

OS/32 SYSTEM LEVEL

Programmer Reference Manual

OS/32 Version 8.2 or Higher

48-040 F00 R05



The information in this document is subject to change without notice and should not be construed as a commitment by Concurrent Computer Corporation. Concurrent Computer Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license, and it can be used or copied only in a manner permitted by that license. Any copy of the described software must include any copyright notice, trademarks, or other legends or credits of Concurrent Computer Corporation and/or its suppliers. Title to and ownership of the described software and any copies thereof shall remain in Concurrent Computer Corporation and/or its suppliers.

The licensed program described herein may contain certain encryptions or other devices which may prevent or detect unauthorized use of the Licensed Software. Temporary use permitted by the terms of the License Agreement may require assistance from Concurrent Computer Corporation.

Concurrent Computer Corporation assumes no responsibility for the use or reliability of the software on equipment that is not supplied by Concurrent Computer Corporation.

Reliance is a trademark of Concurrent Computer Corporation.

© 1981, 1983, 1984, 1986 Concurrent Computer Corporation — All Rights Reserved

Concurrent Computer Corporation, 2 Crescent Place

Oceanport, New Jersey 07757

Printed in the United States of America

TABLE OF CONTENTS

PREFACE		v	
CHAPTERS			
1	OS/32 SUBSYSTEMS		
1.1	INTRODUCTION	1-1	
1.1.1	OS/32 Multiprocessing Support	1-2	
1.2	SOFTWARE SUBSYSTEMS	1-3	
1.2.1	Task Management Subsystem	1-6	
1.2.1.1	Task Scheduling on the 3200MPS Family of Processors	1-11	
1.2.2	Job Accounting Subsystem	1-14	
1.2.3	Memory Management Subsystem	1-14	
1.2.4	Timer Management Subsystem	1-15	
1.2.5	File Management Subsystem	1-16	
1.2.6	Input/Output (I/O) Subsystem	1-17	
1.2.7	Error Recording Subsystem	1-17	
1.2.8	Memory Diagnostics Subsystem	1-18	
1.2.9	Loader and Segmentation Subsystem	1-18	
1.2.10	Basic Data Communications Subsystem	1-19	
1.2.11	Console Monitor Subsystem	1-19	
1.2.12	Command Processor Subsystem	1-20	
1.2.13	System Initialization Subsystem	1-20	
1.2.14	Internal Interrupt Subsystem	1-21	
1.2.15	Optional User Supervisor Call (SVC) Subsystem	1-21	
1.2.16	Floating Point Subsystem	1-21	
2	PRIVILEGED TASKS		
2.1	INTRODUCTION	2-1	
2.2	EXECUTIVE TASKS (E-TASKS)	2-2	
2.2.1	Relocatable Executive Tasks (E-Tasks)	2-2	
2.2.2	Writing Executive Tasks (E-Tasks)	2-2	
2.2.3	Writing Relocatable Executive Tasks (E-Tasks)	2-3	
2.2.4	OS/32 Data Structures Used by Executive Tasks (E-tasks)	2-4	
2.3	PRIVILEGED USER TASKS (U-TASKS)	2-8	
2.4	DIAGNOSTIC TASKS (D-TASKS)	2-9	

CHAPTERS (Continued)

	3	PROGRAMMING IN THE 3200MPS FAMILY OF PROCESSORS MULTIPROCESSING ENVIRONMENT	
	3.1	INTRODUCTION	3-1
	3.2	DESIGNING TASKS TO RUN ON A MULTIPROCESSING SYSTEM	3-2
	3.3	PREPARING AN AUXILIARY PROCESSING UNIT (APU) FOR TASK EXECUTION	3-3
	3.3.1	Queue Assignments	3-3
	3.3.2	Auxiliary Processing Unit (APU) Operating States	3-3
	3.3.3	APU-Only Queue Operating States	3-4
	3.3.4	Logical Processing Unit (LPU) Mapping	3-6
	3.4	ASSIGNING TASKS TO A PROCESSOR QUEUE	3-7
	3.5	CONTROLLING TASK ORDER OF EXECUTION	3-8
	3.5.1	Changing Auxiliary Processing Unit (APU) Task Queue Ordering	3-8
	3.5.2	Monitoring and Preempting Auxiliary Processing Unit (APU) Task Execution	3-10
	3.5.3	Transferring a Task from an Auxiliary Processing Unit (APU) to the Central Processing Unit (CPU)	3-16
	3.5.4	Internal Task Control of Auxiliary Processing Unit (APU) Execution	3-17
	3.5.5	Verifying Task Transfer to an Auxiliary Processing Unit (APU)	3-18
	3.5.6	Customizing Auxiliary Processing Unit (APU) Fault and Supervisor Call (SVC) Handling	3-19
	3.6	PREVENTING MEMORY ACCESS CONFLICTS	3-20
	3.6.1	Avoiding System Deadlock	3-21
	3.7	MEASURING REAL-TIME PERFORMANCE ON THE 3200MPS FAMILY OF PROCESSORS	3-22
	3.8	WHERE TO GO FOR MORE INFORMATION	3-24
	4	SUPERVISOR CALL (SVC) INTERCEPTION	
	4.1	INTRODUCTION	4-1
	4.2	HOW SUPERVISOR CALL (SVC) INTERCEPTION WORKS	4-2
	4.3	PREPARING A TASK FOR SUPERVISOR CALL (SVC) INTERCEPTION	4-3
	4.3.1	Request Descriptor Block (RDB) Buffers	4-4
	4.3.2	Circular List for Request Descriptor Block (RDB) Buffers	4-6

CHAPTERS (Continued)

4.3.3	Task Event Trap	4-8
4.4	CREATING INTERCEPT PATHS (ICREATE)	4-8
4.5	HOW TO CREATE A PSEUDO DEVICE OR TASK WITH ICREATE	4-8
4.6	USE OF GENERIC NAMING FOR PSEUDO DEVICES AND TASKS	4-9
4.7	FUNCTIONAL SUMMARY OF SUPERVISOR CALL (SVC) INTERCEPTION	4-10
4.8	FULL AND MONITOR CONTROL INTERCEPT PATHS	4-11
4.9	HOW INTERCEPT PATHS HANDLE SUPERVISOR CALLS (SVCs) OCCURRING AT END OF TASK	4-13
4.10	TERMINATING THE INTERCEPTED SUPERVISOR CALLS (SVCs)	4-13
4.11	HOW TO REMOVE INTERCEPT PATHS	4-14
4.12	ERROR HANDLING	4-14
4.13	MACROS USED WITH SUPERVISOR CALL (SVC) INTERCEPTION	4-16
4.13.1	ICREATE Macro	4-16
4.13.2	IREMOVE Macro	4-24
4.13.3	IGET Macro	4-25
4.13.4	IPUT Macro	4-27
4.13.5	ICONT Macro	4-28
4.13.6	IPROCEED Macro	4-29
4.13.7	IROLL Macro	4-30
4.13.8	ITERM Macro	4-31
4.13.9	ITRAP Macro	4-32
4.13.10	IERRTST Macro	4-34
4.13.11	\$RDB Macro	4-36
4.14	SAMPLE SUPERVISOR CALL (SVC) INTERCEPTION PROGRAMS	4-36

APPENDIXES

A	SUPPORTED VERTICAL FORMS CONTROL (VFC) CHARACTER SET	A-1
---	--	-----

FIGURES

1-1	Typical 3200MPS Family of Processors Configuration	1-3
3-1	Valid APU Operating States	3-4
3-2	Valid APU Queue Operating States	3-5
4-1	Request Descriptor Block	4-4
4-2	System Task Buffer List (Standard Circular List)	4-7

TABLES

1-1	OS/32 SOFTWARE SUPPORT	1-4
2-1	OS/32 DATA STRUCTURES MACRO LIBRARY	2-5
2-2	MTM DATA STRUCTURES MACRO LIBRARY	2-8
3-1	QUEUE PRIORITY ASSIGNMENTS	3-9
3-2	TIMER MACROS	3-23
3-3	ADDITIONAL INFORMATION SOURCES FOR THE 3200MPS FAMILY OF PROCESSORS PROGRAMMING	3-25
4-1	SYSTEM MACROS FOR SVC INTERCEPTION	4-1
4-2	ERROR CODES RETURNED FOR INTERCEPT MACROS	4-15
4-3	VALID COMBINATIONS FOR SVC, MODE AND NAME PARAMETERS	4-20

INDEX

IND-1

PREFACE

This manual describes operating system features intended for use by system programmers, system analysts, designers, engineers and training instructors.

Chapter 1 presents an overview of the operating system and the software subsystems it supports. Chapter 2 describes the privileged task types supported by OS/32. Chapter 3 describes the techniques used in writing system level control programs that take advantage of the increased throughput offered by the 3200MPS Family of Processors. Chapter 4 contains a functional description of the supervisor call (SVC) interception feature. The vertical forms control (VFC) feature is described in Appendix A.

The R05 revision of this manual includes two new operands in the ICREATE macro. The NOTIFY operand distinguishes an intercepting task as either an entry or exit intercept for use with the AUDIT32 security disk/tape audit utility. The RDB.RID field in the request descriptor block (RDB) has been utilized for this task. The IOPT operand resolves system task deadlocks by removing a task from wait state. Also the SVC= operand of the ICREATE macro has been enhanced to permit an extended option to the SVC7 parameter. This enhancement includes the addition of an SVC7 extended block (RDB.EXT) after the SVC7 parameter block in the RDB. The RDB.EXT field is utilized to hold the name of a file assigned to a logical unit (lu) that is to be passed to the AUDIT32 utility.

Revision F00 R05 is intended for use with the OS/32 R08.2 software release or higher.

CHAPTER 1 OS/32 SUBSYSTEMS

1.1 INTRODUCTION

OS/32 is a general-purpose, event-driven operating system for 32-bit computer systems. Custom versions of OS/32 are created through the use of a system generation program (Sysgen/32) that provides parameters for tailoring OS/32 to a specific installation. The combined hardware and software capabilities of a 32-bit computer system provide support for all phases of program and system development. OS/32 supports concurrent multiprogramming, with up to 252 user programs written in any of the supported languages. The program development facilities are designed to minimize the time and effort needed to test, debug and integrate application programs and systems. In addition, the OS/32 command language allows complex jobs to be performed with minimum operator intervention.

OS/32 incorporates a powerful interrupt handling capability at the task level. This capability permits a task to be interrupted during its normal execution sequence by a variety of hardware and software conditions.

The OS/32 virtual task manager (VTM) allows the memory requirements of a task running under OS/32 to exceed available task memory. VTM has a virtual memory capability that allows tasks consisting of up to 16MB of code and data to execute in as little as 128kB of memory. This feature is provided by the OS/32 linkage editor. See the OS/32 Link Reference Manual for more information.

The roll function allows segments of a task to be rolled out to disk until enough memory is available for the entire task. In real-time applications, rolling is commonly used to queue low priority tasks while tasks of higher priority are active. The roll eligibility of a task is established when the task is link-edited. However, a task option is provided to prevent rolling of a task when necessary (e.g., when the task must be able to respond to real-time events).

A basic data communications facilities package is supplied with OS/32. This package also provides support for higher level data communications products.

The scope and power of the operating system can be extended through the following OS/32 companion products:

- Multi-terminal monitor (MTM)
- Reliance™

MTM is a subsystem monitor that uses the subtasking capabilities of OS/32 to provide a time-sliced, interactive program development environment for up to 64 concurrent terminal users. MTM simultaneously supports both on-line terminal users and batch background tasks. MTM terminal users are also provided with an input/output (I/O) spooler for use with slow speed devices.

Reliance is a transaction software system consisting of the integrated transaction controller (ITC), data management system (DMS/32) and industry standard COBOL. ITC allocates system resources, develops screen formats and controls terminals. DMS/32 supervises disk allocation and data access.

1.1.1 OS/32 Multiprocessing Support

| OS/32 provides a transparent multiprocessing capability for use
| with the 3200MPS Family of Processors. This system consists of
| one central processing unit (CPU) and any combination of one to
| nine auxiliary processing units (APUs) or input/output processors
| (IOPs) (see Figure 1-1). A task can execute on an APU unless it
| is going to take advantage of certain features specific to the
| multiprocessing system (e.g., APU assignment, APU control, etc.).
| IOPs are processors specifically designed to handle the I/O
| portion of I/O intensive tasks. These processors cut down the
| overhead incurred by the CPU when handling I/O requests,
| therefore, enhancing system performance.

OS/32 defines a set of logical processing units (LPUs) that are used to schedule tasks to APU queues for execution on an APU. Tasks are assigned to an LPU that is mapped to an APU queue. The logical processor mapping table (LPMT) defined by OS/32 contains the mapping arrangement between the LPUs and APU execution queues. (More than one LPU can be mapped to an APU queue.)

| Each APU and/or IOP in the 3200MPS Family of Processors is
| assigned a unique identifying number. Each APU is assigned to an
| APU queue. (More than one APU can be assigned to a queue.)
| Tasks on an APU queue execute on an APU assigned to the queue.

If a task is mainly computation intensive, executing that task on an APU increases overall system performance. If a task is I/O-intensive, execution should be performed on the CPU. This

Reliance is a trademark of Concurrent Computer Corporation.

will also increase system performance. IOPs perform physical I/O, they do not perform I/O to the console or to Integrated Telecommunications Access Method (ITAM) devices. An I/O-intensive task executing on an APU is transferred to the CPU when an I/O request is encountered. The task is then transferred to an IOP for I/O servicing. The IOP does the physical I/O and sends it back to the APU via the CPU. An I/O-intensive task, if directed solely to an APU, decreases system performance since each I/O request requires the task to be transferred back to the CPU for OS/32 I/O support services.

The main performance advantage of a multiprocessing system is achieved when a problem is broken down into parts so that several tasks on several processors can work on the problem at the same time or when multiple independent tasks can be executed simultaneously. See Chapter 3 for more information on programming within the 3200MPS Family of Processors environment.

