

**Systems**

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

**VS2 Release 3**

*p 50, 31, 70*

**IBM**

### **First Edition (February, 1975)**

This edition, with **OS/VS2 System Programming Library: Job Management**, GC28-0627, and **OS/VS2 System Programming Library: TSO**, GC28-0629, obsoletes **OS/VS2 System Programming Library: Job Management, Supervisor, and TSO**, GC28-0682-0.

This edition applies to release 3 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 **Virtual Storage Supplement to IBM System/360 and System/370 Bibliography**, GC20-0001, for the editions that are applicable and current.

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

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Programming Systems Publications, Department D58, Building 706-2, PO Box 390, Poughkeepsie, N.Y. 12602. Comments become the property of IBM.

© Copyright International Business Machines Corporation 1974,1975

## 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 and installation-approved personnel, and 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.

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 Scheduler and Supervisor Logic**, SY28-0624, SY28-0625, SY28-0626.  
(3 volumes)

**OS/VS - DOS/VS - VM/370 Assembler Language**, GC33-4010.

**OS/VS2 System Programming Library: System Generation Reference**, GC26-3792.



# Contents

<b>Summary of Amendments</b> . . . . .	9
<b>Part I: Supervisor Services</b> . . . . .	11
<b>Subtask Creation and Control</b> . . . . .	13
Creating a New Task . . . . .	13
User Modifications . . . . .	14
Operator Communication With a Problem Program . . . . .	15
Providing an EXTRACT Answer Area . . . . .	16
<b>Program Management</b> . . . . .	17
Synchronous Exits . . . . .	17
<b>Resource Control</b> . . . . .	19
Locking . . . . .	19
Classes of Locks . . . . .	20
Types of Locks . . . . .	20
Locking Hierarchy . . . . .	21
Locking Considerations . . . . .	21
SETLOCK Macro Instruction . . . . .	22
The Must Complete Function . . . . .	22
Characteristics of the Must Complete Function . . . . .	23
Programming Notes . . . . .	23
Shared Direct Access Storage Devices (Shared DASD) . . . . .	24
Devices That Can be Shared . . . . .	24
Volume/Device Status . . . . .	24
System Configuration . . . . .	25
Volume Handling . . . . .	25
Macro Instructions Used With Shared DASD . . . . .	25
Releasing Devices . . . . .	25
Preventing Interlocks . . . . .	25
Volume Assignment . . . . .	26
Program Libraries . . . . .	26
Finding the UCB Address . . . . .	26
Authorized Program Facility (APF) . . . . .	30
Restricting Services/Resources . . . . .	30
Assigning Authorization . . . . .	30
Defining Authorized Libraries . . . . .	31
Changing System Status . . . . .	32
Generating an SVC . . . . .	32
Generating Inline Code . . . . .	32
Interprocessor Communications (IPC) . . . . .	32
Service Classes . . . . .	33
Status Conditions . . . . .	34
Event Completion . . . . .	35
Cross Memory POST . . . . .	35
Bypassing the POST Routine . . . . .	35
Waiting for Event Completion . . . . .	36
<b>Interruption, Termination, and Dumping Services</b> . . . . .	37
Recovery/Termination . . . . .	37
Task Termination . . . . .	37
Address Space Termination . . . . .	38
Processing Functional Recovery Routines . . . . .	38
SETFRR Macro Instruction . . . . .	39
SPIE Processing . . . . .	39
Dumping Virtual Storage . . . . .	39
Using the SDUMP Macro Instruction . . . . .	39
SQA Buffer . . . . .	40
Using the CHNGDUMP Command . . . . .	40
Task Recovery . . . . .	41
STAE/STAI Exit Routines . . . . .	41
Interface to a STAE/STAI Exit . . . . .	42
STAE/STAI Retry Routines . . . . .	42
ESTAE Extended Capabilities . . . . .	43

The BRANCH Option . . . . .	44
The TERM Option . . . . .	44
The PURGE and ASYNCH Parameters . . . . .	44
ESTAE/ESTAI Exit Routines . . . . .	45
Interface to an ESTAE/ESTAI Exit . . . . .	47
ESTAE/ESTAI Retry Routines . . . . .	49
Clean-Up Routines . . . . .	50
Support for Installation-Written Clean-Up Routines . . . . .	50
Programming Considerations . . . . .	50
<b>Virtual Storage Management . . . . .</b>	<b>53</b>
The BRANCH Parameter . . . . .	53
The KEY Parameter . . . . .	53
<b>Real Storage Management . . . . .</b>	<b>55</b>
Fixing/Freeing Virtual Storage Contents . . . . .	55
Completion Considerations . . . . .	56
Virtual Subarea List (VSL) . . . . .	56
<b>Miscellaneous Services . . . . .</b>	<b>59</b>
Writing Operator Messages . . . . .	59
Routing the Message . . . . .	59
Writing a Multiple-Line Message . . . . .	60
Message Routing Exit Routines . . . . .	60
Characteristics of MCS . . . . .	60
Programming Conventions for WTO/WTOR Routines . . . . .	61
Messages Not Using Routing Codes . . . . .	62
Writing a WTO/WTOR Exit Routine . . . . .	62
Adding a WTO/WTOR Exit Routine to the Control Program . . . . .	62
Service Management . . . . .	63
Scheduling Service Requests . . . . .	64
Service Request Blocks (SRBs) . . . . .	64
Priorities . . . . .	64
Characteristics of Service Requests . . . . .	65
Purging Service Requests . . . . .	65
PURGEDQ Parameters . . . . .	66
Creating Interruption Request Blocks . . . . .	66
Writing SVC Routines . . . . .	67
Characteristics of SVC Routines . . . . .	67
Programming Conventions for SVC Routines . . . . .	67
Inserting SVC Routines Into the Control Program . . . . .	70
Specifying SVC Routines . . . . .	70
Inserting SVC Routines During the System Generation Process . . . . .	71
Missing Interruption Handler . . . . .	71
Establishing a Time Interval . . . . .	71
Adding Code to the Power Warning Feature Support . . . . .	72
Writing Code for the Machine Check Handler Appendage . . . . .	72
Coding Considerations . . . . .	73
Writing Code for the Master Scheduler Initialization Routine . . . . .	74
Coding Considerations . . . . .	75
Control Track Record . . . . .	76
<b>Part II: Reference — Macro Instructions . . . . .</b>	<b>79</b>
Macro Instruction Forms . . . . .	80
Coding the Macro Instructions . . . . .	80
Continuation Lines . . . . .	82
<b>Descriptions of the Macro Instructions . . . . .</b>	<b>83</b>
ATTACH -- Create a New Task . . . . .	84
CALLRTM -- Call Recovery/Termination Manager . . . . .	88
CHNGDUMP -- Change Dump Options . . . . .	90
CIRB -- Create Interruption Request Block . . . . .	94
DEQ -- Release a Serially Reusable Resource . . . . .	96
DEQ (List Form) . . . . .	99
DEQ (Execute Form) . . . . .	100
DSGNL -- Issue Direct Signal . . . . .	101
ENQ -- Request Control of a Serially Reusable Resource . . . . .	103
ENQ (List Form) . . . . .	106
ENQ (Execute Form) . . . . .	107
ESTAE -- Extended STAE . . . . .	108

ESTAE (List Form)	110
ESTAE (Execute Form)	111
EVENTS -- Wait for Events	112
EXTRACT -- Extract TCB Information	113
EXTRACT (List Form)	115
EXTRACT (Execute Form)	116
FREEMAIN -- Free Virtual Storage	117
FREEMAIN (List Form)	119
FREEMAIN (Execute Form)	120
GETMAIN -- Allocate Virtual Storage	121
GETMAIN (List Form)	123
GETMAIN (Execute Form)	124
MODESET -- Change System Status	125
MODESET (List Form)	128
MODESET (Execute Form)	129
NIL -- Provide a Lock Via an AND IMMEDIATE (NI) Instruction	130
OIL -- Provide a Lock Via an OR IMMEDIATE (OI) Instruction	132
PGFIX -- Fix Virtual Storage Contents	134
PGFIX (List Form)	137
PGFREE -- Free Virtual Storage Contents	138
PGFREE (List Form)	140
POST -- Signal Event Completion	141
POST (List Form)	142
POST (Execute Form)	143
RGEDQ -- Purge SRB Activity	144
PURGEDQ (List Form)	145
PURGEDQ (Execute Form)	146
QEDIT -- Link to SVC 34	147
RESERVE -- Reserve a Device	148
RESERVE (List Form)	151
RESERVE (Execute Form)	152
RISGNL -- Issue Remote Immediate Signal	153
RPSGNL -- Issue Remote Pendable Signal	155
SCHEDULE -- Schedule System Services for Asynchronous Execution	157
SDUMP -- Dump Virtual Storage	158
SDUMP (List Form)	161
SDUMP (Execute Form)	162
SETFRR -- Set Up Functional Recovery Routines	164
SETLOCK -- Control Access to Serially Reusable Resources	166
SETRP -- Set Return Parameters	173
SPIE -- Specify Program Interruption Exit	177
SPIE (List Form)	178
SPIE (Execute Form)	179
SPOST -- Synchronize POST	180
STAE -- Specify Task Abnormal Exit	181
STAE (List Form)	184
STAE (Execute Form)	185
STATUS -- Change Subtask Status	186
SYNCH -- Take a Synchronous Exit to a Processing Program	190
TESTAUTH -- Test Authorization of Caller	191
WTO -- Write to Operator	193
WTO (List Form)	196
WTOR -- Write to Operator With Reply	197
WTOR (List Form)	199
<b>Index</b>	<b>201</b>

## Figures

Figure 1.	EXTRACT ECB and CIB Pointers . . . . .	15
Figure 2.	Command Input Buffer Contents . . . . .	15
Figure 3.	EXTRACT Answer Area Fields . . . . .	16
Figure 4.	Summary of Locking Characteristics . . . . .	21
Figure 5.	Valid Volume Characteristic and Device Status Combinations . . . . .	24
Figure 6.	Example of an Interlock Environment . . . . .	26
Figure 7.	Example of Subroutine Issuing RESERVE and DEQ . . . . .	29
Figure 8.	Assigning Authorization via SETCODE . . . . .	31
Figure 9.	Bypassing the POST Routine . . . . .	36
Figure 10.	ESTAE Environment . . . . .	46
Figure 11.	Virtual Subarea List Entries . . . . .	56
Figure 12.	Programming Conventions for SVC Routines . . . . .	68
Figure 13.	Changing the Missing Interruption Handler Time Interval . . . . .	71
Figure 14.	Logical Placement of Your Code in the Machine Check Handler Appendage . . . . .	73
Figure 15.	Logical Placement of Your Code in the Master Scheduler Initialization Module . . . . .	75
Figure 16.	Control Track Record . . . . .	76
Figure 17.	Storage Assignments on MP Systems . . . . .	77
Figure 18.	Macro Instruction Coverage . . . . .	79
Figure 19.	Sample Macro Instruction . . . . .	81
Figure 20.	Continuation Coding . . . . .	82
Figure 21.	Return Code Area Used by DEQ . . . . .	97
Figure 22.	Return Code Area Used by ENQ . . . . .	104
Figure 23.	Return Code Area Used by RESERVE . . . . .	150
Figure 24.	MCSFLAG Fields . . . . .	194



## **Summary of Amendments for GC28-0628-0 OS/VS2 Release 3**

This edition, with **OS/VS2 System Programming Library: Job Management**, GC28-0627, and **OS/VS2 System Programming Library: TSO**, GC28-0629, obsoletes **OS/VS2 System Programming Library: Job Management, Supervisor, and TSO**, GC28-0682-0.

### **Power Warning Feature Support**

This edition incorporates system programmer information previously contained in **OS/VS Power Warning Feature (PWF) Support**, ICR Guide, GC28-0686-0.

### **Virtual Storage Management**

New Subpools 227 and 228 are discussed.

### **Event Completion**

The EVENTS macro instruction, and its use, are discussed.

### **SPIE Macro Instruction**

The syntax of the SPIE macro instruction is presented.

### **Miscellaneous Corrections**

ESTAE processing  
APF  
Locking  
MODESET  
WTO  
Command input buffers

### **Examples**

The examples of the standard, list, and execute forms of the macro instructions are more properly placed at the ends of their appropriate sections.



## 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 will help 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 may be described 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:** The supervisor can be used 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 which 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.

## ***User Modifications***

Rather than accepting the default values, 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 attaching 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 (for example, TCBPKF can be reset to a key other than zero or the attaching task's key). 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=jscb 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.

## 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 will return 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.

0	Address of next CIB		Verb code	CIB length	Reserved
8	Reserved	TSO terminal ID	Console ID	Reserved	Length of data field
10	Variable length data specified on the command.				

```
Verb code  X'04' START
           X'40' STOP
           X'44' MODIFY
```

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
```

The CIB counter should then be set to allow CIBs to be chained and MODIFY commands will 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 will be accepted for the job. STOP commands, however, will be accepted for the job regardless of the value set for CIBCTR.

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 will be cleared when no more CIBs remain. Care should be taken if multiple subtasks are examining these fields. Any CIBs not freed by the task will be 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.

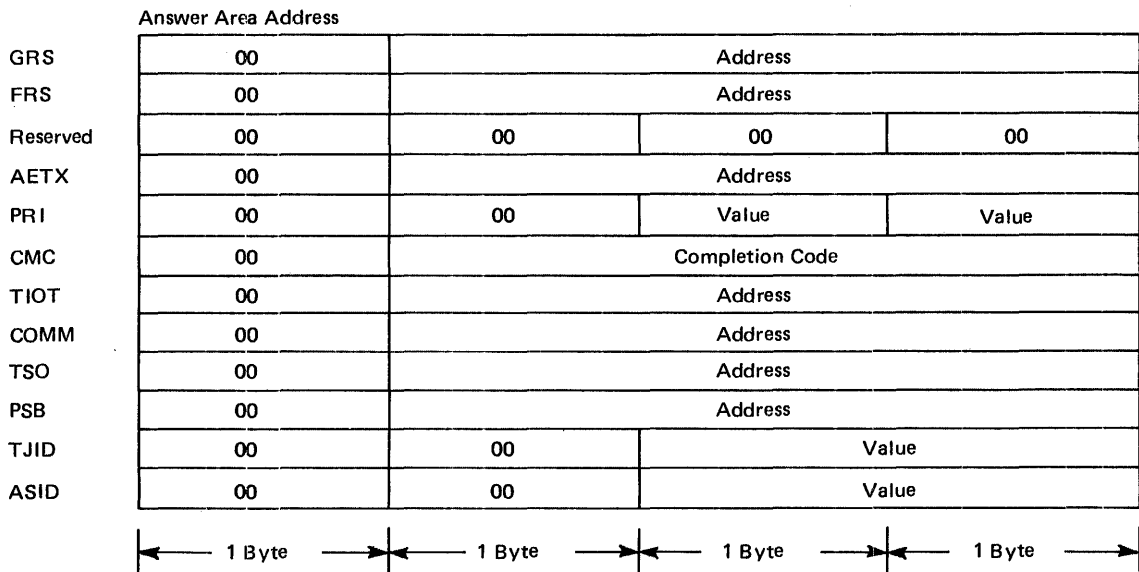


Figure 3. EXTRACT Answer Area Fields

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 address or value has been 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 will appear in the same relative order as shown in Figure 3. (That is, GRS will be 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 will appear 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 only be used by system programmers or other installation-approved personnel.

### Synchronous Exits

In general, the SYNCH macro instruction is used when a control program in the supervisor state is to give temporary control to a processing program routine, and when the processing program is expected to return control to the supervisor state. 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 storage protection 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.

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 entry name address of the subroutine required for the user's private label processing.



## Resource Control

Planning is required when more than one program uses a serially reusable resource. 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)

### Locking

Adequate control of serial resources is a significant problem in a multiprocessing environment. Since some uniprocessor serialization techniques (for example, disablement), are no longer effective, there is the possibility of multiple tasks, even multiple tasks in the same job step, using the same serial resource while running concurrently on different CPUs.

To eliminate this problem, a locking mechanism is provided to control access to serially reusable resources. The lock manager will handle all functions related to the locks (for example, obtaining or releasing locks, or checking the status of a CPU with respect to a particular lock). However, a locking technique can not be effective unless all routines that have the potential for changing the protected resource, or that depend on its status remaining unchanged for a given period, make use of the locking mechanism.

In MVS, a locking manager controls a new hierarchical locking structure with multiple types of locks and synchronizes the use of serially reusable resources. Use of the locking manager is restricted to key 0 programs running in supervisor state, which prevents unauthorized problem programs from interfering with the serialization process.

To enhance performance, each kind of serially reusable resource is assigned a separate lock. In this way, a lock held by a CPU on one resource does not prevent the other CPU from using a different resource.

The locks provided in MVS are:

- Global dispatcher lock (DISP) -- serializes all functions associated with the dispatching process.
- Auxiliary storage manager lock (ASM) -- serializes use of the global ASM control blocks.
- Real storage manager and virtual storage manager space allocation lock (SALLOC) -- serializes the global functions of real storage management and virtual storage management.
- IOS synchronization lock (IOSYNCH) -- serializes global IOS functions.
- IOS channel availability table lock (IOSCAT) -- serializes access and updates to the channel availability table.
- IOS unit control block lock (IOSUCB) -- serializes access and updates to the unit control blocks. There is one lock per UCB.
- IOS logical channel queue lock (IOSLCH) -- serializes access and updates to the IOS logical channel queues. There is one lock per channel queue.

- System resources manager lock (SRM) -- 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 local address space. There is one lock per address space.

### ***Classes of Locks***

Two classes 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 a CPU holds a local lock, the queues and control blocks associated with that address space can be manipulated only by the CPU holding the lock. (For example, an address space-related GETMAIN for a user subpool requires a local lock.)

All of the locks described above, with the exception of the local storage 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 locks, on the other hand, do not serialize across address spaces, but serialize functions executing within the address space.

### ***Types of Locks***

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

- Spin locks -- prevent the requesting CPU from doing any work until the lock is cleared by the other CPU. The requesting CPU enters a loop that keeps testing the lock until the other CPU releases it. As soon as the resource is free, the first CPU can obtain the resource and continue processing.
- Suspend locks -- prevent the requesting program from doing work until the lock is available, but allow the CPU to continue doing other work. The request is queued and the requesting CPU is dispatched to do other work. Upon release of the lock, the highest priority queued requestor will be given control of the lock.

All of the locks described above, with the exception of the local storage and cross memory services locks, are spin locks. The LOCAL and CMS locks run enabled and can be interrupted to run higher priority work. If there is another request for the lock while it is held, the requestor is suspended and other work is dispatched.

The 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 logically disabled 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 normally the functions which run under the locks are of short duration, and 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 thus have to remain disabled.

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

Figure 4 summarizes the characteristics of the locks.

lock	global	local	spin	suspend
DISP	X		X	
ASM	X		X	
SALLOC	X		X	
IOSYNCH	X		X	
IOSCAT	X		X	
IOSUCB	X		X	
IOSLCH	X		X	
SRM	X		X	
CMS	X			X
LOCAL		X		X

Figure 4. Summary of Locking Characteristics

### ***Locking Hierarchy***

To prevent a deadlock between CPUs, the locks are arranged in a hierarchy, and a CPU may unconditionally request only locks higher in the hierarchy than locks that it currently holds. The locking hierarchy is the order in which the locks are listed in Figure 4, with DISP being the highest lock in the hierarchy.

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 which have more than one lock, a CPU may only hold one lock of each type. For example, if a CPU holds a IOSUCB lock, it may not request a different IOSUCB lock.

The LOCAL lock must be held by the caller when requesting the CMS lock. Also, the LOCAL lock cannot be released while holding the CMS lock.

It is not necessary to obtain all locks in the hierarchy up to the highest lock needed. Although only the needed locks have to be obtained, they do, however, have to be obtained in hierarchy sequence.

### ***Locking Considerations***

In MVS the locking function is provided to replace disabling in the control program with a form of logical disabling which works across both CPUs. Locking, however, is not completely equivalent to disablement in VS2 Release 1, and required some changes in the system. The significant differences are:

- In VS2 Release 1, a disabled program could issue an SVC and receive control back disabled. In MVS, a locked routine is not allowed to issue an SVC.
- All user-written functions which disable using the MODESET macro instruction, the SSM 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 CPU even though a global spin lock is held on one CPU. To prevent page stealing, it is necessary to either fix the pages or hold the SALLOC lock.

### ***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.

**Disabled/Enabled State for OBTAIN:** When a global spin lock is successfully obtained, control is returned to the SETLOCK caller with the CPU 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 CPU in the physically enabled or disabled state that existed prior to the SETLOCK request. This is also true for unsuccessful attempts to obtain any of the locks where control is returned to the caller.

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

When the LOCAL lock or CMS lock is released, control is returned with the CPU in the enabled or disabled state that existed prior to the RELEASE request.

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

For multiple RELEASE requests (via SPIN, ALL, or (reg) subparameters), the final state is the same as what 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 engage in updating and/or manipulation of 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 a 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.

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. Thus, in this case, enqueue on the VTOC and use the must complete function (to suspend processing of other tasks) when updating a VTOC.

Just as the ENQ function serializes use of a resource requested by many different tasks, the must complete function serializes execution of tasks.

### ***Characteristics of the Must Complete Function***

The must complete function can be used only at the step level, where only the current task in an address space is allowed to execute. All other tasks, including the initiator task, are placed in a wait state.

When the must complete function is requested, the requesting task is marked as being in the 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. Thus, 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 wanted 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 will result if a serially-resuable routine requested by one of these macro instructions has been requested by one of tasks made nondispatchable by the use of the SMC parameter or was requested by another task and has been only partially fetched.
5. The time a routine is in the must complete mode should be kept as short as possible -- enter at the last moment and leave as soon as possible. One suggested way is to:
  - a. ENQ (on desired resource(s))
  - b. ENQ (on same resource(s)), RET=HAVE, SMC=STEP

Item 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 Auxilliary 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 Series devices may also be configured for shared use via the string switch feature on the IBM 3333 device.

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 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 or reserved

---

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 off line 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 will 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 will free the device when the task is terminated.

## ***Preventing Interlocks***

The more often device reservations occur 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 program use of the RESERVE macro instruction and enabled wait states will occur on the system or systems.

## 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.

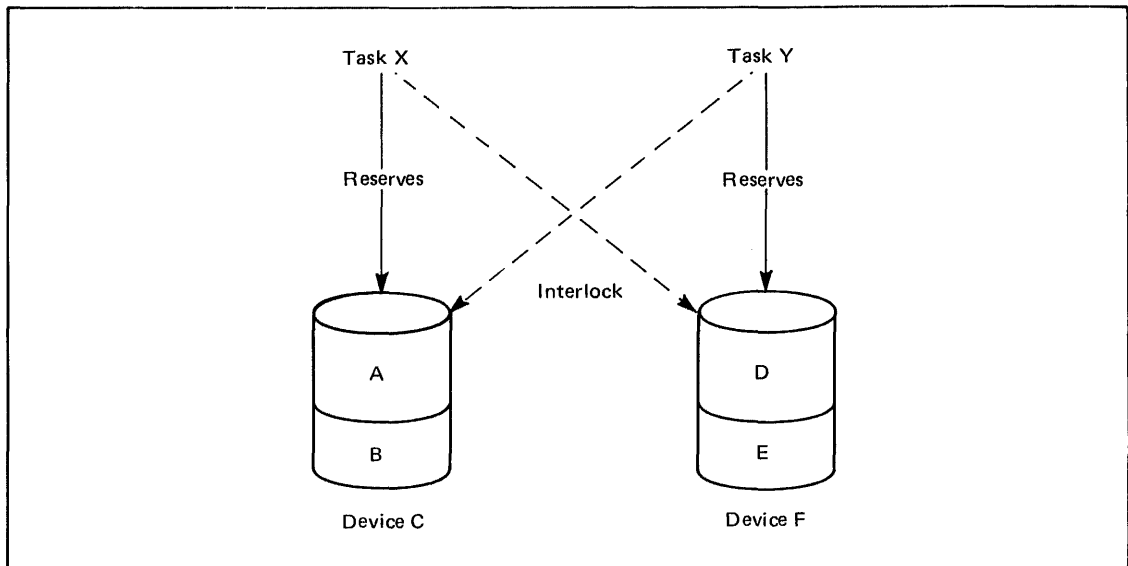


Figure 6. Example of an Interlock Environment

## 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 MBBCCHHR.
3. Load the DEB address from the DCB field labeled DCBDEBAD.
4. Load the address of the field labeled DEBDVMOD in the DEB.
5. Multiply the "M" of the direct access by 16.
6. The sum of steps 4 and 5 is the address of the correct extent entry in the DEB for the next READ/WRITE operation. The sum is also a pointer to the UCB address for this extent.
7. Issue the RESERVE macro specifying the address obtained in step 6 as the parameter of the UCB keyword.

If the data set is an ISAM data set, QISAM in the load mode should be used only at system update time. Further, if it is a multivolume ISAM data set, it must be assumed that all jobs will access the data set through the highest level index. The indexes should never reside in virtual storage when the data set is being shared. In this case, by issuing a RESERVE macro for the volume on which the highest level index resides, the user effectively reserves the volumes on which the prime data and independent overflow areas reside. The following procedures can be used to achieve this:

1. Open the data set.
2. Locate the actual device address (MBBCCHHR) of the highest level index. This address can be obtained from the DCB.
3. Load the DEB address from the DCB field labeled DCBDEBAD.
4. Load the address of the field labeled DEBDVMOD in the DEB.
5. Multiply the "M" of the actual device address located in step 2 by 16.
6. The sum of steps 4 and 5 is the address of the correct extent entry in the DEB for the next READ/WRITE operation. The sum is also a pointer to the UCB address for this extent.
7. Issue the RESERVE macro specifying the address obtained in step 6 as the parameter of the UCB keyword.

**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
                                      SAVE AREA
            ST        11,8(13)
            LR        13,11        ADDRESS OF MY SAVE AREA
            LR        9,1          ADDRESS OF PARAMETER LIST
            L         3,0(9)        DDNAME PARAMETER OR WORD OF
                                      ZEROS
            CLC      0(4,3),=F'0'   WORD OF ZEROS IF DEQ IS
                                      REQUESTED
            BE        WANTDEQ
*PROCESS FOR DETERMINING THE UCB ADDRESS USING THE TIOT
            XR        11,11        REGISTER USED FOR DD ENTRY
            EXTRACT  ADDR TIOT, FIELDS=TIOT
            L         7,ADDR TIOT   ADDRESS OF TASK I/O TABLE
            LA        7,24(7)      ADDRESS OF FIRST DD ENTRY
NEXTDD      CLC      0(8,3),4(7)    COMPARE DDNAMES
            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) limits the use of sensitive system and (optionally) user services and resources to authorized system and user programs. If an installation has its own control program extensions and special SVCs that allow the bypass of normal system security or integrity checks (for example, an SVC that returned control in key 0), and if such SVCs are not currently restricted from use by an unauthorized program, then the APF facility should be used to restrict them and to authorize the control program extensions that use them.

*Note:* The user of APF should be aware that APF authorization is defined such that any SVC that is restricted by the APF mechanism, can be executed by any system key (0-7) or privileged mode routine. However, this not true in reverse. An APF-authorized program will not automatically be allowed system services that are restricted by system key or privileged mode tests.

### *Restricting Services/Resources*

MVS provide two methods of restricting sensitive system SVCs:

- Sensitive SVCs may be restricted via a parameter on the SCVTABLE macro instruction issued at system generation time. The parameter will restrict the use of the SVC to authorized job steps. If the parameter is not specified, the SVC will be unrestricted. The SVC first level interruption handler ensures that SVCs restricted in this manner are accessible only to programs having APF authorization, or executing under supervisor state or system key 0-7.
- When only a part of the SVC's function is sensitive, the TESTAUTH macro instruction can restrict particular paths through the SVCs. The TESTAUTH macro instruction, inserted at appropriate locations in the SVC, returns an authorized or unauthorized indication, and the SVC must take appropriate action based on this return.

TESTAUTH can also test for system key 0-7 and for supervisor state. The authorization checks are made in the following order: state, key, and APF. If any checks are successful, further checks are unnecessary.

### *Assigning Authorization*

The linkage editor permits an installation to establish authorization for programs either through a new parameter in the linkage edit step or through a new linkage editor control statement.

To assign an authorization code via the new parameter, AC=1 should be coded in the PARM field, as follows:

```
//LKED EXEC PGM=HEWL, PARM='AC=1', . . .
```

If no authorization code is assigned in the linkage editor job step, the default is nonauthorization. The authorization code for a given output module can be overridden with the SETCODE control statement.

The SETCODE statement is used to establish authorization for a specific output load module. If it is used, it must be placed 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 will be 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        DSNNAME=SYS1.LINKLIB,DISP=OLD
//SYSLIN    DD        DSNNAME=&&LOADSET,DISP=(OLD,PASS)
//          DD        UNIT=SYSDA
//          DD        *
      SETCODE      AC(1)
      NAME          MOD1(R)
/*

```

---

Figure 8. Assigning Authorization via SETCODE

### ***Defining Authorized Libraries***

This program authorization is accepted by the system only for certain "authorized" system libraries. The libraries automatically considered authorized are SYS1.SVCLIB and SYS1.LINKLIB. However, any library may be designated as an authorized library by inclusion of its name in the SYS1.PARMLIB member IEAAPF00 or IEAPFxx prior to IPL. Such designated libraries are the equivalent of the above system libraries since nonspecific system requests for routines that will execute with system key 0-7, supervisor, or APF authorization can be satisfied by a module loaded from any authorized library. It is the installation's responsibility to control the contents of such libraries.

**Job Step Authorization:** Authorization of the first program executed in the job step determines the authorization of the job step:

- If the first program is not APF-authorized, no part of the job step is APF-authorized.
- During the execution of an APF-authorized job step, if a load request is satisfied from an unauthorized library, the task will be abnormally terminated. It is the responsibility of the authorized program not to recover from the ABEND in a way that would allow the unauthorized module to execute.

**Note:** Because an authorized program normally executes as a job step, executable by any user, it cannot control the identity of JOBLIBS, STEPLIBS, and so forth, since these libraries are identified via JCL. Therefore, if two modules of the same name exist on different authorized libraries, an authorized program attempting to load one of these modules could get the other if the user executing the authorized program were to misuse the two libraries in question. Thus, the installation must ensure that duplicate module names are not be permitted across authorized libraries.

## 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 CPU is in the problem state. When bit 15 is zero, the CPU 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 CPU 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 type 1 SVC issuing MODESET
- Key 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 CPUs sharing the same control program. Those functions executing on one or more of the CPUs which require a CPU or program action on one or more of the CPUs 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 CPUs.



## ***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 which require the modification or sensing of the physical state of one of the configured CPUs. Ten of the twelve SIGP hardware functions are defined as IPC direct services, and are accessible via the DSGNL macro instruction.

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

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

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

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

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

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

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

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

**CPU Reset:** The specified CPU performs CPU reset. The execution of the reset does not affect other CPUs and does not cause any channels, including those configured to the specified CPU, 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 CPUs. The two remaining SIGP functions are defined as remote services:

**External Call:** An "external call" external-interruption condition is generated at the specified CPU. The interruption condition becomes pending during the execution of the SIGP instruction.

The associated interruption occurs when the CPU is interruptible for that condition. Only one external-call condition can be kept pending in a CPU 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 CPU. The interruption condition becomes pending during the execution of the SIGP instruction. The associated interruption occurs when the CPU is interruptible for that condition. At any one time the receiving CPU can keep pending one emergency-signal condition for each CPU of the multiprocessing system, including the receiving CPU itself. The emergency-signal function is accessible via the RISGNL macro instruction.

### ***Status Conditions***

Eight status conditions are defined whereby the issuing and specified CPUs 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 CPU 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 CPU because of a previously issued SIGP instruction. The condition exists from the time an external-call function is accepted until the resultant external interruption has been completed. The condition may be due to the issuing CPU or another CPU. The condition, when present, is indicated only in response to sense and to external call.

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

**Operator Intervening:** This condition exists when the specified CPU 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 CPU 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 function.

**Not Ready:** This condition exists when the specified CPU 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 CPU decodes an unassigned or uninstalled function code.

**Receiver Check:** This condition exists when the specified CPU 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 CPU.

For more details on the SIGP instruction, see **IBM System/370 Principles of Operation**.

## 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 that is given control when the error condition is detected will receive 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 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 Scheduler and Supervisor Logic** for details.) Since SPOST invokes the PURGEDQ SVC see 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, it is necessary to 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) of the ECB to the posted condition should be issued. At this time, if the wait bit is on, 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 9 demonstrates an example of this use of POST.

---

	L	RX,ECB	Get contents of ECB.
	L	RY,='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.
	CS	RX,RY,ECB	Compare and swap completion bit and code.
	BZ	POSTDONE	Branch if CS is successful.
DOPOST	POST	ECB	
POSTDONE	EQU	X	

---

Figure 9. 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 will perform 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. The local lock will be freed by EVENTS. 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 will store 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=YES is not specified, control will later be returned to the caller. The local lock will not be freed by EVENTS.

## Interruption, Termination, and Dumping Services

The supervisor offers many services to assist in the detection and processing of abnormal conditions during the execution of the system. Certain types of abnormal conditions are detected by the hardware and cause program interruptions to occur (for examples, if an attempt is made 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 which is not defined to the system causes an ABEND to be issued by the Open routine).

Although the supervisor provides facilities for writing exit routines to handle specific types of interruptions and abnormal conditions, the supervisor 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.

The services discussed in this chapter include:

- recovery/termination (CALLRTM macro instruction)
- functional recovery routines (SETFRR macro instruction)
- program interruption processing (SPIE macro instruction)
- virtual storage dumping (SDUMP macro instruction, CHNGDUMP command)
- task recovery (STAE and ESTAE macro instructions)
- installation-written clean-up routines

### Recovery/Termination

The recovery/termination manager (RTM) monitors the flow of control through the recovery and termination processes for system, address space, and task failures. Its purpose is to select the appropriate process according to the status of the system and the requests of its invokers.

Some RTM processing is performed on behalf of the caller via the ABEND macro instruction. Other processing is directed to tasks or routines other than the caller via the CALLRTM macro instruction. The ABEND macro instruction is available for use to anyone and is documented in **OS/VS2 Supervisor Services and Macro Instructions**. The CALLRTM macro instruction is restricted to key 0 privileged routines and is described in this publication.

### Task Termination

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 save 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).

### ***Address Space Termination***

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 save 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 save area in register 13.

Since the MEMTERM process circumvents all task recovery and task resource manager processing, its use is restricted to a select group of routines which 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 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.

### **Processing Functional Recovery Routines**

The system recovery routines of the RTM control the software recovery for supervisor functions, SRB functions, and all other locked or physically disabled functions. It provides the interface between the control program and the functional recovery routines (FRRs) defined to recover the control program.

The FRRs are established to protect these global, SRB, and local functions, and are identified to the recovery termination manager via the SETFRR macro instruction. When an FRR is scheduled, it will run in the state of the system (enabled or disabled) as governed by the locks that were held at the time of the error, or as modified by previously scheduled FRRs.

Each FRR is defined on one of several recovery stacks in the system, each stack providing a unique recovery environment. When errors occur in a path through supervisory functions, the system recovery manager uses the recovery stack associated with that path to initiate recovery by routing control to those FRRs defined on that recovery stack. Through this routing capability, the system recovery manager provides a critical function in providing recovery from errors which formerly necessitated re-IPLs of the system.

## ***SETFRR Macro Instruction***

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 processing. The LIFO stack is maintained by the recovery termination manager and contains the addresses of the FRRs that protect a single path through supervisor control and SRB code. The stack is considered part of the path's operating environment and is therefore saved and restored by the supervisor control function on interruptions and redispaching.

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.

**Note:** The size of the FRR stack 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 will be 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.

## **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 the programmer can control his own page fault processing.

If interruption 17 is to be 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.

## **Dumping Virtual Storage**

The SNAP and ABEND macro instructions can be used to request dumps, and can be issued by any user. (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.

The SVC DUMP routine can schedule a dump in the address space specified by the ASID parameter of SDUMP. If the user cannot issue an SVC, this service can be initiated by a branch entry, if desired. If the branch entry is used, the branch entry caller must be key 0, supervisor state, and must be in SRB mode, or own a lock, or be disabled (with supervisor bit on). The branch entry interface is by standard linkage conventions. Branch entry callers must issue the CVT mapping macro instruction with the PREFIX=YES parameter.

The SVC DUMP routine with entry by SVC can also schedule a dump from the address space from which it is issued. The caller of this service only must be authorized via APF or have a control program protection key.

### **SQA Buffer**

A 4K buffer is reserved in the system queue area for the callers of SVC DUMP. A user may 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 which 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. Since the buffer is for use for all callers of SVC DUMP, it must be treated as a serially reusable resource. The high order bit of CVTSDBF must be checked prior to using the buffer. 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.

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.

### **Using the CHNGDUMP Command**

The CHNGDUMP operator command is normally not scheduled by the operator, but is scheduled by the system programmer to override the current system dump options for SYSABEND, SYSUDUMP, and SDUMP dumps. The system obtains its normal dump options as follows:

- *SDUMP requests* -- the options in the system are the ones indicated in the SDUMP parameter list.
- *SYSABEND or SYSUDUMP user dump requests* -- the system obtains its dump options by merging the dump options in the IEAABD00 or IEADMP00 member of SYS1.PARMLIB with the options indicated on the DUMPOPT parameter of the ABEND, SETRP, and CALLRTM macro instruction.

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 recover 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 merge with the parmlib options to produce the combination of options for the dump the system actually takes.

The CHNGDUMP command can override any or all of these options, regardless of their level of origin. There are two major parameters for this command, as follows:

- SET -- to override the existing dump options.
- DEL -- to delete any or all overriding dump options set by a previous CHNGDUMP SET command.

The overriding options indicated with the SET parameter are the only ones that the system will use for all subsequent jobs and tasks for the life of the IPL, or until they are nullified with the DEL parameter in a subsequent CHNGDUMP command.



## Task Recovery

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 specify perform pre-termination processing, or specify a retry address if he wishes to prevent the termination.

This section begins by explaining the processing of the STAE/STAI exit and retry routines. It then continues by providing the extended capabilities of the ESTAE routines. It concludes with a discussion of ESTAE processing.

### ***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 is given control and executes under a program request block created by the SYNCH service routine. Only one STAE exit routine will receive 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 will receive control. This will continue 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 will become active again when the current STAE environment is canceled.

A STAE environment is canceled when the RB which created it goes away (unless it issues XCTL L 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 will have to reissue the STAE macro instruction if it wants continued STAE protection.

A STAI environment is canceled only if the task completes or if it requests that termination continue and no further STAI processing be done. In the latter case, all STAI exits for the task are canceled.

## 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 which 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 parameter list specified on the STAE macro instruction or the STAI parameter of the ATTACH macro instruction.

Upon entry to the STAE routine, parameter registers are as follows:

If an SDWA was obtained:

register 2	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 the control program 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 will not be 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 the control program is preparing to give control to the STAE routine.

### ***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 the 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 closing 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 will occur. However, certain effects on the system will not be canceled by this RB purge. Example of these effects are as follows:

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

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.

The retry routine should use the FREEMAIN macro instruction to free the 512 bytes of storage occupied by the work area when the storage is no longer needed. This storage should be freed from subpool 0 which is the default subpool for the FREEMAIN macro instruction.

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 Extended Capabilities***

The ESTAE service routine provides the following increased capabilities over STAE:

- The BRANCH option to improve performance for SVCs by providing a means of establishing an ESTAE environment without the overhead of an SVC interruption.
- The TERM option to allow ESTAE exits to be scheduled for clean-up processing under certain instances for which STAE exits did not get control.
- Changes in PURGE and ASYNCH defaults to the most commonly used options.

## The BRANCH Option

ESTAE (or SVC 60) can be branched to by SVCs by specifying BRANCH=YES on the ESTAE macro instruction. Following are some guidelines that must be followed when specifying BRANCH=YES:

- Only an SVC running under an SVRB in key 0 supervisor state may specify BRANCH=YES.
- BRANCH=YES may be used by a given SVC only once to create an ESTAE exit. That is, if an SVC wishes more than one ESTAE exit, it must issue ESTAE with BRANCH=NO to get additional exits.
- If BRANCH=YES is used a second time with the create CT option, the ESTAE exit that was previously established by the earlier BRANCH=YES entry will be overlaid by the new options. This exit will be overlaid even if subsequent exits have been created with the ESTAE SVC.  
If subsequent exits had been created with the ESTAE SVC, the exit that was just overlaid will become the latest exit.
- If BRANCH=YES is used with the OV option, the most recently created exit at that RB will be overlaid. If the most recent exit is not an ESTAE exit or if there are not exits in effect for the user's SVRB, the create option will be assumed and the rules governing it will be in effect if an exit address is passed. If no exit address is passed, no exit will be created. In either case, a nonzero return code will be returned to the caller.
- If BRANCH=YES is specified and cancel is requested, the most recently created exit at that B will be canceled. If no exits exist at that RB or if the most recent exit is not an ESTAE exit, a nonzero return code will be returned to the caller.
- The user of BRANCH=YES must be running in key 0, supervisor state, under an SVRB, and must be holding the local lock and only the local lock. The user must also pass the address of a 72-byte savearea in register 13. Also, an expansion of the CVT must be in each module that issues ESTAE with BRANCH=YES.

