records are not dumped.

- Each RTM2 work area (RTM2WA) associated with the failing task, and 2K of storage before and after every valid, unique address in the PSW and registers for the failing task. (The PSW and registers for the failing task are saved in each RTM2WA.)

Notes:

The following system control blocks are dumped in all SVC Dumps:

- The communications vector table (CVT)
- The global data area (GDA)
- The prefixed storage area (PSA) for the current CPU
- Unit control blocks (UCBs)
- The address space vector table (ASVT)
- The address space control block (ASCB) for each address space being dumped

Executing the SDUMP macro results in the ALLPSA, SQA, and SUMDUMP storage areas being dumped unless excluded by the NOALLPSA, NOSQA, or NOSUMDUMP parameter.

The system dump options specified by the CHNGDUMP operator command can add to or override the SDATA options specified with the SDUMP macro instruction.

,STORAGE=(*strt addr,end addr*)

,LIST=*list addr*

specifies one or more pairs of starting and ending addresses or a list of starting and ending addresses of areas to be dumped. (Each starting address must be less than its corresponding ending address.) The storage list must contain an even number of addresses, and each address must occupy one fullword. In the list, the high order bit of the fullword containing the last ending address of the list must be set to 1; all other high order bits must be set to 0. The list specified must be addressable from at least one of the address spaces in which the SVC Dump routine will execute.

The STORAGE parameter is not recommended for programs issuing the ASID or ASIDLST parameters and executing in a non-common area of storage.

,BUFFER=NO

,BUFFER=YES

specifies that the contents of the SQA buffer is (YES) or is not (NO) to be included in the dump. (The SQA buffer does not include the SDUMP parameter list or any data pointed to by the parameter list.) Using BUFFER=YES requires special serialization. Refer to the topic "SQA Buffer" in Part 1.

,QUIESCE=YES

,QUIESCE=NO

specifies that the system is to be set nondispatchable until the contents of the SQA and the CSA are dumped (YES), or that the system is to be left dispatchable (NO). If the SDATA parameter does not specify SQA or CSA, the QUIESCE=YES request is ignored.

Note: Summary dumps (SUMDUMP) for branch entries (BRANCH=YES) always cause the system to be set non-dispatchable until the summary dump is written.

,BRANCH=NO

,BRANCH=YES

specifies that a branch entry is to be used for interfacing with SVC DUMP to schedule a dump (YES), or that an SVC 51 instruction is to be generated for interfacing with SVC DUMP. This parameter can only be used by key 0, supervisor state routines that are in SRB mode, locked, or disabled to schedule a dump.

Routines that issue SDUMP with BRANCH=YES must also issue the CVT mapping macro instruction with the PREFIX=YES parameter.

,SUMLIST = list addr

specifies a list of starting and ending addresses of areas to be included in a summary dump. (SUMDUMP must be specified as an SDATA parameter and each starting address must be less than its corresponding ending address.)

The storage list must contain an even number of addresses, and each address must occupy one fullword. In the list, the high order bit of the fullword containing the last ending address of the list must be set to 1, and all other high order bits must be set to 0.

When BRANCH=YES is also specified, each address in the list must specify a paged-in area. Each specified area must be addressable with the page and segment tables of the current address space.

If the ASID or ASIDLST parameter was not specified, register 15 contains one of the following return codes when control is returned:

Hexadecimal Code	Meaning
00	A complete dump was taken.
04	A partial dump was taken because the dump data set did not have sufficient space.
08	The system was unable to take a dump.

If the ASID or ASIDLST parameter was specified, register 15 contains one of the following return codes when control is returned:

Hexadecimal Code	Meaning
00	A dump was scheduled. If an ECB was supplied, it will be posted on completion of the dump.
08	The system was unable to schedule a dump.

If an ECB was supplied, one of the following codes is returned in the ECB:

Hexadecimal Code	Meaning
00	A complete dump was taken.
04	A partial dump was taken.
08	The system was unable to take a dump.

Example 1

Operation: This example shows how SVC DUMP may be branch entered to initiate a dump in an address space by callers who cannot issue an SVC. Areas to be dumped are requested via three parameters (BUFFER, SDATA, and STORAGE). The dump has the title indicated in the HDR parameter, and the caller requests to be notified of the completion of the scheduled dump via the ECB parameter.

```
SDUMP HDR='USER DATA FOR TEST A',DCB=TESTADCB,BUFFER=YES,
ASID=TSTAASID,ECB=( 8 ),QUIESCE=YES,BRANCH=YES,
STORAGE=( A, B, C, D, ( 9 ), E ),SDATA=( ALLPSA, PSA, NUC, SQA, LSQA,
RGN, LPA, SWA, CSA )
```

Example 2

Operation: This example shows how SVC DUMP can be invoked via a branch entry to initiate a dump of several address spaces by callers who cannot issue an SVC. Areas to be dumped are requested via four parameters (BUFFER, SDATA, LIST, and SUMLIST). The address spaces to be dumped are described by the ASIDLST parameter. Note that areas specified by SUMLIST only apply to the current address space. The LIST addressed by the LIST keyword must be addressable from any address space. The dump has the title indicated in the HDR parameter, and the caller requests to be notified of the completion of the scheduled dump via the ECB parameter.

```
SDUMP   HDR='USER DATA FOR TEST B',  
        BUFFER=YES,ASIDLST=TSTALIST,ECB=( 8 ),  
        QUIESCE=YES,BRANCH=YES,LIST=( 9 ),  
        SDATA=( ALLPSA,PSA,NUC,SQA,RGN,SUMDUMP,LPA,CSA ),  
        SUMLIST=TSTSLIST  
.  
.  
TSTALIST DC X'0000000A800B'  
TSTSLIST DC X'0000000080400000'
```

SDUMP (List Form)

Use the list form of the SDUMP macro instruction to construct a control program parameter list. You can specify any number of storage addresses using the STORAGE parameter. Therefore, the number of starting and ending address pairs in the list form of SDUMP must be equal to the maximum number of addresses specified in the execute form of the macro instruction.

The list form of the SDUMP macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede SDUMP.
SDUMP	
␣	One or more blanks must follow SDUMP.

<i>HDR='dump title'</i>	<i>dump title</i> : from 1 to 100 characters.
<i>,HDRAD=dump title addr</i>	<i>dump title addr</i> : A-type address.
<i>,DCB=dcb addr</i>	<i>dcb addr</i> : A-type address.
<i>,SDATA=(data code)</i>	<i>data code</i> : any combination of the following, separated by commas: PSA CSA NUC SWA LSQA ALLPSA or NOALLPSA (abbreviated NOALL) RGN SQA or NOSQA LPA SUMDUMP (abbreviated SUM) or TRT NOSUMDUMP (abbreviated NOSUM) Default: SDATA=(ALLPSA,SQA,SUMDUMP) ALLPSA, SQA, and SUMDUMP are the defaults even if other parameters are specified for SDATA. NOALLPSA, NOSQA, and NOSUMDUMP must be specified to suppress these defaults.
<i>,STORAGE=(strt addr,end addr)</i>	<i>strt addr</i> : A-type address.
<i>,LIST=list addr</i>	<i>end addr</i> : A-type address. Note: One or more pairs of addresses may be specified, separated by commas. For example: <i>,STORAGE=(strt addr,end addr,strt addr,end addr)</i>
<i>,BUFFER=NO</i>	Default: BUFFER=NO
<i>,BUFFER=YES</i>	
<i>,QUIESCE=YES</i>	Default: QUIESCE=YES
<i>,QUIESCE=NO</i>	
<i>,MF=L</i>	

The parameters are explained under the standard form of the SDUMP macro instruction, with the following exception:

,MF=L

specifies the list form of the SDUMP macro instruction.

The list form expansion creates a 40-byte parameter list.

SDUMP (Execute Form)

A remote control program parameter list is referred to and can be modified by the execute form of the **SDUMP** macro instruction.

If you code one or more of the **SDATA** parameters on the execute form of the macro instruction, any **SDATA** parameters coded on the list form are lost.

The execute form of the **SDUMP** macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede SDUMP.
SDUMP	One or more blanks must follow SDUMP.
␣	
HDR= <i>'dump title'</i>	<i>dump title</i> : from 1 to 100 characters.
HDRAD= <i>dump title addr</i>	<i>dump title addr</i> : RX-type address, or register (2) - (12).
,DCB= <i>dcb addr</i>	<i>dcb addr</i> : RX-type address, or register (2) - (12).
,ASID= <i>ASID addr</i>	<i>ASID addr</i> : RX-type address, or register (2) - (12).
,ASIDLST= <i>list addr</i>	<i>list addr</i> : RX-type address, or register (2) - (12).
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : RX-type address, or register (2) - (12).
,SDATA=(<i>data code</i>)	<i>data code</i> : any combination of the following, separated by commas: PSA CSA NUC SWA LSQA ALLPSA or NOALLPSA (abbreviated NOALL) RGN SQA or NOSQA LPA SUMDUMP (abbreviated SUM) or TRT NOSUMDUMP (abbreviated NOSUM) Default: SDATA=(ALLPSA,SQA,SUMDUMP) ALLPSA, SQA, and SUMDUMP are the defaults even if other parameters are specified for SDATA. NOALLPSA, NOSQA, and NOSUMDUMP must be specified to suppress these defaults.
,STORAGE=(<i>strt addr,end addr</i>)	<i>strt addr</i> : RX-type address, or register (2) - (12).
,LIST= <i>list addr</i>	<i>end addr</i> : RX-type address, or register (2) - (12). <i>list addr</i> : RX-type address, or register (2) - (12). Note: One or more pairs of addresses may be specified, separated by commas. For example: ,STORAGE=(<i>strt addr,end addr,strt addr,end addr</i>)
,BUFFER=NO	
,BUFFER=YES	
,QUIESCE=YES	
,QUIESCE=NO	
,BRANCH=NO	Note: If BRANCH=YES is specified, ASID or ASIDLST must also be specified.
,BRANCH=YES	
,SUMLIST= <i>list addr</i>	<i>list addr</i> : RX-type address, or register (2) - (12).
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the SDUMP macro instruction, with the following exception:

,MF=(E, *ctrl addr*)
specifies the execute form of the SDUMP macro instruction using a remote control program parameter list.

Example 1

Operation: The execute form is used to add (SDATA areas) and change (BUFFER and QUIESCE) options in the SDUMP parameter list. The list form of SDUMP was previously used to create the basic SDUMP parameter list located by register 1.

```
SDUMP  SDATA=( SQA, LPA ), BUFFER=NO, QUIESCE=NO, MF=( E, ( 1 ) )
```

SETFRR — Set Up Functional Recovery Routines

The SETFRR macro instruction provides control program functions with the ability to define their recovery in the FRR (functional recovery routine) LIFO stack, which is used during processing of the system recovery manager. Each branch-entered control program function can use SETFRR to define its own unique recovery environment.

The SETFRR macro instruction can be used to add, delete, or replace FRRs in the LIFO stack, or to purge all FRRs in the stack. The macro instruction also optionally returns to the user the address of a parameter area that is eventually passed to the FRR when an error occurs. The parameter area can be initialized with information such as tracking data that may be useful to the FRR.

SETFRR has no external linkages and expands directly inline. To issue a SETFRR the user must be in supervisor state key zero. If the user is not in SRB or physically disabled mode, the user must also hold a lock. These conditions must be maintained for the duration of the SETFRR environment. All SETFRR users must include the DSECTs for the FRR stack (via the IHAFRRS mapping macro instruction) and the PSA (via the IHAPSA mapping macro instruction) prior to using the SETFRR macro instruction. Note that it is necessary to copy IHAPSA from AMODGEN into MACLIB unless AMODGEN is concatenated to SYSLIB.

When the FRR is entered:

- Register 0 points to a 200-byte work area on a double word boundary for use by the FRR.
- Register 1 points to the SDWA, the first word of which (SDWAPARM) points to the same 24-byte area that was set up for the user at the time the SETFRR macro was issued if he coded the PARMAD parameter.
- Register 13 points to a variable recording area in the SDWA (SDWAVRA).
- Register 14 contains the return address.
- Register 15 points to the FRR entry point.

Note: FRRs need not restore registers upon return.

The SETFRR macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
<i>b</i>	One or more blanks must precede SETFRR.
SETFRR	
<i>b</i>	One or more blanks must follow SETFRR.

A,FRRAD= <i>FRR addr</i>	<i>FRR addr</i> : A-type address, or register (2) - (12).
R,FRRAD= <i>FRR addr</i>	
D	
P	
,WRKREGS=(<i>reg1,reg2</i>)	<i>reg1</i> : decimal digits 1-15. <i>reg2</i> : decimal digits 1-15.
,PARMAD= <i>parm area addr</i>	<i>parm area addr</i> : A-type address, or register (2) - (12). Note: This parameter may only be specified with A or R above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The explanation of the parameter is as follows:

A,FRRAD=*FRRAD addr*
R,FRRAD=*FRRAD addr*
D
P

specifies the operation to be performed on the FRR LIFO stack:

A — an FRR address is to be added to the stack.

R — the FRR address last added to the stack is to be replaced by another FRR address.

D — the FRR address last added to the stack is to be deleted.

P — all entries in the stack are to be purged.

FRRAD specifies the address of a fullword containing the FRR address that is to be added or replaced. The parameter specifies the FRR address in a register or specifies the address of a storage location containing the FRR address.

,WRKREGS=(*reg1,reg2*)

specifies two unique general purpose registers to be used as work registers in the code generated by the SETFRR macro expansion.

,PARMAD=*parm area addr*

specifies the address of a fullword to receive the address of the 24-byte parameter area provided by the system to the issuer of SETFRR. If a register is specified, the address of the 24-byte parameter area is placed in the register. This parameter area is associated with the FRR address that has either been added to or has replaced an FRR address on the stack. This parameter area is passed to the FRR when an error occurs.

,RELATED=*value*

specifies information used to self-document macro instructions by “relating” functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R,LV=4096,RELATED=(FREE1,'GET STORAGE')
FREE1   FREEMAIN   R,LV=4096,A=(1),RELATED=(GET1,'FREE STORAGE')
```

Example 1

Operation: Add an FRR to the FRR stack, and return the address of the parameter list to the issuer of the SETFRR. The FRR address contained in register (R1) is placed on the FRR stack in the next available FRR entry. On return, register (R2) contains the address of the parameter list associated with this FRR entry. Registers R3 and R4 are work registers used in the code generated by SETFRR in performing its operations.