040-1

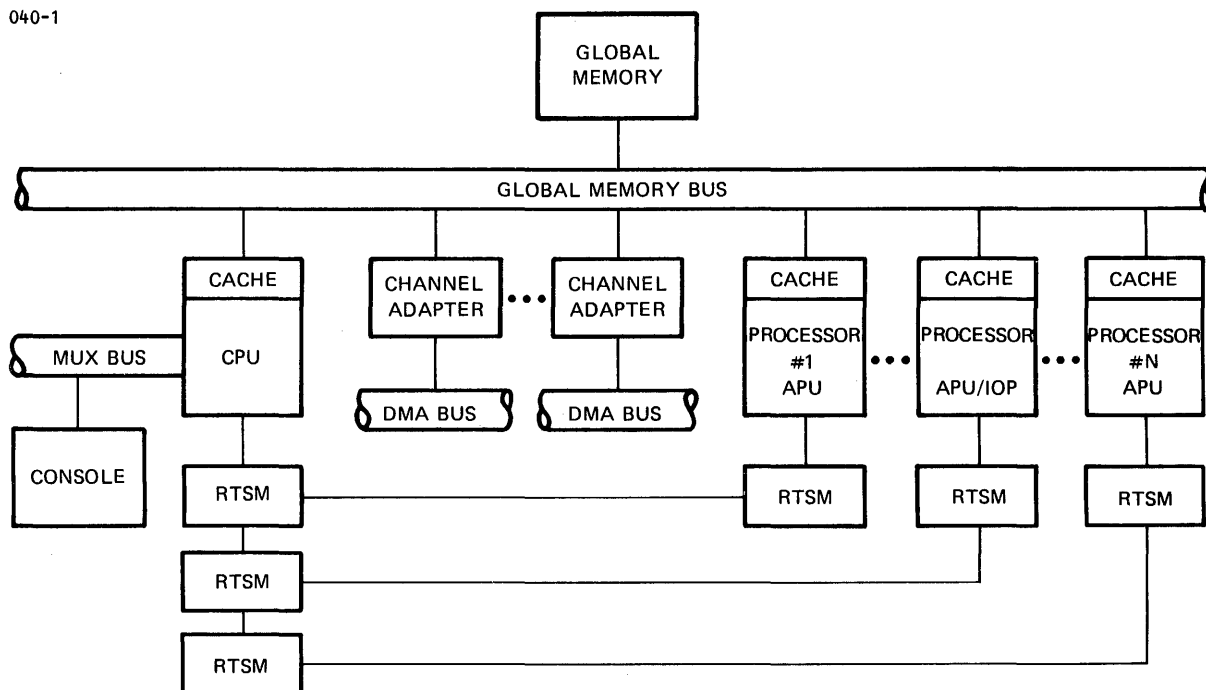


Figure 1-1 Typical 3200MPS Family of Processors Configuration

1.2 SOFTWARE SUBSYSTEMS

OS/32 consists of the following subsystems:

- Task management
- Job accounting
- Memory management
- Timer management
- File management

- I/O management
- Error recording and reporting
- Memory diagnostics
- Loader and segmentation
- Basic communications
- Console monitor
- Command processor
- System initialization
- Internal interrupt
- Optional user supervisor call 14 (SVC14)
- Floating point

Table 1-1 summarizes the software supported by OS/32.

TABLE 1-1 OS/32 SOFTWARE SUPPORT

TYPE	SOFTWARE PRODUCT	STANDARD	OPTIONAL
Program development	Task management	X	
	Job accounting	X	
	Memory management	X	
	Timer management	X	
	File management	X	
	I/O management	X	
	Error recording and reporting	X	
	Memory diagnostics	X	
	Loader and segmentation	X	
	Console monitor	X	
	Command processor	X	
	Floating point	X	
	Internal interrupt subsystem	X	
	ITC*		X
Writable control store (WCS)		X	
MTM		X	
Program debugging	Automatic interactive debugging system (AIDS)	X	
	DEBUG/32		X
Data base management	DMS*		X
Data communications	Asynchronous data communications	X	
	Character synchronous communications	X	
	Bit synchronous communications		X
	2780/3780 RJE emulation		X
	3270 emulation		X
	HASP/32		X
	Ethernet communications	X	

TABLE 1-1 OS/32 SOFTWARE SUPPORT (Continued)

TYPE	SOFTWARE PRODUCT	STANDARD	OPTIONAL
Languages	Common microcode assembler (MICROCAL)		X
	Common assembly language/32 (CAL/32)	X	
	CAL macro/32	X	
	FORTTRAN VII development (D) compiler		X
	FORTTRAN VII global optimizing (O) compiler		X
	FORTTRAN VII universal optimizing (Z) compiler		X
	COBOL*		X
	BASIC Level II		X
	CORAL 66		X
	RPG II		X
	Pascal		X
	C on OS/32		X
	Utilities	Link	X
Edit		X	
Medit			X
Text			X
Source Updater		X	
Copy		X	
Library Loader		X	
Macro Library		X	
Sort/Merge II			X
Patch		X	
OS/32 Spooler		X	
SPL/32		X	
Fastchek		X	
Fastback		X	
Fastcopy		X	
Account Reporting		X	
Backup		X	
Error Reporting		X	
Disk Dump		X	
Dump Print		X	
Mirror Disk Synchronization			
AUDIT32		X	
Virtual Console Facility (VCF)		X	
System Performance Monitor (SPM)	X		

* ITC, COBOL and DMS/32 comprise the Reliance software system designed for transaction processing.

1.2.1 Task Management Subsystem

The task management subsystem consists of five major parts:

- Self-initialization
- Task dispatcher
- Machine state control
- Task queueing
- Task trap support

Self-initialization is invoked during system initialization. It sets up the subsystem for those aspects of machine architecture pertinent to the task manager (number of register sets and floating point support) and system configuration (number of tasks). This process also defines console monitor and command processor subsystems as tasks.

The task dispatcher allocates processor time and activates memory space allocation for the tasks executing in an OS/32 multitasking memory multiplexing environment. The task dispatcher is first invoked upon completion of system initialization to dispatch the command processor and console monitor tasks. It is subsequently invoked upon conclusion of every scheduling event; i.e., (every interrupt signal).

The Task Dispatcher:

- selects the highest priority task from the current task, CPU ready tasks and ready to roll-in tasks. When no current CPU ready or ready to roll-in task exists, the task dispatcher selects a task off the top of APU queue 0 (for 3200MPS Family of Processors only).
- enforces any task status set while the task was outside the control of the dispatcher; i.e., (cancel, pause, roll-in, etc.).
- transfers the task to an APU queue (3200MPS Family of Processors only) based on its LPU and existing LPU to queue mapping arrangement.
- sets up the task's time slice.
- switches task context (program status word (PSW), registers) and transfers control to the selected task.

Several task queues are maintained under task management:

- CPU ready queue, a sequence of the CPU ready tasks in priority order.
- Roll-in queue, a sequence of the ready to roll-in tasks in priority order.
- APU queues, (3200MPS Family of Processors only) numbered 0 through n, where n is the number of APUs (maximum of 9). Each APU queue is a sequence of APU ready tasks in a priority or first-in/first-out (FIFO) order.

The machine state control portion of the task management subsystem is invoked every time there is a need in the operating system to change the PSW to one of the following established values:

- USER TASK (UT) state; memory relocation protection (MAC/MAT/VAT) is enabled, privileged instructions are illegal, SVCs are legal, system queue service (SQS) is enabled and register set X'F' is selected.
- EXECUTIVE TASKS (ET) state; memory relocation protection (MAC/MAT/VAT) is enabled, privileged instructions and SVCs are legal, SQS is enabled and the user register set X'F' is selected.
- DIAGNOSTIC TASKS (DT) state; memory relocation protection (MAC/MAT/VAT) is enabled, privileged instructions and SVCs are legal, SQS is enabled and the register set X'F' is selected.
- REENTRANT SYSTEM (RS) state; memory relocation protection (MAC/MAT/VAT) is disabled, privileged instructions and SVCs are legal and SQS is enabled, except when accessing system structures. On a 2-register set machine, register X'F' of register set 6 is selected.
- REENTRANT SYSTEM ALTERNATE (RSA) state; memory relocation protection (MAC/MAT/VAT) is disabled, privileged instructions and SVCs are legal and SQS is enabled except when accessing system structures. RSA is similar to RS state except an alternate context block is utilized; it is normally used to access supervisor services from the supervisor itself.
- NONREENTRANT SYSTEM (NS) state; memory relocation protection (MAC/MAT/VAT) is disabled, privileged instructions are legal, SVCs are illegal, SQS is disabled and registers 8 through X'F' of register set 0 are used. This is used to execute noninterruptible portions of the operating system.

- NONREENTRANT USER (NSU) state; memory relocation protection (MAC/MAT/VAT) is disabled, privileged instructions are legal, SVCs are illegal, SQS is disabled and registers 0 through X'F' of register set X'F' are used. This is utilized by the OS/32 command processor.
- EVENT SERVICE (ES) state; memory relocation protection (MAC/MAT/VAT) are disabled, privileged instructions are legal, SVCs are illegal and SQS is disabled. On a 2-register set machine, register X'F' of register set 5 is selected.
- INTERCEPT SERVICES (IS) state; memory relocation protection is disabled, privileged instructions are legal, SVCs are illegal, SQS is disabled and registers 0 through 7 of register set 0 are used. This state is used by drivers.
- INTERCEPT SERVICE USER (ISU) state; relocation protection (MAC/MAT/VAT) are disabled, privileged instructions are legal, SVCs are illegal, SQS is disabled and registers 0 through X'F' of register set X'F' are used.

NOTE

For more information regarding PSW settings, see the Processor User Manuals or Instruction Set Reference Manuals.

OS/32 allows a task to specify private interrupt processing or trap routines. The task supplies a task status word (TSW) for each supported trap in a dedicated area known as the user-dedicated location (UDL). If a trap condition occurs, the operating system stores the current TSW, loads the respective new TSW from the UDL and transfers control accordingly. Upon completion, the trap routine restores the current TSW and returns control back to the main routine.

The task management subsystem allocates processor time for each of the tasks executing in an OS/32 multitasking environment. The task manager determines the order in which each task gains processor control on a user-defined priority basis. Task priority levels range from 0 to 254 (0 being the highest priority level). Of these 255 priority levels, 10 through 249 are available for user-written tasks, while 1 through 9 and 250 through 254 are reserved for system use.

The task manager maintains four priority levels for each task:

- Maximum
- Task

- Run
- Dispatch

Maximum priority, set by Link, is the highest priority level (i.e., the smallest number) that can be assigned to a task. Task priority is the priority that is currently assigned to a task. Initially, task priority is set when the task is linked, but this priority can be changed after the task is loaded. Task priority can never be set higher than the maximum priority set by Link.

Run priority can be set dynamically to a value ranging from the task priority to task priority plus n. The value of n is based on the behavior of the task. Run priority can only be set for tasks that have dynamic time-slice/priority scheduling enabled. If dynamic scheduling is not enabled, a task's run priority is equal to its task priority. Currently, only MTM enables dynamic time-slice/priority scheduling.

A dispatched task usually has a priority level equal to its task priority even if dynamic scheduling is enabled. Nevertheless, if a higher priority task requires specific system resources (e.g., a disk directory or bit map) that are currently controlled by a lower priority task, the dispatch priority of this lower priority task is raised to the priority of the higher priority task waiting for the resource. When a task releases control of a system resource, its dispatch priority is reset to its run priority.

Tasks competing for processor time are maintained in priority order on a task control block (TCB) queue known as the ready queue. Tasks competing for both memory and processor time are maintained in priority order on the roll-in queue. Tasks at the same priority level are serviced on a round-robin basis; i.e., tasks are added to the ready queue or roll-in queue behind tasks of the same priority.

In the absence of time-slicing, once a task gains control of the processor, it continues executing until it voluntarily relinquishes that control or is preempted by a higher priority task.

A task will relinquish control of the processor to another task when one of the following occurs:

- It is paused by the system operator.
- It is cancelled by the system operator, user or another task.
- A higher priority task becomes ready due to an external event, such as the completion of an outstanding I/O request.
- It executes an SVC that places it in a wait, paused or dormant state.

- It initiates I/O to a device.
- Its time-slice expires.

After the task relinquishes control of the processor, it is returned to the ready queue where its TCB is placed behind the TCBs of tasks of equal priority. This allows the other tasks on the queue to be given a turn on the processor.

To determine which task should have control of the processor, the task manager chooses the highest priority task among those on the ready queue, the roll-in queue and any currently executing task. If a task is chosen from one of the queues, the currently executing task is placed back on the ready queue and the chosen task becomes the current task.

The task manager supports two types of time-slicing:

- System time-slice
- Dynamic time-slice

System time-slicing limits the execution of a task so that round-robin scheduling of priority tasks can take effect. Time-slicing allows tasks of equal priority to receive equal shares of processing time.

At sysgen, system time-slicing can be enabled through the use of the SLICE command. This allows time-slice scheduling to be activated automatically by the system. Thereafter, the operator SET SLICE command can be used to override the SLICE value set at sysgen.

Dynamic time-slicing is enabled only for MTM subtasks. The dynamic time-slice is calculated as:

$$\text{slice} = 1 + 2^{**}m$$

Where:

$$m = \text{task priority} - \text{run priority} + 1$$

The slice is measured in units of line frequency clock (LFC) ticks (one LFC tick = 8.333ms, 60Hz).

Run priority is adjusted whenever a task uses up a time-slice, is removed from a wait state or has its priority modified by the operator SET PRIORITY command. When a task uses up a time-slice, its run priority is adjusted as follows:

New run priority = run priority + 1 or
task priority + k (whichever is smaller)

Where:

k = number of dynamically scheduled tasks or
12 (whichever is smaller)

Because a task that is placed in a wait state does not use up its last assigned time-slice, the run priority of the task, when it is removed from suspension, is adjusted as follows:

Run priority = run priority - 1 or
task priority (whichever is larger)

The task manager also performs intratask context switches to allow tasks to receive and handle task traps in response to synchronous and asynchronous trap-causing events. Synchronous events include task-initiated faults (e.g., arithmetic, memory access, illegal instruction, etc.) and SVC14 traps. Asynchronous events originate outside of a task and include the task queue traps (e.g., I/O and timer completion, SVC6 send message/data and queue parameter, etc.) and the task event traps currently associated only with SVC intercept support.

In addition to task scheduling and task trap support, the task manager handles the state of a task during execution. Task execution state is determined by the setting of the PSW. The task manager switches or exits tasks from one execution state to another.

1.2.1.1 Task Scheduling on the 3200MPS Family of Processors |

The OS/32 task manager uses four different types of queues to facilitate task scheduling:

- CPU ready queue
- CPU receive queue
- CPU roll-in queue
- APU execution queue

An APU execution queue can be one of four different types:

- APU idle queue, not serviced by any processor
- APU private queue, serviced by a single APU
- APU shared queue, serviced by several APUs
- CPU/APU shared queue, serviced by the CPU whenever its own ready queue is empty and/or by one or more APUs; this is always APU execution queue 0.

Each APU execution queue can be designated either no-priority or priority-ordered. Priority-ordered queues can be either enforced or not enforced. For a priority queue that is not enforced, a new task entry is placed on the queue according to its priority. The currently executing task is not affected. An entry with a higher priority than another entry is placed at the front of the queue to be serviced when the current task relinquishes control. For a priority-enforced queue, if the priority of the new task entry is higher than that of the currently executing task, the executing task is interrupted and the new entry becomes the current task. With enforced priority, any entry of high enough priority will preempt the currently executing task. After system initialization, the CPU ready queue is priority-enforced and the CPU receive queue is no-priority.

When a task requests processor time, the task manager adds the TCB for that task to the CPU ready queue. The task manager selects a task for execution from the queue on a strict priority basis. After selecting a task, the task manager then decides whether the task is to be executed on the CPU or placed on one of the APU queues in the system. A task is transferred to an APU queue for processing only when all of the following conditions are true:

- The task must be executing in the user state, not in the system state.
- The task's "LPU-directed" status must be set. (In MTM, when the load-leveling executive (LLE) is active, subtasks of MTM cannot be LPU-directed unless the user has SVC6 privileges.)
- The TSW does not specify CPU-override status. (If the CPU-override status bit of the TSW is set, the task is executed on the CPU.)

When all of the above conditions are true for the highest priority task on the CPU ready queue, the task manager transfers the TCB for that task from the CPU ready queue to an APU queue.