## The TERM Option

TERM specifies that the exit routine associated with ESTAE will or will not be scheduled in certain instances, in addition to normal ESTAE processing. The instances to which TERM applies are enumerated under the explanation of the TERM parameter of the ESTAE macro instruction in *OS/VS2 Supervisor Services and Macro Instructions*, GC28-0683.

## The PURGE and ASYNCH Parameters

If the PURGE parameter is not specified on the ESTAE macro instruction, PURGE=NONE will be assumed (rather than PURGE=QUIESCE as on the STAE macro instruction).

If the ASYNCH parameter is not specified on the ESTAE macro instruction, ASYNCH=YES will be assumed (rather than ASYNCH=NO as on the STAE macro instruction).

ESTAE processing at time of errors provides the following increased capabilities over STAE:

- Each ESTAE exit established by a task is eligible to receive control. If an error occurs, the most recently created ESTAE will be entered. If it requests that termination continue or if it fails itself, the next ESTAE exit, if any, will be entered. This will continue 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 will receive control.
- STAI and ESTAE exits will receive control after all ESTAE exits and one STAE exit, if any, have been processed.

- The work area (SDWA) passed to an ESTAE exit contains additional information including the resume PSW and general purpose registers of the program whose ESTAE exit is now in control at the time the program was interrupted.
- An ESTAE exit may indicate via an ESTAE parameter what areas of storage should be dumped to the SYSABEND or SYSUDUMP data set.
- An ESTAE exit may request via an ESTAE parameter that the control program free the SDWA instead of freeing it itself in a retry routine.
- An ESTAE exit may specify via an ESTAE parameter what the contents of its registers should be on entry to the retry routine.
- An ESTAE exit remains in effect when its retry routine receives control. It need not reissue the ESTAE to reestablish itself.

### ***ESTAE/ESTAI Exit Routines***

The ESTAE macro instruction identifies recovery routine address to the control program. This recovery routine is associated with the task and the RB which issued ESTAE. 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.

Figure 10 provides an example of an ESTAE environment. In this example, the queueing structure of the ESTAE routines is demonstrated, and the propagation of ESTAI to lower-level subtasks is shown.

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 address should be given control 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 will be performed only for the first exroutinet selected. Subsequent routines will receive an indication of the I/O processing previously performed, but no additional I/O processing will be performed. The asynchronous processing request, however, will be performed for each routines.

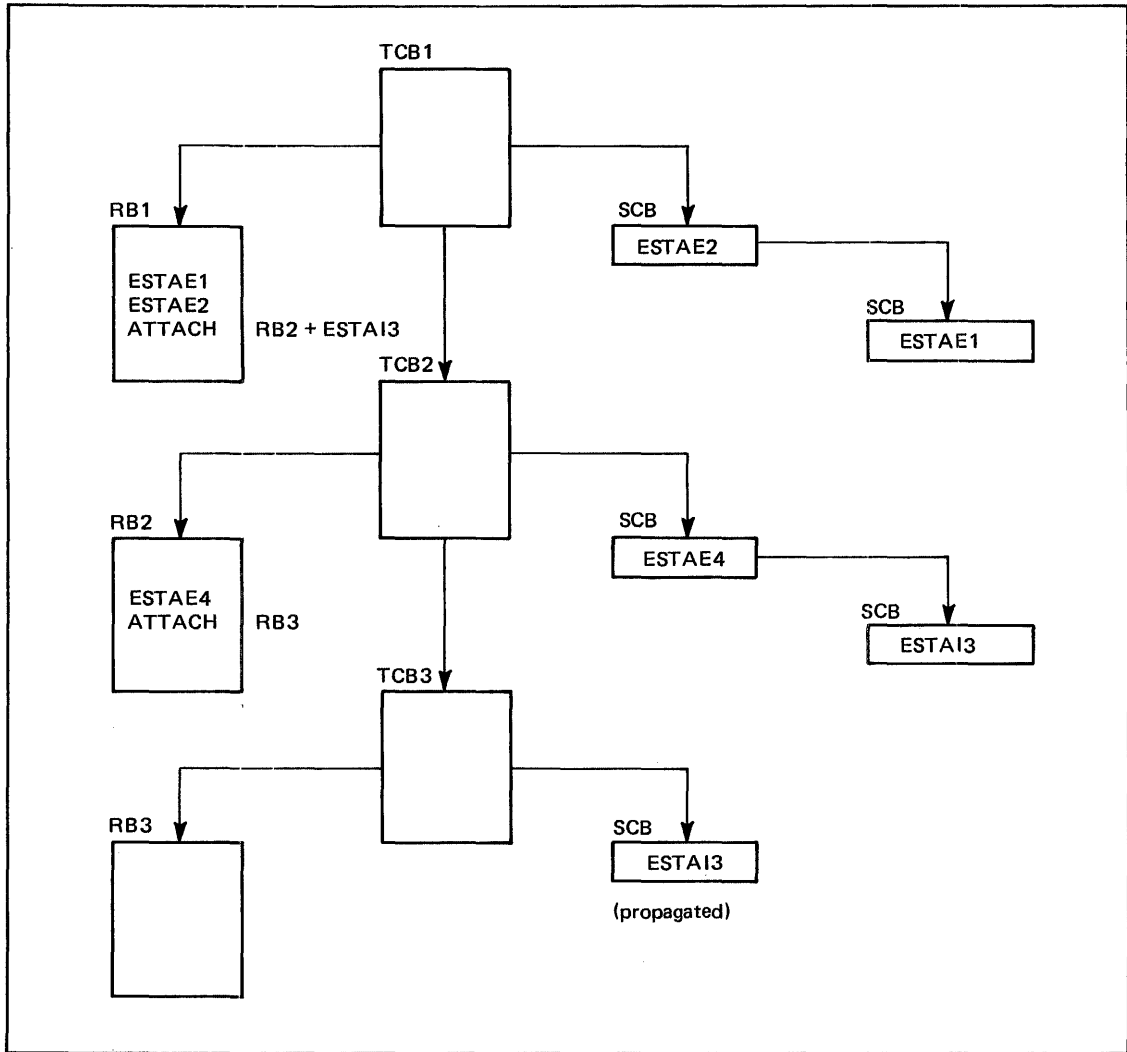


Figure 10. ESTAE Environment

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

- ESTAE or ATTACH may be issued by the ESTAE/ESTAI exit routine.
- ESTAE exits are removed only in the following situations:
  - A return code from the exit routine indicates termination.
  - The exit routine fails.
  - The RB which created the ESTAE exit is removed.
  - ESTAE 0 is issued by a program running at the same RB as the ESTAE exit to be removed.
  - An RB undergoing XCTL processing has XCTL=NO specified for the exit.
- If an ESTAE/ESTAI exit requests termination or fails, the following occurs:
  - All other ESTAE exits established by that task are scheduled in LIFO order. When these are exhausted, the ESTAI exits are scheduled.
  - Dump options will be accumulated, if they exist.

- The asynchronous exit indicator will be reset according to the new exit's request.
- I/O options for the new exit will be ignored.
- A new SDWA will be initialized.
- The new exit will be 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 will be entered.
- ESTAE exits receive control with the same status (supervisor or problem 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 will occur:

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

### **Interface to an ESTAE/ESTAI exit**

Prior to going to an ESTAE/ESTAI recovery routine, the control program attempts to obtain and initialize a work area which 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 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.

When the ESTAE/ESTAI routine has completed, it should use the SETRP macro instruction to inform the control program of the actions it desires. This macro instruction will initialize the SDWA with these options.

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

When the ESTAE or ESTAI routine has completed, its processing, it should rerun to the control program via the contents of register 14. Register 15 should contain one of the following return codes:

return code	action
0	Continue the termination. Any ESTAE exits established prior to this one will receive control.
4	A retry routine is to be scheduled, and its address is placed in register 0.
16	Termination should be continued. No further ESTAE/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 which 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 will occur. However, certain effects on the system will not be canceled by this RB purge. Example 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.



## ***ESTAE/ESTAI Retry Routines***

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

- A dump will be given if requested.
- FREEMAIN of SDWA will be performed, if requested.
- Registers will be updated with user-supplied values, if requested.
- The RB queue will be 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 SWDA was obtained, the user has a choice of interfaces to his retry address. The user can set (in the SWDA) 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 most often used if a special retry routine is to get control.

The parameter registers are as follows:

If no SWDA 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 SWDA was obtained and the exit did not request register update, or freeing of SWDA:

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 SWDA 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

## 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 IGC000IC in SYS1.LPALIB.

Initially, the CSECT consists of four 12-byte entries of all zeros. Each of the first three 12-byte entries is to 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 is to 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.



## Virtual Storage Management

The GETMAIN and FREEMAIN macro instructions allocate and free one or more areas of virtual storage. Although most of the functions of these macro instructions are available to any user, several new functions are available only to programs executing in supervisor state under protection key zero.

One of the new functions permits a branch entry to the GETMAIN and FREEMAIN routines, rather than an SVC entry. Although use of this function requires additional work on the part of the user, the branch entry is significantly more efficient than the SVC entry and does save some system overhead. This function is provided via the BRANCH parameter.

Another function, available only with the branch entry function, allows the user, executing in key zero, to specify the actual key in which the requested storage is to be obtained. This function is provided via the KEY parameter.

### *The BRANCH Parameter*

Branch entry is accomplished by specifying BRANCH=YES on the GETMAIN or FREEMAIN macro instruction. 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 will be 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 will be considered an error.

### *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 otherwise 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.) 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 will be performed before the PGFIX function.

If you specify PGFREE with RELEASE=Y, the PGFREE function will be performed and those pages of the virtual subarea with zero fix counts will be 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 will sometimes complete 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 service functions: PGFIX, PGFREE, PGLOAD, PGRLSE, and PGOUT. The list consists of one or more doubleword entries, each entry describing an area of virtual storage. The list must be nonpageable (for example, in SQA or LSQA) and contained in the address space to be processed.

Each parameter list entry has the format shown in Figure 11.

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

Figure 11. 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. 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.
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.



## Miscellaneous Services

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 will be routed as specified in the ROUTCDE parameter.

The bit definitions for MSGTYP=Y are:

Bit 0	MONITOR JOBNAMES
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 will be 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 will not be sent anywhere; however, a WTOR message will be sent to the master console. The routing codes and REG0 MCSFLAG field, if present, will be 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 multiple-line message with more than one WTO request.

The first WTO request supplies the first lines of the message. Other WTO requests can then add lines to the message. The additional lines would appear at the end of the message, and would continue until an 'END' line is added.

For the first request in a multiple-line sequence, the leftmost three bytes of register zero must be zero; you must ensure that this is done.

The first request receives a message identifier back in register 1. To add more lines, the next multiple-line request must have this identifier in the leftmost three bytes of register zero.

If the conditions on register zero are not met, it would appear to SVC 35 (WTO) that the multiple-line request is adding lines to an existing message, and the new message will not be created.

## **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 which are used to determine the destination of the message. There are fifteen routing codes that can be used. When the message is ready to be routed, 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 prior to the routing of messages so users can examine the message text and modify the message's routing and descriptor codes. The system will use the modified routing codes to route the message. Descriptor codes provide a mechanism for message presentation and deletion, and are explained later in this chapter.

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 may wish to record a selected group of messages, or it may wish to 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 routine 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 will receive 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.
- IECEVXIT 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.
- Exit routine is part of WTO SVC. If the exit routine terminates abnormally, the WTO request will be terminated.
- Exit from the routine is via the RETURN macro instruction.
- Format of text and codes is:

Message flagging and 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 Usage
Bit 5	Routing code 14	Available for Customer Usage
Bit 6	Routing code 15	Available for Customer Usage
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 MONITOR command response
Bit 1	Descriptor code 10	Dynamic Status Displays
	11 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, that is, a response to a DISPLAY A command. Routing of these messages is 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 will be part of the control program. If a message's routing code field is used by the operating system to route the message, the routine will receive 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 may examine but not modify the message text.

A message will be sent to only 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 since 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 will be 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 IEVCVXIT 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 IEVCVXIT in SYS1.AOSC5 with the WTO/WTOR exit routine. The linkage editor should be instructed to load IEVCVXIT 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 IEECVXIT in the SYS1.LPALIB with the user-written WTO/WTOR exit routine.

## Service Management

Service Management facilities provide a basic set of system services which 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 new support are:

- A new control block, called a service request block (SRB), which represents a service request. This block, like a TCB, identifies a unit of work to the dispatcher. The block is smaller than a TCB, and requires less information to be specified about each request.
- The SCHEDULE macro service, which enters SRBs into the dispatchable work queue with a minimum of overhead.
- Changes to the dispatcher to operate under this new service request control structure in addition to the old task structure. Maximum performance is provided when dispatching service requests. 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.

Service management provides a system-wide dispatching facility which can be used to increase parallel processing on multiple CPUs. The programs benefiting the most from these facilities are those which 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 CPUs at a higher priority than the task work.

The service management facilities also provide a mechanism which is used for almost all communications between address spaces, and 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 interruption and schedules a service request block (SRB). The interruption handler can then start I/O requests which were waiting for the I/O path and accept any additional pending interruptions. By delaying complete processing of the interruption, this approach 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 CPU, not just the one which took the interruption
- to process the interruption enabled except where specific serialization through locks is used

Yes!

- 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 will dispatch it in the address space specified by the SRB.

### **Service Request Blocks (SRBs)**

Service requests are represented by service request blocks (SRBs). These control blocks are supplied by the function requesting a service. The basic information contained in an SRB is:

- The address space in which the asynchronous routine will be dispatched.
- The protection key assumed by the routine.
- The entry point of the routine.
- The address of the resource manager termination routine responsible for cleaning up the SRBs which have been scheduled but not yet dispatched.
- The address of a status save area.
- A fullword parameter to be loaded into register 1 on entry to the routine.
- CPU affinity requirements.
- The level (system or nonquiesceable) at which the request is to be scheduled.
- The ASID/TCB associated with the routine.
- A forward chaining field used to queue the request.

### **Priorities**

Services may be scheduled for execution at either global or local priorities. Service requests queued at the global level are given a priority which 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 will be 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; the other (called the nonquiesceable level) is for specialized functions necessary to perform a quiesce of SRBs.

Service requests at the nonquiesceable level continue to be dispatched while the address space is in the process of being quiesced. Requests queued at the system level in an address space will not be dispatched while the quiesced status is in effect. Usage 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. On entry, the address of the SRB is in register 0 and the exit address is in register 14. The routines cannot enter problem program mode and must establish a recovery environment.
- The routines may not issue SVCs. (However, ABEND may be issued.)
- The routines are nonpreemptible -- 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 control is given up 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 will be redispached when the problem is resolved.
- The routines may take page faults. Page faults encountered in the unlocked state will be handled per service request; page faults encountered in the locked state will prevent other processing which requires the lock from proceeding until the page fault is resolved.
- The routines may request any lock through the SETLOCK interface.
- 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 abnormally, 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. 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 will search 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 which have not yet been dispatched will be affected as PURGEDQ will not wait for SRBs already dispatched but not completed. If this parameter is omitted, the current address space will be 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 will occur 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 will produce unpredictable results.

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

## Creating Interruption Request Blocks

The CIRB macro instruction causes the Stage 1 Exit Effector routine to create and initialize an interruption request block (IRB). The IRB is used to control an asynchronous user exit routine requested by the caller.

The Stage 1 routine obtains a work area in which the caller may construct interruption queue elements (IQEs), and obtains a register save area in which the user exit routine may later save the registers of the requesting program. The routine obtains space for the IRBs (and IQEs) and the work area from local supervisor queue space. The work area follows and is contiguous to the IRB. The register save area, if requested, is obtained from subpool zero of the user program's region, and is therefore not contiguous to either the IRB or its work area.

After the storage for the IRBs and optional work and storage areas is obtained, the Stage 1 Exit Effector routine initializes the IRB. The initialization is accomplished according to the flag bits passed to the routine in register 1.

The information initialized in the IRB includes the save area address, the size of the IRB, the entry point of the user exit, the PSW to be loaded to start execution, and a series of flags communicating actions to be taken when the asynchronous exit routine terminates.

For details on the Stage 1 exit effector, and for information on the Stage 2 and 3 exit effectors, see *OS/VS2 Scheduler and Supervisor Logic*.

## 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. Four types of SVC routines are supplied with VS2, and the programming conventions for each type differ. The general characteristics of the four 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 the 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 and 2 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.
- Loads -- Type 3 SVCs have one load; type 4 SVCs may have multiple loads. Since type 4 SVCs require XCTL overhead, and since there are no size limitations for type 4 SVCs, type 4 SVCs have no advantage over type 3 SVCs. Unless an SVC is to be used in another system as well as in MVS, type 4 SVC routines should never be used.
- Serialization -- In MVS, locking has generally replaced CPU disablement as the technique for serializing multiple CPU 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. You must define which locks are to be acquired for your SVC routines during system generation. (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 four types of SVC routines are summarized in Figure 12. 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	Reference Code
Part of resident control program	Yes	Yes	No	No	
Size of routine	Any	Any	Any	Any	
Reenterable routine	Yes	Yes	Yes	Yes	1
Refreshable routine	No*	No*	Yes	Yes	2
Locking requirements	Yes	No	No	No	3
Entry point	Must be the first byte of the routine or load module, and must be on a doubleword boundary				
Number of routine	Numbers assigned to your SVC routines should be in descending order from 255 through 200				
Name of routine	IGCnnn	IGCnnn	IGC00nnn	IGCssnnn	4
Register contents at entry time	Registers 3, 4, 5, 6, 7, and 14 contain communication pointers; registers 0, 1, 13, and 15 are parameter registers				5
May contain relocatable data	Yes*	Yes*	No	No	6
Supervisor request block (SVRB) size	No SVRB exists	200	200	200	7
May issue WAIT macro instruction	No*	Yes	Yes	Yes	8
May issue XCTL macro instruction	No*	No*	No*	Yes	9
May pass control to what other types of SVC routines	None	Any	Any	Any	10
Type of linkage with other SVC routines	Not Applicable	Issue supervisor call (SVC) instruction			11
Exit from SVC routine	Branch using return register 14				12
Method of abnormal termination	ABTERM	ABEND			
Recovery	FRR	ESTAE or FRR			13

Note: Reference code does not apply to items marked with \*.

Figure 12. 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"> <li>• 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.</li> <li>• Types 2, 3, and 4 may also request locks via the SVCTABLE macro instruction or may obtain them dynamically.</li> <li>• 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.</li> <li>• 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 0, enabled mode.</li> <li>• If disabled spin locks are held or obtained, the SVC routine executes in supervisor state, key 0, disabled mode. No SVCs may be issued.</li> <li>• 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.</li> <li>• An FRR may be defined for any SVC routine that holds or obtains locks to provide for abnormal termination (see reference note 9).</li> </ul>
4	all	You must use the following conventions when naming SVC routines: <ul style="list-style-type: none"> <li>• Types 1 and 2 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.</li> <li>• Type 3 must be named IGC00nnn; nnn is the signed decimal number of the SVC routine.</li> <li>• 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.</li> </ul>
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 register is as follows:</p> <ul style="list-style-type: none"> <li>• Register 0 and 1 contain the same information as when the SVC routine was called.</li> <li>• Register 2 contains unpredictable information.</li> <li>• Register 3 contains the starting address of the communication vector table.</li> <li>• Register 4 contains the address of the task control block (TCB) of the task that called the SVC routine.</li> <li>• 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 SVC routine is in control, register 5 contains the address of the last active request block.</li> <li>• Register 6 contains the entry point address.</li> <li>• Register 7 contains the address of the address space control block (ASCB).</li> <li>• Registers 8 through 12 contain unpredictable information.</li> <li>• Register 13 contains the same information as when the SVC routine was called.</li> <li>• Register 14 contains the return address.</li> <li>• Register 15 contains the same information as when the SVC routine was called.</li> </ul> <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	3,4	Since relocatable address constants are not relocated when a type 3 or 4 SVC routine is loaded into virtual storage, you cannot use them in coding these routines; nor can you use macro instructions whose expansions contain relocatable address constants. Types 1 and 2 are not affected by this restriction since they are part of the resident control program.
7	2,3,4	The SVRB is no longer extendable, but is a fixed size of 200 bytes. When a type 2, 3, or 4 SVC routine receives control, register 5 contains the address of the SVRB within this 200-byte area. This SVRB contains a 40-byte "extended save area." In addition, an area is provided for a STAE control block (SCB); this SCB will be used by the recovery termination manager when an ESTAE or ESTAI is issued within an SVC routine.
8	2,3,4	You can issue the WAIT macro instruction if you hold no locks. You can issue WAIT macro instructions that await either single or multiple-events. The event control block (ECB) for single-event waits or the ECB list and ECBS for multiple-event waits must be in virtual storage.
9	4	When you issue an XCTL macro instruction in a routine under control of a type 4 SVRB, the new load module must be located in the fixed or pageable link pack area.  The contents of registers 2 through 13 are unchanged when control is passed to the load module; register 15 contains the entry point of the called load module.
10	all	No SVC routines except ABEND may be called if locks are held. ABEND may be called at any time.
11	all	No locks may be held. If locks are held, branch entry to SVCs is acceptable, or the locks may be freed, the SVC issued, and the locks reobtained.
12	all	Branch using return register 14 should be used. Otherwise, if locks are held, SVC 3 results in abnormal termination.
13	all	If an SVC routine is entered with a lock held or if an SVC routine obtains a lock, it should specify a functional recovery routine (FRR) for as long as the lock is held (see SETFRR macro instruction). The FRR receives control if an error occurs, and ensures the validity of the data being serialized by the lock; the FRR either recovers or releases the lock and continues with termination.  If no FRR is specified, the recovery termination manager will release 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.