```
SETFRR  A,FRRAD=(R1),PARMAD=(R2),WRKREGS=(R3,R4)
```

Example 2

Operation: Delete the last FRR added to the FRR stack.

```
SETFRR  D,WRKREGS=(1,6)
```

SETLOCK — Control Access to Serially Reusable Resources

The SETLOCK macro instruction is used to control access to serially reusable resources. Each kind of serially reusable resource is assigned a separate lock. To use SETLOCK, you must be executing in supervisor state with protection key zero. Also, SETLOCK users must include the DSECT for the PSA (via IHAPSA mapping macro) prior to using the SETLOCK macro instruction. Note that it is necessary to copy IHAPSA from AMODGEN into MACLIB unless AMODGEN is concatenated to SYSLIB.

SETLOCK can be used to:

- Obtain a specified lock or set of locks.
- Release a specified lock or set of locks.
- Test a specified lock or set of locks to determine if the lock is held on the requestor's processor.

There are two classes of locks: global and local. There are two types of locks: spin and suspend. The descriptions of these locks and the hierarchy structure in which these locks are arranged are described under locking in this publication.

Note: In MVS a locked routine is not allowed to issue an SVC, or invoke a routine that would issue an SVC on the locked routine's behalf.

Caution should be used if SETLOCK is invoked and register 11, 12, 13 or 14 is used as the program's base register. Some options of the SETLOCK macro cause branch instructions to be generated after setting registers 11-14 to the required values.

The OBTAIN option of SETLOCK macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede SETLOCK.
SETLOCK	
b	One or more blanks must follow SETLOCK.

OBTAIN

,TYPE=IOSCAT,ADDR=(11)
,TYPE=IOSUCB,ADDR=(11)
,TYPE=IOSLCH,ADDR=(11)
,TYPE=IOSYNCH,ADDR=(11)
,TYPE=ASM,ADDR=(11)
,TYPE=DISP
,TYPE=SALLOC
,TYPE=SRM
,TYPE=CMS
,TYPE=LOCAL

,MODE=COND
,MODE=UNCOND
,MODE=UNCOND,DISABLED

,REGS=SAVE
,REGS=USE

,RELATED=*value*

Note: DISABLED may not be specified if TYPE=CMS or TYPE=LOCAL is specified above.
Note: Registers 11-14 will be destroyed if this parameter is omitted.

value: any valid macro keyword specification.

The parameters are explained as follows:

OBTAIN

specifies that the lockword is to be obtained or locked on the caller's behalf.

,TYPE = IOSCAT,ADDR = (11)
,TYPE = IOSUCB,ADDR = (11)
,TYPE = IOSLCH,ADDR = (11)
,TYPE = IOSYNCH,ADDR = (11)
,TYPE = ASM,ADDR = (11)
,TYPE = DISP
,TYPE = SALLOC
,TYPE = SRM
,TYPE = CMS
,TYPE = LOCAL

specifies the type of lock that is to be obtained on the caller's behalf.

ADDR = (11) specifies that the address of the lockword indicated by the TYPE parameter has been loaded into register 11 prior to the SETLOCK request.

IOSCAT IOS channel availability table lock. It is a global spin lock used by IOS to serialize access and updates to the channel availability table.

IOSUCB IOS unit control block lock. These locks (one per UCB) are global spin locks used to serialize access and updates to UCBs.

IOSLCH IOS logical channel queue lock. These locks (one per channel queue) are global spin locks used to serialize access and updates to the IOS logical channel queues.

IOSYNCH IOS synchronization lock. It is a global spin lock used to serialize the global IOS functions.

ASM Auxiliary storage manager lock. It is a global spin lock used to serialize use of the global ASM control blocks.

DISP Global dispatcher lock. It is a global spin lock used to serialize all functions associated with the dispatching process.

SALLOC Real storage manager and virtual storage manager space allocation lock. It is a global spin lock used to serialize the global functions of RSM and VSM.

SRM Systems resource manager lock. It is a global spin lock used to serialize use of the SRM control algorithms and associated data.

CMS Cross memory services lock. It is a global suspend lock used to serialize on more than one virtual storage where this serialization is not provided by one or more of the global locks.

LOCAL Storage lock of the storage the SETLOCK caller is executing. It is a local suspend lock used by supervisor functions which require serialization within that particular storage only.

,MODE = COND

,MODE = UNCOND

,MODE = UNCOND,DISABLED

specifies whether the lock is to be conditionally or unconditionally obtained.

COND specifies that the lock is to be conditionally obtained. That is, if the lock is not owned on another processor, it is acquired on the caller's behalf. If the lock is already held, control is returned to the caller indicating that the lock is held and that either the caller already owns the lock or that another processor or storage owns the lock.

UNCOND specifies that the lock is to be unconditionally obtained. That is, if the lock is not owned on another processor, it is acquired on the caller's behalf. If the lock is already held by the caller, control is returned to the caller indicating that he already owns the lock. If the lock is held on another processor, the caller's processor spins on the lock until it is released or suspends the SETLOCK caller until the lock is released.

DISABLED specifies that the caller is already in a physically disabled state.

,REGS = SAVE

,REGS = USE

specifies the use of registers 11 through 14.

SAVE specifies that register contents are to be saved. Registers 11 through 14 are saved in the area pointed to by register 13, and are restored upon completion of the SETLOCK request. The savearea consists of at least 5 words. Register 15 contains the return code.

Note: The save area used for the REGS=SAVE parameter must be a different area than the standard linkage save area used by the program.

USE specifies that registers 14, 15, 0, and 1 are available for use. Registers 11, 12, and 13 are saved in registers 15, 0, and 1, respectively, and are restored upon completion of the SETLOCK request. Register 14 is used as a link register; register 15 contains the return code.

Note: If neither SAVE nor USE is specified, registers 11-14 are destroyed and register 13 contains the return code.

,RELATED = value

specifies information used to self-document macro instructions by "relating" functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```

GETLL  SETLOCK  OBTAIN,TYPE=LOCAL,MODE=UNCOND,
          RELATED=(FREELL,'free local lock')
FREELL SETLOCK  RELEASE,TYPE=LOCAL,
          RELATED=(GETLL,'get local lock')

```

When control is returned, register 15 (register 13, if neither SAVE nor USE is specified) contains one of the following return codes:

Hexadecimal Code	Meaning
00	The lock was successfully obtained. The lock was free and is now held on the caller's behalf.
04	The lock was already held by the caller. The lockword id matches the caller's id.
08	The obtain process was unsuccessful. The lockword id does not match the caller's id.

Example 1

Operation: The global dispatcher lock DISP is to be conditionally requested. The RELATED parameter indicates that the DISP lock serializes the TCB resource, and the lock is either freed at the location represented by NAME or SYM1 in module IEFVHA or by SYM2 in module IEFVFA.

```

SETLOCK OBTAIN,TYPE=DISP,MODE=COND,RELATED=(TCB,IEFVHA(NAME,
          SYM1),IEFVFA(SYM2))

```

The RELEASE option of the SETLOCK macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede SETLOCK.
SETLOCK	
b	One or more blanks must follow SETLOCK.

RELEASE	
,TYPE=IOSCAT,ADDR=(11)	<i>reg</i> : decimal digit 2 - 10.
,TYPE=IOSUCB,ADDR=(11)	
,TYPE=IOSLCH,ADDR=(11)	
,TYPE=IOSYNCH,ADDR=(11)	
,TYPE=ASM,ADDR=(11)	
,TYPE=DISP	
,TYPE=SALLOC	
,TYPE=SRM	
,TYPE=CMS	
,TYPE=LOCAL	
,TYPE=SPIN	
,TYPE=ALL	
,TYPE=(<i>reg</i>)	
,DISABLED	Note: DISABLED may not be specified if TYPE=CMS or TYPE=LOCAL is specified above.
,REGS=SAVE	
,REGS=USE	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained under the OBTAIN option of the SETLOCK macro instruction, with the following exceptions:

RELEASE

specifies that the lockword is to be released.

,TYPE=SPIN

,TYPE=ALL

,TYPE=(reg)

specifies the type of lock that is to be released.

SPIN All spin locks currently held on the processor are to be released.

ALL All locks currently held on the processor are to be released.

(reg) The specified register contains a bit string identifying the locks to be released. A value of 1 indicates that the lock held is to be released; a value of 0 indicates that the status of the lock will not change. The bit meanings are:

- Bit 19 DISP
- Bit 20 ASM
- Bit 21 SALLOC
- Bit 22 IOSYNCH
- Bit 23 IOSCAT
- Bit 24 IOSUCB
- Bit 25 IOSLCH
- Bit 26 Reserved
- Bit 27 Reserved
- Bit 28 Reserved
- Bit 29 SRM
- Bit 30 CMS
- Bit 31 LOCAL

,DISABLED

specifies that control is to be returned to the caller with the processor in a physically disabled state (except for machine check) when a lock is successfully released. This form should be used only by those routines which do not have the disabled supervisor indicator on when they are executing and which, upon release of a global spin lock, must remain physically disabled due to noninterruptibility or no recursion restraints.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal

Code	Meaning
00	The lock was successfully released.
04	The lock was not owned. The lock was free when the release request was issued.
08	The release process was unsuccessful. The lockword id does not match the caller's id.

Note: No return codes are supported for multiple releases. That is, return code register contents are unpredictable.

Example 1

Operation: The local lock is requested to be released.

```
SETLOCK RELEASE,TYPE=LOCAL,RELATED=(TCBRQ,MOD1(NAME1),  
MOD2(NAME2))
```

Example 2

Operation: The IOSUCB lock whose address is in register 11 is requested to be released.

```
SETLOCK RELEASE,TYPE=IOSUCB,ADDR=(11),RELATED=(XYZ,MOD1(LABEL))
```

The TEST option of the SETLOCK macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
‡	One or more blanks must precede SETLOCK.
SETLOCK	
‡	One or more blanks must follow SETLOCK.

TEST

,TYPE=IOSCAT	<i>reg</i> : decimal digit 2 - 12
,TYPE=IOSUCB	
,TYPE=IOSLCH	
,TYPE=IOSYNCH	
,TYPE=ASM	
,TYPE=DISP	
,TYPE=SALLOC	
,TYPE=SRM	
,TYPE=CMS	
,TYPE=LOCAL	
,TYPE=SPIN	
,TYPE=ALL	
,TYPE=(<i>reg</i>)	
,ADDR=(<i>reg</i>)	<i>reg</i> : decimal digit 2 - 12
	Note: ADDR may not be specified if any of the following was specified above:
	TYPE=DISP TYPE=LOCAL
	TYPE=SALLOC TYPE=SPIN
	TYPE=SRM TYPE=ALL
	TYPE=CMS TYPE=(<i>reg</i>)
,BRANCH=(HELD, <i>addr</i>)	<i>addr</i> : RX-type address.
,BRANCH=(NOTHELD, <i>addr</i>)	
,REGS=(<i>reg</i>)	<i>reg</i> : decimal digit 2 - 12.
	Note: REGS may only be specified if any of the following was specified above:
	TYPE=SPIN TYPE=ALL TYPE=(<i>reg</i>)
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained under the OBTAIN or RELEASE option of the SETLOCK macro instruction, with the following exceptions:

TEST

specifies that the lockword is to be checked to determine if it is currently held on the requesting processor.

,BRANCH = (HELD,*addr*)

,BRANCH = (NOTHELD,*addr*)

specifies that the return code setting output of the macro instruction is to be suppressed and replaced by a direct branch to the specified address.

If HELD is specified, the address is branched to; if the specified lock, or at least one lock for TYPE=ALL or TYPE=SPIN, or all the specified locks for TYPE=(reg) are held on the requesting processor.

If NOTHELD is specified, the address is branched to; if the specified lock is not currently held on the requesting processor, or if not all the locks specified for TYPE=(reg) are held, or if no lock for TYPE=ALL or TYPE=SPIN is held.

,REGS=(reg)

specifies the register containing a bit string identifying which locks are held. If the bit string is partially correct (that is, one of the locks specified is not held), the connected string is returned in the register specified.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	The lock was held by the requestor, or all the locks were held (if the request was for several locks via a register), or at least one lock was held (if TYPE=ALL or TYPE=SPIN was specified).
04	The lock was not held by anybody, or not all the locks were held (if the request was for several locks via a register), or no lock was held (if TYPE=ALL or TYPE=SPIN was specified).

Example 1

Operation: If the local lock is not held, a branch to DSROLLINT is to be performed; otherwise, the next sequential instruction is to be executed.

```
SETLOCK TEST, TYPE=LOCAL, BRANCH=( NOTHELD, DSROLLINT )
```

SETRP — Set Return Parameters

The SETRP macro instruction is used to indicate the various requests that a recovery exit can return.

The macro instruction is valid only for exits established via functional recovery exits and ESTAE/ESTAI/ESTAR exits. The table following the description of the macro instruction indicates which parameters are valid for each situation.

The SETRP macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the **RECORD**, **RECPARM**, **FRELOCK** and **CPU** parameters. These parameters are restricted in use and should be used only by programs executing in supervisor state or under protection key 0-7 and executing as a functional recovery routine.

The syntax of the complete SETRP macro instruction is shown below. However, only the explanation of the restricted parameters is presented. Explanation of the other parameters can be found in **OS/VS2 Supervisor Services and Macro Instructions**.

Note: This macro instruction requires that the **IHSDWA** mapping macro be assembled as a **DSECT** into the caller's program.