If the APU is waiting for the task (i.e., APU processing has been suspended until the task arrives), the TCB becomes the current TCB and execution begins immediately. If the APU is not waiting for the task, the TCB is placed on the APU queue.

Whenever it is not processing a task, the APU continually checks its APU queue. If the APU finds entries on the queue, it will execute the task at the top of the queue.

Once the APU starts a task, the task will execute until it:

- relinquishes control voluntarily (reschedules itself),
- encounters a fault,
- issues an SVC, or
- is returned to the CPU via an operating system request on behalf of a monitoring task, operator command, etc.

The task may reschedule itself to the rear of the APU queue or to the CPU. In a no-priority APU queue, the task is placed at the bottom of the queue. In a priority APU queue, the task is placed behind all tasks of equal or higher priority or at the queue top if there is no task of equal or higher priority on the queue. In a priority-enforced APU queue, the task is placed on the queue in the same manner as for a priority queue. In addition, whenever the task happens to be placed at the queue top, the operating system executes the preempt procedure to ensure execution of the highest priority tasks even if a lower priority task is currently executing.

The task is returned to the CPU receive queue if it is rescheduled to the CPU, if a fault occurs or if an SVC or operating system request occurs. The task waits on the receive queue until the CPU places the task on the CPU ready queue.

If the task is placed on the receive queue as a result of a fault, the task is moved to the CPU ready queue. If the appropriate bits in the TSW are set, the task's TSW location is set to the address of the task trap handler. The task can then be dispatched back to the APU queue.

If the task is placed on the receive queue as a result of issuing an SVC, the task is moved to the CPU ready queue and executed on the CPU until SVC processing is complete. The task can then automatically move back to the APU queue.

Rollable tasks are moved from the roll-in queue to the CPU ready queue and are processed in the same manner as any other currently executing tasks. Rollable tasks may be dispatched to an APU.

Tasks running under MTM will run on APUs as determined by the LLE at a priority scheduled by the priority scheduling mechanism (PSM). When the LLE is active, MTM controls whether the task or request will be assigned to one of the APUs, IOPs or to the CPU.

1.2.2 Job Accounting Subsystem

The job accounting subsystem reports CPU usage and time elapsed, memory and disk space utilized and number and length of I/O transfers by device class. The OS/32 job accounting subsystem now reports APU usage, IOP usage and time elapsed in the 3200MPS Family of Processors. The job accounting subsystem contains the:

- Data Collection Facility
- Account Reporting Utility

The Data Collection Facility collects accounting data on all user activities and stores this information in the accounting transaction file (ATF) when the task terminates.

The Account Reporting Utility is designed to accommodate specific customer site requirements. The performance information gathered by the Data Collection Facility is prepared by the Account Reporting Utility for use by system maintenance personnel. Reports can be requested for individual user accounts, summaries of user accounts and system usage. For more information, see the OS/32 System Support Utilities Reference Manual.

Through the DISPLAY ACCOUNTING command, the system operator has access to accounting data for one or all tasks in the system. This command also gives MTM users access to accounting data for a task monitored by MTM on their behalf.

1.2.3 Memory Management Subsystem

When a task is loaded, the memory management subsystem dynamically allocates necessary space in memory. OS/32 supports three types of memory:

- Local
- Shared
- System

Local memory is physically contiguous starting from location 0 and contains the operating system, task space and system space. Local memory is allocated on a first-fit basis when sufficient memory is available for a specific task. Free segments are allocated in ascending address order. When no space is available for a task, the memory manager determines which tasks are to be rolled out to ensure that higher priority tasks take precedence. When memory becomes free, adjacent areas are merged together to minimize search time and to provide large free blocks of memory for bigger tasks.

Shared memory is located above local memory and is not required to be contiguous. Global task common (TCOM) segments, located in shared memory can be used by more than one processor.

System task space is maintained by the memory manager and is dynamically allocated when a task or device structure is built. System memory is shared by all processors in the 3200MPS Family of Processors. System memory contains both local and shared areas with local memory areas used by the CPU and all APUs.

The memory manager maintains task space through free and allocated lists. Segments are allocated dynamically on a first-fit basis by searching the free lists. When free task space is allocated to a segment, it is removed from the free list and connected to the allocated list. This list is called the segment control list (SCL). Similarly, whenever a segment is released, its memory space is removed from the allocated list and connected to (or merged into) the free list.

1.2.4 Timer Management Subsystem

The timer management subsystem provides tasks with a set of timer management/maintenance services. These services control all time-dependent functions (e.g., time-slicing, I/O, job accounting and file dating) through the universal clock (UCLOCK).

The following timer queues are maintained by the timer management subsystem:

- Time of day
- Device time-out
- Communications device time-out
- Interval timer

There are several timer routines that service these queues. Entries are placed on the time of day queue and the interval timer queue as a result of SVC2 timer requests. The control blocks on the time of day queue are referred to as timer queue elements (TMQs). The interval timer queue has the same format as the time of day queue but is maintained as a separate queue.

The UCLOCK consists of an LFC and a precision interval clock (PIC). In a 60Hz system, the LFC generates an interrupt every 8.3ms or 120 times per second. In a 50Hz system, the LFC generates an interrupt every 10ms or 100 times per second. The PIC interrupts when a task's requested time interval has expired or at intervals of 4,096ms, whichever is shorter. If the interval terminates or the time of day is reached, the TMQ is removed from system space and a trap is generated or the task is removed from timer wait.

In the 3200MPS Family of Processors configuration, the real-time support module (RTSM) provides each processor, having only APU's and IOP's, with a 32-bit real-time counter for timing program execution. These counters are incremented every microsecond by a 1MHz on-board oscillator. The read real time clock (RRTC) instruction allows tasks to read the counters. See the appropriate instruction set reference manual for more information.

1.2.5 File Management Subsystem

The OS/32 file management subsystem stores and retrieves information for a task on secondary storage devices (disks, floppy disks, etc.). The file manager partitions this storage into smaller areas, called files, that can be used by tasks for data and program storage. In addition, the file manager provides tasks with the following support services for file management:

Allocate	creates a file by allocating space on a secondary storage device.
Delete	removes a file from a secondary storage device.
Rename	changes the name of a file.
Open	assigns a logical unit (lu) to a file.
Close	cancels the lu assignment.
Fetch attributes	examines the attributes of a file.
Checkpoint	ensures that all data in an output buffer is written to a secondary storage device.

Software selects recording density for 6,250 bits per inch (bpi) magnetic tape drives.
density
selection

1.2.6 Input/Output (I/O) Subsystem

The I/O subsystem provides a uniform programming interface between the task and external devices. I/O operations can occur in the following task modes:

Wait	halts execution until data transfer is completed.
Proceed I/O	continues task execution during data transfer.
Halt I/O	allows the task to cancel previous proceed I/O requests.
Queued I/O	allows a task to queue I/O requests to a busy device and continue execution until the device is free.

A task trap mechanism can be used to report each completed I/O operation. Wait-only and test I/O functions allow the task to synchronize its execution with the completed I/O operations.

1.2.7 Error Recording Subsystem

The error recording subsystem logs all memory errors, file manager errors, system milestones and system detected errors on a disk for the Error Reporting Utility, which analyzes the data and generates reports.

OS/32 memory error recording software supports the memory error log hardware of the Series 3200 processors. Error log hardware keeps a history of the single-bit corrected memory errors. The operating system reads the error log hardware and stores the error information into an internal error log buffer. On a 3280 System, the operating system requests error log information through the Control Diagnostic System (CDS). When the error log buffer is full, its contents are stored on an error recording file with the date and time of the last error recorded. When the error recording file is almost full, a warning message is displayed on the system console indicating that a new error recording file should be allocated or that the Error Reporting Utility should be initiated. The Error Reporting Utility provides reports on the previously recorded error information in the error recording file.

The current error status can be displayed to the system console by using the DISPLAY ERRORS command. The internal error log read-out period can be changed by the system operator.

1.2.8 Memory Diagnostics Subsystem

The memory diagnostics subsystem eliminates inoperable memory areas from the system without affecting task execution. It allows the operating system to execute when portions of real memory have been removed (holes) or when a part of the system is powered down for maintenance. Memory can be tested and marked on and off through the operator MEMORY command or when the operating system is initialized.

The marked-off areas are noted as allocated in the memory map. Memory is marked on when previously marked-off memory is to be used again. If an irrecoverable memory error occurs during task execution on a Series 3200 processor, the operating system automatically marks off the area occupied by the task.

1.2.9 Loader and Segmentation Subsystem

The OS/32 resident loader is responsible for loading tasks, reentrant libraries, TCOM segments and partial images. These tasks and segments must have been built by Link. Each task image generated by Link contains information related to the task (e.g., task options, size, libraries referenced) in a record called the loader information block (LIB). The OS/32 resident loader uses this information to generate data areas, set the task options, create segment tables for the tasks and map the task segments into the memory space of the processor.

All user tasks (u-tasks) in OS/32 are built as though they were loaded at physical address 0 in memory. The relocation/protection hardware automatically relocates the task relative addresses at run-time by using the task segment table. This process is totally transparent to the user.

The loader is also responsible for creating the task environment, allocating roll files, creating, maintaining and deleting segment tables, maintaining a segment control list and mapping and unmapping partial images.

The task image can be divided into pure and impure segments. A pure segment is one that is read and execute only, while an impure segment is one that is read, write and excute. Pure and impure segments are defined by specifying the SEGMENTED task option when the task is built by Link. Regardless of the number of times a task is loaded, the loader will allow only one copy of the task's pure segment in memory at any one time. A separate copy of the task's impure segment is loaded each time the task is loaded. The relocation/protection hardware ensures the integrity of pure segments by allowing read-only and execute-only access privileges to those segments.

FORTRAN and assembly programs utilize TCOM segments. Access to TCOM is achieved mnemonically, that is by reference to the name of the segment. Linkages to these TCOM segments are resolved by Link. Link commands are also used to request read, write and execute privileges for TCOM blocks. See the OS/32 Link Reference Manual for more information.

1.2.10 Basic Data Communications Subsystem

The basic data communications subsystem provides a software interface between tasks and common carrier facilities. Basic data communications facilities allow the user to access remote terminals or computers as though they were locally attached peripherals. For example, with OS/32 Data Communications software, a task performs I/O to a remote terminal in the same manner as I/O to a local device.

In addition to providing device-independent (logical I/O) access to the task, the subsystem provides a device-dependent I/O capability that allows the systems programmer to tailor a communications package to a particular installation. Such a package can include device-dependent and device-independent support of asynchronous line devices as well as device-dependent support of binary synchronous lines.

The OS/32 Basic Data Communications software support package is required for all 32-bit communications products; e.g., HASP, 2780/3780 Remote Job Entry, the zero-bit data link control (ZDLC) Channel Terminal Manager and the Ethernet Data Link Controller (EDLC), which support the synchronous data logic control (SDL), high-level data link controller (HDLC) and advanced data communications control procedure (ADCCP) protocols.

1.2.11 Console Monitor Subsystem

The console monitor subsystem processes all I/O requests directed to the system console device and the system log device from all tasks including the command processor task. The console driver is responsible for intercepting system console I/O requests and for directing them to the console monitor or to another monitor task such as MTM. All I/O operations between the system console and tasks running under MTM are routed to the user's terminal through MTM and not through the console monitor.

When a command is issued from the system console, the console monitor issues an SVC6 to the command processor notifying it of a command to be processed. The command processor interprets the command and issues an SVC6 to the console monitor indicating that it is ready to accept another command.

The console driver is a part of the OS/32 I/O subsystem and is the module that intercepts I/O requests to the system console, processes them and gives them to MTM or to the console monitor to perform the actual I/O.

The console monitor is the first task dispatched at OS/32 initialization. The console monitor initializes both itself and the "dummy" device control block (DCB) used by the console driver to intercept requests from the system console. The monitor then issues an SVC6 to start the command processor.

1.2.12 Command Processor Subsystem

The command processor subsystem accepts commands from the system console monitor, decodes them and calls the appropriate executor. Commands can be input to the command processor by entering them directly through the system console or issuing them through a foreground task that uses the system console as an interactive I/O device. Commands input from a foreground task are executed by the command processor in the same manner as commands entered from the system console. If an error occurs during execution of a command, the command processor outputs an error message to the console.

An extension to the command processor, the command substitution system (CSS), allows commonly performed sequences of operations to be executed with one command. The CSS routines provide the user with the ability to build, execute and control files of operator and MTM commands. The user establishes command files that are called from the user console and executed in the user-defined sequence. In this way, complex operations can be carried out by the user with few commands. These commands are analogous to macro instructions in assembly language.

The CSS provides a set of logical CSS commands to conditionally control the precise sequence of commands to be executed. Parameters can be passed as part of a CSS call so that general sequences can be written that take on specific meaning only when the parameters are substituted. Other calls to CSS files can be imbedded within a CSS file (nested calls).

The command processor normally runs at the second highest priority level after the console monitor in OS/32. This task is strictly trap driven and responds to the SVC6 task queue parameter calls from the console monitor to service a command request. When the command is processed, it signals the console monitor for a new command read via an SVC6 queue parameter call and then enters into a trap wait state. The command processor priority can be decreased by the operator ATTN command. This command can be used in a real-time application environment to allow a task to run at a higher priority than the command processor.

1.2.13 System Initialization Subsystem

After the operating system is loaded, the system initialization subsystem initializes the memory diagnostics subsystem, error recording subsystems and system control blocks and tables in memory. It then dispatches the console monitor, which readies the command processor to accept commands from the system console.

1.2.14 Internal Interrupt Subsystem

The internal interrupt subsystem is responsible for:

- handling illegal instruction faults,
- handling arithmetic faults,
- detecting memory faults,
- handling SQS interrupts,
- handling relocation/protection hardware faults,
- handling data format/alignment faults,
- handling power fail and power restore conditions,
- restoring an interrupted task to its previous program state upon resumption of the task,
- handling parameter block errors,
- handling illegal SVCs and SVC interrupts,
- handling machine malfunction interrupts, and
- performing memory image dumps.

Processor-dependent interrupt handlers comprise the internal interrupt subsystem. This subsystem does not support external I/O interrupts; they are handled by the appropriate device drivers.

On the 3200MPS Family of Processors, the CPU handles all fault conditions or interrupts that occur during execution of a task on an APU. Thus, the APU can execute another task while the CPU is handling the fault or interrupt.

1.2.15 Optional User Supervisor Call (SVC) Subsystem

SVCl4 is provided as an optional SVC that can be defined by the user. On execution, the task receives a task trap for SVCl4. See the Supervisor Call (SVC) Reference Manual for information on how to implement the SVCl4 trap feature.

1.2.16 Floating Point Subsystem

A task has optional access to single and/or double precision floating point instructions under OS/32. Floating point instructions can be executed through hardware or simulated by software. Systems that do not support floating point options handle all floating point instructions as illegal instructions.

CHAPTER 2 PRIVILEGED TASKS

2.1 INTRODUCTION

In a multiuser system, improper use of certain machine instructions, called privileged instructions, can have a detrimental effect on system integrity. Privileged instructions include storage protection setting, interrupt handling, timer control, input/output (I/O) and some processor status-setting instructions. To prevent accidental or intentional misuse of these instructions, OS/32 provides a privileged operating state in which tasks can execute these instructions. In addition to the privileged operating state, OS/32 provides a privileged task state in which tasks can access the file account and bare disk OS/32 supervisor routines.