## Missing Interruption Handler

The missing interruption handler checks whether expected I/O interruptions occur within a specified period of time. If the interruptions do not occur, the operator is notified so that steps can be taken to correct the situation before system status is harmed.

The missing interruption handler checks for pending device ends, channel ends, DDR swaps, and MOUNT commands. When a pending condition is found, the condition is indicated in the UCB of the device. After the specified time elapses, another check is made for the pending condition. If the condition is still pending, a message is issued informing the operator what condition is pending and what operator action is required.

### *Establishing a Time Interval*

The IBM-supplied CSECT IGFINTVL provides a time interval of 3 minutes. If the interval is not changed or is incorrectly modified, an interval of 3 minutes is assumed.

To change the time interval, specify the desired interval in a modification of the CSECT IGFINTVL as indicated in Figure 13.

---

```
//MODIFY    JOB      MSGLEVEL=( 1, 1 )
//STEP      EXEC     PGM=AMASPZAP
//SYSPRINT  DD       SYSOUT=A
//SYSLIB    DD       DSNAME=SYS1.LPALIB,DISP=OLD
//SYSIN     DD       *
              NAME    IGFIDIO  IGFINTVL
              VERIFY  0000  F0F0F0F3
              REP     0000  F0F0FXFX
```

---

**Figure 13. Changing the Missing Interruption Handler Time Interval**

In the REP statement in Figure 13, FXFX represents the time interval, and must be replaced within the range F0F1 to F9F9. If F0F0 is specified, the time interval defaults to 3 minutes.

If a new time interval is specified, it does not take effect immediately. The new time interval becomes effective only after the next IPL of the system, and only if the CLPA parameter is specified at the IPL.

**Note:** The CSECT IGFINTVL is eight bytes in length, and initially contains the character string C'00030000', where 3 indicates the time interval.

## **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 will be executed 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. Insert your code in the machine check handler appendage ICFBDF00.

**Adding code to the master scheduler initialization module --** You can add code that will execute when a warn data set contains an image of real storage and the system operator chooses FORM during system IPL. Your code will execute just before the warn data sets are erased and reformatted. Insert your code in the Power Warning Feature Support Initialization module ICFBIF00.

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

**Note:** Your code must replace instructions, in either of the routines, that are bracketed with asterisks and specifically identified with comments. 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 which will be executed 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 transfer control to the dump routine or have control returned to the machine check handler for the system to continue processing.

As shown in Figure 14, 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 your code:

- Obtain the source code for ICFBDF00.
- Replace the code that is bracketed with asterisks with your code.
- Reassemble the modified ICFBDF00.
- Re-linkedit the nucleus data set SYS1.NUCLEUS (IEANUC01), replacing the old ICFBDF00, with your modified one.
- Re-IPL your system.

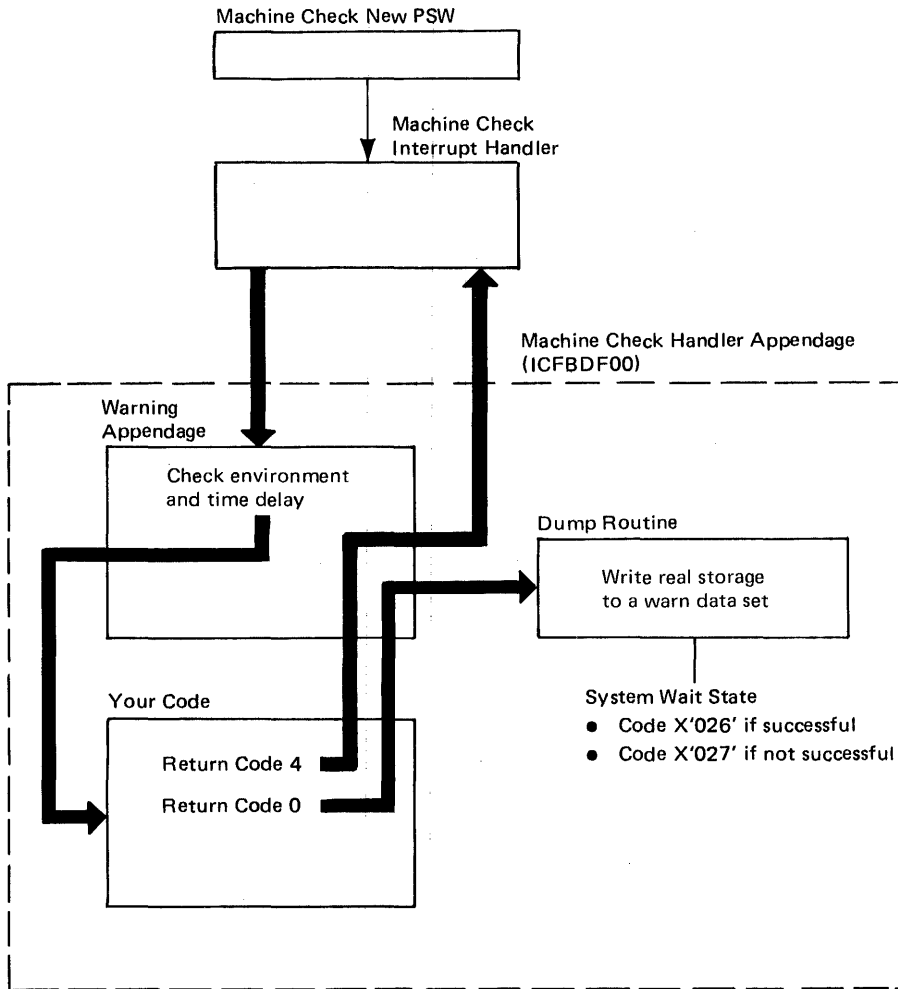


Figure 14. Logical Placement of Your Code in the Machine Check Handler Appendage

### 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 will be entered:

- Key zero.
- Disabled for all interrupts.
- In supervisor state.
- In extended control mode with Dynamic Address Translation (DAT)  
(No DAT in MVS)

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 will be executed by only one CPU.

### ***Writing Code for the Master Scheduler Initialization Routine***

You can write a routine that will execute, 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 15, during IPL, if the warn data set contains information from real storage, the system operator can choose to respond either REST or FORM. REST will cause real storage to be refreshed with the contents of the warn data set. FORM will cause a transfer of control to your code; after your code returns control, the warn data sets will be erased and reformatted by the Power Warning Feature Support. If you have not inserted any code, FORM will immediately cause the warn data set to be erased and reformatted.

To include you code:

1. Obtain the source code for ICFBIF00.
2. Replace the code bracketed with asterisks, with your code.
3. Reassemble the modified ICFBIF00.
4. Inserted your modified ICFBIF00 as follows:
  - If you have not generated your system -- replace the old ICFBIF00 on DLIB with your modified ICFBIF00 and generate your system.
  - If you have generated your system -- re-linkedit your modified ICFBIF00 found in DLIB, with the master scheduler initialization module found in SYSI.LINKLIB in the system library. The master scheduler initialization module is: IEEMB860 for MVS, or IEEVIPL for VS2 Release 1.

### Master Scheduler Initialization Routine

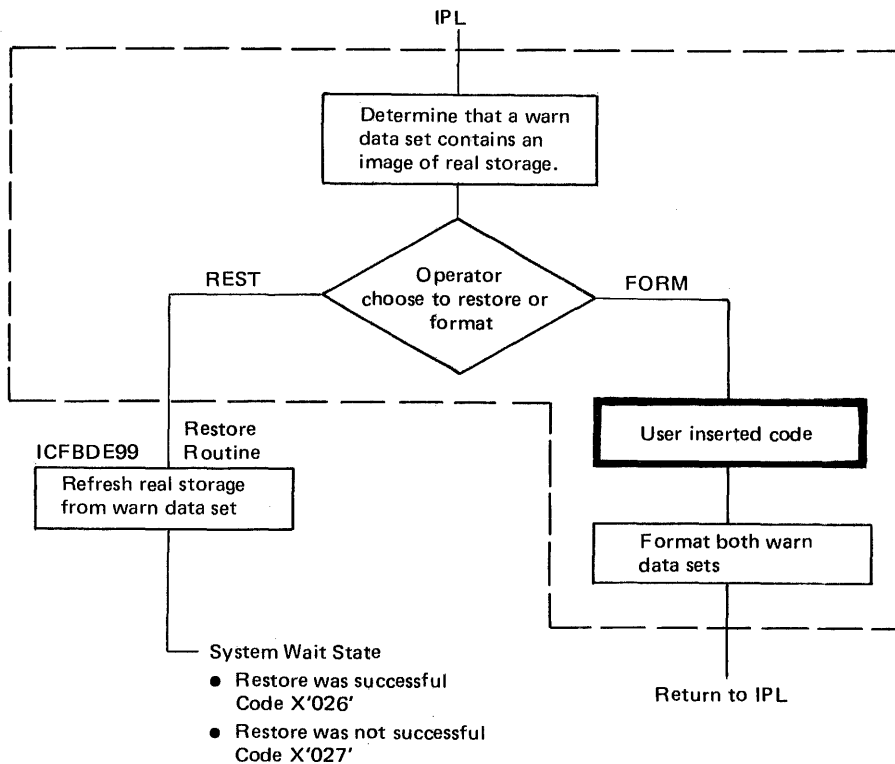


Figure 15. Logical Placement of Your Code in Master Scheduler Initialization Module

### 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 16.

When your code gets control the environment is as follows:

- Your code is pageable.
- Your code is transient.
- Your code can not 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.

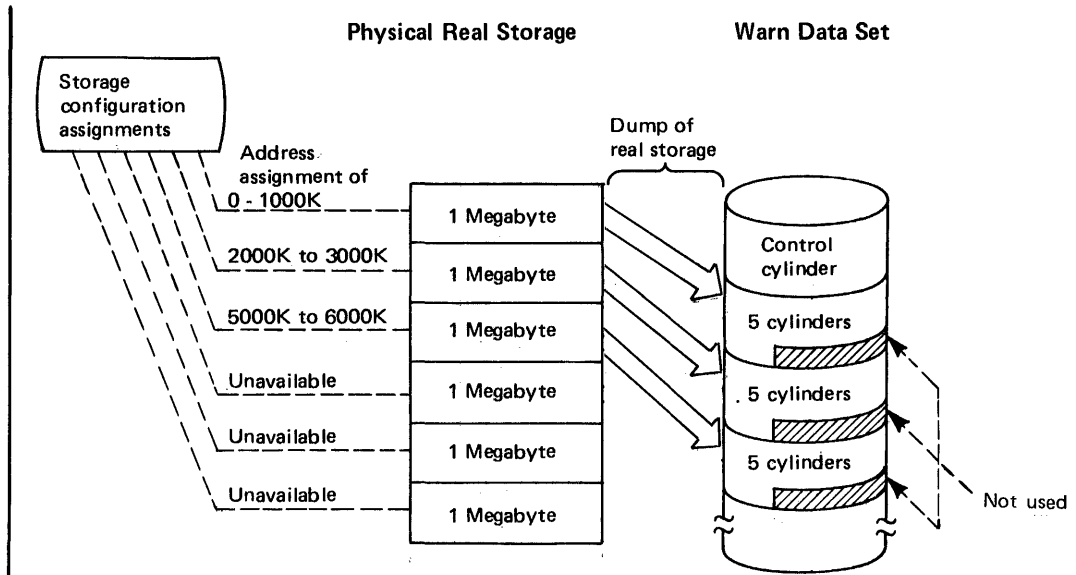
If your system is an M158 or M168 multiprocessor, you must consider that the first address after the last unused address starts a new cylinder on the warn data set. As shown in Figure 17, addresses 0, 2000K, and 5000K start new cylinders. An unused address refers to a location in real storage that is 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 PFW 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	ICFTTPC	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 16. 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 17. Storage Assignments on MP Systems



## Part II: Reference — Macro Instructions

You can communicate service requests to the control program using a set of macro instructions provided by IBM. Whereas most of the macro instructions have no restrictions on use by applications programmers some of the macro instructions should be restricted in use 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 18 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
CALLRTM	X	
CHNGDUMP	X	
CIRB	X	
DSGNL	X	
EXTRACT	X	
MODESET	X	
NIL	X	
OIL	X	
PGFIX	X	
PGFREE	X	
PURGEDQ	X	
QEDIT	X	
RESERVE	X	
RISGNL	X	
RPSGNL	X	
SCHEDULE	X	
SDUMP	X	
SETFRR	X	
SETLOCK	X	
SPOST	X	
STAE	X	
SYNCH	X	
TESTAUTH	X	
ATTACH		X
DEQ		X
ENQ		X
ESTAE		X
EVENTS		X
FREEMAIN		X
GETMAIN		X
POST		X
SETRP		X
SPIE		X
STATUS		X
WTO		X
WTOR		X

Figure 18. 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 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.

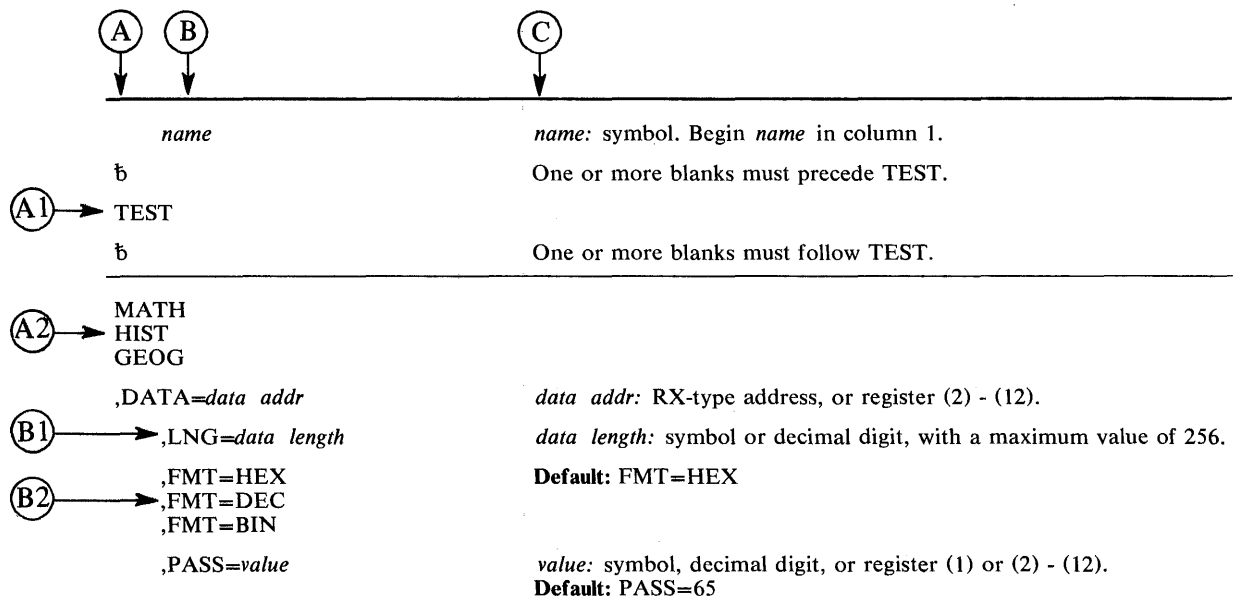
The descriptions of the following macro instructions assume that the standard begin, end, and continue columns are used -- for example, column 1 is assumed as the begin column. To change the begin, end, and continue columns, code the ICTL instruction to establish the coding format you wish to use. If 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 19 presents a sample macro instruction, TEST, and summarizes all the coding information that is available for it. The table is divided into three columns.





**Figure 19. 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), the parameter appearing on one and only one of the lines may be coded if desired.
- 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 alphanumeric 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 a 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 20 shows an example of each method.

---

NAME1	OP1	OPERAND1, OPERAND2, OPERAND3, OPERAND4, OPERAND5, OPERAND6	THIS IS ONE WAY
NAME2	OP2	OPERAND1, OPERAND2, OPERAND3, OPERAND4	THIS IS ANOTHER WAY X

---

**Figure 20. Continuation Coding**

## Descriptions of the Macro Instructions

## **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 parameters. These parameters are restricted in use and should only be used with tasks in supervisor mode, 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 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 ATTACH 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 ATTACH.
ATTACH	
␣	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	<b>Note</b> : <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).
,G SPL= <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	<b>Default</b> : 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> )	
,PURGE=QUIESCE	<b>Note</b> : PURGE may be specified only if STAI or ESTAI is specified.
,PURGE=NONE	<b>Default for STAI</b> : PURGE=QUIESCE
,PURGE=HALT	<b>Default for ESTAI</b> : PURGE=NONE
,ASYNCH=NO	<b>Note</b> : ASYNCH may be coded only if STAI or ESTAI is specified.
,ASYNCH=YES	<b>Default for STAI</b> : ASYNCH=NO
	<b>Default for ESTAI</b> : ASYNCH=YES
,TERM=NO	<b>Note</b> : TERM may be specified only if ESTAI is specified.
,TERM=YES	<b>Default</b> : TERM=NO
,JSTCB=NO	<b>Default</b> : JSTCB=NO
,JSTCB=YES	
,SM=PROB	<b>Default</b> : SM=PROB
,SM=SUPV	
,SVAREA=YES	<b>Default</b> : SVAREA=YES
,SVAREA=NO	
,KEY=PROP	<b>Default</b> : KEY=PROP
,KEY=ZERO	
,DISP=YES	<b>Default</b> : 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).
	<b>Default</b> : 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).
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained below. 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 will obtain 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.