The SETRP macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede SETRP.
SETRP	
b	One or more blanks must follow SETRP.
WKAREA=(<i>reg</i>)	<i>reg</i> : decimal digits 1-12. Default: WKAREA=(1)
,REGS=(<i>reg1</i>)	<i>reg1</i> : decimal digits 0-12, 14, 15.
,REGS=(<i>reg1,reg2</i>)	<i>reg2</i> : decimal digits 0-12, 14, 15. Note: If <i>reg1</i> and <i>reg2</i> are both specified, order is 14, 15, 0-12.
,DUMP=IGNORE	Default: DUMP=IGNORE
,DUMP=YES	
,DUMP=NO	
,DUMPOPT=<i>parm list addr</i>	<i>parm list addr</i> : RX-type address, or register (2) - (12). Note: This parameter may be specified only if DUMP=YES is specified above.
,RC=0	Default: RC=0
,RC=4	
,RC=16	
,RETADDR=<i>retry addr</i>	<i>retry addr</i> : RX-type address, or register (2) - (12). Note: This parameter may be specified only if RC=4 is specified above.
,RETREGS=NO	<i>info addr</i> : RX-type address, or register (2) - (12). Default: RETREGS=NO
,RETREGS=YES	Note: This parameter may be specified only if RC=4 is specified above. If RETREGS=YES is specified for a FRR, all registers are restored from SDWASRSV with the exception of register 15. Register 15 always contains the entry point of the retry routine.
,RETREGS=YES,RUB=<i>info addr</i>	
,FRESDDWA=NO	Default: FRESDDWA=NO
,FRESDDWA=YES	Note: This parameter may be specified only if RC=4 is specified above.
,COMPCOD=<i>code</i>	<i>code</i> : symbol, decimal digit, or register (2) - (12). Default: COMPCOD=(<i>code</i> ,USER)
,COMPCOD=(<i>code</i>,USER)	
,COMPCOD=(<i>code</i>,SYSTEM)	
,FRELOCK=(<i>locks</i>)	<i>locks</i> : any combination of the following, separated by commas: DISP IOSCAT(<i>lockword</i>) SRM IOSUCB(<i>lockword</i>) SALLOC IOSLCH(<i>lockword</i>) CMS IOSYNCH(<i>lockword</i>) LOCAL ASM(<i>lockword</i>) <i>lockword</i> : RX-type address. Note: This parameter may be specified only if RC=0 is specified above.
,CPU=<i>reg</i>	<i>reg</i> : decimal digits 2-12.
,RECORD=IGNORE	Default: RECORD=IGNORE
,RECORD=YES	
,RECORD=NO	
,RECPARM=<i>record list addr</i>	<i>record list addr</i> : RX-type address, or register (2) - (12). Note: This parameter may be specified only if RECORD=IGNORE or RECORD=YES is specified above.

The parameters restricted in use are explained as follows. (The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.)

,FRELOCK = (locks)

specifies the locks to be freed and the corresponding lockwords:

DISP — Global dispatcher lock.

SRM — Systems resource manager lock.

SALLOC — Real storage manager and virtual storage manager space allocation lock.

CMS — Cross memory services lock.

LOCAL — Storage lock of the storage the caller is executing in.

IOSCAT — IOS channel availability table lock.

IOSUCB — IOS unit control block lock.

IOSLCH — IOS logical channel queue lock.

IOSYNCH — IOS synchronization lock.

ASM — Auxiliary storage manager lock.

,CPU = (reg)

specifies the register that contains the logical processor identification of the processor holding the resource that this processor is waiting for.

,RECORD = IGNORE

,RECORD = YES

,RECORD = NO

specifies that the entire SDWA, both fixed and variable areas, is to be written on SYS1.LOGREC (YES), is not to be written on SYS1.LOGREC (NO), or is to be written as indicated prior to the SETRP macro instruction (IGNORE).

,RECPARM = record list addr

specifies the address of a user-supplied record parameter list. The parameter list consists of three 8-byte fields:

- The first field contains the module name (microfiche name).
- The second field contains the CSECT name.
- The third field contains the FRR identification.

The three fields are left-justified, and padded with blanks.

Notes:

1. The variable information record, containing two 2-byte length fields at the beginning of the record consists of:
 - The first field, filled in by the system, specifies the total length available to the user (exclusive of the two length fields).
 - The second field, filled in by the user, defines the format of the data to be dumped in the first byte and the actual length of the data in the second byte.
2. The FRESDDWA parameter cannot be specified or defaulted for a functional recovery routine (FRR). The SDWA is always released before an FRR's retry routine gets control.

The following table indicates which parameters are available to functional recovery routines (FRRs) and which parameters are available to ESTAE/ESTAE exits.

Parameter	FRR	ESTAE
WKAREA	x	x
REGS	x	x
DUMP	x	x
DUMPOPT	x	x
RC=0	x	x
RC=4	x	x
RC=16		x
RETADDR	x	x
RETREGS	x	x
RUB	x	x
FRESDDWA		x
COMPCOD	x	x
FRELOCK	x	
CPU	x	
RECORD	x	x
RECPARM	x	x

Example 1

Operation: Request continue with termination and freeing of the IOSCAT and SRM locks. The IOSCAT lockword is label X.

```
SETRP RC=0, FRELOCK=( IOSCAT( X ), SRM )
```

Example 2

Operation: Cause a restart interruption on the processor identified by the contents of register 7. In this example, the interrupted function is spinning on a lock currently being held by the processor identified in register 7.

```
SETRP CPU=( 7 )
```

SPIE — Specify Program Interruption Exit

The SPIE macro instruction is described in the OS/VS2 Supervisor Services and Macro Instructions, with the exception of interruption type 17. This interruption type designates page faults and may be specified by an installation-authorized system programmer.

Each succeeding SPIE macro instruction completely overrides the previous SPIE macro instruction specification for that task. The specified exit routine is given 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 PIE is reset (ZEROED). Thus a SPIE exit routine should save any required PIE data prior to issuing a SPIE. 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.

For more information on the SPIE macro instruction, see the section on "Interruption, Recovery/Termination, and Dumping Services".

The syntax of the complete SPIE macro instruction is shown below.

The standard form of the SPIE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
 	One or more blanks must precede SPIE.
SPIE	
 	One or more blanks must follow SPIE.

<i>exit addr,(interrupts)</i>	<i>exit addr</i> : A-type address, or register (2) - (12). <i>interrupts</i> : decimal digits 1-15, or 17 expressed as single values: (2,3,4,7,8,9,10) ranges of values: ((2,4),(7,10)) combinations: ((2,4),6,8,(10,13),15)
-------------------------------	--

The parameters are explained as follows:

exit addr,(interrupts)

specifies the address of the exit routine to be given control when a program interruption of the type specified occurs. The interruption types are:

Number	Interruption Type
1	Operation
2	Privileged operation
3	Execute
4	Protection
5	Addressing
6	Specification
7	Data
8	Fixed-point overflow (maskable)
9	Fixed-point divide
10	Decimal overflow (maskable)
11	Decimal divide
12	Exponent overflow
13	Exponent underflow (maskable)
14	Significance (maskable)
15	Floating-point divide
17	Page fault

Note:

- If a specified program interruption type is maskable, the corresponding bit is set to 1. Interruption types not specified above are handled by the control program.
- The control program returns the address of the previous PICA in register 1. If no previous SPIE environment existed, the control program returns zeros in register 1.
- If an exit address is zero or no parameters are specified, the SPIE environment is cancelled.

Example 1

Operation: Give control to an exit routine for interruptions 17. DOITSPIE is the address of the SPIE exit routine.

```
SPIE    DOITSPIE,(17)
```

SPIE (List Form)

Use the list form of the SPIE macro instruction to construct a control program parameter list in the form of a program interruption control area.

The list form of the SPIE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
‡	One or more blanks must precede SPIE.
SPIE	
‡	One or more blanks must follow SPIE.

<i>exit addr</i>	<i>exit addr</i> : A-type address.
<i>, (interrupts)</i>	<i>interrupts</i> : decimal digits 1-15, or 17, expressed as single values: (2,3,4,7,8,9,10) ranges of values: ((2,4),(7,10)) combinations: ((2,4),6,8,(10,13),15)
<i>,MF=L</i>	

The parameters are explained under the standard form of the SPIE macro instruction, with the following exception:

,MF=L
specifies the list form of the SPIE macro instruction.

Note: If SPIE is coded with a 0 as the control address, the SPIE environment is cancelled.

SPOST — Synchronize POST

The SPOST macro instruction is used in a cross-memory post environment to ensure that all outstanding cross-memory post requests to the current address space have completed. SPOST resolves a synchronization problem that arises when it becomes necessary to free an ECB that is a potential target for a cross-memory post request. Before issuing SPOST, you must stop any new posts from being initiated.

For explanation of the parameters in a cross-memory post request, see the POST macro instruction.

SPOST invokes the PURGEDQ SVC. For details, see the PURGEDQ macro instruction.

The SPOST macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
<i>b</i>	One or more blanks must precede SPOST.

SPOST

Note: SPOST contains no optional or required parameters.

Example 1

Operation: Execute the SPOST macro instruction, with a comment.

```
SPOST    ,ISSUE SPOST
```

STAE — Specify Task Abnormal Exit

The STAE macro instruction enables the user to intercept a scheduled ABEND and to have control returned to him at a specified exit routine address. The STAE macro instruction operates in both problem program and supervisor modes.

Note: The STAE macro instruction is available for compatibility with release 1 of VS2 and with MFT and MVT. However, it is recommended that ESTAE be used.

The standard form of the STAE macro instruction is written as follows:

<i>name</i>	<i>name:</i> symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede STAE.
STAE	
␣	One or more blanks must follow STAE.

0	<i>exit addr:</i> A-type address, or register (2) - (12).
<i>exit addr</i>	
0	
,CT	Default: CT
,OV	
,PARAM= <i>list addr</i>	<i>list addr:</i> A-type address, or register (2) - (12).
,XCTL=NO	Default: XCTL=NO
,XCTL=YES	
,PURGE=QUIESCE	Default: PURGE=QUIESCE
,PURGE=HALT	
,PURGE=NONE	
,ASYNCH=NO	Default: ASYNCH=NO
,ASYNCH=YES	
,RELATED= <i>value</i>	<i>value:</i> any valid macro keyword specification.

The parameters are explained as follows:

0	
<i>exit addr</i>	specifies the address of a STAE exit routine to be entered if the task issuing this macro instruction terminates abnormally. If 0 is specified, the most recent STAE request is canceled.
,CT	
,OV	specifies the creation of a new STAE exit (CT) or indicates that the parameters passed in this STAE macro instruction are to overlay the data contained in the previous STAE exit (OV).
,PARAM= <i>list addr</i>	specifies the address of a user-defined parameter list containing data to be used by the STAE exit routine when it is scheduled for execution.

,XCTL=NO
,XCTL=YES

specifies that the STAE macro instruction will be canceled (NO) or will not be canceled (YES) if an XCTL macro instruction is issued by this program.

,PURGE=QUIESCE
,PURGE=HALT
,PURGE=NONE

specifies that all outstanding requests for I/O operations are not saved when the STAE exit is taken (HALT), that I/O processing is allowed to continue normally when the STAE exit is taken (NONE), or that all outstanding requests for I/O operations are saved when the STAE exit is taken (QUIESCE). For QUIESCE, at the end of the STAE exit routine, the user can code a retry routine to handle the outstanding I/O requests.

Note: If any IBM-supplied access method, except EXCP, is being used, the PURGE=NONE option is recommended. If you use PURGE=NONE, all control blocks affected by input/output processing can continue to change during STAE exit routine processing.

If PURGE=NONE is specified and the ABEND was originally scheduled because of an error in input/output processing, an ABEND recursion develops when an input/output interruption occurs, even if the exit routine is in progress. Thus, it appears that the exit routine failed when, in reality, input/output processing caused the failure.

ISAM Notes: If ISAM is being used and PURGE=HALT is specified or PURGE=QUIESCE is specified but I/O is not restored:

- Only the input/output event on which the purge is done is posted. Subsequent event control blocks (ECBs) are not posted.
- The ISAM check routine treats purged I/O as normal I/O.
- Part of the data set may be destroyed if the data set is being updated or added to when the failure occurred.

,ASYNCH=NO
,ASYNCH=YES

specifies that asynchronous exit processing is allowed (YES) or is not allowed (NO) while the STAE exit is executing.

ASYNCH=YES must be coded if:

- The STAE exit routine requests any supervisor services that require asynchronous interruptions to complete their normal processing.
- PURGE=QUIESCE is specified for any access method that requires asynchronous interruptions to complete normal input/output processing.
- PURGE=NONE is specified and the CHECK macro instruction is issued in the STAE exit routine for any access method that requires asynchronous interruptions to complete normal input/output processing.

Note: If ASYNCH=YES is specified and the ABEND was originally scheduled because of an error in asynchronous exit handling, an ABEND recursion develops when an asynchronous interruption occurs. Thus, it appears that the exit routine failed when, in reality, asynchronous exit handling caused the failure.

,RELATED = value

specifies information used to self-document macro instructions by “relating” functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R,LV=4096,RELATED=( FREE1, 'GET STORAGE' )
FREE1   FREEMAIN   R,LV=4096,A=( 1 ),RELATED=( GET1, 'FREE STORAGE' )
```

Control returns to the instruction following the STAE macro instruction; register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	Successful completion of STAE request.
04	STAE was unable to obtain storage for STAE request.
08	Attempt was made to cancel or overlay a nonexistent STAE request.
0C	Exit routine or parameter list address was invalid, or STAI request was missing a TCB address.
10	Attempt was made to cancel or overlay a STAE request of another user, or an unexpected error was encountered while processing this request.

Example 1

Operation: Request an overlay of the existing STAE recovery exit with the following options: new exit address is ADDR, parameter list is at PLIST, halt I/O, do not take asynchronous exits, transfer ownership to the new request block resulting from any XCTL macro instructions.

```
STAE    ADDR,OV,PARAM=PLIST,XCTL=YES,PURGE=HALT,ASYNCH=NO
```

STAE (List Form)

The list form of the STAE macro instruction is used to construct a remote control program parameter list.

The list form of the STAE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede STAE.
STAE	
␣	One or more blanks must follow STAE.

<i>exit addr</i>	<i>exit addr</i> : A-type address.
,PARAM= <i>list addr</i>	<i>list addr</i> : A-type address.
,PURGE=QUIESCE	Default: PURGE=QUIESCE
,PURGE=HALT	
,PURGE=NONE	
,ASYNCH=NO	Default: ASYNCH=NO
,ASYNCH=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=L	

The parameters are explained under the standard form of the STAE macro instruction, with the following exception:

,MF=L

specifies the list form of the STAE macro instruction.

STAE (Execute Form)

A remote control program parameter list is used in, and can be modified by, the execute form of the STAE macro instruction. The control program parameter list can be generated by the list form of the STAE macro instruction. If you want to dynamically change the contents of the remote STAE parameter list, he may do so by coding a new exit address and/or a new parameter list address. If exit address or PARM= is coded, only the associated field in the remote STAE parameter list is changed. The other field remains as it was before the current STAE request was made.

The execute form of the STAE macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede STAE.
STAE	
b	One or more blanks must follow STAE.

<i>exit addr</i>	<i>exit addr</i> : RX-type address, or register (2) - (12).
0	
,CT	
,OV	
,PARAM= <i>list addr</i>	<i>list addr</i> : RX-type address, or register (2) - (12).
,XCTL=NO	
,XCTL=YES	
,PURGE=QUIESCE	
,PURGE=HALT	
,PURGE=NONE	
,ASYNCH=NO	
,ASYNCH=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address, or register (1) or (2) - (12).

The parameters are explained under the standard form of the STAE macro instruction, with the following exception:

,MF = (E, *ctrl addr*)

specifies the execute form of the STAE macro instruction using a remote control program parameter list.

Example 1

Operation: Provide the pointer to the recovery code in the register called EXITPTR, and the address of the STAE exit parameter list in register 9. Register 8 points to the area where the STAE parameter list (created with the MF=L option) was moved.