Only privileged tasks can execute in a privileged operating or task state. OS/32 allows three types of privileged tasks:

- Executive tasks (e-tasks)
- Privileged user tasks (u-tasks)
- Diagnostic tasks (d-tasks)

A task can be linked as a privileged task by specifying one or more of the following task options in the Link OPTION command:

ETASK, ACPRIVILEGE, DISC, DTASK

See the OS/32 Link Reference Manual.

This chapter describes the privileges that are available to each type of privileged task through the Link OPTION command.

2.2 EXECUTIVE TASKS (E-TASKS)

E-tasks run with the memory address relocation/protection hardware disabled and are allowed to execute all instructions provided by the hardware. E-tasks always have file account and bare disk privileges. In addition, e-tasks can execute code that modifies or enhances the OS/32 system software. For example, an e-task can modify one of the system modules to enhance an existing OS/32 feature. E-tasks can also function as drivers that support nonstandard peripheral devices. A task can be linked as an e-task by specifying the ETASK task option in the Link OPTION command. The following sections detail the programming considerations that must be taken into account when writing e-tasks.

2.2.1 Relocatable Executive Tasks (E-tasks)

A relocatable e-task is a method by which a programmer is freed from the traditional restrictions associated with writing an e-task. Within a relocatable e-task, a programmer can specify address constants and (RX3) instructions without having to relocate them manually from within the program. The programmer can also write programs in modules, which means that overlays can also be developed. All relocation is user-transparent so that the programmer does not have to worry about any special housekeeping or additional memory requirements.

2.2.2 Writing Executive Tasks (E-Tasks)

Because e-tasks execute in a privileged state, certain precautions must be taken when e-tasks are programmed.

When an e-task is executing, no memory address protection or relocation is provided and all interrupts are enabled. The task has access to all machine instructions and memory address space in the system. In addition, the e-task can access system tables and control information via the system pointer table (SPT). The address of the SPT is stored in the halfword at location X'62' in memory.

Link builds the image for an e-task as if it were loaded at absolute location zero. The loader, however, is free to load the e-task into any available memory location. Therefore, an e-task must be coded as if it were positionally independent; an e-task must not contain relocatable code unless the the RELOCATE option is specified.

Because Link relocates e-task addresses to absolute zero, e-tasks cannot assemble code containing address constants, as shown in the following example.

Example:

```
SVC7BLK  DB    X'80',7
          DAC    ADDR
```

An e-task must dynamically set the addresses required by the task. To reference addresses in the 16kB range, use the following technique:

```
LA      UE, BUFSTART
LA      UF, BUFEND
LA      U3, SVC1PBK
STM     UE, SVC1.SAD(U3)
SVC     1,0(U3)
```

References to addresses exceeding the 16kB range can be made in the following manner.

Example:

```
BASE    LA      U4, BASE
        LA      UE, BUFSTART-BASE(U4)
        LA      UF, BUFEND-BASE(U4)
        LA      U3, SVC1BLK-BASE(U4)
        STM     UE, SVC1.SAD(U3)
        SVC     1,0(U3)
```

E-tasks smaller than 16kB must use the no RX3 (NORX3) (CAL/32) instruction to force all RX instructions to the RX1 or RX2 format. The tasks must not contain any RX1 or RX3 instructions with relocatable addresses. See the Common Assembly Language/32 (CAL/32) Programming Reference Manual.

2.2.3 Writing Relocatable Executive Tasks (E-Tasks)

When writing relocatable e-tasks, any 3- or 4-byte instruction can be used within the task; this includes address constants such as RX3 and RI2 instructions along with external instructions. Halfword address constants and absolute data will not be relocated. Due to previous restrictions, only programs written in Common Assembly Language (CAL) may be used with relocation features since some of the object code generated by the compilers may reference absolute address locations. All other restrictions associated with e-tasks apply (e.g., shared segments can not be specified); on the other hand, a program may be linked with overlays. Common blocks can also be referenced, as long as the common block is not linked as a sharable segment.

To declare an e-task as relocatable, the user must specify it to be relocatable at link time via the OPTION command.

Example:

```
OPTION    ETASK,RELOCATE
```

Both options must be specified in order for the relocation tables to be built. If ETASK is omitted, the RELOCATE option is ignored; if RELOCATE is omitted, an e-task is established with no relocation tables being built.

NOTE

No attempt should be made to run any utilities as relocatable e-tasks. This is due to the fact that some of these utilities refer to absolute address locations.

For information on linking a relocatable e-task, see the OS/32 Link Reference Manual.

A relocatable e-task requires no additional memory to run, since the relocation tables are only used at load time. However, an extra 256-byte buffer is reserved within the task control block (TCB) for a task established with overlays; this buffer is used to read the relocation table every time an overlay is loaded.

2.2.4 OS/32 Data Structures Used by Executive Tasks (E-Tasks)

OS/32 provides two macro libraries that contain OS/32 and multi-terminal monitor (MTM) data structures. The OS/32 data structure macro library is stored in file SYSSTRUC.MLB. Table 2-1 contains a list of the macros and corresponding data structures in this library. Data structures specific to the MTM subsystem are stored in file MTMSTRUC.MLB. The contents of this library are listed in Table 2-2.

Using the OS/32 e-task capability and the data structures available to e-tasks, the system level programmer can incorporate changes or add user-written modules to the source of the OS/32 system modules.

TABLE 2-1 OS/32 DATA STRUCTURES MACRO LIBRARY

MACRO	DATA STRUCTURE
\$ACB	Directory access control block (ACB)
\$ACTCD	Input/output processor (IOP) action code
\$AOPT	Auxiliary processing unit (APU) options
\$APB	Auxiliary processor block (APB)
\$APB\$	\$APB, \$APRC, \$APS, \$AOPT
\$APRC	Passback reason codes and equates
\$APS	APB status codes and equates
\$APST	APU status codes, error codes and equates
\$ATF	Account transaction file (ATF)
\$CABINET	Cabinet configuration (3280 System only)
\$CCB	Channel control block (CCB)
\$CDS	Control Diagnostic Subsystem (CDS) (3280 System only : unsolicited message format)
\$CDSHDR	CDS message header (3280 System only)
\$CDSTIME	CDS time (3280 Systems only; request time response format)
\$CMMADDR	Composite memory module (CMM) address (3280 System only)
\$CMMCONF	CMM configuration (3280 System only)
\$CMMMSG	CMM message (3280 System only)
\$CTX	U-task context block
\$DATB	Device attributes equates
\$DCB\$	\$PDCB, \$DDCB, DCB EQUATE, \$DFLAG, \$DATB, \$DXFL
\$DDCB	Device-dependent device control block (DCB)
\$DDE	Error log data structure
\$DFLG	DCB flags
\$DIR	Primary directory entry
\$DXFL	Disk-extended flags
\$EMIL	System milestone recording entries
\$EFMG	Bulk device error recording entries
\$EREGS	16 executive registers (E1=register 1)
\$ERRC\$	\$GERC, \$EFMG, \$ESYS, \$EMIL, \$MERC
\$ESYS	System error recording entries
\$EVN	Event node
\$FCB	File control block (FCB)
\$FCB\$	FCB and FCB flags
\$FDE	Free block descriptor entry
\$FFLG	FCB flags
\$FD	File descriptor (fd)
\$GERC	General error recording information
\$HB	Help subroutine argument block
INTCPARM	Supervisor call (SVC) intercept information

TABLE 2-1 OS/32 DATA STRUCTURES MACRO LIBRARY (Continued)

MACRO	DATA STRUCTURE
\$ICB	Intercept control block
\$IOB	I/O block
\$IOB\$	I/O and I/O flags
\$IOBF	I/O block flags
\$IOH	I/O handler list
\$IPB	IOP parameter block (IPB)
\$IPCB	Intercept path control block
\$IRCB	Intercept control block
\$IVT	Initial value table
\$LIB	Loader information block (LIB)
\$LIB\$	LIB and loader options
\$LLE	Load-leveling executive (LLE)
\$LPMT	Logical processor mapping table (LPMT)
\$LOPT	Loader options
\$LSG	Load segment
\$LTCB	Loader TCB redefinitions
\$MAGDCB	Magnetic tape DCB
\$MERC	Memory error recording entry
\$OCB	Overlay control block
\$ODT	Overlay descriptor table (ODT) structure
\$ORT	Overlay reference table
\$PDCB	Primary (device-independent) DCB
\$PFCB	Private FCB
\$PSCMSG	Power supply cabinet (3280 System only)
\$PSDCB	Pseudo DCB structure (device-dependent)
\$PSTCB	Pseudo TCB
\$PSW	Program status word (PSW)
\$QCB	Queue control block (QCB)
\$QCHEAD	Queue control block header (QCHEAD)
\$QH	SVC intercept queue handler structure
\$QPB	Queue parameter block (QPB)
\$QPB\$	\$QPB, \$QPSTAT
\$QPSTAT	QPB status
\$RCTX	RS/RSA context block
\$REGS\$	\$SOPT, \$UREGS, \$EREGS, \$RREGS, \$PSW
\$REL	Read error logger (REL) (3280 System only)
\$REQSTAT	Request status (3280 System only)
\$RLST	Roll selection list
\$RREGS	16 general user registers (R1 = register 1)
\$RREL	Response for read error logger (3280 System only)
\$RSARCPY	Reentrant system state alternate save area

TABLE 2-1 OS/32 DATA STRUCTURES MACRO LIBRARY (Continued)

MACRO	DATA STRUCTURE
\$\$SXO	SVC1 extended options masks
\$\$DCB	Pseudo print DCB structure
\$\$SD	Send data message block
\$\$SDE	Segment descriptor element
\$\$SENDREL	Send read error logger (3280 System only)
\$\$SOPT	System options
\$\$SPLMSG	Spooler message structures
\$\$SPT	System pointer table (SPT)
\$\$SPTE	SPT extern definitions
\$\$SPOL	Spooler message
\$\$SSUB	Shared segment usage block
\$\$STE	Segment table entries (STEs) memory address translators ((MAT) processors)
\$\$SPR	Segment privilege flags
\$\$SVC1	SVC1
\$\$SVC1\$	SVC1 and SVC1 error codes
\$\$SVC1ERR	SVC1 error codes
\$\$SVC1MTE	SVC1 magnetic tape specific error codes
\$\$SVC4	System SVC - reserved
\$\$SVC5	SVC5 parameter block
\$\$SVC6	SVC6 parameter block
\$\$SVC7	SVC7 parameter block
\$\$SVC7EXT	Extended SVC7 functions
\$\$SVC7SPL	Spooler SVC7 parameter block
\$\$SVC13	SVC13 parameter block
\$\$SVC13\$	\$\$SVC13, \$APST
\$\$SVT	System value tab
\$\$SYP	System space structure
\$\$SPT	SPT table definitions
\$\$STABL\$	Structure/extern generating macro
\$\$TCB	TCB, \$\$SDE, \$IOB\$, \$TCB, \$CTX
\$\$TCB\$	\$TCB, \$TOPT, \$TSTT, \$TWT, \$TLFL, \$PSTCB, \$OCB, \$TQE, \$TFL, \$TPRC, TQH
\$\$TFL	TCB flags
\$\$TKQ	Task queue head
\$\$TLFL	Logical unit (lu) table of flags
\$\$TMQ	Timer queue entry
\$\$TOPT	Task options flags
\$\$TPRC	Task passback codes
\$\$TQE	Task event queue entry
\$\$TQH	Task event queue header
\$\$TQ27	SVC2 code 27 parameter block
\$\$TSTT	Internal task status flags
\$\$TSW	Task status word (TSW)
\$\$TTB	APU trap block
\$\$TWT	Task wait status flags
\$\$UDL	User-dedicated locations (UDLs)

TABLE 2-1 OS/32 DATA STRUCTURES MACRO LIBRARY (Continued)

MACRO	DATA STRUCTURE
\$UDL\$	UDL and TSW
\$UREGS	16 general user registers (U1 = register 1)
\$VD	Volume descriptor
\$VFCHARS	Vertical forms control (VFC) characters
\$VFDCB	Common VFC DCB structure
\$VSTE	Virtual address translator (VAT) segment table entry (VSTE) (3280 System only)
\$SWAP	Read/write access matrix header structure

TABLE 2-2 MTM DATA STRUCTURES MACRO LIBRARY

MACRO	DATA STRUCTURE
\$TERMUSR	Terminal user block
\$AUF	Authorized user file (AUF) record
\$MTMSTE	Terminal state definitions
\$PRIV	User privileges
\$VAR	CSS variable flags and structures
\$BTQ	Batch queue header and entry structures
\$CMB	Command buffer structure
\$LMB	Log/broadcast message buffer structures
\$CBH	Common buffer header structure
\$CSTK	Command substitution system (CSS) pointer stack structure

2.3 PRIVILEGED USER TASKS (U-TASKS)

Privileged u-tasks run with the memory address relocation/protection hardware enabled and are restricted to a subset of instructions known as nonprivileged instructions. If a u-task attempts to execute a privileged instruction, it causes an illegal instruction fault. However, unlike nonprivileged u-tasks, privileged u-tasks have file account and bare disk privileges. File account privileges allow tasks to specify an account number in the file account/class field of an fd. Bare disk privileges allow tasks to perform I/O operations to a bare disk device. See the OS/32 Supervisor Call (SVC) Reference Manual.

A u-task acquires file account and bare disk privileges by specifying the ACPRIVILEGE and DISC task options, respectively, in the Link OPTION command when the task is built.

2.4 DIAGNOSTIC TASKS (D-TASKS)

D-tasks, like e-tasks, can execute all instructions provided by the hardware. However, like u-tasks, d-tasks run with the memory address relocation/protection hardware enabled and execute in the nonprivileged task state. D-tasks are designed for use in diagnostic applications, loading writable control store (WCS), and direct execution of I/O instructions.

A task can be linked as a d-task by specifying the DTASK task option in the Link OPTION command. To execute in the privileged task state, a d-task must be built with the ACPRIVILEGE and DISC task options enabled.

CHAPTER 3
PROGRAMMING IN THE 3200MPS FAMILY OF PROCESSORS
MULTIPROCESSING ENVIRONMENT

3.1 INTRODUCTION

Programming in a multiprocessing environment is similar to programming in a uniprocessing environment. However, due to the nature of the hardware configuration, the multiprocessing environment offers one major programming advantage: increased throughput. For efficient use of this expanded computing ability, the system level programmer should take the following into consideration:

- The selection of tasks or input/output (I/O) requests that are to be executed on a central processing unit (CPU), auxiliary processing units (APUs), or I/O Processors (IOPs).
- The preparation of the APUs for task execution.
- The assignment of tasks to processor queues.
- The establishment and control over the order of task execution.
- The prevention of invalid data variables, caused when two tasks running on different processors are allowed to concurrently read and modify a common data structure or memory location.
- The measurement of real-time performance of the individual tasks in the system.
- The customization of APU fault handling.