**,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	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.

### **Ignored Parameters**

The following parameters, available with Release 1 of OS/VS2, are ignored if coded in MVS:

```
GIVEJPQ=YES
GIVEJPQ=NO
TSLOGON=YES
TSLOGON=NO
LSQA=n
```

### **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 )
```

## 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 0 supervisor state routines. After execution of the macro instruction, control is returned to the caller.

The CALLRTM 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 CALLRTM.
CALLRTM	
b	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 addr</i>	<i>ASID addr</i> : decimal digits 0 - 32,765 or register (2) - (15).
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : 0, or register (0) or (2) - (12). <b>Note:</b> This parameter may only be specified with TYPE=ABTERM.
,DUMP=YES	<b>Default:</b> DUMP=YES
,DUMP=NO	<b>Note:</b> This parameter may only be specified with TYPE=ABTERM.
,STEP=NO	<b>Default:</b> STEP=NO
,STEP=YES	<b>Note:</b> 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 below:

TYPE = ABTERM

TYPE = MEMTERM

specifies that the services of the recovery/termination manager is 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.

,COMPCOD = *comp code*

specifies the 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 addr*

specifies the address space identifier of the address space to be terminated (for MEMTERM) or the address space identifier 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 is defined by the //SYSABEND or //SYSDUMP 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 by using the list form of the SNAP macro instruction, or a compatible list may be created. The TCB and DCB options available on SNAP will be ignored if they appear in the parameter list; the TCB used will be for the task that received ABEND, the DCB used will be provided by the ABDUMP routine. If a //SYSABEND or //SYSDUMP 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 will take place in the address space identified by the ASID specified in register 5, and will have a completion code of 123.

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

## CHNGDUMP — Change Dump Options

When the system or a user requests a SYSABEND or SYSUDUMP dump, the system uses the dump options specified in the SYS1.PARMLIB member or on the ABEND, CALLRTM, or SETRP macro instructions. When a SDUMP dump is requested, the options are obtained from the SDUMP parameter list.

The CHNGDUMP command is used to override these options.

The SET option of the CHNGDUMP command is written as follows:

---

CHNGDUMP  
CD

␣ . . . . . One or more blanks must follow CHNGDUMP or CD.

---

SET

,NODUMP  
,SDUMP  
,SYSABEND  
,SYSUDUMP

**Note:** If NODUMP is specified, no other parameters may be specified.

,NODUMP  
,Q=*code*  
=*(dump codes)*  
=*(dump codes)*,Q=*code*

**Note:** These parameters are specified only when SDUMP is specified above.  
*code:* YES OR NO.  
*dump codes:* Any combination of the following, separated by commas:

PSA	LSQA	TRT
ALLPSA	RGN	CSA
NUC	LPA	SWA
SQA		

,NODUMP  
,SDATA=*(sys data codes)*  
,PDATA=*(prob data codes)*  
,SDATA=*(sys data codes)*  
 ,PDATA=*(prob data codes)*

**Note:** These parameters are specified only when SYSABEND or SYSUDUMP is specified above.  
*sys data codes:* Any combination of the following, separated by commas:

NUC	SWA	ENQ
SQA	CB	TRT
LSQA	ALLSDATA	

*prob data codes:* Any combination of the following, separated by commas:

PSW	ALLPDATA	JPA
REGS	SA	LPA
SPLS	SAH	ALLPA

---

The parameters are explained below:

SET

specifies that the current system dump options are to be overridden by those options specified in this command.

,NODUMP  
,SDUMP  
,SYSABEND  
,SYSUDUMP

specifies whether SDUMP, SYSABEND, or SYSUDUMP dumps are to be taken. If NODUMP is specified, no dumps are to be taken. If SDUMP, SYSABEND, or SYSUDUMP is specified, see the explanation of the parameters that follow.

,NODUMP

,Q=*code*

=(*dump codes*)

=(*dump codes*),Q=*code*

specifies the parameters that may be specified with SDUMP.

NODUMP specifies that no SDUMP dumps are to be taken.

Q specifies that the system should be placed (YES) or should not be placed (NO) in a quiesced state while dumping the system queue area and the common service area.

The *dump codes* specify the SDUMP areas to be included in the SDUMP dump:

ALLPSA — prefix storage area for all CPUs

PSA — prefix storage area for the dumping CPU

NUC — resident nuclues

SQA — system queue area

LSQA — local system queue area

RGN — address space being dumped

LPA — link pack area for dumping task

TRT — GTF incore trace buffers (when GTF is active)

CSA — common service area

SWA — scheduler work area

,NODUMP

,SDATA=(*sys data codes*)

,PDATA=(*prob data codes*)

,SDATA=(*sys data codes*),PDATA=(*prob data codes*)

specifies the parameters that may be specified with SYSABEND or SYSUDUMP.

NODUMP specifies that no SYSABEND or SYSUDUMP dumps are to be taken.

SDATA specifies the system areas to be included in the SYSABEND or SYSUDUMP dump:

ALLSDATA — all of the following SDATA areas

NUC — resident nuclues

SQA — system queue area

LSQA — local system queue area

SWA — scheduler work area

CB — task control blocks

ENQ — enqueued control blocks

TRT — GTF incore trace buffers (when GTF is active)

PDATA specifies the problem program areas to be included in the SYSABEND or SYSUDUMP dump:

ALLPDATA — all of the following PDATA areas

PSW — program status word

REGS — general registers

SPLS — virtual storage subpools

SA — save area linkage information and back-trace through the save area

SAH — save area linkage information

JPA — job pack area

LPA — link pack area associated with the task

ALLPA — job pack area and link pack area associated with the task

The DEL option of the CHNGDUMP command is written as follows:

---

CHNGDUMP  
CD

␣ One or more blanks must follow CHNGDUMP or CD.

---

DEL

,ALL  
,SDUMP  
,SYSABEND  
,SYSUDUMP

**Default:** ALL

**Note:** If ALL is specified or assumed, no other parameters may be specified.

,ALL  
,Q=*code*  
=(*dump codes*)  
=(*dump codes*),Q=*code*

**Default:** ALL

**Note:** These parameters are specified only when SDUMP is specified above.

*code:* YES or NO

*dump codes:* any combination of the following, separated by commas:

ALLPSA		
PSA	LSQA	TRT
NUC	RGN	CSA
SQA	LPA	SWA

,ALL  
,SDATA=(*sys data codes*)  
,PDATA=(*prob data codes*)  
,SDATA=(*sys data codes*),  
PDATA=(*prob data codes*)

**Default:** ALL

**Note:** These parameters are specified only when SYSABEND or SYSUDUMP is specified above.

*sys data codes:* any combination of the following, separated by commas:

ALLSDATA	LSQA	
NUC	SWA	ENQ
SQA	CB	TRT

*prob data codes:* any combination of the following, separated by commas:

ALLSDATA	SPLS	JPA
PSW	SA	LPA
REGS	SAH	ALLPA

---

The parameters are explained under the SET option of the CHNGDUMP command, with the following exceptions:

DEL

specifies that the dump overrides previously set in the CHNGDUMP command are to be discontinued, and that the system is to use the options obtained from the SYS1.PARMLIB data set, the ABEND, CALLRTM, and SETRP macro instructions, and the SDUMP parameter list. Any options specified on a CHNGDUMP SET command and not specified on a CHNGDUMP DEL command remain in effect.

ALL

specifies that all dump option overrides specified on the CHNGDUMP SET command are to be discontinued. If SDUMP, SYSABEND, or SYSUDUMP is specified, all SDUMP, SYSABEND, or SYSUDUMP overrides, respectively, are to be discontinued.

### Example 1

**Operation:** Place the system in a quiesced state while dumping the system queue area and the common service area. This command applies to all occurrences of SDUMP.

```
CHNGDUMP SET,SDUMP,Q=YES
```

## Example 2

**Operation:** All system areas and only three problem program areas (program status word, general registers, and virtual storage subpools) are to be dumped to the data set defined by the SYSUDUMP DD statement.

```
CHNGDUMP    SET, SYSUDUMP, PDATA=( PSW, REGS, SPLS ), SDATA=ALLSDATA
```

## 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) if one does not already exist for the requesting task. 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 the scheduling of the execution of user exit routines.

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

The parameters are explained below:

*EP=entry name addr*

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

,KEY = PP

,KEY = SUPR

specifies whether the asynchronous exit routine will operate 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 will be executed in problem program (PP) or supervisor (SUPR) mode.

,SVAREA = NO

,SVAREA = YES

specifies whether a 72-byte register save area is to be obtained 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) will be freed by EXIT.

,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 will (YES) or will not (NO) be returned to the available queue when the asynchronous exit terminates.

### ***Ignored Parameters***

The following parameters, available with Release 1 of OS/VS2, are ignored if coded in MVS:

TYPE = IRB

ENABLE = YES

STAB = (RE)

The following parameters are no longer acceptable:

TYPE = TIRB

ENABLE = NO

### **Example 1**

***Operation:*** Create an IRB to be used in scheduling an asynchronous exit. The exit will be scheduled via the IQE interface to Stage 2 Exit Effector, and will receive control in the supervisor state. The IRB will be freed when it terminates. The exit will receive 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 will be used to schedule the routine. The exit will get 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 syntax of the complete DEQ 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 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>	<b>Note:</b> <i>rname length</i> must coded if a register is specified for <i>rname addr</i> .
,	<b>Default:</b> 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	<b>Default:</b> RMC=NONE
,RMC=STEP	<b>Default:</b> GENERIC=NO
,GENERIC=NO	<b>Note:</b> 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). <b>Note:</b> 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 below. 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 will be dequeued, regardless of whether they have control of the resource.

,TCB=*tcb addr*

specifies the address of a fullword on a fullword boundary that contains the address of a TCB on whose behalf the DEQ is to be done. The caller (not the directed task) will be 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 the device.

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 21.




Figure 21. 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.
<b>b</b>	One or more blanks must precede DEQ.
DEQ	
<b>b</b>	One or more blanks must follow DEQ.

---

(	
<i>qname addr</i>	<i>qname addr</i> : A-type address.
,	
<i>rname addr</i>	<i>rname addr</i> : A-type address.
,	
<i>rname length</i>	<i>rname length</i> : symbol or decimal digit.
,	
,STEP	<b>Default: STEP</b>
,SYSTEM	
,SYSTEMS	
, <i>var1234</i>	<i>var1234</i> : The preceding 4 parameters may be repeated up to 65,535 times.
)	
,RET=HAVE	<b>Default: RET=NONE</b>
,RET=NONE	
, RMC=NONE	<b>Default: RMC=NONE</b>
, RMC=STEP	<b>Default: GENERIC=NO</b>
,GENERIC=NO	<b>Note:</b> If GENERIC=YES is specified, you must also specify
,GENERIC=YES	RET=HAVE above.
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : A-type address. <b>Note:</b> 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.

---

(	<b>Note:</b> ( 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.
)	<b>Note:</b> See note opposite ( above.
,RET=HAVE	
,RET=NONE	
,RMC=NONE	
,RMC=STEP	
,GENERIC=NO	
,GENERIC=YES	<b>Note:</b> If GENERIC=YES is specified, you must also specify RET=HAVE above.
,TCB= <i>tcb addr</i>	<i>tcb addr</i> : RX-type address, or register (2) - (12). <b>Note:</b> 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**.

## DSGNL — Issue Direct Signal

The DSGNL macro instruction uses the signal processor (SIGP) instruction to modify or sense the physical state of one of the CPUs 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 two SIGP functions are accessible via the RISGNL and RPSGNL macro instructions.

The DSGNL 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 DSGNL.
DSGNL	
␣	One or more blanks must follow DSGNL.

---

SENSE  
 START  
 STOP  
 RESTART  
 IPR  
 PR  
 SSS  
 IMPL  
 ICPUR  
 CPUR  
 (0)

,CPU=*PCCA addr*                      *PCCA addr*: RX-type address, or register (1).

---

The parameters are explained below:

SENSE  
 START  
 STOP  
 RESTART  
 IPR  
 PR  
 SSS  
 IMPL  
 ICPUR  
 CPUR  
 (0)

specifies the action to be performed. If (0) is specified, the code indicating the desired function has already been loaded into bits 24-31 of register 0. The actions and codes are:

	Code	Action
SENSE	01	State of specified CPU 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 CPU reset function
CPUR	0C	CPU 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 CPU whose PCCA is located at PCCA is requested. If the CPU is executing or is in a wait state, a return code of 0 in register 15 will be provided; otherwise, a return code of 8 with status indicators in register 0 will be returned.

```
DSGNL SENSE,CPU=PCCA
```

### Example 2

**Operation:** The CPU 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).
,	<b>Default:</b> E
,E	
,S	
,	
, <i>rname length</i>	<i>rname length</i> : symbol, decimal digit, or register (2) - (12). <b>Default:</b> assembled length of <i>rname</i> <b>Note:</b> <i>rname length</i> must be coded if a register is specified for <i>rname addr</i> .
,	<b>Default:</b> STEP
,STEP	
,SYSTEM	
,SYSTEMS	
, <i>var12345</i>	<i>var12345</i> : The preceding 5 parameters may be repeated up to 65,535 times.
)	
,RET=CHNG	<b>Default:</b> 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>	<b>Default:</b> SMC=NONE
,TCB= <i>tcb addr</i>	<b>Note:</b> ECB cannot be specified with RET above. <b>Note:</b> TCB cannot be specified with RET=HAVE or RET=NONE above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

---

The parameters restricted in use are explained below. 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.

TCB specifies the address of a fullword on a fullword boundary that contains the address of 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, or RET=HAVE; 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 22.




Figure 22. 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 shared. 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	<b>Default:</b> E
,S	
,	
<i>rname length</i>	<i>rname length</i> : symbol or decimal digit. <b>Default:</b> assembled length of <i>rname</i>
,	
,STEP	<b>Default:</b> STEP
,SYSTEM	
,SYSTEMS	
, <i>var12345</i>	<i>var12345</i> : The preceding 5 parameters may be repeated up to 65,535 times.
)	
,RET=CHNG	<b>Default;</b> RET=NONE
,RET=HAVE	
,RET=TEST	
,RET=USE	
,RET=NONE	
,SMC=NONE	<i>ecb addr</i> : A-type address.
,SMC=STEP	<b>Default:</b> SMC=NONE
,ECB= <i>ecb addr</i>	<b>Note:</b> ECB cannot be specified with RET above.
,TCB=0	<b>Note:</b> TCB cannot be specified with RET=HAVE or RET=NONE above, and must be specified on the list form if used on the execute form.
,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.
(	<b>Note:</b> ( 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.
)	<b>Note:</b> 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>tcb addr</i> : A-type address, or register (2) - (12).
,ECB= <i>ecb addr</i>	<b>Note:</b> ECB cannot be specified with RET above.
,TCB= <i>tcb addr</i>	<b>Note:</b> TCB cannot be specified with RET=HAVE or RET=NONE 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=0 (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. These parameters are restricted in use and should only be used by types 2, 3, and 4 SVCs executing in supervisor state, under protection key 0, and owning the local lock.

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	<b>Default:</b> CT
,OV	
,PARAM= <i>list addr</i>	<i>list addr</i> : A-type address, or register (2) - (12).
,XCTL=NO	<b>Default:</b> XCTL=NO
,XCTL=YES	
,PURGE=NONE	<b>Default:</b> PURGE=NONE
,PURGE=QUIESCE	
,PURGE=HALT	
,ASYNCH=YES	<b>Default:</b> ASYNCH=YES
,ASYNCH=NO	
,TERM=NO	<b>Default:</b> TERM=NO
,TERM=YES	
,BRANCH=NO	<b>Default:</b> BRANCH=NO
,BRANCH=YES, SVEAREA= <i>save area</i>	<i>save addr</i> : A-type address, or register (2) - (12) or (13).
,RECORD=NO	<b>Default:</b> RECORD=NO
,RECORD=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained below. The explanation of the other parameters is as explained in **OS/VS2 Supervisor Services and Macro Instructions**.

,BRANCH=NO

,BRANCH= YES,SVEAREA=*save addr*

specifies that an SVC 60 entry to the ESTAE service routine is not to be performed (NO) or that a branch entry is to be performed (YES). (The branch entry is for type 2, 3, or 4 SVCs only.) The save area is a 72-byte area used to save the general registers. If BRANCH=YES is specified, the caller must be in key 0 and own the local lock.

,RECORD=NO

,RECORD= YES

specifies that the SDWA workarea will not be written to SYS1.LOGREC (NO) or that the entire SDWA workarea, both fixed and variable, will be written to SYS1.LOGREC (YES).

Control is returned to the instruction following the ESTAE macro instruction. When control is returned, 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.
08	BRANCH=YES was issued for the current SVRB with a create request; the previous BRANCH=YES exit is canceled and the new exit is made the current 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	<b>Default:</b> PURGE=NONE
,PURGE=QUIESCE	
,PURGE=HALT	
,ASYNCH=YES	<b>Default:</b> ASYNCH=YES
,ASYNCH=NO	
,TERM=NO	<b>Default:</b> TERM=NO
,TERM=YES	
,RECORD=NO	<b>Default:</b> 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 changed 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 will be changed. The other field will remain 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 0, 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.
␣	One or more blanks must precede EVENTS.
EVENTS	
␣	One or more blanks must follow EVENTS.

---

ENTRIES= <i>nbr</i>	<i>nbr</i> : 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>	<b>Note:</b> 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).
	<b>Note:</b> If LAST is specified, WAIT must also be specified.
,WAIT=YES	
,WAIT=NO	
,BRANCH=NO	<b>Default:</b> BRANCH=NO
,BRANCH=YES	

---

The parameters restricted in use are explained below. 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.



## EXTRACT — Extract TCB Information

The EXTRACT macro instruction causes the control program to provide information from specified fields of the task control block or a subsidiary control block for either the active task or one of its subtasks. The information is placed in an area provided by the problem program.

The standard 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, or register (2) - (12).
,'S'	<i>tcb addr</i> : A-type address, or register (2) - (12).
, <i>tcb addr</i>	Default: 'S'
,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

---

The parameters are explained below:

### *answer addr*

specifies the address of the answer area to contain the requested information. The address is of one or more fullwords, starting on a fullword boundary. The number of fullwords required is the same as the number of fields specified in the FIELDS parameter, unless ALL is coded. If ALL is coded, seven fullwords are required.

### ,'S'

### ,*tcb addr*

specifies the address of a fullword on a fullword boundary containing the address of a task control block for a subtask of the active task. If 'S' is coded or assumed, no address is specified and the active task is assumed.

### ,FIELDS=(*tcb info*)

specifies the task control block information requested:

ALL requests information from the GRS, FRS, reserved, AETX, PRI, CMC, and TIOT fields.  
(If ALL is specified, 7 words are required just for ALL.)

GRS is the address of the save area used by the control program to save the general registers 0-15 when the task is not active.

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. The high-order bit of the last pointer is set to one to indicate the end of the list.

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 which 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>	Default: '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 exceptions:

*,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 cancelled 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      TIOD     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 exceptions:

,MF=(E, *ctrl addr*)

specifies the execute form of the EXTRACT macro instruction using a remote control program parameter list.

## 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 0.

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.
␣	One or more blanks must precede FREEMAIN.
FREEMAIN	
␣	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) - (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.
VC	<b>Note:</b> 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.
VU	
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	<b>Note:</b> BRANCH=(YES,GLOBAL) may only be specified with RC or RU above.
,BRANCH=(YES,GLOBAL)	
,KEY= <i>nmb</i>	<i>nmb</i> : decimal digits 0-15, or register (2) - (12). <b>Note:</b> This parameter may be specified only if BRANCH above is also specified.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained below. 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 memory lock prior to entry.

If `BRANCH=(YES,GLOBAL)` is specified, the `SP` parameter may only designate subpools 227, 228, 231, 239, 241, or 245.

`,KEY=key nmb`

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 is returned, 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 will contain 0; otherwise, register 15 will contain a nonzero value.

```
FREEMAIN    EC, LV=400, A=( 1 ), BRANCH=YES
```

### Example 2

**Operation:** Free all storage in subpool 239. Register 3 has been preset to contain the storage key of the storage to be released. If the request is unsuccessful, the caller will be abnormally terminated.

```
FREEMAIN    RU, SP=239, 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.
␣	One or more blanks must precede FREEMAIN.
FREEMAIN	
␣	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. <b>Note:</b> LA may only be specified with LC, LU, or L above. <b>Note:</b> 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*.

## FREEMAIN (Execute Form)

The execute form of the FREEMAIN 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 FREEMAIN.
FREEMAIN	
␣	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> : RX-type address or register (2) - (12).
,LV= <i>length value</i>	<i>length value</i> : symbol, decimal digit, or register (2) - (12). <b>Note:</b> LA may only be specified with LC, LU, or L above. <b>Note:</b> LV may only be specified with EC, EU, or E above.
,A= <i>addr</i>	<i>addr</i> : RX-type address, or register (2) - (12)
,SP= <i>subpool nmb</i>	<i>subpool nmb</i> : symbol, decimal digit, or register (0) or (2) - (12).
,BRANCH=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E , <i>ctrl prog</i> )	<i>ctrl prog</i> : RX-type address, or register (1) or (2) - (12).

---



## 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 0.

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> LU,LA= <i>length addr</i> ,A= <i>addr</i> VC,LA= <i>length addr</i> ,A= <i>addr</i> 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>	<i>length addr</i> : A-type address, or register (2) - (12). <i>length value</i> : symbol, decimal digit, or register (2) - (12). If R is specified, register (0) may also be specified. <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). <b>Note</b> : If R,LV=(0) is specified above, SP may not be specified.
,BNDRY=DLWD ,BNDRY=PAGE	<b>Default</b> : BNDRY=DLWD <b>Note</b> : This parameter may not be specified with R above.
,BRANCH=YES ,BRANCH=(YES,GLOBAL)	<b>Note</b> : BRANCH=(YES,GLOBAL) may only be specified with RC or RU above.
,KEY= <i>key number</i>	<i>key nmb</i> : decimal digits 0-15, or register (2) - (12). <b>Note</b> : This parameter may be specified only if BRANCH above is also specified.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained below. 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 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 memory lock prior to entry.

If BRANCH=(YES,GLOBAL) is specified, the SP parameter may only designate subpools 227, 228, 231, 239, 241, or 245.