```
STAE      (EXITPTR),PARAM=(9),MF=(E,(8))
```

STATUS — Change Subtask Status

The STATUS macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of those parameters that are restricted in use and available only to supervisor state, key zero callers. These restricted parameters allow the caller to manipulated the dispatchability of TCBs, SRBs, ASCBs, a STEP, or the SYSTEM.

The description of the STATUS macro instruction has been divided into two parts: the START/STOP option, and the SET/RESET option.

The START/STOP options of the STATUS macro instruction are written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede STATUS.
STATUS	
␣	One or more blanks must follow STATUS.

START	
STOP	
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : RX-type address, or register (2) - (12), or 0.
,SRB	<i>ASID addr</i> : RX-type address, or register (2) - (12).
,SRB,ASID= <i>ASID addr</i>	Note: ASID may only be specified with START.
,SYNCH	Note: SYNCH may only be specified with STOP.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are as follows:

START
STOP

specifies that the appropriate START/STOP count is to be adjusted and the dispatchability bits are to be set/reset.

,TCB=*tcb addr*
,SRB
,SRB,ASID=*ASID addr*
,SYNCH

specifies the status of the stop/start function:

TCB specifies the address of a fullword on a fullword boundary containing the address of the TCB that is to have its START/STOP count adjusted.

SRB specifies that the STOP function affects the dispatchability of system-level SRBs only; all other tasks in the address space area set/reset nondispatchable. For START, the ASID *addr* specifies the address of a halfword containing the address space identifier.

SYNCH specifies that the STOP function affects all the subtasks of the caller. If any of the subtasks are deferring STOPs when the request is issued, the caller is placed in a WAIT condition. The issuer is taken out of the wait state when all deferred stops are complete.

,RELATED=value

specifies information used to self-document macro instructions by “relating” functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter can be used, for example, as follows:

```
GET1    GETMAIN    R, LV=4096, RELATED=( FREE1, 'GET STORAGE' )
FREE1   FREEMAIN   R, LV=4096, A=( 1 ), RELATED=( GET1, 'FREE STORAGE' )
```

The SET/RESET options of the STATUS macro instruction are written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede STATUS.
STATUS	
b	One or more blanks must follow STATUS.

SET	
RESET	
,MC	Note : If MC or MC,STEP is specified, no other parameters can be specified.
,MC,STEP	
,SD	
,ND	
,SYSTEM	<i>mask</i> : for SD, any of decimal digits 1-32 (except 18), separated by commas; for ND, any of decimal digits 1-16 (except 14), separated by commas.
,STEP	
,STEP,(mask)	
,tcb addr,(mask)	<i>tcb addr</i> : RX-type address, or register (2) - (12).
.,(mask)	Default : STEP
,E	Note : This parameter can only be specified with <i>tcb addr,(mask)</i> .
,ASID=ASID addr	<i>ASID addr</i> : RX-type address, or register (2) - (12). Note : For SET, this parameter can only be specified with <i>tcb addr,(mask)</i> . For RESET, this parameter may <i>not</i> be specified with SYSTEM.
,RELATED=value	<i>value</i> : any valid macro keyword specification.

The parameters are explained as follows:

SET

RESET

specifies that the TCBs or ASCBs are to be set or reset nondispatchable.

,MC

,MC,STEP

,SD

,ND

specifies the nondispatchability status:

ND specifies that the primary nondispatchability bits are affected by this request.

SD specifies that the secondary nondispatchability bits are affected by this request.

MC and MC,STEP specifies that the initiator and all TCBs in the job step TCBs (except the issuer's TCB) are to be set/reset nondispatchable.

,SYSTEM

,STEP

,STEP,(*mask*)

,*tcb addr*,(*mask*)

„(*mask*)

specifies more information on the nondispatchability status:

SYSTEM specifies that all ASCBs are to be set/reset nondispatchable except for certain exempt ones (for examples, the master scheduler or the issuer).

STEP specifies that all job step TCBs (except the issuer's TCB) are to be set/reset nondispatchable.

tcb addr specifies that the specified TCB and all its subtasks are to be set/reset nondispatchable.

(*mask*) specifies the nondispatchability bits that are to be set/reset.

,E

specifies that only the specified TCB is to be set/reset nondispatchable.

,ASID=*ASID addr*

specifies the address of a halfword containing the address space identifier.

,RELATED=*value*

specifies information used to self-document macro instructions by "relating" functions or services to corresponding functions or services. The format and contents of the information specified are at the discretion of the user, and may be any valid coding values.

The RELATED parameter is available on macro instructions that provide opposite services (for example, ATTACH/DETACH, GETMAIN/FREEMAIN, and LOAD/DELETE, and on macro instructions that relate to previous occurrences of the same macro instructions (for example, CHAP and ESTAE).

The parameter may be used, for example, as follows:

```
GET1    GETMAIN    R,LV=4096,RELATED=(FREE1,'GET STORAGE')
FREE1   FREEMAIN   R,LV=4096,A=(1),RELATED=(GET1,'FREE STORAGE')
```

Example 1

Operation: Set primary nondispatchability bit 3 for the specified TCB and all its subtasks.

```
STATUS SET,ND,TCBADDR,(3)
```


SYNCH — Take a Synchronous Exit to a Processing Program

The SYNCH macro instruction makes it possible for a supervisor routine to take a synchronous exit to a processing program. The SYNCH routine initializes a PRB (program request block) and schedules execution of the requested program. After the processing program has been executed, the supervisor routine that issued the SYNCH macro instruction regains control.

The SYNCH macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
b	One or more blanks must precede SYNCH.
SYNCH	
b	One or more blanks must follow SYNCH.
<i>entry point addr</i>	<i>entry point addr</i> : RX-type address, or register (2) - (12) or (15).
,RESTORE=NO	Default: RESTORE=NO
,RESTORE=YES	
,KEYADDR= <i>addr</i>	<i>address</i> : RX-type address, or register (2) - (12)
,KEYADDR=NOKEYADDR	
,STATE=PROB	Default: STATE=PROB
,STATE=SUPV	
,MF=(E, <i>ctrl addr</i>)	<i>ctrl addr</i> : RX-type address or register (1), (2) - (12)
,MF=L	

The parameters are explained as follows:

entry point addr
 specifies the address of the entry point of the processing program to receive control.

,RESTORE=NO
 ,RESTORE=YES
 specifies whether registers 2-13 are to be restored when control is returned to the caller.

,KEYADDR=*addr*
 ,KEYADDR=NOKEYADDR
addr specifies the address of a one-byte area that contains the key in which the exit is to receive control. The key must be in bits 0-3; bits 4-7 must be zero. If KEYADDR=*addr* is not specified, the key in the TCB is used as the default.
 NOKEYADDR is used with the execute form of the macro to specify that the default should be used instead of the key in the parameter list previously defined by a list form of the macro.
 To use this keyword, the issuer must be in supervisor state, system key, or APF-authorized.

,STATE=PROB
 ,STATE=SUPV
 specifies the state in which the requested program receives control.
 To use this keyword, the issuer must be in supervisor state, system key, or APF-authorized.

,MF(E,ctrl addr)

specifies the execute form of the SYNCH macro instruction using a list generated by the list form of SYNCH.

,MF=L

specifies the list form of the SYNCH macro instruction.

Example 1

Operation: Take a synchronous exit to a processing program whose entry point address is specified in register 8.

```
SYNCH (8)
```

Example 2

Operation: Take a synchronous exit to a processing program labeled SUBRTN and restore registers 2-13 when control is returned.

```
SYNCH SUBRTN,RESTORE=YES
```

TCTL — Transfer Control from an SRB Process

The TCTL (transfer control) macro instruction allows an SRB process to exit from its processing and to pass control directly to a task without going through the dispatcher.

The TCTL macro instruction is coded as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede TCTL.
TCTL	
␣	One or more blanks must follow TCTL.

TCB=(4)	Default: TCB address contents of register (4)
TCB= <i>tcbaddr</i>	<i>tcbaddr</i> : A-type address or registers (2) - (12).

The parameters are explained as follows:

TCB=(4)

TCB=*tcbaddr*

specifies the task designated for dispatching. Register (4) is the default; it is assumed to contain the appropriate TCB address.

The TCTL macro instruction uses registers as follows:

Register	Use
0-3	Work registers
4	TCB address
5-14	Work registers
15	EPA of the TCTL routine

Example 1

Operation: From SRB mode processing, terminate the SRB and give control to the task specified in register 4.

```
TCTL    TCB=(4)
```

TESTAUTH — Test Authorization of Caller

The TESTAUTH macro instruction is used on behalf of a privileged or sensitive function to verify that its caller is appropriately authorized.

TESTAUTH supports the authorized program facility (APF) - a facility that permits the identification of programs that are authorized to use restricted functions. In addition, TESTAUTH provides the capability for testing for system key 0-7 and supervisor state.

The TESTAUTH macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede TESTAUTH.
TESTAUTH	
␣	One or more blanks must follow TESTAUTH.

FCTN= <i>fcn</i>	<i>fcn</i> : decimal digit 0 or 1 or register (2) - (12).
FCTN= <i>fcn</i> ,AUTH= <i>auth</i>	<i>auth</i> : decimal digit 0 or 1, or register (2) - (12).
	Default: FCTN=0
,STATE=NO	Default: STATE=NO
,STATE=YES	
,KEY=NO	Default: KEY=NO
,KEY=YES	
,RBLEVEL=2	Default: RBLEVEL=2
,RBLEVEL=1	
,BRANCH=NO	Default: BRANCH=NO
,BRANCH=YES	

The parameters are explained as follows:

FCTN = *fcn*

FCTN = *fcn*,AUTH = *auth*

specifies the authorization (via APF) of a program.

FCTN = 0 specifies that APF-authorization is not checked.

FCTN = 1 specifies that APF-authorization is checked.

AUTH = 0 specifies that the job step is not authorized to perform any restricted function.

AUTH = 1 specifies that the job step is authorized to perform restricted functions.

Note: If FCTN=1 is specified by itself (that is, without the AUTH parameter), the JSCB is used to check for authorization. AUTH should only be coded when it is not possible for TESTAUTH to acquire the code from the JSCB. In MVS, AUTH is generally not be coded except in a testing environment.

,STATE = NO

,STATE = YES

specifies whether or not (YES or NO) a check is to be made for supervisor/problem program state. (Supervisor state is authorized, problem program state is not authorized.)

,KEY = NO
,KEY = YES

specifies whether or not (YES or NO) a check is to be made of the protection keys.
(Protection keys 0-7 are authorized, protection keys 8-15 are not authorized.)

Note: TESTAUTH Is used to test one or more of three conditions FCTN, STATE, or KEY. If any of the requested conditions are tested favorably, a return code of 0 is returned in register 15. If all of the requested conditions are tested unfavorably, the return code is set to 4.

,RBLEVEL = 2
,RBLEVEL = 1

specifies whether the TESTAUTH caller is a type 2, 3, or 4 SVC (RBLEVEL=2), or a type 1 SVC (RBLEVEL=1).

,BRANCH = NO
,BRANCH = YES

specifies a branch entry (YES) or an SVC entry (NO). If BRANCH=YES is specified, registers 2 and 3 are modified by the TESTAUTH routine.

When control is returned, register 15 contains one of the following return codes:

Hexadecimal Code	Meaning
00	Task is authorized.
04	Task is not authorized.

Example 1

Operation: Test jobstep for APF authorization.

TESTAUTH FCTN=1

Example 2

Operation: Test for APF authorization and supervisor state, and do not check protection keys.

TESTAUTH STATE=YES, KEY=NO, FCTN=1

T6EXIT — Type 6 Exit

The T6EXIT macro instruction returns control from a Type 6 SVC routine to the SVC first level interrupt handler (FLIH). This exit macro can only be used in a Type 6 SVC.

The T6EXIT macro instruction is written as follows:

<i>name</i>	<i>name</i> : any valid assembler symbol
<i>b</i>	One or more blanks must precede T6EXIT.
T6EXIT	
<i>b</i>	One or more blanks must follow T6EXIT.

RETURN= <i>CALLER</i>	CALLER is the default.
RETURN=DISPATCH	
RETURN=SRB	

The explanation of the RETURN parameter is as follows:

RETURN =

specifies how the Type 6 SVC has chosen to exit.

CALLER specifies that the return is directly to the caller or issuer of the SVC without going through the dispatcher. The operand is the default return option.

DISPATCH specifies that the return should be through the dispatcher. This function is for the use of routines that have suspended the current task.

Note: No registers are returned to the caller.

SRB specifies that the SVC FLIH should immediately dispatch an SRB. This SRB must:

- be initialized by the Type 6 SVC
- be pointed to by register 1
- execute in the same address space as the SVC. The SRB has the same format as the SCHEDULE SRB.

Note: No registers are returned to the caller.

Example 1

Operation: Terminate Type 6 SVC processing and return control from the Type 6 SVC to the caller of the SVC.