This chapter focuses on some techniques that can be used by an assembly language programmer in solving unique programming problems in a multiprocessing environment. For additional programming information, see Table 3-3.

3.2 DESIGNING TASKS TO RUN ON A MULTIPROCESSING SYSTEM

The main performance advantage of designing an application to run on a multiprocessing system is that a job can be broken down into several parts that can be run on different processors simultaneously.

A job can be divided among a number of tasks that control individual operations, such as processing I/O, performing calculations resulting from a particular action and providing an operator interface for reporting and responding to the results of the calculations.

The individual APUs running these tasks can transmit all status information regarding the components of the system to another task, called the supervisor monitor. The supervisor monitor can then output messages to a console or printer as the status is received. Another function of the supervisor monitor is to store a code in a status word in memory that can be accessed by a stand-by task. The stand-by task would then be able to periodically check the status of the system and adjust task execution accordingly.

Once the programmer has divided a job into several tasks that can be run simultaneously, the next step should be to assign each task to an APU for execution. It should be remembered, that generally, execution of a computation-intensive task on an APU increases overall system performance, while an I/O-intensive task running on an APU decreases system performance. Because the operating system executes exclusively on the CPU, each physical I/O request made by an APU task causes the task's execution to transfer back to the CPU for operating system support. For this reason, all I/O-intensive tasks should be assigned to the CPU for execution. An I/O request from an APU will be passed to the CPU, which in turn passes it to an IOP, if one or more are available.

Tasks that perform extensive supervisor calls (SVCs) (other than SVCls to an indexed file with large blocking factors), are best left for the CPU to execute. Tasks that perform I/O to indexed files or proceed I/O to contiguous files should be executed on APUs. This takes advantage of an APU's ability to efficiently perform logical I/O. Computational-intensive tasks are most efficiently executed on APUs leaving the CPU free to service other tasks. Tasks that perform I/O to indexed files with one to one blocking factors or to contiguous files that are not proceed I/O (or proceed I/Os that are not followed by computational-intensive routines) should be scheduled to the CPU.

3.3 PREPARING AN AUXILIARY PROCESSING UNIT (APU) FOR TASK EXECUTION

OS/32 supports a multiprocessing configuration consisting of one CPU and any combination of one to nine APUs or IOPs. The operating system schedules tasks for execution by arranging them in queues. These queues consist of a CPU ready queue and APU execution queues.

3.3.1 Queue Assignments

The CPU ready queue is intended for SVC I/O-intensive tasks and is serviced by the CPU. The APU execution queues are numbered 0 through n, where n represents the number of APUs in the system. They are intended for the computation-intensive tasks and are serviced by APUs assigned to them. APU queue 0 is serviced by the CPU when the CPU ready queue is empty. The APU execution queues numbered 1 and above are also referred to as APU-only queues.

The APU-only queues may have the following assignment possibilities:

- The queue is idle with no APUs assigned.
- The queue is private and has one APU assigned.
- The queue is shared with two or more APUs assigned.

When the operating system is loaded, each of the APU-only queues is designated as a private queue and is assigned to one APU. The number of the queue will correspond to the number of the APU to which it is assigned. Subsequently, the APUs may be reassigned using a corresponding SVC13 control function or the operator command APC. To employ an SVC13 control function, a task must be linked using the OPTION APCONTROL command of LINK.

3.3.2 Auxiliary Processing Unit (APU) Operating States

OS/32 maintains two operating states for an APU. Each differs in the degree of APU readiness for task execution. These states are:

DISABLED	APU is unavailable for all purposes except running the power-up link check procedure.
ENABLED	APU has successfully passed the power-up link check procedure and is ready for task execution.

All APUs are put into the DISABLED state upon operating system load or power restore. On a power fail restart, an attempt is made to upgrade each APU not disabled prior to the power fail to the ENABLED state.

The transition from one APU state to another can be accomplished along the paths shown in Figure 3-1. These transitions are executed by the corresponding SVCL3 control functions or the operator APC command. The APU firmware logic requires resetting the APU state after it is disabled in order to be enabled again. The resetting is done using the appropriate button on the APU board or by powering down the APU cabinet.

040-2




Figure 3-1 Valid APU Operating States

3.3.3 APU-Only Queue Operating States

OS/32 maintains three operating states for each APU-only queue. Each differs in the degree of queue availability for task scheduling. These states are:

- | | |
|--------------|---|
| OFF | APU queue is not available for task scheduling. |
| ON EXCLUSIVE | APU queue has only a designated task scheduled to it. (Only an idle or private APU queue can be marked ON EXCLUSIVE.) |
| ON | APU queue is fully available for task scheduling. |

All the APU-only queues are put into the OFF state upon operating system load. Upon a power fail restart, the load power fail monitor (LPFM) restores the queue states. The queue 0 is always maintained in the ON state.

The transition from one APU queue state to another can be accomplished along the paths shown in Figure 3-2. These transitions are executed by the corresponding SVC13 mapping functions or the operator QUEUE command. To use an SVC13 mapping function, a task must be linked using an OPTION APMAPPING command of LINK.

040-3

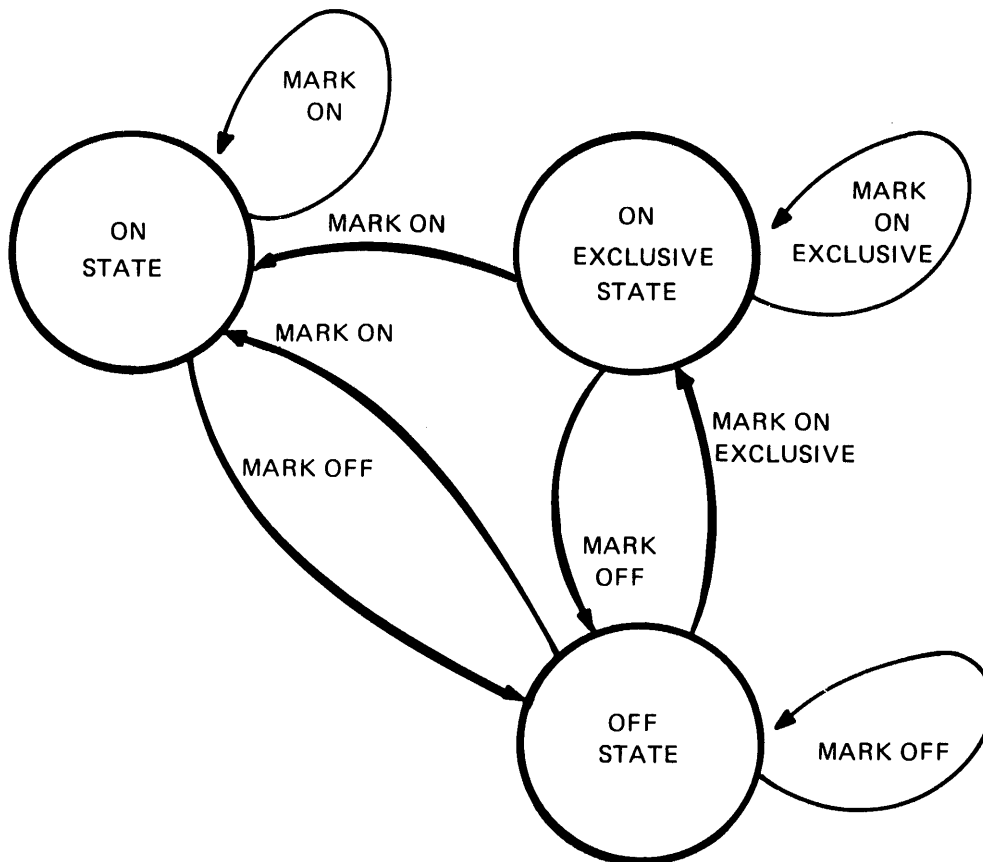


Figure 3-2 Valid APU Queue Operating States

```

*****
*       The following code demonstrates how SVC13          *
*       is used to enable an APU and mark on the queue    *
*       for task scheduling. This example does not        *
*       check for SVC13 execution errors. A task incor-  *
*       porating this code must be linked using a LINK   *
*       command OPTION APMAPPING, APMAPPING.              *
*****
  
```

Example:

```

$SVC13
ALIGN 4
ENABLE DS SVC13 ALLOCATE STORAGE FOR SVC 13 PARBLK
ENABLEE EQU *
*GAIN CONTROL RIGHTS, ENABLE APU, START APU, ASSIGN TO QUEUE,
*RELEASE CONTROL RIGHTS
ORG ENABLE+SV13.OPT
DB X'CD'
ORG ENABLE+SV13.FUN
DB X'03' FUNCTION CODE=3
ORG ENABLE+SV13.DOP
DB X'01' SEND START APU COMMAND
ORG ENABLE+SV13.APN
DB 2 APU NUMBER
ORG ENABLE+SV13.USE
DCX 3 ASSIGN APU TO QUEUE
ORG ENABLEE
***BUILD SVC 13 PARAMETER BLOCK FOR MARKING QUEUE
ALIGN 4
MARK DS SVC13. ALLOCATE STORAGE FOR SVC 13 PARBLK
MARKE EQU *
*GAIN MAPPING RIGHTS, MARK QUEUE ON, MAP LPU,
*RELEASE MAPPING RIGHTS
ORG MARK+SV13.OPT
DB X'B1'
ORG MARK+SV13.FUN
DB 2 FUNCTION CODE=2
ORG MARK+SV13.DOP
DB 2 LPU NUMBER TO BE MAPPED
ORG MARK+SV13.APN
DB 3 QUEUE TO MAP LPU TO
ORG MARKE
*****ISSUE SVC 13 TO ENABLE APU AND MARK QUEUE*****
SVC 13,ENABLE ENABLE APU
SVC 13,MARK MARK QUEUE ON
```

3.3.4 Logical Processing Unit (LPU) Mapping

For the purpose of directing tasks to the queues, OS/32 defines LPUs ranging from 0 to 255. LPUs are mapped into the APU queues while each task is associated with a particular LPU.

All the LPUs are initially mapped to queue 0 at operating system load time. LPUs 1 through 255 can later be mapped to other queues using a corresponding SVC13 mapping function or the operator LPU command. LPU 0 always remains mapped to queue 0.

3.4 ASSIGNING TASKS TO A PROCESSOR QUEUE

As mentioned previously, each task in the 3200MPS Family of Processors is associated with an LPU. The initial LPU value is established at task link editing time to be either LPU=0 by default or a value specified in the OPTION LPU command of Link. The LPU value may be changed at task load time or whenever the task is paused via a corresponding SVC6 function or with an operator OPTION LPU command.

Each task's LPU mapping (association with an APU queue) is enabled or disabled in a task status either by default or via a corresponding SVC6 function or via operator OPTION LPU and OPTION NLPU commands. By default, LPU mapping is disabled if the task is linked with LPU=0 and is enabled if the task is linked with a nonzero LPU. The operator OPTION LPU command, however, sets the specified LPU and also enables mapping even if LPU=0 was previously specified. The operator OPTION NLPU command enables mapping without changing the LPU number.

All tasks with mapping enabled are called LPU-directed tasks. OS/32 places the LPU-directed tasks onto corresponding APU queues. Tasks with mapping disabled and CPU-directed tasks are placed onto the CPU ready queue.

```
*****
*
*       This example loads and starts a copy of a task
*       and makes it LPU-directed via the SVC6 function.
*
*****
```

```

          $SVC6
          ALIGN 4
PARBLK   DS      SVC6           ALLOCATE STORAGE FOR PARBLK
ENDBLK   EQU     *
*SET LOAD, ASSIGN LPU, LPU-DIRECTED, & START FUNC CODES
          ORG    PARBLK+SVC6.FUN
          DC     SFUN.DOM!SFUN.LM!SFUN.LPM!SFUN.XLM!SFUN.SIM
          ORG    PARBLK+SVC6.LU
          DB     5                LU OF DIRECTED.TSK (IMAGE)
          ORG    PARBLK+SVC6.SAD
          DC     0                TASK EXECUTION START ADDR
          ORG    PARBLK+SVC6.SOP
          DC     0                START OPTIONS (none)
          ORG    PARBLK+SVC6.SEG
          DC     Y'40'           TASK WORKSPACE
          ORG    ENDBLK
START    EQU     *
*SETUP NAME OF TASK TO BE LOADED
          LI     R1,C'APU1'
          ST     R1,PARBLK
          LI     R1,C'TASK'
          ST     R1,PARBLK+4
```

```

*ASSIGN LPU NUMBER
      LIS   R1,2
      STB   R1,PARBLK+SVC6.LPU
*ISSUE SVC6 TO LOAD TASK FROM LU5
      SVC   6,PARBLK
      END   START

```

After the SVC6 in the previous example is executed, the task will be loaded into memory from the file (DIRECTED.TSK) with a workspace of 64 (X'40') bytes. When the task is started, the task manager dispatches it to the APU queue into which LPU2 is mapped.

3.5 CONTROLLING TASK ORDER OF EXECUTION

In a uniprocessor system, priority scheduling determines the execution flow of the tasks in the system. In order to affect task scheduling, a programmer must change the priority of the tasks in the system. In a multiprocessing environment, there is a choice of options to control the order of task execution as described in the following sections.

3.5.1 Changing Auxiliary Processing Unit (APU) Task Queue Ordering

Each of the APU queues can be set to handle its assigned tasks through the following priority disciplines.

- The no-priority queue services tasks in a first-in/first-out (FIFO) order, regardless of task priority.
- The priority queue services its highest priority tasks first and its equal priority tasks in a FIFO order. No preemption of currently executing tasks by higher priority tasks will occur.
- The priority-enforced queue services its tasks in the same manner as the priority queue; however, higher priority tasks are allowed to preempt lower priority tasks being executed on the processor assigned to the queue.

At operating system load time, the queues are initially set with the following priority assignments or disciplines. See Table 3-1.

TABLE 3-1 QUEUE PRIORITY ASSIGNMENTS

QUEUE	PRIORITY DISCIPLINE
CPU Ready Queue	Priority-enforced
APU Queue 0	Priority-enforced
APU Queue 1 to n	No-priority

These initial settings can be subsequently altered via a corresponding SVC13 mapping function or with an operator QUEUE command.

```

*****
*           The following example uses SVC13 to change a           *
*           queue priority discipline.  If the discipline           *
*           of this queue prior to the SVC13 call was no-         *
*           priority, OS/32 will reorder the queue according     *
*           to the task priorities subsequent to the SVC13         *
*           call.  A task incorporating this code must be         *
*           linked using the LINK command OPTION APMAPPING.       *
*****

```

Example:

```


          $SVC13
          ALIGN 4
DISCIP   DS   SVC13      ALLOCATE STORAGE FOR SVC13 PARBLK
SETD    EQU   *
*GAIN MAPPING RIGHTS, SET QUEUE DISCIPLINE,
*RELEASE MAPPING RIGHTS
          ORG   DISCIP+SVC13.OPT
          DB    X'85'
          ORG   DISCIP+SVC13.FUN
          DB    2              FUNCTION CODE=2
          ORG   DISCIP+SVC13.APN
          DB    3              QUEUE NUMBER
          ORG   DISCIP+SVC13.USE
          DC    H'1'          PRIORITY DISCIPLINE
          ORG   SETD
*ISSUE SVC13 TO CHANGE QUEUE DISCIPLINE
          SVC13, DISCIP

```

3.5.2 Monitoring and Preempting Auxiliary Processing Unit (APU) Task Execution

The 3200MPS Family of Processors provide facilities to monitor APU operation via the mechanism of task trap service. The APU reports its significant events to OS/32 by issuing asynchronous status signals.