,KEY=*key nmbr*

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 will be used; otherwise, subpool 0 will be used. If the storage cannot be obtained, the caller will be 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) will be 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	
,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.
	<b>Note:</b> LA may be specified with EC or EU above.
	<b>Note:</b> LV may not be specified with LC, LU, VC or VU 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.
,BNDRY=DBLWD	<b>Default:</b> BNDRY=DBLWD
,BNDRY=PAGE	
,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 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	
,LA= <i>length addr</i>	<i>length addr</i> : RX-type address or register (2) - (12).
,LV= <i>length value</i>	<i>length value</i> : symbol, decimal digit, or register (2) - (12). <b>Note:</b> LA may not be specified with EC or EU above. <b>Note:</b> LV may not be specified with LC, LU, VC, or VU above.
,A= <i>addr</i>	<i>addr</i> : RX-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).
,BNDRY=DBLWD	<b>Default:</b> BNDRY=DBLWD
,BNDRY=PAGE	
,BRANCH=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.
,MF=(E , <i>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 will be 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>	
	SCHED    SRM    ZERO
	JES       SUPR    TCB
	RSM       DATAMGT RBT1
	VSM       TCAM    RBT234
	<i>key addr</i> : A-type address or register (2).
	<b>Note:</b> The WORKREG parameter is required if the following are specified:
	EXTKEY=TCB                    EXTKEY=RBT234
	EXTKEY=RBT1                   KEYADDR=A-type address
,SAVEKEY= <i>old key addr</i>	<i>old key addr</i> : A-type address or register (2).
	<b>Note:</b> If KEYADDR=(2) is specified above, then SAVEKEY=(2) cannot be specified.
	<b>Note:</b> 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 below:

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.

,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.
␣	One or more blanks must precede MODESET.
MODESET	
␣	One or more blanks must follow MODESET.
KEY=ZERO KEY=NZERO	<b>Note:</b> KEY is required only if no other parameter is specified.
,MODE=PROB ,MODE=SUP	<b>Note:</b> MODE is required only if no other parameter is specified.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained below:

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.
<code>  b</code>	One or more blanks must precede MODESET..
MODESET	
<code>  b</code>	One or more blanks must follow MODESET.

---

<code>  RELATED=<i>value</i>,</code>	<i>value</i> : any valid macro keyword specification.
<code>  MF=L</code>	

---

The parameters are explained under the standard form of the MODESET macro instruction, with the following exceptions:

- 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 exceptions:

MF=(E ,*list addr*)

specifies the execute form of the MODESET macro instruction, using a parameter list address.

### ***Incompatible Parameters***

The ENABLE and SYSMASK parameters, available with Release 1 of VS2, are no longer supported on MODESET. The functions formerly available via ENABLE and SYSMASK are now provided by the SETLOCK macro instruction and the STNSM and STOSM instructions.

## 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>	<b>Default for reg1: 0</b>
<i>,WREGS=(reg1)</i>	<b>Default for reg2: 1</b>
<i>,WREGS=(,reg2)</i>	<b>Default for reg3: 2</b>
<i>,WREGS=(,,reg3)</i>	

The parameters are explained below:

*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 AND 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>	<b>Default for reg1: 0</b>
<i>,WREGS=(reg1)</i>	<b>Default for reg2: 1</b>
<i>,WREGS=(,reg2)</i>	<b>Default for reg3: 2</b>
<i>,WREGS=(,,reg3)</i>	

---

The parameters are explained below:

*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
```

## PGFIX — Fix Virtual Storage Contents

The PGFIX macro instruction makes virtual storage areas resident in real storage and ineligible for page-out while the requesting task's address space occupies real storage. The PGFIX function is available only to authorized system functions and users.

PGFIX does not prevent pages from being paged out when an entire address space is swapped out of real storage. Consequently, when using the PGFIX macro instruction, you can not assume a constant real address mapping for fixed pages that are susceptible to swapping.

The standard form of the PGFIX 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 PGFIX.
PGFIX	
␣	One or more blanks must follow PGFIX.

---

R	
,A= <i>start addr</i>	<i>start addr</i> : A-type address, or register (1) or (2) - (12).
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (0) or (2) - (12).
,EA= <i>end addr</i>	<i>end addr</i> : A-type address, or register (2) - (12) or (15). <b>Default:</b> <i>start addr</i> + 1
,LONG=Y	<b>Default:</b> LONG=Y
,LONG=N	
,RELEASE=N	<b>Default:</b> RELEASE=N
,RELEASE=Y	<b>Note:</b> RELEASE=Y may only be specified with EA above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

---

The parameters are explained below:

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 fixed.

,ECB=*ecb addr*

specifies the address of the ECB that is used to signal event completion.

,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' )
```

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

Hexadecimal Code	Meaning
00	Operation completed normally; ECB posted complete.
04	Operation abnormally terminated. Operation incomplete because of invalid address in virtual subarea list entry; ECB posted complete.
08	Operation proceeding; ECB will be posted when all requested pages are fixed in real storage.
10	Operation abnormally terminated. Virtual subarea list entry or ECB address invalid; no ECB is posted.

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.

### ***Incompatible Parameters***

The following parameters were valid in Release 1 of OS/VS2, but are not supported in MVS:

```
SUSPEND = N, will be ignored.
SUSPEND = Y, will be ignored.
ECBIND = address, will probably cause errors.
```

### **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 will actually be fixed. The storage will be long fixed.

```
PGFIX   R, A=( R3 ), ECB=( R5 )
```

### **Example 2**

**Operation:** Fix virtual storage without using a virtual subarea list. Storage will be 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 to be forfeited before fixing the area specified by registers 3 and 4.

```
PGFIX    R,A=(R3),EA=(R4),ECB=(R5),LONG=N,RELEASE=Y
```



## PGFIX (List Form)

The list form of the PGFIX macro instruction uses a virtual subarea list.

The list form of the PGFIX 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 PGFIX.
PGFIX	
␣	One or more blanks must follow PGFIX.

---

L	
,LA= <i>list addr</i>	<i>list addr</i> : A-type address, or register (1) or (2) - (12).
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (0) or (2) - (12).
,LONG=N	<b>Default:</b> LONG=N
,LONG=Y	
,RELEASE=N	<b>Default:</b> RELEASE=N
,RELEASE=Y	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

---

The parameters are explained under the standard form of the PGFIX macro instruction, with the following exceptions:

- 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.

### Example 1

**Operation:** Fix virtual storage, providing a virtual subarea list addressed by register 5.

```
PGFIX L, LA=(R5), ECB=(R6)
```

## 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 the issuance of 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.

---

R	
,A= <i>start addr</i>	<i>start addr</i> : A-type address, or register (1) or (2) - (12).
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (0) or (2) - (12).
,EA= <i>end addr</i>	<i>end addr</i> : A-type address, or register (2) - (12) or (15). <b>Default:</b> <i>start addr</i> + 1
,RELEASE=N	<b>Default:</b> RELEASE=N
,RELEASE=Y	<b>Note:</b> RELEASE=Y may only be specified with EA above.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

---

The parameters are explained below:

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.

### ***Incompatible Parameters***

The following parameters were valid in Release 1 of OS/VS2, but are not supported in MVS:

ECBIND = address, will probably cause errors.

#### **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
```

## PGFREE (List Form)

The list form of the PGFREE macro instruction uses a virtual subarea list.

The list of the PGFREE 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 PGFREE.
PGFREE	
b	One or more blanks must follow PGFREE.

---

L	
,LA= <i>list addr</i>	<i>list addr</i> : A-type address, or register (1) or (2) - (12).
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (0) or (2) - (12).
,RELEASE=N	<b>Default:</b> RELEASE=N
,RELEASE=Y	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

---

The parameters are explained under the standard form of the PGFREE macro instruction, with the following exceptions:

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.

### Example 1

**Operation:** Free the storage in Example 1 of list-form PGFIX.

```
PGFREE L, LA=( R5 )
```

## 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.
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). <b>Range of values:</b> 0 - 2 <sup>30</sup> -1 <b>Default:</b> 0
<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>,RELATED=value</i>	<i>value</i> : any valid macro keyword specification.

The parameters restricted in use are explained below. 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 begin posted, and the address of the routine to be given control when an error condition resulting from a POST failure is detected.

### 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.
,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.

## 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.
␣	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). <b>Range of values:</b> 0 - 2 <sup>30</sup> -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>,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 exceptions:

*,MF=(E ,prob addr)*

specifies the execute form of the POST macro instruction using a remote control program parameter list.

## PURGEDQ — Purge SRB Activity

The PURGEDQ macro instruction provides the facility for a task to purge particular SRB activity. Because an SRB routine is dispatched asynchronously to the actual issuance of a SCHEDULE macro instruction, the conditions that existed in the system at the time the SCHEDULE was issued may have totally changed by the time the routine is dispatched. If, in this time interval, the environment that the asynchronous routine requires to run successfully has been changed, the results would be unpredictable. For this reason, the PURGEDQ macro instruction is available to:

- Dequeue SRBs not yet dispatched.
- Dequeue or allow completed processing for SRBs previously scheduled.
- 'Clean up' each dequeued SRB.

The standard 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).

---

The parameters are explained below:

RMTR=*RMTR addr*

specifies the address of the resource manager termination routine.

,ASID=*ASID addr*

specifies the address of a halfword containing the address space identifier. PURGEDQ will search for SRBs scheduled to be dispatched into the address space specified by ASID.

,ASIDTCB=*TCB addr*

specifies the address of a doubleword in the following format:

bytes 0-1	Reserved
bytes 2-3	ASID or zero
bytes 4-7	TCB address or zero

### Example 1

**Operation:** All SRBs scheduled into the current address space and related to the current (terminating) task are to be purged by the RMTR routine IEAVRSPG.

```
PURGEDQ RMTR=IEAVRSPG
```



## PURGEDQ (List Form)

The list form of the PURGEDQ macro instruction is used to construct a remote program parameter list.

The list 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> : A-type address.
,ASID= <i>ASID addr</i>	<i>ASID addr</i> : A-type address.
,ASIDTCB= <i>TCB addr</i>	<i>TCB addr</i> : A-type address.
,MF=L	

---

The parameters are explained under the standard form of the PURGEDQ macro instruction, with the following exceptions:

,MF=L  
specifies the list form of the PURGEDQ macro instruction.

## 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=(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 PURGEDQ macro instruction, with the following exceptions:

,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 — Link to SVC 34

The QEDIT macro instruction generates the required entry parameters and the linkage to SVC 34 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 below:

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 will be freed.

,BLOCK=*CIB addr*

specifies the address of the CIB that is 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.

### 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 )
```

## RESERVE — Reserve a Device

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. If a task issues two RESERVE instructions for the same resource without an intervening DEQ an abnormal termination will result 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 will not restore the RESERVE; the user's program must reissue the RESERVE.) Even if a DEQ is not issued for a particular device, termination routines will 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.
␣	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, or register (2) - (12).
, <i>rname addr</i>	<i>rname addr</i> : A-type address, or register (2) - (12).
,E	<b>Default:</b> 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 below:

- (
- 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 name 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 resource 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.

,UCB = *ucb addr*

specifies the address of a fullword that contains the address of the UCB for the device to be reserved.

,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 23.




Figure 23. 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>	<i>rname addr</i> : A-type address.
,	
,E	
,S	
,	
<i>rname length</i>	<i>rname length</i> : symbol or decimal digit. <b>Note:</b> <i>rname length</i> must be coded if a register is specified for <i>rname 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 exceptions:

,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.
␣	One or more blanks must precede RESERVE.
RESERVE	
␣	One or more blanks must follow RESERVE.

---

(	<b>Note:</b> ( 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 exceptions:

,MF=(E, *ctrl addr*)  
 specifies the execute form of the RESERVE macro instruction using a remote control program parameter list.



## 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 CPUs 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 below:

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 CPU.

,CPU=*PCCA addr*

specifies the address of the physical configuration communication area (PCCA) of the CPU 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 CPU.

,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 CPU not online.
08	Function unsuccessful. Emergency signal could not be generated on CPU. Status returned in register 0.
12	Function unsuccessful. Specified CPU is either not installed, not configured into system, or powered off.
16	CPU is a uniprocessor and does not have signal processor sending and receiving capabilities.
20	CPU alive bit was turned off during the remote immediated 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 CPU 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 CPU 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 CPUs 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.
␣	One or more blanks must precede RPSGNL.
RPSGNL	
␣	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 below:

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 TQE check function, to ensure that TQE 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 CPU to which channels are attached.

,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	Specified CPU is online and has been notified that the specified service is to be executed.
04	Function not initiated because addressed CPU not online.
08	Function unsuccessful. External call signal could not be generated on CPU. Status returned in register 0.
12	Specified CPU is either not installed, not configured into 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
1-25	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 CPU 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 CPU 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 will have a priority that is greater than, and independent of, any address space priority. Services scheduled at a local priority will have the priority of the specific address space they execute in, but will still have a priority greater than that of any task within the address space.

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	<b>Default:</b> SCOPE=LOCAL.
,SCOPE=GLOBAL	

The parameters are explained below:

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 0, 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 an address space 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 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.
‡	One or more blanks must follow 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> : A-type address, or register (2) - (12).
,DCB= <i>dcb addr</i>	<i>dcb addr</i> : A-type address, or register (2) - (12).
,ASID= <i>ASID addr</i>	<i>ASID addr</i> : A-type address, or register (2) - (12).
,ECB= <i>ecb addr</i>	<i>ecb addr</i> : A-type address, or register (2) - (12).
,SDATA=( <i>data code</i> )	<i>data code</i> : any combination of the following, separated by commas: SQA      RGN ALLPSA   LPA PSA      TRT NUC      CSA LSQA     SWA
,STORAGE=( <i>strt addr,end addr</i> )	<i>strt addr</i> : A-type address, or register (2) - (12).
,LIST= <i>list addr</i>	<i>end addr</i> : A-type address, or register (2) - (12). <i>list addr</i> : A-type address, or register (2) - (12). <b>Note:</b> 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	<b>Default:</b> BUFFER=NO
,BUFFER=YES	
,QUIESCE=YES	<b>Default:</b> QUIESCE=YES
,QUIESCE=NO	
,BRANCH=NO	<b>Default:</b> BRANCH=NO
,BRANCH=YES	<b>Note:</b> If BRANCH is specified, ASID must also be specified.

---

The parameters are explained below:

**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 appear 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.

**,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 will be used.

**,ASID = ASID addr**

specifies the address of a halfword containing the address space identifier of the address space to be dumped. If this parameter is omitted, the current address space will be dumped. If 0 is specified, the dump is scheduled in the current address space.

**,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.

**,SDATA = (data code)**

specifies the system control program information to be dumped:

SQA — The system queue area.

ALLPSA — All of the prefixed storage areas in the system.

PSA — The prefixed storage area for the current CPU.

NUC — The nucleus.

LSQA — The local system queue area for the address space being dumped.

RGN — The allocated pages in the private area of the address space being dumped. This includes the LSQA and the SWA.

LPA — The active link pack area modules and SVCs for the 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 will be made to include trace information.

CSA — The common service area subpools.

SWA — The scheduler work area subpools in the address space being dumped.

**,STORAGE = (strt addr, end addr)**

**,LIST = list addr**

specifies one or more pairs of starting and ending address 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.

**,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.)

,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 SDATA parameter does not specify SQA or CSA, the QUIESCE=YES request is ignored.

,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.

If the ASID 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.
08	The system was unable to take a dump.

If the ASID 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:

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 for callers who cannot issue an SVC. Areas to be dumped are requested via three parameters (BUFFER, SDATA, and STORAGE). The dump will have 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)
```



## 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.
HDR= <i>'dump title'</i> ,HDRAD= <i>dump title addr</i> ,DCB= <i>dcb addr</i> ,SDATA=( <i>data code</i> )	<i>dump title</i> : from 1 to 100 characters. <i>dump title addr</i> : A-type address. <i>dcb addr</i> : A-type address. <i>data code</i> : any combination of the following, separated by commas: SQA           RGN ALLPSA       LPA PSA           TRT NUC           CSA LSQA          SWA
,STORAGE=( <i>strt addr,end addr</i> ) ,LIST= <i>list addr</i>	<i>strt addr</i> : A-type address. <i>end addr</i> : A-type address. <b>Note:</b> 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	<b>Default:</b> BUFFER=NO
,QUIESCE=YES ,QUIESCE=NO	<b>Default:</b> QUIESCE=YES
,MF=L	

The parameters are explained under the standard form of the SDUMP macro instruction, with the following exceptions:

,MF=L  
specifies the list form of the SDUMP macro instruction.

## 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 will be 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).
,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: SQA          RGN ALLPSA     LPA PSA         TRT NUC         CSA LSQA         SWA
,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). <b>Note</b> : 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	<b>Note</b> : If BRANCH is specified, ASID must also be specified.
,BRANCH=YES	
,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 exceptions:

,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. Users of SETFRR must be key 0 supervisor state and must hold a lock or be in SRB mode.

The SETFRR macro instruction is written as follows:

---

<i>name</i>	<i>name</i> : symbol. Begin <i>name</i> in column 1.
<b>␣</b>	One or more blanks must precede SETFRR.
SETFRR	
<b>␣</b>	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). <b>Note:</b> 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 containing the address of where the parameter area is to be placed after the FRR address has either been added to the stack or has replaced an FRR address on the stack. The parameter area is a 24-byte area that is eventually 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. Register (R2) will contain 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 0.

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 by the requesting CPU.

Two classes of locks exist: global and local. Two types of locks exist: 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.

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* *value*: any valid macro keyword specification.

---

The parameters are explained below:

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 of storage.

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 in which 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 by another CPU, it will be acquired on the caller's behalf. If the lock is already held, control will be returned to the caller indicating that the lock is held and that either the caller already owns the lock or that another CPU or storage owns the lock.

UNCOND specifies that the lock is to be unconditionally obtained. That is, if the lock is not owned by another CPU, it will be acquired on the caller's behalf. If the lock is already held by the caller, control will be returned to the caller indicating that he already owns the lock. If the lock is held by another CPU, the caller's CPU will either spin on the lock until it is released or suspend 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 1.

SAVE specifies that register contents are to be saved. Registers 11 through 14 will be saved in the area pointed to by register 13, and will be restored upon completion of the SETLOCK request. The savearea will consist of at least 5 words. Register 15 will contain the return code.

USE specifies that registers 14, 15, 0, and 1 are available for use. Registers 11, 12, and 13 will be saved in registers 15, 0, and 1, respectively, and will be restored upon completion of the SETLOCK request. Register 14 will be used as a link register; register 15 will contain the return code.

**Note:** If neither SAVE nor USE is specified, registers 11-14 are destroyed and register 13 contain 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:

```
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	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 will either be 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.
↳	One or more blanks must precede SETLOCK.
SETLOCK	
↳	One or more blanks must follow SETLOCK.
<hr/>	
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	<b>Note:</b> 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 by the CPU are to be released.

ALL All locks currently held by the CPU 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 CPU 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 <b>Note:</b> 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. <b>Note:</b> 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 by the requesting CPU.

,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 will be 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 by the requesting CPU.

If NOTHELD is specified, the address will be branched to if the specified lock is not currently held by the requesting CPU, 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 are 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 may 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**.

The SETRP 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 SETRP.
SETRP	
␣	One or more blanks must follow SETRP.

---

WKAREA=( <i>reg</i> )	<i>reg</i> : decimal digits 1-12. <b>Default:</b> WKAREA=(1)
,REGS=( <i>reg1</i> ) ,REGS=( <i>reg1,reg2</i> )	<i>reg1</i> : decimal digits 0-12, 14, 15. <i>reg2</i> : decimal digits 0-12, 14, 15. <b>Note:</b> If <i>reg1</i> and <i>reg2</i> are both specified, order is 14, 15, 0-12.
,DUMP=IGNORE ,DUMP=YES ,DUMP=NO	<b>Default:</b> DUMP=IGNORE
,DUMPOPT= <i>parm list addr</i>	<i>parm list addr</i> : RX-type address, or register (2) - (12). <b>Note:</b> This parameter may be specified only if DUMP=YES is specified above.
,RC=0 ,RC=4 ,RC=16	<b>Default:</b> RC=0
,RETADDR= <i>retry addr</i>	<i>retry addr</i> : RX-type address, or register (2) - (12). <b>Note:</b> This parameter may be specified only if RC=4 is specified above.
,RETREGS=NO ,RETREGS=YES ,RETREGS=YES,RUB= <i>info addr</i>	<i>info addr</i> : RX-type address, or register (2) - (12). <b>Default:</b> RETREGS=NO <b>Note:</b> This parameter may be specified only if RC=4 is specified above.
,FRESWA=NO ,FRESWA=YES	<b>Default:</b> FRESWA=NO <b>Note:</b> This parameter may be specified only if RC=4 is specified above.
,COMPCOD= <i>code</i> ,COMPCOD=( <i>code,USER</i> ) ,COMPCOD=( <i>code,SYSTEM</i> )	<i>code</i> : symbol, decimal digit, or register (2) - (12). <b>Default:</b> COMPCOD=( <i>code,USER</i> )
,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. <b>Note:</b> 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 ,RECORD=YES ,RECORD=NO	<b>Default:</b> RECORD=IGNORE
,RECPARM= <i>record list addr</i>	<i>record list addr</i> : RX-type address, or register (2) - (12). <b>Note:</b> This parameter may be specified only if RECORD=IGNORE or RECORD=YES is specified above.