```
T6EXIT      RETURN=CALLER
```

WTO — Write to Operator

The WTO macro instruction is described in the *OS/VS2 Supervisor Services and Macro Instructions*, with the exception of the DESC, AREAID, MSGTYP, and MCSFLAG parameters. The MSGTYP and MCSFLAG parameters should only be used by system programmers familiar with MCS, because using the parameters improperly could impede the entire message routing scheme. The AREAID parameter can only be used by authorized (supervisor state, protection key 0-7, or APF authorized) users.

The syntax of the complete WTO macro instruction is shown below. However, only the explanation of the DESC, AREAID, MSGTYP and MCSFLAG parameters are presented. Explanation of the other parameters can be found in *OS/VS2 Supervisor Services and Macro Instructions*. An authorized user can issue a multiple line WTO (MLWTO) of up to 255 lines with one WTO macro instruction. Additional lines can be added as discussed in "Writing a Multiple-Line Message" in "Miscellaneous Services".

The standard form of the WTO macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin name in column 1.
b	One or more blanks must precede WTO.
WTO	
b	One or more blanks must follow WTO.

<i>'msg'</i>	<i>msg</i> : Up to 124 characters.																		
<i>('text')</i>	<i>text</i> : Up to 124 characters.																		
<i>('text',line type)</i>	The permissible <i>line types</i> , <i>text</i>																		
lengths and maximum numbers are	shown below:																		
	<table border="0"> <thead> <tr> <th style="padding-right: 20px;">line type</th> <th style="padding-right: 20px;">text</th> <th>maximum number</th> </tr> </thead> <tbody> <tr> <td>C</td> <td>34 char</td> <td>1 C type</td> </tr> <tr> <td>L</td> <td>70 char</td> <td>2 L type</td> </tr> <tr> <td>D</td> <td>70 char</td> <td>10 D type</td> </tr> <tr> <td>DE</td> <td>70 char</td> <td>1 DE type</td> </tr> <tr> <td>E</td> <td>None</td> <td>1 E type</td> </tr> </tbody> </table>	line type	text	maximum number	C	34 char	1 C type	L	70 char	2 L type	D	70 char	10 D type	DE	70 char	1 DE type	E	None	1 E type
line type	text	maximum number																	
C	34 char	1 C type																	
L	70 char	2 L type																	
D	70 char	10 D type																	
DE	70 char	1 DE type																	
E	None	1 E type																	
	The maximum total number of line types that can be coded in one instruction is 255.																		
,ROUTCDE=(<i>route code</i>)	<i>route code</i> : decimal digit from 1 to 16. The <i>route code</i> is one or more codes, separated by commas.																		
,DESC=(<i>desc code</i>)	<i>desc code</i> : decimal digit from 1 to 16. The <i>desc code</i> is one or more codes, separated by commas.																		
,AREAID= <i>id char</i>	<i>id char</i> : alphabetic character A - Z.																		
,MSGTYP=(<i>msg type</i>)	<i>msg type</i> : any of the following N SESS,JOBNAMES Y SESS,STATUS SESS JOBNAMES,STATUS JOBNAMESESS,JOBNAMES,STATUS STATUS																		
,MCSFLAG=(<i>field name</i>)	<i>field name</i> : any combination of the following, separated by commas: REG0 HRDCPY RESP QREG0 REPLY NOTIME BRDCST NOCPY																		

The parameters restricted in use are explained as follows. (The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instruction*.)

,DESC=(desc code)

specifies the message descriptor code(s) to be assigned to the message. Descriptor codes 1 through 6 and descriptor code 11 are mutually exclusive. Codes 7 through 10 can be assigned in combination with any other code.

The descriptor codes are:

1 System failure	7	Application program processor
2 Immediate action required	8	Out-of-line message
3 Eventual action required	9	Operator request
4 System status	10	Dynamic status displays
5 Immediate command response	11	Critical eventual action required
6 Job status	12-16	Reserved for future use

Note: All WTO messages with descriptor codes of 1, 2, or 11 are action messages. An asterisk appears before the first character of an action message to indicate a need for operator action.

If both the ROUTCDE and DESC parameters are omitted and the message is not a MLWTO message, the routing code specified in the OLDWTOR parameter of the system generation CONSOLE macro instruction is assigned; and no descriptor code is assigned. If the OLDWTOR parameter is omitted, no descriptor or routing codes are assigned. Routing codes should be used with MLWTO messages. If DESC is specified with no ROUTCDE, a default routing code of zero is generated, causing a message to be queued to the master console by default.

,AREAD=id char

specifies a display area of the console screen on which a multiple-line message is to be written. This parameter is meaningful only for out-of-line (descriptor code 8 and 9) MLWTO messages that are to be sent to CRT consoles.

The character Z designates the message area (the screen's general message area, rather than a defined display area); it is assumed nothing is specified.

Note: You must be an authorized (supervisor state, key 0-7, or APF-authorized) user to use this parameter. Otherwise, Z is used. Also, if this parameter specifies an area, the area could be overlaid by a currently running dynamic display. Support for queuing messages with descriptor code 8 is by console id only.

,MSGTYP=(msg type)

specifies how the message is to be routed.

For SESS, JOBNames, or STATUS, the message is to be routed to the console and TSO terminal in operator mode that issued the MONITOR SESS, MONITOR JOBNames, or MONITOR STATUS command, respectively. When the message type is identified by the operating system, the message is routed only to those consoles that requested the information.

For Y or N, the message type specifies whether flags are to be set in the WTO macro expansion to describe what functions (MONITOR SESS, MONITOR JOBNames, and MONITOR STATUS) are desired. N, or omission of the MSGTYP parameter, indicates that the message is to be routed as specified in the ROUTCDE parameter.

,MCSFLAG=(field name)

specifies that the macro expansion should set bits in the MCSFLAG field as indicated by each name coded. The names and corresponding bit settings are shown in Figure 35.

Name	Bit	Meaning
—	0	Invalid entry.
REG0	1	Queue the message to the console whose source ID is passed in register 0.
RESP	2	The WTO is an immediate command response.
—	3	Invalid entry.
REPLY	4	The WTO macro instruction is a reply to a WTOR macro instruction.
BRDCST	5	Broadcast the message to all active consoles.
HRDCPY	6	Queue the message for hard copy only.
QREG0	7	Queue the message unconditionally to the console whose source ID is passed in Register 0.
NOTIME	8	Do not append time to the message.
—	9-12	Invalid entry.
NOCPY	13	If the WTO or WTOR macro instruction is issued by a program in the supervisor state, do not queue the message for hard copy. Otherwise, this parameter is ignored.
—	14-15	Invalid entry.

Note: Invalid specifications are ignored and produce an appropriate error message from the assembler.

Figure 35. MCSFLAG Fields

Example 1

Operation: Send a WTO message to the hardcopy log only.

```
WTO      'THIS MSG IS TO HARDCOPY ONLY WITH RC=ALL',MCSFLAG=HRDCPY,
        ROUTCDE=( 1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16 )
```

Example 2

Operation: Send a WTO message to all active consoles and broadcast it to all consoles or terminals which have issued MONITOR commands.

```
WTO      'THIS MSG IS BROADCAST WITH RC=ALL',MCSFLAG=BRDCST,
        ROUTCDE=( 1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16 )
```

Example 3

Operation: Send a WTO message to all consoles and TSO terminals which have issued a MNJOB NAMES command.

```
WTO      'WTO BY MSGTYP=JOB NAMES WITH RC=ALL,NO CONSOLE MONITORING
        JOB NAMES',MSGTYP=JOB NAMES,ROUTCDE=( 1,2,3,4,5,6,7,8,9,10,11,
        12,13,14,15,16 )
```

WTO (List Form)

The list form of the WTO macro instruction is described in the OS/VS2 Supervisor Services and Macro Instructions, with the exception of the DESC, AREAID, MSGTYP and MCSFLAG parameters. These parameters are restricted in use, and are described below.

The list form of the WTO macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede WTO.
WTO	
␣	One or more blanks must follow WTO.

<i>'msg'</i>	<i>msg</i> : up to 124 characters.																		
<i>('text')</i>	The permissible <i>line types</i> , <i>text</i> lengths and maximum numbers are shown below:																		
<i>('text',line type)</i>																			
	<table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">line type</th> <th style="text-align: left;">text</th> <th style="text-align: left;">maximum number</th> </tr> </thead> <tbody> <tr> <td>C</td> <td>34 char</td> <td>1 C type</td> </tr> <tr> <td>L</td> <td>70 char</td> <td>2 L type</td> </tr> <tr> <td>D</td> <td>70 char</td> <td>10 D type</td> </tr> <tr> <td>DE</td> <td>70 char</td> <td>1 DE type</td> </tr> <tr> <td>E</td> <td>None</td> <td>1 E type</td> </tr> </tbody> </table>	line type	text	maximum number	C	34 char	1 C type	L	70 char	2 L type	D	70 char	10 D type	DE	70 char	1 DE type	E	None	1 E type
line type	text	maximum number																	
C	34 char	1 C type																	
L	70 char	2 L type																	
D	70 char	10 D type																	
DE	70 char	1 DE type																	
E	None	1 E type																	
	The maximum total number of line types that can be coded in one instruction is 255.																		
<i>,ROUTCDE=(route code)</i>	<i>route code</i> : decimal digit from 1 to 16. The <i>route code</i> is one or more codes, separated by commas.																		
<i>,DESC=(desc code)</i>	<i>desc code</i> : decimal digit from 1 to 16. The <i>desc code</i> is one or more codes, separated by commas.																		
<i>,AREAID=id char</i>	<i>id char</i> : an alphabetic character A-Z.																		
<i>,MSGTYP=(msg type)</i>	<i>msg type</i> : any one of the following: N SESS,JOBNAMES Y SESS,STATUS SESS JOBNAMES,STATUS JOBNAMES SESS,JOBNAMES,STATUS STATUS																		
<i>,MCSFLAG=(field name)</i>	<i>field name</i> : any combination of the following, separated by commas: REG0 HRDCPY RESP QREG0 REPLY NOTIME BRDCST NOCPY																		
<i>,MF=L</i>																			

The restricted parameters are explained under the standard form of the WTO macro instruction. The other parameters are explained in OS/VS2 Supervisor Services and Macro Instructions.

WTOR — Write to Operator with Reply

The WTOR macro instruction is described in the *OS/VS2 Supervisor Services and Macro Instructions*, with the exception of the MSGTYP and MCSFLAG parameters. These parameters should be used only by system programmers familiar with MCS, since using the parameters improperly could impede the entire message routing scheme.

The syntax of the complete WTOR macro instruction is shown below. However, only the explanation of the MSGTYP and MCSFLAG parameters are presented. Explanation of the other parameters can be found in *OS/VS2 Supervisor Services and Macro Instructions*.

The standard form of the WTOR macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
 	One or more blanks must precede WTOR.
WTOR	
 	One or more blanks must follow WTOR.

<i>msg</i>	<i>msg</i> : up to 121 characters.
, <i>reply addr</i>	<i>reply addr</i> : A-type address, or register (2) - (12).
, <i>reply length</i>	<i>reply length</i> : symbol, decimal digit, or register (2) - (12). The minimum length is 1; the maximum length is 115 when the operator enters REPLY id, 'reply' and 119 when the operator enters R id, 'reply'.
, <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (2) - (12).
,ROUTCDE=(<i>route code</i>)	<i>route code</i> : decimal digit from 1 to 16. The <i>route code</i> is one or more codes, separated by commas.
,DESC=(<i>desc code</i>)	<i>desc code</i> : decimal digit from 1 to 16. The <i>desc code</i> is one or more codes, separated by commas.
,MSGTYP=(<i>msg type</i>)	<i>msg type</i> : any one of the following: N SESS,JOBNAMES Y SESS,STATUS SESS JOBNAMES,STATUS JOBNAMES SESS,JOBNAMES,STATUS STATUS
,MCSFLAG=(<i>field name</i>)	<i>field name</i> : any combination of the following, separated by commas: REG0 HRDCPY RESP QREG0 REPLY NOTIME BRDCST NOCPY

The parameters restricted in use are explained as follows. (The other parameters are explained in *OS/VS2 Supervisor Services and Macro Instructions*.)

,MSGTYP=(*msg type*)

specifies how the message is to be routed.

For SESS, JOBNames, or STATUS, the message is to be routed to the console or TSO terminal in operator mode which issued the MONITOR SESS, MONITOR JOBNames, or MONITOR STATUS command, respectively. When the message type is identified by the operating system, the message is routed only to those consoles that had requested the information.

For Y or N, the message type specifies whether flags are to be set in the WTO macro expansion to describe what functions (MONITOR SESS, MONITOR JOBNames, and MONITOR STATUS) are desired. N, or omission of the MSGTYP parameter, indicates that the message is to be routed as specified in the ROUTCDE parameter.

,MCSFLAG=(*field name*)

specifies that the macro expansion should set bits in the MCSFLAG field as indicated by each name coded. The names and corresponding bit settings are shown in Figure 30 that appears in the description of WTO.

Example 1

Operation: Send a WTOR message to the hardcopy log only.

```
WTOR      'THIS MSG IS TO HARDCOPY ONLY WITH RC=ALL',MCSFLAG=HRDCPY,  
          ROUTCDE=(1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16)
```

Example 2

Operation: Send a WTOR message to all active consoles and broadcast it to all consoles or terminals which have issued MONITOR commands.

```
WTOR      'THIS MSG IS BROADCAST WITH RC=ALL',MCSFLAG=BRDCST;  
          ROUTCDE=(1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16)
```

Example 3

Operation: Send a WTOR message to all consoles and TSO terminals which have issued a MN JOBNames command.