An APU status signal is a byte with the following format:



Bits:
 0 1 2 3 4 5 6 7

- PAR is the parity bit which is adjusted to maintain odd parity for the byte.
- RUN is set to 1 if the APU is currently not idle. The APU may be executing task instructions or performing servicing functions (NON-TASK) such as: selecting a task from the queue, releasing a task back to the queue, processing a task fault, etc.
- NON-TASK is set to 1 if the APU is performing any of the servicing functions.
- WAIT is set to 1 if the APU is idle and in the wait state imposed upon it by the last executed task.
- RESP is set to 1 if the APU is returning the status in response to a command from the CPU (normally an SVC13 read APU status function); this bit is always reset in an APU asynchronous status signal.
- ERROR is set to 1 if the APU has detected an error in system data structures and subsequently entered an idle state.
- MOD1, MOD2 are set or reset to supply some additional status information.

When the last three bits are set to 1, they have the values of 4, 2 and 1, respectively. The actual value of these three bits reflects the APU status condition as follows:

- 0 Undefined.
- 1 The APU has detected that the queue to which it is assigned is empty (contains no tasks).
- 2 The APU has rescheduled (released) a task to the APU queue as a result of an RSCH1 instruction in the task or as a result of an appropriate command issued via an SVC13 control function.
- 3 The APU has rescheduled a task to the CPU ready queue as a result of an RSCH 0 instruction in the task, an appropriate command issued via an SVC13 control function, a nonexecutable APU instruction (normally an SVC) or a task fault.
- 4 The APU has detected an error in data structures at an arbitrary moment.
- 5 The APU has detected an error in data structures while attempting to select a task from its queue.
- 6 The APU has detected an error in data structures while attempting to lock its queue. Each processor locks its assigned queue prior to manipulating it.
- 7 Undefined.

Detailed information regarding the data structure errors detected by the APU can be obtained using SVC13 read APU status function.

NOTE

In the absence of an APU monitor task, OS/32 reports APU errors via the operator console.

This section will examine the methods used by an APU monitor task to:

- receive status signals from an APU, and
- preempt the current task executing on an APU with another task after a certain time interval has elapsed.

To receive a status signal from an APU, the APU must be connected and thawed (via SVC6) as a trap-generating device to the monitor task. In addition, the monitor task must have the appropriate bits in its task status word (TSW) set, a task queue to receive the status signal and a task queue trap-handling routine to service the trap.

```
*****
*       The following example demonstrates how to code      *
*       a typical APU monitor program to receive and       *
*       handle task queue traps from an APU. For more     *
*       information on task trap handling, see the OS/32   *
*       Application Level Programmer Reference Manual.      *
*****
```

Example:

```
**** Define a task queue to receive APU signals *****
*
*       ALIGN 4
*       TASKQ DLIST 100 DEFINE TASK Q OF 100 ELEMENTS.
*
* Put the address of task queue in UDL (UDL.TSKQ)
*
*       LA R14,TASKQ
*       ST R14,UDL.TSKQ
*
* Set TSW bits to enable the applicable task traps.
*
*       LI R14,TSW.TSKM+TSW.APTM
*
* TSW.TSKM enables task queue service traps
* TSW.APTM enables signals from APU
* Save TSW values to enable APU signals and task Q entries
*       ST R14,ENTRIES SAVE TSW VALUES
* TO ENABLE APU SIGNALS AND TASK Q ENTRIES
*
* SET UP TSW FOR TRAPS IN UDL
*
*       LA R15,QSERVICE
*       STM R14,UDL.TSKN SET UP TSW ON TRAPS IN UDL
*       SVC 9,UDL.TSKN ENABLE TASK QUEUE ENTRIES
```

For information on writing a task queue trap handling routine that removes the APU status entries from the task queue, see the OS/32 Application Level Programmer Reference Manual.

```
*****
*       The following code demonstrates a method of      *
*       connecting the APUs as trap-generating devices     *
*       to the APU monitor task.                          *
*****
```

Example:

```

* Enable each APU in the system if it is not enabled and then
* connect to each APU, but first
* read APU assignment information to obtain the
* number of APUs in the system.
*
START    SVC    13,APUASGN
        LB     R1,BUFFER1+1          LOAD MAX APU NO.INTO R1
* SET UP SVC13 PARAMETER BLOCK TO
* FETCH APU STATUS
        LIS    R3,X'80'
        STB    R3,FETAPU+SV13.OPT    SET APU STATUS OPTION
        LIS    R3,1
        STB    R3,FETAPU+SV13.FUN    SET UP FUNCTION CODE 1
        LA     R4,APUBUF
        ST     R4,FETAPU+SV13.BUF    SET UP BUFFER ADDR.
        LHI   R3,40
        STH   R3,FETAPU+SV13.LEN    SET UP BUFFER LENGTH
*
* SET UP SVC13 PARAMETER BLOCK TO ENABLE THE APU
*
        LIS    R3,3                  SET UP SVC 13 FUNC CODE 3
        STB    R3,ENABAPU+SV13.FUN
        LIS    R3,X'C1'              SET UP CONTROL OPTIONS
        STB    R3,ENABAPU+SV13.OPT  GAIN,ENABLE,RELEASE
*
* GET THE APU STATUS.  IF APU IS DISABLED,
* ATTEMPT TO ENABLE IT.  IF APU CAN'T BE
* ENABLED, LOG MESSAGE TO CONSOLE AND
* CONNECT TO IT ANYWAY JUST IN CASE IT IS
* ENABLED LATER.
*
APULOOP EQU *
* GET APU STATUS
        STB    R1,FETAPU+SV13.APN    SET UP APU NO.
        SVC    13,FETAPU             ISSUE SVC 13
        LH     R4,FETAPU+SV13.ERR    GET SVC 13 ERROR STATUS
        BZ     GETSTAT                IF NO ERROR-GET APU STATUS
        BNE   ER.ROUTE               IF ERROR, BRANCH TO ER.ROUTE
        LB     R5,APUBUF+5           GET 2ND BYTE OF APU S-STATUS
        BNZ   CONNECT                NOT DISABLED, GO CONNECT
*APU IS DISABLED; ISSUE SVC 13 TO ENABLE IT.
        STB    R1,ENABAPU+SV13.APU  SAVE APU NUMBER
        SVC    13,ENABAPU            ENABLE THE APU
        LH     R3,ENABAPU+SV13.ERR  GET SVC 13 ERROR STATUS
        BNZ   ENAB.ERR               BRANCH TO ERROR ROUTINE ON ERROR
CONNECT EQU *
*
*SAVE APU NO. AS PART OF APU'S TGD MNEMONIC
        STB    R1,SVC6.DEV
*ISSUE SVC6 TO CONNECT AND THAW THE APU
*

```

```

        SVC      6,APUTRAPS
        LH       R6,SVC6.STA      GET SVC 6 ERROR STATUS
        BZ       NEXT.APU        NO ERROR-GO CONNECT TO NEXT
*
        STB     R1,CONB.ERR+24    SAVE APU NO. IN MESSAGE
        SVC     2,LOGMSG         LOG MESSAGE: COULD NOT
*
NEXT.APU SIS    R1,1             MOVE ON TO NEXT LOWEST APU NO.
        BP     APULOOP          GO HANDLE NEXT APU.

```

The parameter blocks used in the previous example are defined as follows:

*SVC 13 Read APU Assignment Parameter Block and Buffer

```

        ALIGN 4
APUASGN DS      SVC13
ENDPBK  EQU     *
        ORG    APUASGN+SV13.FUN
        DB     X'00'             SET FUNC CODE 0
        ORG    APUASGN+SV13.BUF
        DAC    BUFFER 1         DATA BUFFER ADDR
        ORG    APUASGN+SV13.LEN
        DC     H'50'           MAX LENGTH OF BUFF
        ORG    ENDPBK
        ALIGN 4
BUFFER  DS      50

```

```

*
* SVC13 Fetch APU Status Parameter Block & Buffer
*
* FETAPU DS      SVC13

```

```

        ALIGN 4
APUBUF  DS      40

```

*** SVC13 Enable APU Parameter Block

```

        ALIGN 4
ENABAPU DS      SVC13.

```

* SVC 6 Connect & Thaw APU Parameter Block

```

        $$SVC6
        ALIGN 4
APUTRAPS DS      SVC6.
ENDAPUTB EQU     *
        ORG    APUTRAPS+SVC6.FUN
        DS     Y'C000 C000'     SVC6 FUNC CODE-
*                               SELF-DIRECTED, CONNECT & THAW
*
        ORG    APUTRAPS+SVC6.DEV
        DC     C'APU'           TRAP-GENERATING DEVICE MNEMONIC
        ORG    ENDAPUTB

```

**** SVC 2 Log Message Parameter Block

```

*
LOGMSG      DB      0,7
            DCZ     CONE.ERR-CONB.ERR
CONB.ERR    DB      C'UNABLE TO CONNECT TO APU'
CONE.ERR    EQU     *

```

The code in the previous example allows the monitor to receive traps from the APUs. Status returned from these traps can be reported to the console (via SVC1 or SVC2 code 7) or to a file designated for the APU output (via SVC1). In addition, this monitor program can be coded to run a certain task (TASK1) every ten minutes on a specific APU. To do this, the monitor sets an interval timer via SVC2 code 23. Upon expiration of the timer, the monitor task issues an SVC13 code 3 to preempt the current executing task on the APU, as shown below. This preemption mechanism is only allowed on no-priority queues. It is used when the overhead associated with maintaining a priority queue is to be avoided.

Example:

```

                SVC  13,PREQ
                .
                .
                .
                ALIGN 4
* PREEMPT TASK EXECUTION, RESTART APU
PREQ.OPT DB     X'B9'           SET SVC 13 OPTIONS:
PREQ.FUN DB     X'03'           SET FUNCTION CODE-
*                               CONTROL FUNCTION
PREQ.DOP DB     X'01'           DIRECTIVE OPTION-
*                               START APU
PREQ.APN DB     X'01'           APU NO. - APU 1
PREQ.APS DS     2               APU HARDWARE STATUS
PREQ.ERR DS     2               SVC 13 ERROR STATUS
PREQ.BUF DAC    BUF2           DATA BUFFER ADDRESS
PREQ.USE DS     2               LENGTH OF BUFFER USED
PREQ.LEN DC     H'8'           MAX LENGTH OF BUFFER
                ALIGN 4
BUF2         DC     C'TASK1 000'    TASK ID BUFFER

```

Execution of the previous SVC13 will cause the monitor to gain control rights to the specified APU (APU1), provided that the task has been link-edited with the APCONTROL task option and no other task has control rights to the APU. The control options, specified in the SVC13 parameter block, will then cause the following actions:

- Execution of the current executing task on the APU will be stopped.

- The current task will be rescheduled to the end of the APU queue.
- The APU's queue pointer will be repositioned to point to TASK1. (This will cause TASK1 to be selected as the next task to be executed on the APU.)
- The APU will be restarted for execution of TASK1.
- The monitor task will release the control rights to the APU.

The remaining code in the monitor program should check the PREQ.ERR field of the PREQ parameter block for errors as follows.

Example:

```
LH    R2,SV13.ERR
BNZ   ERR.PREQ
```

If an error has occurred, ERR.PREQ can log a message to the console.

Finally, to reexecute TASK1 in ten minutes, the interval timer (via SVC2 code 23) should be reset so that the SVC13 code 3 to preempt the current APU task can be reissued when ten minutes have elapsed.

See the OS/32 Supervisor Call (SVC) Reference Manual for more information on SVC13, SVC6 and SVC2 code 23.

3.5.3 Transferring a Task from an Auxiliary Processing Unit (APU) to the Central Processing Unit (CPU)

Under certain conditions, a monitor task may need to transfer some other task back to the CPU ready queue. The task to be transferred may be executing on an APU or waiting on its queue. The monitor task can transfer a task back to the CPU ready queue by issuing an SVC6, specifying the following function codes:

- Suspend (SFUN.SM)
- Transfer to CPU (SFUN.XCM)
- Release (SFUN.RM)

The suspend will transfer the task back to the CPU ready queue and then the LPU-directed task status is reset. Upon release, the task will stay on the CPU ready queue and not be dispatched according to its LPU assignment.

Example:

```
SVC    6,CPUDIR
      .
      .
      .
ALIGN  4
$SVC6
CPUDIR DS    SVC6.
CPUDIRE EQU  *
      ORG    CPUDIR+SVC6.ID
      DC    C'TASK1'          ID OF TASK TO BE
*                                     TRANSFERRED
      ORG    CPUDIR+SVC6.FUN
* SET OTHER-DIRECTED, SUSPEND, TRANSFER TO CPU, & RELEASE FUNC CODES
* FOR TASK1
      DB    SFUN.DOM!SFUN.SM!SFUN.XCM!SFUN.RM
      ORG    CPUDIRE
```

Execution of this SVC6 causes TASK1 to be suspended (if it is not already in a wait state) and transferred to the CPU ready queue. Resetting the LPU-directed status directs the task manager to ignore its LPU mapping and to schedule this task for execution on the CPU ready queue. When released, the task will execute on the CPU at the location following the instruction that was executed before the task was suspended. If the SVC6 in the previous example did not reset the LPU-directed status bit, the task will again be dispatched to the APU queue into which its LPU is mapped upon release from the suspended state.

3.5.4 Internal Task Control of Auxiliary Processing Unit (APU) Execution

A task can exercise control over its own execution on an APU through the SVC6 mechanism described previously since SVC6 can be made self-directed; however, there are more efficient mechanisms achieving the same result that are particularly valuable for real-time and APU diagnostic applications.

1. A task wishing to relinquish use of an APU while remaining on the same processor queue may issue the following instruction:

```
RSCH R1,1
```

The APU places the task at the queue tail and immediately picks up the task residing at the queue head. OS/32 will restore the queue order according to the queue discipline, if necessary.

2. A task wishing to transfer to the CPU indefinitely, may issue the following instruction:

RSCH R1,0

The APU sends the task to the CPU and then immediately picks up the task residing at its queue head. OS/32 resets the tasks LPU-directed status, which prevents the task from going anywhere but the CPU receive queue.

3. A task wishing to transfer to an APU indefinitely, according to its LPU mapping, may issue the following instruction:

RSCH R1,2

OS/32 insures that the task is scheduled to the appropriate APU queue according to the task's priority and the queue discipline.

4. A task may manipulate its TSW CPU-override status to enable or disable its transfer to the APU, for a given reason, to which the TSW corresponds. Bit TSW.CPOB (currently bit 8) prevents task scheduling to an APU queue when set to 1. This is necessary when a particular task fault, not a single instruction, should be executed on the APU.