---

The parameters restricted in use are explained below. 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 CPU identification of the CPU holding the resource that this CPU 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.

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, contains the actual length of the record.

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
DUMP		x
DUMPOPT	x	x
RC=0	x	x
RC=4	x	x
RC=16		x
RETADDR		x
RETREGS	x	x
RUB	x	x
FRESWA	x	x
COMPCOD	x	x
FRELOCK	x	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 CPU identified by the contents of register 7. In this example, the interrupted function is spinning on a lock currently being held by the CPU 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.

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 <b>single values</b> : (2,3,4,7,8,9,10) <b>ranges of values</b> : ((2,4),(7,10)) <b>combinations</b> : ((2,4),6,8,(10,13),15)
-------------------------------	--

---

The parameters are explained below:

*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.

**Note:** As shown in the table above, interruption types can be designated as one or more single numbers, as one or more pairs of numbers (designating ranges of values), or as any combination of the two forms. For example, (4,8) indicates interruption types 4 and 8; ((4,8)) indicates interruption types 4 through 8.

### Example 1

**Operation:** Give control to an exit routine for interruptions 1,5,7,8,9, and 10. DOITSPIE is the address of the SPIE exit routine.

```
SPIE    DOITSPIE,(1,5,7,(8,10))
```

## **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 follow 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 vaues: ((2,4),(7,10)) combinations: ((2,4),6,8,(10,13),15)
,MF=L	

---

The parameters are explained under the standard form of the SPIE macro instruction, with the following exceptions:

,MF=L  
specifies the list form of the SPIE macro instruction.

## SPIE (Execute Form)

A remote control program parameter list (program interruptions control area) is used in, and can be modified by, the execute form of the SPIE macro instruction. The PICA (program interruptions control area) can be generated by the list form of SPIE, or you can use the address of the PICA returned in register 1 following a previous SPIE macro instruction. If this macro instruction is being issued to reestablish a previous SPIE environment, code only the MF parameter.

The address of the remote control program parameter list associated with any previous SPIE environment is returned by the SPIE macro instruction.

The execute 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> : RX-type address, or register (2) - (12).
<i>,(interrupts)</i>	<i>interrupts</i> : decimal digits 1-15, or 17, expressed as <b>single values</b> : (2,3,4,7,8,9,10) <b>ranges of values</b> : ((2,3),(7,10)) <b>combinations</b> : ((2,4),6,8,(10,13),15)
<i>,MF=(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 SPIE macro instruction, with the following exceptions:

*,MF=(E,ctrl,addr)*

specifies the execute form of the SPIE macro instruction using a remote control program parameter list.

## SPOST — Synchronize POST

The SPOST macro instruction is used in a cross-memory post environment to ensure that all outstanding cross-memory post requests for the ECB specified have completed. SPOST resolves a synchronization problem that arises when it becomes necessary to eliminate an ECB which is a potential target for a cross-memory post request.

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.
␣	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.
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> : A-type address, or register (2) - (12).
0	
,CT	<b>Default:</b> CT
,OV	
,PARAM= <i>list addr</i>	<i>list addr</i> : A-type address, or register (2) - (12).
,XCTL=NO	<b>Default:</b> XCTL=NO
,XCTL=YES	
,PURGE=QUIESCE	<b>Default:</b> PURGE=QUIESCE
,PURGE=HALT	
,PURGE=NONE	
,ASYNCH=NO	<b>Default:</b> ASYNCH=NO
,ASYNCH=YES	
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained below:

*exit addr*

0

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=*list addr*

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 will not be saved when the STAE exit is taken (HALT), that I/O processing will be allowed to continue normally when the STAE exit is taken (NONE), or that all outstanding requests for I/O operations will be 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 this is done, all control blocks affected by input/output processing may 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 will develop when an input/output interruption occurs, even if the exit routine is in progress. Thus, it will appear that the exit routine failed when, in reality, input/output processing was the cause of 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 will be posted. Subsequent event control blocks (ECBs) will not be posted.
- The ISAM check routine will treat 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 will be allowed (YES) or will not be allowed (NO) while the STAE exit is executing.

ASYNCH=YES must be coded if:

- Any supervisor services that require asynchronous interruptions to complete their normal processing are going to be requested by the STAE exit routine.
- 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 will develop when an asynchronous interruption occurs. Thus, it will appear that the exit routine failed when, in reality, asynchronous exit handling was the cause of 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 is returned to the instruction following the STAE macro instruction. When control is returned, 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, I/O will be halted, no asynchronous exits will be taken, ownership will be transferred 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	<b>Default:</b> PURGE=QUIESCE
,PURGE=HALT	
,PURGE=NONE	
,ASYNCH=NO	<b>Default:</b> 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 exceptions:

,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 the user desires to dynamically changed 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 will be changed. The other field will remain 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.
␣	One or more blanks must precede STAE.
STAE	
␣	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 exceptions:

,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, contain 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 authorized 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>tcn addr</i>	<i>tcn 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>	<b>Note:</b> ASID may only be specified with START.
,SYNCH	<b>Note:</b> SYNCH may only be specified with STOP.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

The parameters are explained below:

START

STOP

specifies that the START/STOP count is to be adjusted and the dispatchability bits are to be set/reset.

,TCB=*tcn 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 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')
```

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.
␣	One or more blanks must precede STATUS.
STATUS	
␣	One or more blanks must follow STATUS.

---

SET	
RESET	
,MC	<b>Note:</b> If MC or MC,STEP is specified, no other parameters may 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, <i>mask</i>	<i>tcb addr</i> : RX-type address, or register (2) - (12).
, <i>tcb addr</i> , <i>mask</i>	<b>Default:</b> STEP
., <i>mask</i>	
,E	<b>Note:</b> This parameter may only be specified with <i>tcb addr</i> , <i>mask</i> .
,ASID=ASID <i>addr</i>	<i>ASID addr</i> : RX-type address, or register (2) - (12). <b>Note:</b> For SET, this parameter may only be specified with <i>tcb addr</i> , <i>mask</i> . For RESET, this parameter may <i>not</i> be specified with SYSTEM.
,RELATED= <i>value</i>	<i>value</i> : any valid macro keyword specification.

---

The parameters are explained below:

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. STEP indicates that the set-must-complete indicator in the issuer's TCB and a count in the ASCB are to be set/reset.

,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>▯</b>	One or more blanks must precede SYNCH.
SYNCH	
<b>▯</b>	One or more blanks must follow SYNCH.

---

<i>entry name addr</i>	<i>entry name addr</i> : RX-type address, or register (2) - (12) or (15).
------------------------	---

---

The parameters are explained below:

*entry name addr*

specifies the address of the entry name of the processing program to receive control.

### Example 1

**Operation:** Take a synchronous exit to a processing program whose entry name address is specified in register 8.

```
    SYNCH   ( 8 )
```

### Example 2

**Operation:** Take a synchronous exit to a processing program labeled SUBRTN.

```
    SYNCH   SUBRTN
```

## 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> FCTN= <i>fcn</i> ,AUTH= <i>auth</i>	<i>fcn</i> : decimal digit 1 or register (2) - (12). <i>auth</i> : decimal digit 0 or 1, or register (2) - (12).
,STATE=NO ,STATE=YES	Default: STATE=NO
,KEY=NO ,KEY=YES	Default: KEY=NO
,RBLEVEL=2 ,RBLEVEL=1	Default: RBLEVEL=2
,BRANCH=NO ,BRANCH=YES	Default: BRANCH=NO

The parameters are explained below:

FCTN=*fcn*

FCTN=*fcn*,AUTH=*auth*

specifies the authorization (via APF) of a program.

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 FUNC=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.

,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.)

,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).

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    AUTH=1 ,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
```



## WTO — Write to Operator

The WTO macro instruction is described in the **OS/VS2 Supervisor Services and Macro Instructions**, with the exception of the AREAID, MSGTYP, and MCSFLAG parameters. The MSGTYP and MCSFLAG parameters should only be used by system programmers familiar with MCS, since using the parameters improperly could impede the entire message routing scheme. The AREAID parameter can only be used by APF-authorized users.

The syntax of the complete WTO macro instruction is shown below. However, only the explanation of the AREAID, 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 WTO macro instruction is written as follows:

---

<i>name</i>	<i>name</i> : symbol. Begin name 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>	<i>text</i> : Up to 124 characters.
<i>('text',line type)</i>	The permissible <i>line types</i> and <i>text</i> lengths are shown below:
	line type    VS2 text
	C            34 char
	L            70 char
	D            70 char
	DE          70 char
	E            -----
	Up to 10 occurrences of the second and/or third formats may be coded.
,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
	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        QREGO
	REPLY       NOTIME
	BRDCST      NOCPY

---

The parameters restricted in use are explained below. The other parameters are explained in **OS/VS2 Supervisor Services and Macro Instruction**.

,ARE Aid=*id char*

specifies a display area of the console screen on which a multiple-line message is to be written. This parameter is useful only for out-of-line (descriptor code 8 and 9) MLWTO messages which 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 APF-authorized to use this parameter. Also, if an area is specified by this parameter, there exists the possibility that this area will be overlaid by a currently running dynamic display.

,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 which issued the MONITOR SESS, MONITOR JOBNAMES, or MONITOR STATUS command, respectively. When the message type is identified by the operating system, the message will be routed to only 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 24.

---

Name	Bit	Meaning
----	0	Invalid entry.
REG0	1	Message is to be queued 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	Message should be broadcast to all active consoles.
HRDCPY	6	Message queued for hard copy only.
QREG0	7	Message is to be queued unconditionally to the console whose source ID is passed in Register 0.
NOTIME	8	Time is not appended 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, the message is not queued 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 24. 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 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> ' ('text') ('text',line type)	<i>msg</i> : up to 124 characters. The permissible <i>line types</i> and <i>text</i> lengths are shown below: line type    VS2 text C            34 char L            70 char D            70 char DE          70 char E            ----- Up to 10 occurrences of the second and/or third formats may be coded.
,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> : an alphabetic character A-Z.
,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
,MF=L	

---

The parameters restricted in use are explained under the standard form of the WTO macro instruction. The explanation of the other parameters is as 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 only be used 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).
<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>,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