```
WTOR      'WTOR BY MSGTYP=JOBNames WITH RC=ALL, NO CONSOLE MONITORING  
          JOBNames',MSGTYP=JOBNames,ROUTCDE=(1,2,3,4,5,6,7,8,9,10,11,  
          12,13,14,15,16)
```

WTOR (List Form)

The list form of the WTOR macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the MSGTYP and MCSFLAG parameters. These parameters are restricted in use, and are described below.

The list form of the WTOR macro instruction is written as follows:

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
␣	One or more blanks must precede WTOR.
WTOR	
␣	One or more blanks must follow WTOR.

'msg'	<i>msg</i> : up to 121 characters.
,	<i>reply addr</i> : an A-type address.
, <i>reply addr</i>	
,	<i>reply length</i> : symbol or decimal digit. The minimum length is 1; the maximum length is 115 when the operator enters REPLY id, 'reply' and 119 when the operator enters R id, 'reply'.
, <i>reply length</i>	
,	<i>ecb addr</i> : A-type address.
, <i>ecb addr</i>	
,ROUTCDE=(<i>route code</i>)	<i>route code</i> : decimal digit from 1 to 16. The <i>route code</i> is one or more codes, separated by commas.
,DESC=(<i>desc code</i>)	<i>desc code</i> : decimal digit from 1 to 16. The <i>desc code</i> is one or more codes, separated by commas.
,MSGTYP=(<i>msg type</i>)	<i>msg type</i> : any one of the following: N SESS,JOBNAMES Y SESS,STATUS SESS JOBNAMES,STATUS JOBNAMES SESS,JOBNAMES,STATUS STATUS
,MCSFLAG=(<i>field name</i>)	<i>field name</i> : any combination of the following, separated by commas: REGO HRDCPY RESP QREG0 REPLY NOTIME BRDCST NOCPY
,MF=L	

The restricted parameters are explained under the standard form of the WTOR macro instruction. The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instructions**.

- A parameter
 - FREEMAIN macro instruction 165
 - GETMAIN macro instruction 169
 - PGFIX macro instruction 182
 - PGFREE macro instruction 185
 - SETFRR macro instruction 241
- A-type address, meaning of 124
- ABEND macro instruction 50
- ABTERM parameter 49
- AC= parameter 23
- ACEE parameter 195-215
- ACTINFO parameter 206
- ADDR parameter 244
- address, finding UCB 19-20
- address space termination 48
- address, subsystem screen table 107
- ADDSD command 201
- ADDVOL parameter 195
- allocate virtual storage (GETMAIN) 169-172
- ALLPSA parameter 232
- ALTDS command 201
- ALTER parameter 200-203
- alternate power down 87,88
- AMDPRDMP service aid 73,74
- AND IMMEDIATE (NI) instruction, providing a lock via 178-179
- answer area (EXTRACT) 6
- APF (authorized program facility) 23-27
- APF authorization code 23
- APFTABLE 24
- APPL parameter 200-215
- AREAID parameter 275
- ASCB parameter 187
- ASCBTCBS count 40
- ASID parameter
 - CALLRTM macro instruction 131
 - PURGEDQ macro instruction 191
 - SDUMP macro instruction 229
 - STATUS macro instruction 264-266
- ASIDLST parameter on the SDUMP macro instruction 231
- ASIDTCB parameter 191
- ASIS parameter 200-203
- ASM (auxiliary storage management) 50
- ASM lock 11
- ASM lock 15 (5740-XE1)
- asymmetric storage feature 88
- ASYNCH parameter
 - ATTACH macro instruction 127
 - ESTAE macro instruction 152
 - STAE macro instruction 260
 - used for serialization of resources 29
 - used for serialization of resources 11 (5740-XE1)
- asynchronous execution, scheduling system services for 228
- Asynchronous Exit Routines 97
- ASXB 206
- ATTACH macro instruction 126-129,51
 - authorization considerations 27
 - parameters ignored in MVS 129
 - use of restricted parameters 4
 - User modifications 4
- ATTR parameter 200,202-203
- AUTH parameter 271
 - authorization 23
 - authorization rules 27
 - authorization, testing 271-272
- authorized job step task
- authorized library
 - general information 23
 - SVC routines 100-103
- authorized program facility (APF) 23-27
- auxiliary storage management (ASM) 49
- auxiliary storage manager lock (ASM) 11
- auxiliary storage manager lock (ASM) 15 (5740-XE1)
- BLOCK parameter 194
- BNDRY parameter 169
- branch entering POST 35
- branch entering WAIT 36
- branch entry
 - CHANGKEY 81
 - EVENTS 31
 - GETMAIN and FREEMAIN 81
 - output, POST 126
 - POST 35
 - SDUMP 76
 - WAIT
 - description 126
 - register contents for 126
- BRANCH parameter
 - CHANGKEY macro instruction 81
 - CIRB macro instruction 136
 - ESTAE macro instruction 152
 - EVENTS macro instruction 157
 - FREEMAIN macro instruction 166,81
 - GETMAIN macro instruction 169,81
 - SDUMP macro instruction 233
 - SETLOCK macro instruction 248
 - TESTAUTH macro instruction 272
- BRANCH=YES option of CALLDISP 41
- BUFFER parameter 229
- build in-storage profiles for RACK 211
- bypassing the POST routine 31-32
- call recovery/termination manager (CALLRTM) 131-132
- CALLDISP macro instruction 131
- CALLRTM macro instruction 131-132
 - general use of 48-50
 - use by FORCE command 50
- CHANGE parameter 204-210
- change subtask status (STATUS) 264-266
- change system status (MODESET) 173-177
- CHANGKEY macro instruction
 - coding description 133-134
 - examples of 134
 - how to use 81
 - parameters of 133-134
- channel ends 108-110
- check RACF authorization macro 200
- checkpoint/restart considerations 9
- check stop status condition 30
- CHNGDUMP, use of 71
- CIB (command input buffer) 5-6
- CIB chain, manipulation of 5-6
- CIBCTR parameter 194
- CIRB macro instruction 135-137
 - branch entry interface 136
 - parameters ignored in MVS 136
 - use in Asynchronous Exit Routines 97-99
- CLASS parameter 195-215
- classes of locks 12

CLEAR parameter 228.1 (5740-XE1)
clock comparator value 78
CMS lock 78
CMS (cross memory services lock)
 general information 12
 SVC routines 102
CMS (general cross memory services lock) 14 (5740-XE1)
CMSEQDQ (ENQ/DEQ cross memory services lock) 14 (5740-XE1)
CMSSMF (SMF cross memory services lock) 14 (5740-XE1)
coding the macro instructions 123-125
command input buffer (CIB) 5-6
command input buffer manipulation (QEDIT) 194
common service area subpools 81
COMPARE AND SWAP (CS) instruction
 NIL macro instruction 178-179
 OIL macro instruction 180-181
 POST macro instruction 31-32,34
COMPAT parameter 228.1 (5740-XE1)
COMPCOD parameter
 CALLRTM macro instruction 131
 SETRP macro instruction 251
console switching 90
continuation lines 125
control access to serially reusable resources (SETLOCK) 243-249
CONTROL parameter 200-203
control program extensions 45
control track record 116
correlating diagnostic material 74-76
CPU parameter
 DSGNL macro instruction 144
 RISGNL macro instruction 224
 RPSGNL macro instruction 226
 SETRP macro instruction 251
CPU
 reset function 29
 RESUME macro instruction influence 40
CPUR parameter 143
create a new task (ATTACH) 3
create interruption request block (CIRB) 135-137
CREATE parameter 204-210
cross-memory post 31
cross memory services lock (CMS) 12
CS (COMPARE AND SWAP) instruction 178,180
CT parameter
 ESTAE macro instruction 152
 STAE macro instruction 259
CTRLPROG macro instruction 87
CURRENT parameter 267
CVT mapping macro 39
CVTQTD00 field 79
CVTTTPC field 76
CVTTQE00 field 78
CVTVWAIT 38
CVTOPT01 37
CVTOPT02 37

D parameter 241,242
DASD, shared 17
DATASET macro instruction 104
DATASET parameter 200-203
DCB parameter
 ATTACH macro instruction 127
 SDUMP macro instruction 230-231
DDR swaps 108
DE parameter 127
DEFINE parameter 196,197
decimal digit, meaning of 124
default, meaning of 124
define a resource to RACF macro 195
DELETE parameter 195
deleting exit routines 33
DEQ macro instruction 138-142,15
 and RESERVE 216
 execute form 142
 list form 141
 standard form 138-140
 use of restricted parameters 15-16
DESC parameter
 WTO macro instruction 274
 WTOR macro instruction 278
descriptor codes 92
device, reserving a 216-221
diagnostic material, correlation of 74-76
DIE 76-78
 (see also disabled interrupt exit)
direct class 29
direct signal, issuing 143-144
disabled interrupt exit (DIE)
 characteristics 76
 clock failure 79
 execution 76-77
 exit from 78
 interval cancellation 79
 RESUME macro instruction restriction 40
 task queue element control
 freeing 78
 obtaining 79
 serializing 78
DISABLED parameter 244-245
DISABLED parameter 228.1,236-240 (5740-XE1) 190.1
disable, enable low address protection (PROTPSA) (5740-XE1)
disabling 14
DISP (global dispatcher lock) 11
DISP (global dispatcher lock) 14 (5740-XE1)
DISPATCH parameter 273
dispatchability tests 41
DISP parameter 126-128
DPMOD parameter 127
DSGNL macro instruction
 general information 143-144
 use of 29
DSTYPE parameter 195
DUMP parameter
 CALLRTM macro instruction 131
 SETRP macro instruction 250
dump virtual storage (SDUMP) 229-239
DUMPOPT parameter
 CALLRTM macro instruction 132
 SETRP macro instruction 251
dumps of virtual storage 68,71-72
duplicate names, restrictions 24

E parameter
 ENQ macro instruction 146
 FREEMAIN macro instruction 165
 RESERVE macro instruction 216
 STATUS macro instruction 266
EA parameter
 PGFIX macro instruction 182
 PGFREE macro instruction 185
EC (external call) function 30
EC parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
ECB extension
 illustration 33-34
 in example 34
 initializing 33

ECB for POST exit 32
 ECB parameter
 ATTACH macro instruction 127
 ENQ macro instruction 147
 EVENTS macro instruction 157
 PGFIX macro instruction 182
 PGFREE macro instruction 185
 RESERVE macro instruction 216
 SDUMP macro instruction 230
 ECBIND parameter
 PGFIX macro instruction 184
 PGFREE macro instruction 186
 emergency signal (ES) function 30
 ENABLE parameter
 CIRB macro instruction 136
 MODESET macro instruction 177
 ENQ macro instruction 145-151
 execute form 150-151
 list form 149
 standard form 146-148
 use of restricted parameters 15-16
 ENQ/DEQ cross memory services lock (CMSEQDQ) 14
 (5740-XE1)
 ENTITY parameter 195
 ENTRIES parameter 157
 ENVIR parameter 204-215
 EP parameter
 ATTACH macro instruction 127
 CIRB macro instruction 135
 RISGNL macro instruction 224
 EPLOC parameter 127
 equipment check status condition 30
 ERRET parameter 187
 error identifiers 75
 errorid 74-75
 ES (emergency signal) function 30
 establishing a timer disabled interrupt exit 76-78
 (see also disabled interrupt exit)
 ESTAE environment 54
 ESTAE exit routine 51-53
 ESTAE extended capabilities 52-54
 ESTAE macro instruction 152-156
 execute form 155
 list form 154
 standard form 152-153
 use of 52-54,56-64
 ESTAE recovery routine 50
 ESTAE retry routine 58
 ESTAI exit routine 52-53
 ESTAI parameter 127
 ESTAI recovery routine 50
 ESTAI retry routine 58
 ETXR parameter 127
 EU parameter
 FREEMAIN macro instruction 168
 GETMAIN macro instruction 169
 event completion 31-32
 EVENTS 32
 PGFIX 84
 POST 31
 SPOST 31
 EVENTS macro instruction
 general information 157
 use of 32
 examples
 ATTACH 129
 CALLDISP 130
 CALLRTM 132
 CIRB 137
 DEQ 140
 DSGNL 144
 ENQ 148
 ESTAE 153,
 EXTRACT 159
 FREEMAIN 166
 GETMAIN 170
 INTSECT 172.1 (5740-XE1)
 MODESET 175
 NIL 179
 OIL 181
 PGFIX 184
 PGFIXA 184.1 (5740-XE1)
 PGFREE 186
 POST 188
 PROTPSA 190.1 (5740-XE1)
 PURGEDQ 191,193
 QEDIT 194
 RESERVE 219
 RISGNL 225
 RPSGNL 227
 SCHEDULE 228
 SDUMP 234,235
 SETFRR 242
 SETLOCK 245,247,249
 SETRP 253
 SPIE 255
 SPOST 243
 STAE 261,263
 STATUS 266
 SYNCH 269
 TESTAUTH 271
 WTO 274
 WTOR 279
 execute form of macro instruction, use of 123
 exit, SNAP/ABEND 72
 explicit serialization 14
 explicit serialization 11 (5740-XE1)
 extended ECB
 illustration 34
 initializing 133
 purpose 132
 extended STAE (ESTAE) 152-156
 external call (EC) function 30
 external call pending status condition 30
 EXTKEY parameter 173-174
 EXTRACT macro instruction 158-161
 execute form 161
 list form 160
 standard form 158-159
 extract TCB information (EXTRACT) 158-161

 fast extended STAE (fast ESTAE) 51-52
 (see also FESTAE)
 fast path to fix virtual storage contents (PGFIXA) 184.1
 (5740-XE1)
 fast path to free virtual storage contents (PGFREEA)
 186.1 (5740-XE1)
 FCTN parameter 271
 FESTAE (fast extended STAE)
 exit routines 51
 macro instruction
 coding description 162
 how to use 52
 RBFEPARM parameter area 52
 task recovery 60
 FIELDS parameter 158
 fix virtual storage contents (PGFIX) 182-184
 FLIH, role in subsystem SVC screening 105
 FORCE command 50
 force dispatcher entry (CALLDISP) 130
 frames, page 83
 free virtual storage (FREEMAIN) 168-171
 free virtual storage contents (PGFREE) 185-186

FREEMAIN macro instruction 164-168
 disabling 14
 execute form 168
 list form 167
 standard form 168-169
 use of restricted parameters 81
 FRELOCK parameter 252
 FRESWA parameter 251
 FRR parameter 228 (5740-XE1)
 FRR stack 50-51
 FRRAD parameter 234
 FRRs (functional recovery routines) 57,60,50
 functional recovery routines (FRRs)
 general information 57,60,50
 setting up 241

 general cross memory services lock (CMS) 14 (5740-XE1)
 GENERIC parameter 138
 GETMAIN macro instruction 169-172
 execute form 172
 list form 171
 standard form 169-170
 use of restricted parameters 81
 GIVEJPQ parameter 129
 GLBRANCH 81
 global branch entry point 81
 global dispatcher lock (DISP) 11
 global dispatcher lock (DISP) 14 (5740-XE1)
 global locks
 general information 12
 user-written SVC routines 100
 global priority 95
 global spin locks 77
 global storage subpools 81
 GROUP parameter 204,208-210
 GSPL parameter 127
 GSPV parameter 127
 GTFCRM parameter 226
 guidelines for using APF 27