3.5.5 Verifying Task Transfer to an Auxiliary Processing Unit (APU)

It may be necessary for a task to verify whether or not it has actually been transferred to an APU queue. For example, suppose a task on the CPU is assigned to LPU3 and executes the following instruction:

RSCH R1,2

Execution of this instruction will cause OS/32 to set the LPU-directed status of the task. The OS/32 task manager will then attempt to transfer the task to the APU queue into which LPU3 has been mapped. Suppose LPU3 is mapped to APU queue 3 and APU4 is assigned to this queue. To verify that the task is indeed executing on APU4, the next instructions executed by the task could be:

```

        LIS    R1,0                GET RTSM PULSE LINE
* TO PULSE
        LI     R2,15              FILL IN APU ID
        GSIG   R1,R2              GENERATE SIGNAL
* HERE THE NO. (15) CAN NEVER MATCH
* THE APU ID IN THE RTSM. NO SIGNAL WILL BE
* SENT. INSTEAD, ONLY THE APU ID IS RETURNED TO R1

```

After execution of GSIG, R1 will contain the number of the APU that the task is currently executing. See the appropriate instruction set reference manual for more information on the RSCH and GSIG instructions.

3.5.6 Customizing Auxiliary Processing Unit (APU) Fault and Supervisor Call (SVC) Handling

OS/32 allows customization of fault and SVC handling by the APUs. When consistently pursued, this route may allow reduction of the task traffic between the APUs and the CPU caused by SVCs or it may provide for APU I/O handling "invisible" to OS/32 and subsequently more efficient.

As an example of this customized handling, an APU can be made to wait for a task return while the task fault or SVC is processed by the CPU. This may be needed to leave private queue orders undisturbed by occasional SVCs.

This feature is not fully supported by OS/32 and therefore, is intentionally made difficult to use. However, software tools may be easily developed to exploit the customization feature.

In order to allow for custom processing of the faults and SVCs in a given task by the APUs, the following actions must be performed:

- An APU trap block has to be allocated in memory. This block will contain pairs of fullwords, each being a program status word (PSW) for a given APU detected reason in this order:
 - arithmetic fault
 - illegal instruction
 - memory controller fault
 - instruction format fault
 - SVC
 - machine malfunction fault

- A single trap block is allocated during OS/32 system generation (sysgen) and is designated in the map under a symbol TBLK1. Any additional blocks can be allocated using the MODULE command at sysgen.

- The trap block has to be patched with zeros for various reasons. If an APU wait is desired, the first word of the pair is set to X'8000' (bit 16 set) and the second word is ignored. However, if custom processing is desired, the first word of the pair is set to the required status and the second word is set to the location where the custom processing begins.
- The APU queue parameter block (QPB) fullword at location QPB.TPTR, (currently X'8' in QPB) has to be patched to the address of the APU trap block after task loading. This patching can be performed using the operator MODIFY command or via a dedicated executive task (e-task) assembled with the appropriate data structures.

When a task executes on an APU assigned to the patched queue and a fault is detected for which the PSW in the trap block is not zero, the APU transfers control according to this PSW. In the case of the bit 16 of the first word set in PSW, the APU transfers the task to the CPU ready queue and awaits the task's return. OS/32 will restart the APU when the task is scheduled back to it. OS/32 also restarts the APU when the task for which the APU is waiting is cancelled or terminated.

If the customized processing needs to be done on a per task rather than per queue basis, this can be arranged by patching out the OS/32 code in module APSV routine TMCKAPU that loads the QPB.TPTR into every TCB scheduled to the queue. Then, instead of patching QPB.TPTR with the selected task's TCB at location TCB.TPTR (currently X'20' in TCB), it can be patched with the address of the trap block. This task level trap block support is used in the operating system to intercept SVC calls from tasks running on APUs. This intercept is used to enhance index file I/O performance. The SVC trap blocks vectors to the routine FLIHAPU in APSV, so any further processing that the user wishes to do with SVC requests should be done by modifying this routine.

3.6 PREVENTING MEMORY ACCESS CONFLICTS

When several processors are executing simultaneously, it is possible for tasks running on two or more processors to require access to the same data. For example, suppose two tasks share a buffer list consisting of 30 buffers defined as follows:

```
BLISTBIT DS      2
BUFLIST  DLIST  30
```

BUFLIST contains the addresses of the buffers. BUFLIST and the actual buffers reside in an area of memory shared by the two tasks. One task collects data, writes it to a buffer and adds the address of that buffer to the bottom of the list. The other task removes an address of a buffer from the top of the list and processes it. Since both tasks can be run simultaneously on different APUs, both tasks may attempt to access the list at the same time. The Test and Set instruction (TS) can be used to ensure that only one task at a time can access the buffer.

To ensure that only one task at a time can access BUFLIST, a test and set operation is performed on BLISTBIT. BLISTBIT acts as a lock-out mechanism that is set and reset. A task can only access BUFLIST if BLISTBIT is not set.

3.6.1 Avoiding System Deadlock

When using the test and set operation, care should be taken to ensure that system deadlock is avoided.

For example, suppose task A uses TS to lock out data structure X while task B is locking out data structure Y. Task A now finds that it needs to access data structure Y, so it waits for Y to be released. Similarly, Task B finds it needs to access data structure X, so it waits for X to be released. Since each task holds the data structure needed by the other, processing stops. Both tasks are deadlocked.

To avoid system deadlock, the Test and Set instruction should be used with a time-out mechanism.

```
*****
*           The following example shows how to prevent           *
*           memory access conflicts without system deadlock.     *
*****
```

Example:

```

      TS      BLISTBIT      TASK CHECKS IF IT CAN GET
*                               ACCESS TO LIST
      BNM     CONTINUE     PROCESS LIST IF FREE
      LI      R2,50        LOAD TIMEOUT VALUE OF 50
*                               MICROSEC IN R2
SETBITLP EQU *           TIMER ROUTINE
      SIS    R2,1         DECREMENT TIMEOUT COUNT
      BM     TIMEOUT      BRANCH TO TIMEOUT ROUTINE
*IF BRANCH TO TIMEOUT IS TAKEN IT MEANS THAT THE
*TASK STILL COULD NOT GET ACCESS TO LIST
*THE TIMEOUT ROUTINE PRINTS A MESSAGE TO THE CONSOLE
*SO OPERATOR CAN TAKE NECESSARY ACTION
*ELSE CONTINUE

```

```

        LH      R4,BLISTBIT          USE APU CACHE TO MATCH LOCKS
        BMS     SETBITLP             BUFLIST NOT AVAILABLE YET;
*                                     TRY AGAIN
        TS      BLISTBIT             BUFLIST IS AVAILABLE SO
*                                     ATTEMPT TO GRAB ACCESS
        BMS     SETBITLP             NOT QUICK ENOUGH, RETRY
**IF SUCCESSFUL, PROCESS LIST      **
*                                     *
CONTINUE EQU      *
*                                     *
*         .                           *
*         .                           *
*         .                           *
*ACCESS BUFLIST EITHER BY ABL (ADD TO BOTTOM OF LIST
*INSTRUCTION) OR  RTL (REMOVE FROM TOP OF LIST INSTRUCTION).
*         .                           *
*         .                           *
*         .                           *
*AFTER PROCESSING BUFFER, UNLOCK BLISTBIT SO OTHER TASK CAN
*ACCESS IT.
*                                     *
        LIS    R4,0
        RBT    R4,BLISTBIT

```

3.7 MEASURING REAL-TIME PERFORMANCE ON THE 3200MPS FAMILY OF PROCESSORS

The OS/32 system macro library provides a set of timer macros that can be used to measure the real-time performance of individual tasks currently executing. These macros allow the programmer to set up a named timer in memory. A named timer can be compared to a stopwatch that measures the amount of time elapsed from the time the watch is started to the time it is stopped. The following example shows the data structure setup in memory for a timer named TIMRNAME. The timer macro, CRTIMERS, is used to set up timer data areas.

Example:

```

        ALIGN 4
TIMRNAME DCF  C'TIMRNAME'          TIMER NAME (8 CHAR MAX)
        DCF  0                      TIMER COUNTER
        DCF  0                      TIMER START VALUE
        DCF  0                      ACCUMULATED TIME
        DCF  0                      REGISTER SAVE AREA

```

The timer macros are used to set the watch and read the accumulated time after a specified interval has elapsed. The timer macros are listed in Table 3-2.

TABLE 3-2 TIMER MACROS

MACRO	FUNCTION
CRTIMERS (NAME1[,NAME2,...])	Creates a data area for each named timer.
STRTIME NAME(,REG)	Starts the named timer.
STOPTIME NAME(,REG)	Stops the named timer.
GETTIME NAME,REG	Gets the total time accumulated by the named timer.
READTCNT NAME,REG	Gets the number of intervals that have been timed by this timer.
RESETIME NAME	Resets accumulated time counts.

```

*****
*       The following example demonstrates how these      *
*       macros can be used to time the execution of a    *
*       program and its subroutine.                       *
*****

```

Example:

```

* Create a data area for the timer                        *
* for MAIN and the timer for SUB                          *
*                                                         *
*       CRTIMERS (MAIN,SUB)                               *
*                                                         *
* Start timer for MAIN.                                  *
*                                                         *
START   EQU   *                                           *
        STRTIME MAIN                                       *
        .                                               *
        .                                               *
        .                                               *
        BAL   R15,SUBPROG                                  *
        .                                               *
        .                                               *
* Stop timer for MAIN                                    *
        STOPTIME MAIN                                       *
* Get total time accumulated by MAIN                     *
* Timer. Load into REG 0                                 *

```

```

        GETIME MAIN,R0
        .
        .
        .
* Log MAIN program execution time.
        .
        .
* Get total time accumulated by SUB
* timer. Load into REG 3
        GETIME SUB,R3
* Get number of intervals timed by
* SUB timer. Load into R0
        READTCNT SUB,R0
        .
        .
        .
* Compute average subroutine execution
* time.
        DR      R2,R0
        .
        .
        .
SUBPROG EQU   *
*
* Start timer for SUB
*
        STRTIME SUB
        .
        .
*
* Stop timer for SUB
*
        STOPTIME SUB
        BR      R15

```

Detailed descriptions of the timer macros can be found in the OS/32 System Macro Library Reference Manual.

3.8 WHERE TO GO FOR MORE INFORMATION

This chapter is intended to demonstrate assembly language programming techniques used in designing system level control programs that take advantage of the 3200MPS Family of Processors capabilities. However, all the programming facilities available for writing system level control programs are not shown. Table 3-3 summarizes additional facilities and lists the manuals in which they are described.

TABLE 3-3 ADDITIONAL INFORMATION SOURCES FOR THE 3200MPS FAMILY OF PROCESSORS PROGRAMMING

MANUAL	PROGRAMMING METHODS DESCRIBED
3260MPS Instruction Set and System Architecture Reference Manual	Describes the machine instructions specific to each processor with a discussion of the APU processor states.
3280MPS Instruction Set and System Architecture Reference Manual	
3230XP Instruction Set and System Architecture Reference Manual	
3230MPS Supplement to the 3230XP Instruction Set and System Architecture Reference Manual	
OS/32 System Support Run Time Library (RTL) Reference Manual	All RTL routines available for the writing of system level control programs that perform the functions described.
OS/32 Operator Reference Manual	Describes the operator commands that can be used to perform SVC13 mapping and control functions. APU-related functions included APC, LPU, OPTION LPU, QUEUE.
OS/32 Supervisor Call (SVC) Reference Manual	Gives details on how to use SVC6, SVC13 and assembly language programming SVCs.
OS/32 System Macro Library Reference Manual	Describes the time and SVC13 macros.
OS/32 Link Reference Manual	Describes the use of OPTION LPU, APCONTROL and APMAPPING at task linkage time.
OS/32 Application Level Programmer Reference Manual	Gives details on writing a task trap handling routine that can be used to handle APU-related events.
OS/32 System Generation (Sysgen/32) Reference Manual	Describes the use of the MODULE command.

CHAPTER 4
SUPERVISOR CALL (SVC) INTERCEPTION

4.1 INTRODUCTION

SVC interception software is used to write programs that can emulate the SVC processing ability of OS/32. This software consists of macros that allow a task (intercepting task) to intercept the SVC of another task before it goes to the operating system for processing. Once intercepted, the SVC can be monitored by the intercepting task and sent to the operating system for processing or it can be processed by the intercepting task. Table 4-1 lists the system macros used for SVC interception.

TABLE 4-1 SYSTEM MACROS FOR SVC INTERCEPTION

MACRO	FUNCTION
ICREATE	Creates an SVC intercept path.
IREMOVE	Removes a previously created path.
IGET	Gets data from a data area of the task that issued an intercepted SVC.
IPUT	Puts data into a data area of the task that issued an intercepted SVC.
ICONT	Continues standard execution of an intercepted SVC by passing control to an OS/32 SVC executor.
IPROCEED	Allows the task that issued the intercepted SVC to proceed with its execution.
IROLL	Makes an intercepted task rollable.
ITERM	Terminates an intercepted SVC after processing.
ITRAP	Sends a task queue trap to a task.
IERRTST	Evaluates errors returned by any of the above macros and branches execution to specific error routines within the intercepting task.

The intercepting task tells the OS/32 SVC executor which SVC it will process or monitor. When the intercepting task is sent an SVC from the executor, the intercepting task handles the intercepted SVC while the task that issued the SVC is placed in a wait state. While executing the intercepted SVC, the intercepting task can read from or write to the address space of the task that issued the SVC.

A task is not aware that its SVC has been intercepted unless it is informed by the intercepting task.

SVC interception software must be configured in OS/32 at the time of system generation (sysgen). See the INTERCEPT configuration statement in the OS/32 System Generation (Sysgen/32) Reference Manual.

A task can intercept SVC calls only after it is linked with the intercept task option enabled (OPTION INTERCEPT). See the OS/32 Link Reference Manual for further details. The task can then be programmed to intercept any of the following SVCs issued by any application task in the system:

- SVC1
- SVC2 code 7
- SVC3
- SVC6
- SVC7

Intercepting tasks can be loaded and executed under the multi-terminal monitor (MTM). However, the intercepting task must be loaded from an account that has executive task (e-task) load privileges. See the OS/32 Multi-Terminal Monitor (MTM) System Planning and Operation Reference Manual for more information regarding e-task privileges.

4.2 HOW SUPERVISOR CALL (SVC) INTERCEPTION WORKS

In general, SVC interception software functions as follows:

1. A task with SVC interception enabled is built by Link. This intercepting task must:
 - reserve memory for a set of request descriptor block (RDB) buffers for each SVC to be intercepted,
 - build a circular list for storing addresses of RDB buffers containing information on intercepted SVCs,

- create, via the ICREATE macro, intercept paths that designate the SVCs to be intercepted, and
 - define, via the ICREATE macro, what control the intercepting task has over the SVCs it intercepts.
2. An application task issues an SVC.
 3. If no intercept path was created for that particular SVC, one of the standard OS/32 executors services the SVC.
 4. If an intercept path has been created for that SVC, the operating system:
 - intercepts the SVC before it reaches the OS/32 executor,
 - removes an RDB address from the circular list of the intercepting task,
 - loads the SVC's parameter block and identifying information into the RDB, and
 - sends a task event trap to the intercepting task to notify the task that an SVC has been intercepted.
 5. Execution of the intercepting task branches to the task event trap-handling routine. The address of this routine is specified when the path is created via the ICREATE macro.
 6. If the intercept path was built to monitor this SVC, the task event trap-handling routine issues an ICONT macro to return the SVC to the OS/32 executor for execution.
 7. If the intercept path was built to service the SVC, the task event trap-handling routine processes the SVC by the intercept macros IGET, IPUT, IROLL and ITRAP. Also, the routine can issue the IPROCEED macro to allow the application task to continue executing during SVC processing.
 8. After the task event trap-handling routine processes the SVC, it issues an ITERM macro that transfers control back to the application task that issued the SVC.
 9. The intercepting task exits the trap handler through the TEXTIT macro.