The parameters restricted in use are explained below. 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 will be routed to only 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 24 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 JOBNAME command.

```
WTOR      'WTOR BY MSGTYP=JOBNAME WITH RC=ALL, NO CONSOLE MONITORING  
          JOBNAME',MSGTYP=JOBNAME,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>	<i>reply addr</i> : an A-type address.
,	
, <i>reply length</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>ecb addr</i>	<i>ecb addr</i> : A-type address.
,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
,MF=L	

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





# Index

Indexes to OS/VS2 publications are consolidated in the OS/VS2 Master Index, GC28-0693, and the OS/VS2 Master Index of Logics, GY28-0694. For additional information about any subject listed below, refer to other publications listed for the same subject in the Master Index.

A parameter  
FREEMAIN macro instruction 117  
GETMAIN macro instruction 121  
PGFIX macro instruction 134  
PGFREE macro instruction 138  
SETFRR macro instruction 164

A-type address  
meaning 81

ABEND macro instruction 37,39

ADDR parameter 167

address space termination 38

ALL parameter 92

allocate virtual storage (GETMAIN) 121-124

AND IMMEDIATE (NI) instruction, providing a lock via 130-131

APF (authorized program facility) 30-31

AREAID parameter 194

ASCB parameter 141

ASID parameter  
CALLRTM macro instruction 88  
PURGEDQ macro instruction 144  
SDUMP macro instruction 159  
STATUS macro instruction 186,189

ASIDTCB parameter 144

ASM (auxiliary storage manager lock) 19

ASYNCH parameter  
ATTACH macro instruction 85  
ESTAE macro instruction 108  
STAE macro instruction 182

asynchronous execution, scheduling system services for 157

ATTACH macro instruction 84-87  
use of restricted parameters 13-14

AUTH parameter 191

authorization 30

authorization, testing 191-192

authorized library 31  
SVC routines 67

authorized program facility (APF) 30-31

auxiliary storage manager lock (ASM) 19

BLOCK parameter 147

BNDRY parameter 121

branch entry  
EVENTS 36  
GETMAIN and FREEMAIN 53  
SDUMP 39

BRANCH parameter  
CIRB macro instruction 95  
ESTAE macro instruction 108  
EVENTS macro instruction 112  
FREEMAIN macro instruction 117  
GETMAIN macro instruction 121  
SDUMP macro instruction 160  
SETLOCK macro instruction 171  
TESTAUTH macro instruction 192

BUFFER parameter 159

call recovery/termination manager (CALLRTM) 88-89  
CALLRTM macro instruction 88-89  
use of 37-38

CD (CHNGDUMP) command 90-93  
change dump options (CHNGDUMP) 90-93  
change subtask status (STATUS) 186-189  
change system status (MODESET) 125-129  
channel ends 71  
check stop status condition 34

CHNGDUMP command 90-93  
use of 40

CIB (command input buffer) 15

CIBCTR parameter 147

CIRB macro instruction 94-95  
use of 66

classes of locks 20

CMS (cross memory services lock) 20  
SVC routines 69

coding the macro instructions 80-82

command input buffer (CIB) 15

common service area subpools 53

COMPARE AND SWAP (CS) instruction  
NIL macro instruction 130  
OIL macro instruction 132-133  
POST macro instruction 35

COMPCOD parameter  
CALLRTM macro instruction 88  
SETRP macro instruction 174

console switching 60

continuation lines 82

control access to serially reusable resources (SETLOCK) 166-172

control track record 76

CPU parameter  
DSGNL macro instruction 102  
RISGNL macro instruction 153  
RPSGNL macro instruction 155  
SETRP macro instruction 175

CPU reset function 33

CPUR parameter 101

create a new task (ATTACH) 84-87

create interruption request block (CIRB) 94-95

cross-memory post 35

cross memory services lock (CMS) 20

CS (COMPARE AND SWAP) instruction 130,132

CT parameter  
ESTAE macro instruction 108  
STAE macro instruction 181

D parameter 164

DCB parameter  
ATTACH macro instruction 85  
SDUMP macro instruction 159

DDR swaps 71

DE parameter 85

decimal digit  
meaning 81

default  
meaning 82

DEL parameter 92

DEQ macro instruction 96-100  
and RESERVE 148  
execute form 100  
list form 99  
standard form 96-98  
use of restricted parameters 22-24

SC parameter  
WTO macro instruction 193  
WTOR macro instruction 197

descriptor codes 62

device ends 34

device, reserving a 148-152

direct class 33

direct signal, issuing 101-102

DISABLED parameter 170

disabling 20

DISP (global dispatcher lock) 19

DISP parameter 86

DPMOD parameter 85

DSGNL macro instruction 101-102  
use of 33

dump options, changing 90-93

DUMP parameter  
  CALLRTM macro instruction 89  
  SETRP macro instruction 174  
dump virtual storage (SDUMP) 158-163  
DUMPOPT parameter  
  CALLRTM macro instruction 89  
  SETRP macro instruction 174  
dumps of virtual storage 39-40

E parameter  
  ENQ macro instruction 103  
  FREEMAIN macro instruction 117  
  RESERVE macro instruction 149  
  STATUS macro instruction 189

EA parameter  
  PGFIX macro instruction 134  
  PGFREE macro instruction 138

EC (external call) function 33

EC parameter  
  FREEMAIN macro instruction 117  
  GETMAIN macro instruction 121

ECB parameter  
  ATTACH macro instruction 85  
  ENQ macro instruction 104  
  EVENTS macro instruction 112  
  PGFIX macro instruction 134  
  PGFREE macro instruction 138  
  RESERVE macro instruction 149  
  SDUMP macro instruction 159

ECBIND parameter  
  PGFIX macro instruction 135  
  PGFREE macro instruction 139

emergency signal (ES) function 34

ENABLE parameter  
  CIRB macro instruction 95  
  MODESET macro instruction 129

ENQ macro instruction 103-107  
  execute form 107  
  list form 106  
  standard form 103-105  
  use of restricted parameters 22-24

ENTRIES parameter 112

EP parameter  
  ATTACH macro instruction 85  
  CIRB macro instruction 94  
  RISGNL macro instruction 153

EPLOC parameter 85

equipment check status condition 34

ERRET parameter 141

ES (emergency signal) function 34

ESTAE environment 46  
ESTAE exit routine 45-48  
ESTAE extended capabilities 43-45  
ESTAE macro instruction 108-111  
  execute form 111  
  list form 110  
  standard form 108-109  
  use of 43-49  
ESTAE retry routine 49  
ESTAI exit routine 45-48  
ESTAI parameter 85  
ESTAI retry routine 49  
ETXR parameter 85

EU parameter  
  FREEMAIN macro instruction 117  
  GETMAIN macro instruction 121

event completion 35-36  
  EVENTS 36  
  PGFIX 56  
  POST 35  
  SPOST 35

EVENTS macro instruction 112  
  use of 36

examples  
  ATTACH 87  
  CALLRTM 89  
  CHNGDUMP 92,93  
  CIRB 95  
  DEQ 98  
  DSGNL 102

ENQ 105  
ESTAE 109,111  
EXTRACT 114  
FREEMAIN 118  
GETMAIN 122  
MODESET 127  
NIL 131  
OIL 133  
PGFIX 135,136,137  
PGFREE 139,140  
POST 141  
PURGEDQ 144,146  
QEDIT 147  
RESERVE 150  
RISGNL 154  
RPSGNL 156  
SCHEDULE 157  
SDUMP 160,163  
SETFR 165  
SETLOCK 168,170,172  
SETRP 176  
SPIE 177  
SPOST 180  
STAE 183,185  
STATUS 189  
SYNCH 190  
TESTAUTH 192  
WTO 195  
WTOR 198

execute form of macro instruction  
  use of 80

extended STAE (ESTAE) 108-111

external call (EC) function 33

external call pending status condition 34

EXTKEY parameter 125

EXTRACT macro instruction 113-116  
  execute form 116  
  list form 115  
  standard form 113-114  
  use of 15-16,27

extract TCB information (EXTRACT) 113-116

FCTN parameter 191

FIELDS parameter 113

fix virtual storage contents (PGFIX) 134-137

frames, page 55

free virtual storage (FREEMAIN) 117-120

free virtual storage contents (PGFREE) 138-140

FREEMAIN macro instruction 117-120  
  disabling 22  
  execute form 120  
  list form 119  
  standard form 117-118  
  use of restricted parameters 53

FRELOCK parameter 175

FRESDDWA parameter 174

FRR stack 38

FRRAD parameter 164

FRRs (functional recovery routines) 38-39  
functional recovery routines (FRRs) 38-39  
  setting up 164-165

GENERIC parameter 96

GETMAIN macro instruction 121-124  
  execute form 124  
  list form 123  
  standard form 121-122  
  use of restricted parameters 53

GIVEJPO parameter 87

GLBRANCH 53

global dispatcher lock (DISP) 19

global locks 20  
  user-written SVC routines 69

global priority 64-65

global storage  
  subpools 53

global subpools 53

GSPL parameter 85

GSPV parameter 85

GTFCRM parameter 155  
  
 hardcopy log 60  
 HDR parameter 159  
 HDRAD parameter 159  
 hierarchy, locking 21  
  
 ICFBDF00 73  
 ICFBIF00 74  
 ICPUR parameter 101  
 ICTL instruction 80  
 IEAVTRML 50  
 IGFINTVL 71  
 IHAFRRS mapping macro instruction 39  
 IHAPSA mapping macro instruction 39  
 IHARMPL mapping instruction 50  
 IMPL parameter 101  
 initial CPU reset function 33  
 initial microprogram load function 33  
 initial program reset function 33  
 installation-written clean-up routines 50-51  
 interlock conditions 26  
 interprocessor communications 32-35  
 interruption handler, missing 71  
 interruption request block, creating 66  
 interruption, termination, and dumping services 37-51  
 interruption request block (IRB) 94,66  
 invalid function status condition 34  
 IOS channel availability table lock (IOSCAT) 19  
 IOS logical channel queue lock (IOSLCH) 19  
 IOS synchronization lock (IOSYNCH) 19  
 IOS unit control block lock (IOSUCB) 19  
 IOSCAT (IOS channel availability table lock) 19  
 IOSLCH (IOS logical channel queue lock) 19  
 IOSUCB (IOS unit control block lock) 19  
 IOSYNCH (IOS synchronization lock) 19  
 IPR parameter 101  
 IRB (interruption request block) 66  
 issue direct signal (DSGNL) 101-102  
 issue remote immediate signal (RISGNL) 153-154  
 issue remote pendable signal (RPSGNL) 155-156  
  
 job step authorization 31  
 JSCB parameter 86  
 JSTCB parameter 86  
  
 KEY parameter  
   ATTACH macro instruction 86  
   CIRB macro instruction 94  
   FREEMAIN macro instruction 118  
   GETMAIN macro instruction 122  
   MODESET macro instruction 127  
   TESTAUTH macro instruction 191  
 KEYADDR parameter 125  
  
 L parameter  
   FREEMAIN macro instruction 117  
   PGFIX macro instruction 137  
   PGFREE macro instruction 140  
 LA parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
   PGFIX macro instruction 137  
   PGFREE macro instruction 140  
 label processing 17  
 LAST parameter 112  
 LC parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
 link to SVC 34 (QEDIT) 147  
 linkage editor 30  
 list form of macro instruction  
   use of 80  
 LIST parameter 159  
 LOCAL (local storage lock) 20  
   GETMAIN and FREEMAIN 53  
   SVC routines 69  
 local locks 20  
 local priority 64-65  
 local storage  
   subpools 53  
 local storage lock (LOCAL) 20  
 locking 19-22  
   classes of locks 20  
   considerations 21-22  
   hierarchy 21  
   locks 19-20  
   manager 19  
   types of locks 20-21  
 LONG parameter 134  
 long-term fix 55  
 LPMOD parameter 85  
 LSQA parameter 87  
 LU parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
 LV parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
  
 machine check handler appendage  
   adding code to 72-74  
 macro instruction forms 80  
 master scheduler initialization module  
   adding code to 74-75  
 MC parameter 188  
 MCS (multiple console support) 60  
 MCSFLAG parameter  
   WTO macro instruction 194  
   WTOR macro instruction 198  
 message routing 60  
 message routing exit routines 60-63  
 MF parameter  
   DEQ macro instruction 99,100  
   ENQ macro instruction 106,107  
   ESTAE macro instruction 110,111  
   EXTRACT macro instruction 115,116  
   FREEMAIN macro instruction 119,120  
   GETMAIN macro instruction 123,124  
   MODESET macro instruction 128,129  
   POST macro instruction 142,143  
   PURGEDQ macro instruction 145,146  
   RESERVE macro instruction 151,152  
   SDUMP macro instruction 161,162  
   SPIE macro instruction 178,179  
   STAE macro instruction 184,185  
   WTO macro instruction 196  
   WTOR macro instruction 199  
 MFITCH parameter 155  
 miscellaneous services, supervisor 59-77  
 missing interruption handler 71  
 MODE parameter  
   CIRB macro instruction 94  
   MODESET macro instruction 127  
   RPSGNL macro instruction 155  
   SETLOCK macro instruction 167  
 MODESET macro instruction 125-129  
   disabling 21  
   execute form 129  
   list form 128  
   standard form 125-127  
   use of 32  
 MOUNT command 71  
 MP systems 29  
 MSGTYP parameter  
   WTO macro instruction 194  
   WTOR macro instruction 197  
 multiple console support (MCS) 60  
 multiple-line message 60  
 multiple locks 21  
 must complete function 22-24  
  
 ND parameter 188  
 NI (AND IMMEDIATE) instruction 130  
 NIL macro instruction 130-131

NODUMP parameter 90  
 nonquiesceable priority level 64  
 not ready status condition 34  
 NSHSPL parameter 86  
 NSHSPV parameter 86

OBTAIN parameter 166  
 OI (OR IMMEDIATE) instruction 132  
 OIL macro instruction 132-133  
 operator  
   writing to with reply 193-196  
   writing to without reply 197-199  
 operator intervening status condition 34  
 operator messages 59-60  
 OR IMMEDIATE (OI) instruction, providing a lock via 132-133  
 ORIGIN parameter 147  
 OV parameter  
   ESTAE macro instruction 108  
   STAE macro instruction 181

P parameter 164  
 page fixing 55  
 page frames 55  
 PARALLEL parameter 153  
 PARAM parameter  
   ATTACH macro instruction 85  
   ESTAE macro instruction 108  
   STAE macro instruction 181  
 PARM parameter 153  
 PARMAD parameter 165  
 partitioned data set  
   SVC routines 70  
 PDATA parameter 91  
 PGFIX macro instruction 134-137  
   list form 137  
   standard form 134-136  
   use of 55-56  
   virtual subarea list 56-57  
 PGFREE macro instruction 138-140  
   list form 140  
   standard form 138-139  
   use of 55-56  
   virtual subarea list 56-57  
 PGLOAD macro instruction 55  
   virtual subarea list 56-57  
 PGOUT macro instruction 55  
   virtual subarea list 56-57  
 PGRLSE macro instruction 55  
   virtual subarea list 56-57  
 POST macro instruction 141-143  
   execute form 143  
   list form 142  
   standard form 141  
   use of restricted parameters 35  
 POST, synchronizing 180  
 power warning feature support 72-77  
 PR parameter 101  
 priorities 64-65  
 program management 17  
 program reset function 33  
 provide a lock via an AND IMMEDIATE (NI) instruction (NIL) 130-131  
 provide a lock via an OR IMMEDIATE (OI) instruction (OIL) 132-133  
 PURGE parameter  
   ATTACH macro instruction 85  
   ESTAE macro instruction 108  
   STAE macro instruction 182  
 purge SRB activity (PURGEDQ) 144-146  
 PURGEDQ macro instruction 144-146  
   and SPOST 180  
   execute form 146  
   list form 145  
   standard form 144  
   use of 65-66

Q parameter 91

QEDIT macro instruction 147  
   use of 15  
 QUIESCE parameter 160

R parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
   PGFIX macro instruction 134  
   PGFREE macro instruction 138  
   SETFRR macro instruction 164  
 RBLEVEL parameter 191  
 RC parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
   SETTRP macro instruction 174  
 real storage manager and virtual storage manager space  
   allocation lock (SALLO) 19  
 real storage management 55-57  
 real storage manager 55  
 receiver check status condition 34  
 RECORD parameter  
   ESTAE macro instruction 108  
   SETTRP macro instruction 175  
 recovery environment 37  
 recovery/termination 37-38  
 recovery/termination manager (RTM) 37  
 recovery/termination manager, calling 88-89  
 RECPARM parameter 175  
 reenterable SVC routines 67  
 REF parameter  
   NIL macro instruction 130  
   OIL macro instruction 132  
 reference - macro instructions and commands 79-198  
 refreshable SVC routines 67  
 register (0)  
   meaning 81  
 register (1)  
   meaning 81  
 register (2) - (12)  
   meaning 81  
 REGS parameter  
   SETLOCK macro instruction 168,171  
   SETTRP macro instruction 174  
 RELATED parameter  
   ATTACH macro instruction 86  
   DEQ macro instruction 96  
   ENQ macro instruction 103  
   ESTAE macro instruction 108  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
   MODESET macro instruction 126  
   PGFIX macro instruction 134  
   PGFREE macro instruction 138  
   POST macro instruction 141  
   RESERVE macro instruction 149  
   SETFRR macro instruction 164  
   SETLOCK macro instruction 168  
   STAE macro instruction 182  
   STATUS macro instruction 186,189  
 release a serially reusable resource (DEQ) 96-100  
 RELEASE parameter  
   PGFIX macro instruction 134  
   PGFREE macro instruction 138  
   SETLOCK macro instruction 169  
 remote class 33  
 remote immediate signal, issuing 153-154  
 remote pendable signal, issuing 155-156  
 request control of a serially reusable resource (ENQ) 103-107  
 reserve a device (RESERVE) 148-152  
 RESERVE macro instruction 148-152  
   execute form 152  
   list form 151  
   standard form 148-150  
   use of 25-28  
 RESET parameter 188  
 resource control 19-36  
 restart function 33  
 RESTART parameter 101

RET parameter  
   DEQ macro instruction 96  
   ENQ macro instruction 103  
   RESERVE macro instruction 149  
 RETADDR parameter 174  
 RETIQC parameter 95  
 RETREGS parameter 174  
 RETRN parameter 95  
 return codes  
   ATTACH macro instruction 87  
   DEQ macro instruction 97-98  
   DSGNL macro instruction 102  
   ENQ macro instruction 104-105  
   ESTAE macro instruction 109  
   FREEMAIN macro instruction 118  
   GETMAIN macro instruction 122  
   PGFIX macro instruction 135  
   PGFREE macro instruction 139  
   RESERVE macro instruction 149-150  
   RSIGNL macro instruction 153-154  
   RPSGNL macro instruction 156  
   SDUMP macro instruction 160  
   SETLOCK macro instruction 168,170,172  
   STAE macro instruction 183  
   TESTAUTH macro instruction 192  
 return parameters, setting 173-176  
 RISGNL macro instruction 153-154  
   use of 34  
 RMC parameter 96  
 RMTR parameter 144  
 ROUTCDE parameter  
   WTO macro instruction 193  
   WTOR macro instruction 197  
 routing codes 61  
 RPSGNL macro instruction 155-156  
   use of 34  
 RQCHECK parameter 155  
 RTM (recovery/termination manager) 37  
 RU parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
 RUB parameter 174  
 RX-type address  
   meaning 81

**S** parameter  
   ENQ macro instruction 103  
   EXTRACT macro instruction 113  
   RESERVE macro instruction 149  
 SALLOC (real storage manager and virtual storage  
   manager space allocation lock) 19  
   GETMAIN and FREEMAIN 53  
 SAVEKEY parameter 126  
 SCHEDULE macro instruction 157  
   and PURGEDQ 144  
   use of 63-64  
 schedule system services for asynchronous execution  
   (SCHEDULE) 157  
 SCOPE parameter 157  
 SD parameter 188  
 SDATA parameter  
   CHNGDUMP command 91  
   SDUMP macro instruction 159  
 SDUMP dump requests 40  
 SDUMP macro instruction 158-163  
   execute form 162-163  
   list form 161  
   standard form 158-160  
   use of 39-40  
 SDUMP parameter 90  
 sense function 33  
 SENSE parameter 101  
 SERIAL parameter 153  
 serialization 67  
 serially reusable resources  
   controlling access to 166-172  
   locking 19-22  
   releasing 96-100  
   requesting control of 103-107  
 service classes 33-34  
 service management 63-66  
 service request block (SRB) 63,64  
 SET parameter 188  
   CHNGDUMP command 90  
   STATUS macro instruction 188  
 set return parameters (SETRP) 173-176  
 set up functional recovery routines (SETFRR) 164-165  
 SETCODE statement 30  
 SETFRR macro instruction 164-165  
   use of 38-39  
 SETLOCK macro instruction 166-172  
   use of 22  
 SETRP macro instruction 173-176  
 shared DASD 24-29  
 SHSPL parameter 85  
 SHSPV parameter 85  
 signal event completion (POST) 141-143  
 signal processor (SIGP) instruction  
   DSGNL macro instruction 101  
   RISGNL macro instruction 153  
   RPSGNL macro instruction 155  
 SIGP (signal processor) instruction 32-35  
 single system locks 21  
 SIO parameter 155  
 SM parameter 86  
 SMC parameter 104  
 SP parameter  
   FREEMAIN macro instruction 117  
   GETMAIN macro instruction 121  
   special clean-up processing 50  
   specify program interruption exit (SPIE) 177-179  
   specify task abnormal exit (STAE) 181-185  
 SPIE macro instruction 177-179  
   execute form 179  
   list form 178  
   standard form 177  
   use of restricted interruptions 39  
 spin locks 20  
   SVC routines 69  
 SPOST macro instruction 180  
   use of 35  
 SQA buffer 40  
 SRB (service request block) 63,64  
 SRB activity, purging 144-146  
 SRB parameter  
   SCHEDULE macro instruction 157  
   STATUS macro instruction 186  
 SRM (system resource manager lock) 20  
 SSM instruction 21  
 SSS parameter 101  
 STAB parameter 95  
 STAE environment 41-42  
 STAE exit routines 41-42  
 STAE macro instruction 181-185  
   execute form 185  
   list form 184  
   standard form 181-183  
   use of 41-43  
 STAE retry routine 42-43  
 stage 1 exit effector 66  
 stage 2 exit effector 66  
 stage 3 exit effector 66  
 STAI environment 41-42  
 STAI exit routines 41-42  
 STAI parameter 85  
 STAI retry routine 42-43  
 start function 33  
 START parameter  
   DSGNL macro instruction 101  
   STATUS macro instruction 186  
 STATE parameter 191  
 status conditions 34-35  
 STATUS macro instruction 186-189  
 STEP parameter  
   CALLRTM macro instruction 89  
   DEQ macro instruction 96  
   ENQ macro instruction 103  
   STATUS macro instruction 188  
 stop and store status function 33  
 stop function 33  
 STOP parameter  
   DSGNL macro instruction 101  
   STATUS macro instruction 186

- stopped status condition 34
- STORAGE parameter 159
- subpools
  - common service area 53
  - fetch protected 53
  - global 53
  - not fetch protected 53
  - system queue area 53
  - 227 53
  - 228 53
  - 229 53
  - 230 53
  - 231 53
  - 236 13
  - 237 13
  - 239 53
  - 241 53
  - 245 53
- subtask creation and control 13-16
- subtask status, changing 186-189
- supervisor services 11-78
- suspend locks 20
- SUSPEND parameter 135
- SVAREA parameter
  - ATTACH macro instruction 86
  - CIRB macro instruction 95
- SVC DUMP 39-40
- SVC routines, user-written 67-71
  - characteristics 67
  - inserting 70
  - programming conventions 67-70
  - specifying 70
- SVC 34, linking to 147
- SVCTABLE macro instruction 70
- SVEAREA parameter 108
- SVRB
  - SVC routines 70
- SWITCH parameter 155
- symbol
  - meaning 81
- SYNCH macro instruction 189
  - use of 17
- SYNCH parameter 186
- synchronize POST (SPOST) 180
- synchronous exits, taking 190
- SYSABEND dump requests 40
- SYSABEND parameter 90
- SYSMASK parameter 129
- system generation
  - SVC routines 70
- system log 61
- SYSTEM parameter
  - DEQ macro instruction 96
  - ENQ macro instruction 103
  - STATUS macro instruction 188
- system queue area subpools 53
- system resource manager lock (SRM) 20
- system status 125-129,32
- SYSTEMS parameter
  - DEQ macro instruction 96
  - ENQ macro instruction 103
  - RESERVE macro instruction 149
- SYSUDUMP dump requests 140
- SYSUDUMP parameter 90
- SYS1.LINKLIB
  - power warning feature support 74
  - SVC routines 67
- SYS1.LPALIB
  - clean-up routines 50
  - SVC routines 67
- SYS1.NUCLEUS
  - power warning feature support 73
- SYS1.SVCLIB
  - SVC routines 67
- SZERO parameter 85
- TABLE parameter 112
- take asynchronous exit to a processing program (SYNCH) 190
- task abnormal exit
  - extended 108-111
  - specifying 181-185
- task, creating a new 34-87
- task recovery 41-49
- task termination 37-38
- TASKLIB parameter 85
- TCB information, extracting 113-116
- TCB parameter
  - CALLRTM macro instruction 88
  - DEQ macro instruction 97
  - ENQ macro instruction 104
  - STATUS macro instruction 186
- TERM parameter
  - ATTACH macro instruction 85
  - ESTAE macro instruction 108
- test authorization of caller (TESTAUTH) 191-192
- TEST parameter 171
- TEST UNDER MAST (TM) instruction 35
- TESTAUTH macro instruction 191-192
  - use of 30
- TID parameter 86
- time interval for missing interruption handler 71
- TM (TEST UNDER MASK) instruction 35
- TSOLOGON parameter 87
- TYPE parameter
  - CALLRTM macro instruction 88
  - CIRB macro instruction 95
  - SETLOCK macro instruction 167,169
- types of locks 20-21
- UCB parameter
  - DEQ macro instruction 97
  - RESERVE macro instruction 149
- user-written message routing exit routines 60-63
- user-written SVC routines 67-71
- V parameter 117
- VC parameter
  - FREEMAIN macro instruction 117
  - GETMAIN macro instruction 121
- virtual storage
  - allocating 121-124
  - dumping 158-163,39-40
  - freeing 117-120
- virtual storage contents
  - fixing 134-137,55-56
  - freeing 138-140,55-56
- virtual storage, dumping 39-40
- virtual storage management 53
- virtual subarea list (VSL) 56-57
- VL parameter 85
- volume handling 25
- VSL (virtual subarea list) 56-57
- VU parameter
  - FREEMAIN macro instruction 117
  - GETMAIN macro instruction 121
- | wait for events (EVENTS) 112
- WAIT macro instruction
  - SVC routines 70
- | WAIT parameter 112
- WKAREA parameter
  - CIRB macro instruction 95
  - SETRP macro instruction 174
- WORKREG parameter 126
- WREGS parameter
  - NIL macro instruction 130
  - OIL macro instruction 132
- write-to-operator (WTO) 193-196
- write-to-operator-with-reply (WTOR) 197-199
- WRKREGS parameter 164
- WTO macro instruction 193-196
  - list form 96
  - standard form 193-195
  - use of restricted parameters 59-63

WTOR macro instruction 197-199  
list form 199  
standard form 197-198  
use of restricted parameters 59-63

XCTL macro instruction  
SVC routines 70  
XCTL parameter  
ESTAE macro instruction 108  
STAE macro instruction 181



**International Business Machines Corporation**  
**Data Processing Division**  
1133 Westchester Avenue, White Plains, New York 10604  
(U.S.A. only)

**IBM World Trade Corporation**  
821 United Nations Plaza, New York, New York 10017  
(International)



OS/VS2 System Programming Library:  
Supervisor  
GC28-0628-0

**READER'S  
COMMENT  
FORM**

*Your views about this publication may help improve its usefulness; this form will be sent to the author's department for appropriate action. Using this form to request system assistance or additional publications will delay response, however. For more direct handling of such requests, please contact your IBM representative or the IBM Branch Office serving your locality.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Index Figures Examples Legibility

Cut or Fold Along Line

Are the tables used to describe the macro instructions in this publication an improvement over the brackets-and-braces syntax?

What is your occupation? \_\_\_\_\_

Number of latest Technical Newsletter (if any) concerning this publication: \_\_\_\_\_

Please indicate your address in the space below if you wish a reply.

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.)

Cut or Fold Along Line

**Your comments, please . . .**


This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. Your comments on the other side of this form will be carefully reviewed by the persons responsible for writing and publishing this material. All comments and suggestions become the property of IBM.

Fold

Fold

First Class  
Permit 81  
Poughkeepsie  
New York

**Business Reply Mail**  
No postage stamp necessary if mailed in the U.S.A.



Postage will be paid by:

International Business Machines Corporation  
Department D58, Building 706-2  
PO Box 390  
Poughkeepsie, New York 12602

Fold

Fold

OS/VS2 SPL: Supervisor (S370-36)

Printed in U.S.A.

GC28-0628-0



**International Business Machines Corporation**  
**Data Processing Division**  
1133 Westchester Avenue, White Plains, New York 10604  
(U.S.A. only)

**IBM World Trade Corporation**  
821 United Nations Plaza, New York, New York 10017  
(International)

OS/VS2 System Programming Library:  
Supervisor  
GC28-0628-0

**READER'S  
COMMENT  
FORM**

*Your views about this publication may help improve its usefulness; this form will be sent to the author's department for appropriate action. Using this form to request system assistance or additional publications will delay response, however. For more direct handling of such requests, please contact your IBM representative or the IBM Branch Office serving your locality.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Index Figures Examples Legibility

Cut or Fold Along Line

Are the tables used to describe the macro instructions in this publication an improvement over the brackets-and-braces syntax?

What is your occupation? \_\_\_\_\_

Number of latest Technical Newsletter (if any) concerning this publication: \_\_\_\_\_

Please indicate your address in the space below if you wish a reply.

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.)

Cut or Fold Along Line

**Your comments, please . . .**


This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. Your comments on the other side of this form will be carefully reviewed by the persons responsible for writing and publishing this material. All comments and suggestions become the property of IBM.

Fold

Fold

First Class  
Permit 81  
Poughkeepsie  
New York

**Business Reply Mail**  
No postage stamp necessary if mailed in the U.S.A.



Postage will be paid by:

International Business Machines Corporation  
Department D58, Building 706-2  
PO Box 390  
Poughkeepsie, New York 12602

Fold

Fold

OS/VS2 SPL: Supervisor (S370-36)

Printed in U.S.A.

GC28-0628-0



**International Business Machines Corporation**  
Data Processing Division  
1133 Westchester Avenue, White Plains, New York 10604  
(U.S.A. only)

**IBM World Trade Corporation**  
821 United Nations Plaza, New York, New York 10017  
(International)