 hardcopy log 90
 HDR parameter 230
 HDRAD parameter 230
 hierarchy, locking 13
 Hot I/O
 changing threshold values 110
 intercepting 110

 ICFBDF00 112
 ICFBIF00 114
 ICPUR parameter 143
 ICTL instruction 123
 identifying exit routines 33
 IEAAPF00 24
 IEALIMIT 117
 IEAVTRML 64
 IEAVVTPC mapping macro 76
 IEAVWAIT 38
 IEAOPT01 36
 IEAOPT02 36
 IEA911A system message 74
 IEAOPT0E entry point 35
 IEECVXIT 92
 IGFINTVL 108
 IHAABDPL mapping macro 73
 IHAECB mapping macro 33
 IHAFRRS mapping macro instruction 51,240
 IHAPSA mapping macro instruction 51,240
 IHARMPL mapping instruction 65

 IHASDWA mapping macro instruction 250
 IHASRB 106
 IHAWSAVT 166,173
 IMPL parameter 143
 initial processor reset function 29
 initial microprogram load function 29
 initial program reset function 29
 initializing extended ECBs and ECB extension 33
 installation-written clean-up routines 64-65
 INSTLN parameter 195-210
 integrity 41-45
 intercepting Hot I/O 110
 intercepting system errors 65
 (see also SLIP command)
 interface
 to ESTAE/ESTAI exit 56
 to functional recovery routines 54
 to STAE/STAI exit 55
 interlock conditions 18
 interprocessor communications 28-31
 interruption handler, missing 108
 interruption request block, creating 97
 interruption, recovery/termination, and dumping services
 47-79
 interruption request block (IRB) 135,97
 interrupts, Mass Storage System 108
 interrupts, synchronous 39
 intersect with the dispatcher (INTSECT) 29,175.0
 (5740-XE1)
 INTSECT macro instruction 14,172.1 (5740-XE1)
 invalid function status condition 31
 IOS channel availability table lock (IOSCAT) 11
 IOS channel availability table lock (IOSCAT) 14
 (5740-XE1)
 IOS logical channel queue lock (IOSLCH) 11
 IOS logical channel queue lock (IOSLCH) 14 (5740-XE1)
 IOS synchronization lock (IOSYNCH) 11
 IOS synchronization lock (IOSYNCH) 14 (5740-XE1)
 IOS unit control block lock (IOSUCB) 11
 IOS unit control block lock (IOSUCB) 14 (5740-XE1)
 IOSCAT (IOS channel availability table lock) 11
 IOSCAT (IOS channel availability table lock) 14
 (5740-XE1)
 IOSLCH (IOS logical channel queue lock) 11
 IOSLCH (IOS logical channel queue lock) 14 (5740-XE1)
 IOSUCB (IOS unit control block lock) 11
 IOSUCB (IOS unit control block lock) 14 (5740-XE1)
 IOSYNCH (IOS synchronization lock) 11
 IOSYNCH (IOS synchronization lock) 14 (5740-XE1)
 IPR parameter 143
 IRB (interruption request block) 97
 issue direct signal (DSGNL) 143-144
 issue remote immediate signal (RISGNL) 224-225
 issue remote pendable signal (RPSGNL) 226-227

 job step authorization 23
 JOBNAME parameter 206
 JSCB parameter 128
 JSCBAUTH bit 25
 JSTCB parameter 27,128

 KEY parameter
 ATTACH macro instruction 128
 CIRB macro instruction 135
 FREEMAIN macro instruction 166,81
 GETMAIN macro instruction 169
 MODESET macro instruction 175
 TESTAUTH macro instruction 271
 KEYADDR parameter 173

L parameter
 FREEMAIN macro instruction 165
 RACDEF macro instruction 197
 RACHECK macro instruction 202
LA parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
 label processing 9
LAST parameter 157
LC parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
library, authorized 23
 limiting user region size 117
 list form of macro instruction, use of 123
LIST parameter 233
LLOCK parameter 228 (5740-XE1)
LOCAL (local storage lock) 12
LOCAL (local storage lock) 14 (5740-XE1)
 GETMAIN and FREEMAIN 81
 SVC routines 100
local locks 12
local locks 13 (5740-XE1)
local priority 95
local storage subpools 81
local storage lock (LOCAL) 12
local storage lock (LOCAL) 14 (5740-XE1)
locking 11
locking 13 (5740-XE1)
 categories of locks 12 (5740-XE1)
 classes of locks 12
 classes of locks 27 (5740-XE1)
 considerations 12-15
 considerations 14 (5740-XE1)
 hierarchy 13
 hierarchy 13 (5740-XE1)
 locks 11-12
 locks 12-14 (5740-XE1)
 manager 11
 manager 12 (5740-XE1)
 types of locks 12-13
 types of locks 12 (5740-XE1)
LOG parameter 200-203
long-term fix 83
low address protection 119 (5740-XE1)
LPMOD parameter 127
LSQA for subsystem SVC screening 107
LSQA parameter 129
LU parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
LV parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169

machine check handler appendage, adding code to
 110-112
macro instruction forms 123
Mass Storage System missing interrupts 108
master scheduler initialization module
 adding code to 110
 general information 114-115
MC parameter 265
MCS (multiple console support) 90
MCSFLAG parameter
 WTO macro instruction 274
 WTOR macro instruction 277
MEMSWT parameter 226 (5740-XE1)
MEMTERM parameter 49
MENTITY parameter 196
message IEA911A 74

message routing 89
message routing exit routines 90-93
MF parameter
 DEQ macro instruction 141,142
 ENQ macro instruction 149,150
 ESTAE macro instruction 154,155
 EXTRACT macro instruction 160,161
 FREEMAIN macro instruction 167,168
 GETMAIN macro instruction 171,172
 MODESET macro instruction 176,177
 POST macro instruction 189,190
 PURGEDQ macro instruction 192,193
 RACF macro instructions 198,199,202,203,209,211
 RESERVE macro instruction 220,221
 SDUMP macro instruction 236
 SPIE macro instruction 256,257
 STAE macro instruction 262,263
 WTO macro instruction 277
 WTOR macro instruction 280
MFITCH parameter 226
miscellaneous services, supervisor 89-119
missing interruption handler 108
 choosing time intervals 109
 use of CSECT IGFNTVL 108
MODE parameter
 CIRB macro instruction 135
 MODESET macro instruction 175
 RPSGNL macro instruction 226
 SETLOCK macro instruction 236
MODESET macro instruction 173-177
 disabling 122
 execute form 177
 incompatible parameters 177
 list form 176
 standard form 173-175
 use of 28
MOUNT command 108
MP systems 87-88,9
MSGTYP parameter
 WTO macro instruction 275
 WTOR macro instruction 278
multiple console support (MCS) 90
multiple-line message 90
multiple locks 13
multiprocessing configuration 87-88
multiprocessing considerations 9
must complete function 15-16
MVOLSER parameter 196

N parameter 196-203
ND parameter 265
NEWNAME parameter 196-199
NEWPASS parameter 204-210
NI (AND IMMEDIATE) instruction 178
NIL macro instruction 178-179
NOALL parameter 232
NOALLPSA parameter 232
NOFAIL parameter 200-203
NONE parameter 200-203
non-preemptable SVCs
 description of 106
 restrictions for 106
nonquiesceable priority level 95
NOSQA parameter 232
NOSUM parameter 232
NOSUMDUMP parameter 232
not ready status condition 31
NSHSPL parameter 128
NSHSPV parameter 128

OBTAIN parameter 243
OI (OR IMMEDIATE) instruction 180
OIDCARD parameter 206
OIL macro instruction 180-181
OLDVOL parameter 197
operator
 writing to with reply 278
 writing to without reply 274
operator communication with a problem program 5
operator intervening status condition 30-31
operator messages 89-90
OR IMMEDIATE (OI) instruction, providing a lock via 180-181
ORIGIN parameter 194
OV parameter
 ESTAE macro instruction 152
 STAE macro instruction 259

P parameter 241
page fixing 83
page frames 83
PARALLEL parameter 224
PARAM parameter
 ATTACH macro instruction 127
 ESTAE macro instruction 152
 STAE macro instruction 259
PARM parameter 224
PARMAD parameter 242
partitioned data set, SVC routines 103
PASSCHK parameter 205
PASSWRD parameter 204-210
PER traps 65-70
PGFIX macro instruction 182-184
 incompatible parameters with MVS 184
 standard form 182-184
 use of 83-84
 virtual subarea list 84-85
PGFIXA macro instruction 184.1 (5740-XE1)
 considerations 84 (5740-XE1)
 restrictions 184.2 (5740-XE1)
 standard form 184.1 (5740-XE1)
PGFREE macro instruction 185-186
 incompatible parameters with MVS 186
 standard form 185-186
 use of 83
 virtual subarea list 84-85
PGFREEA macro instruction 186.1 (5740-XE1)
 considerations 84 (5740-XE1)
 restrictions 186.1 (5740-XE1)
 standard form 186.1 (5740-XE1)
PGLOAD macro instruction, virtual subarea list 84-85
PGMNAME parameter 206
PGOUT macro instruction, virtual subarea list 84-85
PGRLSE macro instruction, virtual subarea list 84-85
PICA 47
PIE 47
PIRL (see purged I/O request list)
POST macro instruction 187-190
 branch entry output 37
 branch entry to 35
 cross address space 34
 execute form 190
 exit routines, writing 32
 example of using exit function 34
 initializing extended ECBs and ECB extension 33
 interface with exit routines 34
 list form 189
 re-entry from POST exit 34
 scheduling SRB for 34
 standard form 187
 use of restricted parameters 31

POST, synchronizing 258
power warning feature support 110-112
PR parameter 143
PREVIOUS parameter 267
 priorities 95
 priority considerations 10
PROFILE parameter 200-203
program FLIH 47
program management 9
program reset function 29
PROTPSA macro instruction 190.1 (5740-XE1)
 restrictions 190.1 (5740-XE1)
 provide a lock via an **AND IMMEDIATE (NI)** instruction (NIL) 178-179
 provide a lock via an **OR IMMEDIATE (OI)** instruction (OIL) 180-181
PURGE parameter
 ATTACH macro instruction 127
 ESTAE macro instruction 152
 STAE macro instruction 260
purge SRB activity (PURGEDQ) 191-193
PURGEDQ macro instruction 191-193
 and **SPOST** 258
 execute form 193
 list form 192
 standard form 191
 use of 93-95
 purpose of extended ECB 32

QEDIT macro instruction
 general information 194
 uses of 5-6
QUIESCE parameter 233

R parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
 PGFIX macro instruction 182
 PGFREE macro instruction 185
 SETFRR macro instruction 241
RACDEF macro instruction
 coding descriptions
 execute form 199
 list form 198
 standard form 195
 how to use 46
RACF 44-46
 (see also resource access control facility)
RACHECK macro instruction
 coding descriptions
 execute form 203
 list form 202
 standard form 200
 how to use 46
RACINIT macro instruction
 coding descriptions
 execute form 210
 list form 208
 standard form 204
 how to use 46
RACLIST macro instruction
 coding descriptions
 execute form 215
 list form 214
 standard form 211
 how to use 46
RB parameter 222,267
RBFE Parm parameter area 53
RBLEVEL parameter 271
RBSCF suspend count field 39,40

RBWCF wait count field 38
 RC parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
 SETRP macro instruction 251
 READ parameter 200-203
 real storage manager and virtual storage manager space
 allocation lock (SALLOC) 11
 real storage manager and virtual storage manager space
 allocation lock (SALLOC) 14 (5740-XE1)
 real storage management 83-88
 real storage manager 83
 receiver check status condition 31
 reconfiguration using vary storage command 85
 RECORD parameter
 ESTAE macro instruction 152
 SETRP macro instruction 252
 recovery environment 31
 recovery guidelines 95
 recovery routines 50-64
 recovery/termination 48-64
 recovery/termination manager (RTM) 48
 recovery/termination manager, calling 131,132
 RECPARM parameter 250
 reenterable SVC routines 100
 reentrant modules 10
 REF parameter
 NIL macro instruction 178
 OIL macro instruction 180
 reference - macro instructions and commands 121-280
 refreshable SVC routines 100
 register (0), meaning of 124
 register (1), meaning of 124
 register (2) - (12), meaning of 124
 REGS parameter
 SETLOCK macro instruction 245,248
 SETRP macro instruction 251
 RELATED parameter
 ATTACH macro instruction 127
 DEQ macro instruction 138
 ENQ macro instruction 146
 ESTAE macro instruction 152
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
 MODESET macro instruction 174
 PGFIX macro instruction 182
 PGFREE macro instruction 185
 POST macro instruction 187
 RESERVE macro instruction 216
 SETFRR macro instruction 241
 SETLOCK macro instruction 245
 STAE macro instruction 260
 STATUS macro instruction 264,266
 release a serially reusable resource (DEQ) 138-142
 RELEASE parameter
 PGFIX macro instruction 182
 PGFREE macro instruction 185
 SETLOCK macro instruction 246
 remote class 30
 remote immediate signal, issuing 224-225
 remote pendable signal, issuing 226-227
 reporting interface, using the SRM 120 (5740-XE1)
 request control of a serially reusable resource (ENQ)
 145-151
 reserve a device (RESERVE) 216-221
 RESERVE macro instruction 216-221
 execute form 221
 list form 220
 standard form 216-219
 use of 18-22
 RESET parameter 266
 resource access control facility (RACF)
 data set 45,46
 group name 46
 identification 45-46
 password 46
 profiles 46
 RACDEF macro instruction 46
 RACHECK macro instruction 46
 RACINIT macro instruction 46
 RACLIST macro instruction 46
 resource manager 46
 scope of authorization 46
 userid 46
 verification 45-46
 resource control 11-45
 resource serialization 14
 resource serialization 11 (5740-XE1)
 restart function 29
 RESTART parameter 143
 restricted functions 23
 restricting
 load module access 25
 SVC service routines 23,25
 RESUME macro instruction
 coding description 222-223
 considerations 40
 considerations (5740-XE1)
 conditional 40.1 (5740-XE1)
 unconditional 40 (5740-XE1)
 how to use 40
 lock restrictions 40-41
 restrictions (5740-XE1)
 conditional 223 (5740-XE1)
 unconditional 223 (5740-XE1)
 resumption and suspension of request block tasks 38-41
 resumption and suspension of tasks 39-40.1 (5740-XE1)
 RET parameter
 DEQ macro instruction 138
 ENQ macro instruction 146
 RESERVE macro instruction 217
 RETADDR parameter 251
 RETIOQ parameter 136
 RETREGS parameter 251
 RETRN parameter 136
 retry routines 57-64
 return codes
 ATTACH macro instruction 129
 DEQ macro instruction 140
 DSGNL macro instruction 144
 ENQ macro instruction 148
 ESTAE macro instruction 153
 EXTRACT macro instruction 163
 FREEMAIN macro instruction 166
 GETMAIN macro instruction 170
 PGFIX macro instruction 183
 PGFREE macro instruction 186
 RACDEF macro instruction 197
 RACINIT macro instruction 204
 RACLIST macro instruction 210.1, 210.2
 RESERVE macro instruction 219
 RISGNL macro instruction 224-225
 RPSGNL macro instruction 227
 SDUMP macro instruction 234
 SETLOCK macro instruction 245,247,249
 STAE macro instruction 261
 TESTAUTH macro instruction 272
 RETURN parameter 222
 return parameters, setting 250-253
 RISGNL macro instruction
 general information 224-225
 use of 30
 RMC parameter 138

RMTR parameter 191,96
 ROUTCODE parameter
 WTO macro instruction 274
 WTOR macro instruction 278
 routing codes 91
 RPSGNL macro instruction
 general information 226
 use of 30
 RQCHECK parameter 226
 RSAPF keyword 28
 RTM (recovery/termination manager) 48
 RU parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
 RUB parameter 251
 RX-type address, meaning of 124

S parameter
 ENQ macro instruction 146
 EXTRACT macro instruction 158
 RESERVE macro instruction 217
 SALLOC (real storage manager and virtual storage manager space allocation lock)
 general information 11
 general information 14 (5740-XE1)
 GETMAIN and FREEMAIN 81
 SAVEKEY parameter 174
 SCHEDULE macro instruction 228
 and PURGEDQ 191
 use of 93-94
 schedule system services for asynchronous execution (SCHEDULE) 228
 SCOPE parameter 228
 screen flag bit (TCBSVCS) 107
 screen table address (TCBSVCA2) 107
 screening, subsystem SVC 107
 SD parameter 266
 SDATA parameter on the SDUMP macro instruction 229
 SDUMP dump requests 71
 SDUMP macro instruction 229-239
 execute form 238-239
 list form 236
 standard form 229-235
 use of 71
 SDWA (see system diagnostic work area)
 SDWACLUP 63
 SDWAFMID 63
 SDWAPERC 63
 sense function 29
 SENSE parameter 143
 SERIAL parameter 224
 serialization
 considerations 14
 of POST 33
 of resources 14
 of resources 11 (5740-XE1)
 of SVC routines 100
 requirements 14
 requirements 11 (5740-XE1)
 serially reusable resources
 controlling access to 243-249
 locking 11-14
 locking 12-14 (5740-XE1)
 releasing 138-142
 requesting control of 145-151
 service aids 73
 service classes 29-30
 service management 93-97
 service request block (SRB) 93-97
 serviceability level interception processing 65-67
 (see also SLIP command)

SET parameter on the STATUS macro instruction 265
 set return parameters (SETRP) 250-253
 set up functional recovery routines (SETFRR) 240-242
 SETCODE statement 26
 SETDIE 76-77
 (see also disabled interrupt exit)
 SETDIE entry point 76
 SETFRR macro instruction 240-242
 type 6 SVC recovery 105
 use of 50-51
 SETLOCK macro instruction
 general information 243-249
 use of 15
 SETRP macro instruction 250-253
 shared DASD 17,216
 SHSPL parameter 127
 SHSPV parameter 127
 signal event completion (POST) 187-190
 signal processor (SIGP) instruction
 DSGNL macro instruction 143
 RISGNL macro instruction 224
 RPSGNL macro instruction 226
 SIGP (signal processor) instruction 29-30
 SIO parameter 226
 SLIP command, examples of
 obtaining a dump with queue elements and control blocks 68
 obtaining a dump with SQA control blocks 67
 replacing an SVC dump with a standalone dump 68
 setting interception definitions for an application program 69
 setting a trap using instruction fetch PER event 69
 setting a trap using storage alteration PER event 69
 setting a trap using storage alteration PER event and the ACTION=IGNORE option 70
 setting a trap using successful branch PER event 70
 using 66
 SM parameter 128
 SMC parameter 147,205,207
 SMF cross memory services lock (CMSSMF) 14 (5740-XE1)
 SNAP/ABEND user exit 73
 SP parameter
 FREEMAIN macro instruction 166
 GETMAIN macro instruction 169
 special clean-up processing 64
 specify program interruption exit (SPIE) 254-257
 specify task abnormal exit (STAE) 259-263
 SPIE macro instruction 254-257
 execute form 257
 list form 256
 standard form 254
 use of restricted interruptions 47
 SPIE processing 47
 SPIE 17 routine 47
 spin locks
 general information 12
 general information 12 (5740-XE1)
 SVC routines 102
 SPOST macro instruction
 general information 258
 use of 31
 SQA buffer 71
 SQA for subsystem SVC screening 107
 SRB (service request block) 93-96
 input format 106
 mode processing, type 6 SVC 105
 SRB activity, purging 191-193
 SRB parameter
 SCHEDULE macro instruction 228
 STATUS macro instruction 264
 T6EXIT macro instruction 273

SRBASCBC field 94
 SRBASID field 95
 SRBCPAFF field 94
 SRBEP field 94
 SRBPARM field 94
 SRBPKF field 94
 SRBPRIOR field 94
 SRBPTCB field 95
 SRBRMTR field 94
 SRBSAVE field 94
 SRM (system resource manager) lock 12
 SRM (system resource manager) lock 13 (5740-XE1)
 SRM (system resource manager) reporting interface 120 (5740-XE1)
 SSM instruction 14
 SSS parameter 143
 SSTMASK screening area 107
 SSTSVCN subsystem SVC entry area 107
 STAB parameter 136
 STAE environment 51-52
 STAE exit routines 51-52
 STAE macro instruction 259-263
 execute form 263
 list form 262
 standard form 259-261
 use of 51-52
 STAE recovery routine 50,51
 STAE retry routine 57,58
 stage 1 exit effector 97
 stage 2 exit effector 97
 stage 3 exit effector 97
 STAI environment 51-52
 STAI exit routines 51-52
 STAI parameter 127
 STAI recovery routine 50,51
 STAI retry routine 57-58
 start function 29
 START parameter
 DSGNL macro instruction 143
 STATUS macro instruction 264
 RACINIT macro instruction 206
 STATE parameter 271
 status conditions 30-31
 STATUS macro instruction 264-266
 STCK instruction 78
 STEP parameter
 CALLRTM macro instruction 132
 DEQ macro instruction 138
 ENQ macro instruction 146
 STATUS macro instruction 258
 stimulus for event completion 40
 stop and store status function 29
 stop function 29
 STOP parameter
 DSGNL macro instruction 143
 STATUS macro instruction 264
 stopped status condition 30
 STORAGE parameter 233
 SUBPOOL parameter 205,207,211,212
 subpools
 common service area 82
 fetch protected 82
 global 82
 not fetch protected 82
 system queue area 81
 227 81-82
 228 81-82
 229 81-82
 230 81-82
 231 81-82
 236 81-82
 237 81-82
 239 81-82
 241 81-82
 245 81-82
 subsystem SVC entry area (SSTSVCN) 107
 subsystem SVC screening 107
 subtask creation and control 3
 subtask status, changing 264-266
 SUMDUMP parameter 232
 SUMLIST parameter 232
 supervisor services 121
 suspend count field (RBSCF) 39-40
 suspend locks 12
 SUSPEND macro instruction
 coding description 267
 considerations 39-40
 examples of 39
 how to use 38-40
 SUSPEND parameter 184
 suspension and resumption of request block tasks 38-41
 suspension and resumption of tasks 39-40.1 (5740-XE1)
 SVAREA parameter
 ATTACH macro instruction 128
 CIRB macro instruction 135
 SVC dumps, error correlation 79-80
 SVC routines, user-written 100-108
 calling SVC routines 54
 characteristics 100
 inserting 103
 programming conventions 100-103
 specifying 103
 SVC screening, subsystem 107
 SVC 34, in response to FORCE command 50
 SVC Table entries 104
 SVCTABLE macro instruction 103,
 SVCTABLE options 105
 SVEAREA parameter 152
 SVRB, SVC routines 103
 SWITCH parameter 226
 symbol, meaning of 124
 SYNCH macro instruction
 general information 268
 use of 9
 SYNCH parameter 264
 synchronize POST (SPOST) 258
 synchronous exits, taking 268
 SYSABEND dump requests 72
 SYSEVENT macro instruction 269.1 (5740-XE1)
 return codes 269.2 (5740-XE1)
 examples 269.2 (5740-XE1)
 SYSMASK parameter 177
 system diagnostic workarea 54-64
 system generation, SVC routines 103
 system integrity 40-44
 system log 92
 system message IEA911A 74
 SYSTEM parameter
 DEQ macro instruction 138
 ENQ macro instruction 152
 STATUS macro instruction 264
 system queue area subpools 81
 system resource manager lock (SRM) 12
 system resource manager lock (SRM) 14 (5740-XE1)
 system status 173-178
 SYSTEMS parameter
 DEQ macro instruction 138
 ENQ macro instruction 146
 RESERVE macro instruction 217
 SYS1.AOSCE 114
 SYS1.LINKLIB
 APF authorization 24
 power warning feature support 110
 SVC routines 100

SYS1.LOGREC records 74-75
 SYS1.LPALIB
 APF authorization 24
 clean-up routines 64
 SVC routines 100
 SYS1.NUCLEUS, power warning feature support 110-112
 SYS1.SVCLIB
 APF authorization 24
 SVC routines 100
 SZERO parameter 127

 TABLE parameter 157
 take asynchronous exit to a processing program (SYNCH)
 264
 task abnormal exit
 extended 152-155
 specifying 259-263
 task, creating a new 30
 task recovery routines 50-64
 task termination 49-50
 TASKLIB parameter 127
 TCB information, extracting 158-161
 TCB mode processing, Type 6 SVC 105
 TCB parameter
 CALLRTM macro instruction 131
 DEQ macro instruction 138
 ENQ macro instruction 46
 RESUME macro instruction 222
 STATUS macro instruction 264
 TCBACTIV flag 107
 TCBSVCA2 screen table address 107
 TCBSVCS screen flag bit 107
 TCTL macro instruction
 coding description 270
 how to use 40
 TERM parameter
 ATTACH macro instruction 127
 ESTAE macro instruction 152
 TERMD parameter 206
 test authorization of caller (TESTAUTH) 271-272
 TEST parameter 248
 TEST UNDER MAST (TM) instruction 31
 TESTAUTH macro instruction
 general information 271-272
 use of 23,25
 TID parameter 128
 time interval for missing interruption handler 108-110
 time intervals, DIE 76
 timer disabled interrupt exit 76
 (see also disabled interrupt exit)
 timer SLIH 77
 timer supervision work area 76
 TM (TEST UNDER MASK) instruction 31
 TOD clock 79
 TPCSDIE field 76
 TQE DEQUEUE routine 78
 TQE ENQUEUE routine 78
 TQE AID field initialization 77
 TQE ASCB field 77
 TQE DREGS field 78
 TQE EXIT field initialization 78
 TQE TCB field 78
 TQE VAL field initialization 78
 trapping system errors
 (see SLIP command)
 TRAXERPT parameter 269.1 (5740-XE1)
 TRAXFRPT parameter 269.1 (5740-XE1)
 TRAXRPT parameter 269.1 (5740-XE1)
 TYPE parameter
 CALLRTM macro instruction 131
 CIRB macro instruction 136
 RACDEF macro instruction 196
 SETLOCK macro instruction 243
 types of locks 12-13
 type 5 SVC 104
 type 6 SVC routine
 considerations for use 105
 description 104-105
 exit options 105
 location 100
 modification plans 105
 SRB mode processing 105
 TCB mode processing 105
 T6EXIT macro instruction
 coding description 273
 use of 105

 UCB parameter
 DEQ macro instruction 138
 RESERVE macro instruction 217
 UPDATE parameter 200-203
 user region size 117
 user storage areas 42
 user-written message routing exit routines 90-93
 user-written SVC routines 97-106
 USERID parameter 204

 V parameter 164,199-203
 varying storage 85-86
 VC parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169
 virtual storage
 allocating 169-172
 dumping 229-239,67-75
 freeing 164-168
 virtual storage contents
 fixing 182-184,83-84
 freeing 185
 virtual storage, dumping 67-74
 virtual storage management 81
 virtual subarea list (VSL) 84-85
 VL parameter 127
 VOLSER parameter 195
 volume handling 17
 VSL (virtual subarea list) 84-85
 VU parameter
 FREEMAIN macro instruction 165
 GETMAIN macro instruction 169

 wait for events (EVENTS) 157
 WAIT macro instruction
 SVC routines 103
 WAIT parameter 157
 WKAREA parameter
 CIRB macro instruction 136
 SETRP macro instruction 251
 WORKREG parameter 174
 WREGS parameter
 NIL macro instruction 178
 OIL macro instruction 180
 write-to-operator (WTO) 274-277
 write-to-operator-with-reply (WTOR) 278-280
 writing POST exit routines 32
 WRKREGS parameter 240
 WTO macro instruction 274-277
 list form 277
 standard form 274-276
 use of restricted parameters 89-93

WTOR macro instruction 278-280
list form 280
standard form 278-280
use of restricted parameters 89-93

XCTL macro instruction, SVC routines 103

XCTL parameter
ESTAE macro instruction 152
STAE macro instruction 259

Y parameter 222

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

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 comments are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If comments apply to a Selectable Unit, please provide the name of the Selectable Unit _____.

If you wish a reply, give your name and mailing address:

Note: Staples can cause problems with automated mail sorting equipment. Please use pressure sensitive or other gummed tape to seal this form.

Cut or Fold Along Line

Please circle the description that most closely describes your occupation.

Customer	(Q) Install Mgr.	(U) System Consult.	(X) System Analyst	(Y) System Prog.	(Z) Applica. Prog.	(F) System Oper.	(I) I/O Oper.	(L) Term. Oper.					(O) Other
IBM	(S) System Eng.	(P) Prog. Sys. Rep.	(A) System Analyst	(B) System Prog.	(C) Applica. Prog.	(D) Dev. Prog.	(R) Comp. Prog.	(G) System Oper.	(J) I/O Oper.	(E) Ed. Dev. Rep.	(N) Cust. Eng.	(T) Tech. Staff Rep.	

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.)

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:

POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department D58, Building 706-2
PO Box 390
Poughkeepsie, New York 12602



Fold and tape

Please Do Not Staple

Fold and tape



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

OS/VS2 System Programming Library: Supervisor Printed in U.S.A. GC28-0628-3