4.3 PREPARING A TASK FOR SUPERVISOR CALL (SVC) INTERCEPTION

Before creating an intercept path, an intercepting task must:

- build a set of RDB buffers for each type of SVC to be intercepted,

- build a circular list to store the addresses of the RDB buffers, and
- be prepared to handle a task event trap.

4.3.1 Request Descriptor Block (RDB) Buffers

The size of each RDB buffer built by the intercepting task depends on the size of the parameter block for the particular SVC to be intercepted. For example, a set of buffers allocated for SVC6 interception will be larger than a set of buffers for SVC1 interception. When an intercepting task uses one set of buffers for intercepting two or more SVC types, the buffer size must equal the size of the RDB needed to hold the largest parameter block associated with the SVCs to be intercepted. Figure 4-1 shows the RDB fields. To define a structure containing these fields, use the \$RDB macro.

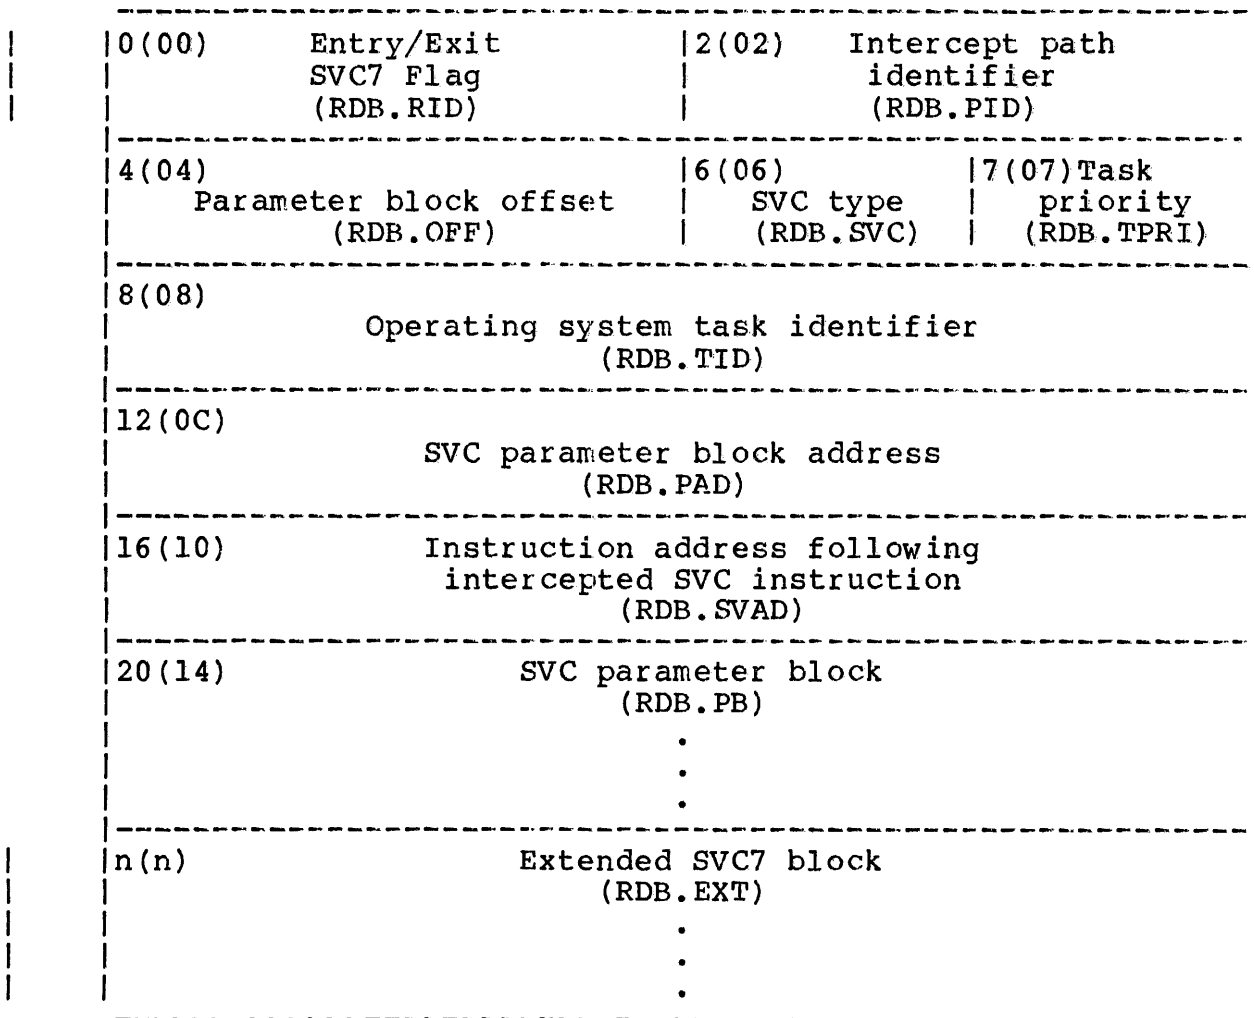


Figure 4-1 Request Descriptor Block

The fields contained within the RDB are described as follows:

Fields:

Exit/Entry
SVC7 Flag
(RDB.RID) is a halfword field containing an SVC7 flag. This flag is utilized exclusively to distinguish entry intercepts from exit intercepts. This field contains '00' on an entry intercept and '01' on an exit intercept. The RDB.RID field is only valid if an SVC7 intercept path is established.

Intercept
path
identifier
(RDB.PID) is a halfword field containing an SVC intercept path identifier exclusively reserved for one particular SVC interception.

NOTE

RDB.PID is not to be mistaken for the PID returned by ICREATE and used in IREMOVE.

Parameter
block
offset
(RDB.OFF) is a halfword field containing the hexadecimal offset value for the parameter block field within the RDB.

SVC
type
(RDB.SVC) is a 1-byte field containing a decimal number specifying the type of SVC that is to be intercepted.

- 01 indicates SVC1.
- 02 indicates SVC2, code 7.
- 03 indicates SVC3.
- 06 indicates SVC6.
- 07 indicates SVC7.

Task
priority
(RDB.TPRI) is a 1-byte field containing a decimal number specifying the priority of the task that issued the intercepted SVC.

OS task
identifier
(RDB.TID) is a 4-byte field containing the operating system task identifier for the task that issued the intercepted SVC.

SVC parameter block address (RDB.PAD) is a 4-byte field containing a hexadecimal number specifying the address of the parameter block for the SVC being intercepted. For SVC3 interceptions, this field contains the end of task code.

Instruction address following intercepted SVC instruction (RDB.SVAD) is a 4-byte field containing a hexadecimal number specifying the address of the instruction following the intercepted SVC instruction. This field is set to 0 for SVC3 interceptions.

SVC parameter block (RDB.PB) is a variable length field containing the parameter block of the intercepted SVC.

Extended SVC7 block (RDB.EXT) is a 32-byte field containing the name of the file assigned to a logical unit (lu). This field is only utilized if the SVC=(7,X) operand is specified in the ICREATE macro.

4.3.2 Circular List for Request Descriptor Block (RDB) Buffers

The intercepting task must have a standard circular list to hold the address of each RDB buffer. Figure 4-2 shows the fields of the standard circular list. When an SVC is sent to the intercepting task for processing, one RDB buffer address is automatically removed from the circular list and the RDB is filled with information identifying the intercepted SVC. The circular list can be created by the assembler instruction DLIST. See the appropriate Series 3200 Processor User's Manual or the Instruction Set Reference Manual for a more detailed explanation of the standard circular list.

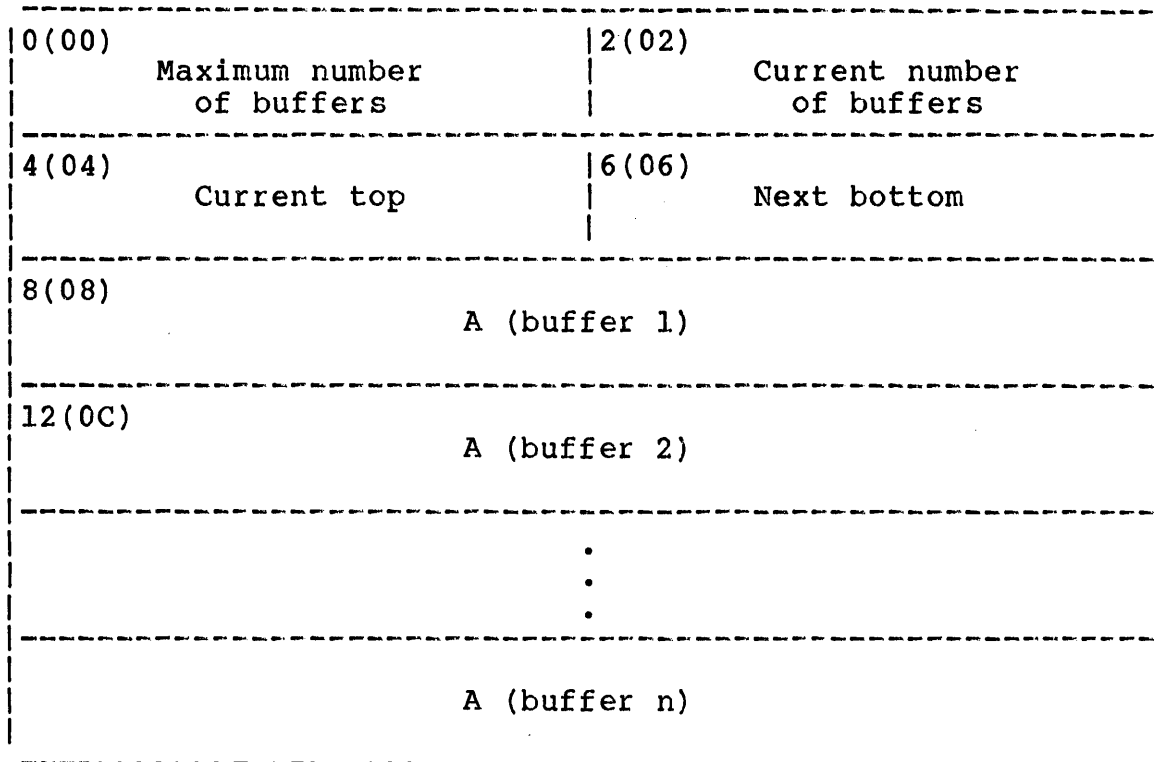


Figure 4-2 System Task Buffer List (Standard Circular List)

Fields:

- Maximum number of buffers is a halfword field indicating the maximum number of fullwords in the entire list.
- Current number of buffers is a halfword field indicating the number of fullwords currently in use. When this field equals zero, the list is empty. When this field equals the number of fullwords in the list, the list is full.
- Current top is a halfword field indicating the address of the RDB buffer that is currently at the top of the list.
- Next bottom is a halfword field indicating the address of the next RDB buffer that is at the bottom of the list.
- A (buffer n) indicates the address of an RDB buffer.

4.3.3 Task Event Trap

To receive a task event trap, an intercepting task must have the TSW.TESB bit in its task status word (TSW) set. See the OS/32 Application Level Programmer Reference Manual for more information on TSW bit settings. If this bit is not set, the task event trap will be queued until a TSW is loaded with this bit set. In addition, a task cannot receive a task event trap or task queue trap during execution of the task event trap-handling routine. These traps will be queued until the task exits from the routine.

Before execution branches to the task event trap-handling routine, the operating system places the address of the RDB in register 1 and a unique intercept identifier in register 0. To prevent the data in these registers from being lost during execution of the task event trap-handling routine, the intercepting task should be link-edited with the TEQSAVE task option. TEQSAVE informs the operating system which register contents should be saved and restored when a task enters or exits the task event trap-handling routine. See the OS/32 Link Reference Manual for more information on TEQSAVE.

4.4 CREATING INTERCEPT PATHS (ICREATE)

Before an intercepting task can intercept an SVC, it must create a path to the application task that contains the SVC to be intercepted. This path is created by executing code built by the ICREATE macro that informs the OS/32 SVC executor which SVC is to be intercepted by this path. The intercepting task also accesses the application task's address space through the intercept path.

An intercept path remains in effect until it is removed by the intercepting task creating it or until the intercepting task terminates. Although only one type of SVC can be intercepted by each path, there is no limit to the number of paths that can be created by one intercepting task.

The mode parameter of the ICREATE macro specifies when an SVC is to be intercepted. Under caller mode, the specified SVC is intercepted every time it is issued by the application task. When the recipient existent mode is specified, the SVC is intercepted only when it is directed toward a specified task, device, pseudo task or pseudo device that exists in the system. Under the recipient nonexistent mode, the SVC is intercepted only when it is directed toward a specified pseudo task or pseudo device created by execution of code built by the ICREATE macro.

4.5 HOW TO CREATE A PSEUDO DEVICE OR TASK WITH ICREATE

A pseudo device consists of a name and the SVC1 or SVC7 intercept paths attached to it. The pseudo device name, which is known to the system but does not actually refer to any system device or file, consists of a device name, filename and extension. A

device name that does not already exist for a real device or disk volume must be used. Pseudo devices ignore the file class/account number field of the file descriptor (fd).

When the operating system cannot find a device or filename in the system, it will search the list of pseudo devices. If a match occurs, the system will continue processing the SVC using the pseudo device.

To create a pseudo device using SVC interception software, the ICREATE macro should be set to specify either an SVC1 or SVC7. The recipient nonexistent mode should also be specified. An SVC1 intercept path must be in effect when an input/output (I/O) operation is attempted to a pseudo device; otherwise, an invalid function (X'C0') error status is returned.

A pseudo task consists of a name attached to one or more SVC6 intercept paths. A pseudo task name is known to the system but does not refer to an actual task existing in the system.

To create a pseudo task, issue the ICREATE macro specifying SVC6 and the recipient nonexistent mode. Because a pseudo task does not refer to a real task, the pseudo task cannot be cancelled. Both pseudo tasks and pseudo devices can be deleted by removing all intercept paths attached to them.

4.6 USE OF GENERIC NAMING FOR PSEUDO DEVICES AND TASKS

A pseudo device or task can be generically named. The following characters can be used for generic naming:

- An asterisk (*) represents any character or blank.
- A backward slash (\) represents any character.

If a pseudo device or task name specifies the filename and extension fields as blanks, the system substitutes filename and extension fields filled with asterisks. This has the effect of generically naming the filename and extension fields so that they will always match the input filename and extension.

If the operands of an ICREATE macro specify the recipient existent mode and a generic pseudo device or task name, a pseudo device or task must exist with its name exactly matching the one specified by ICREATE. An error will result if the names do not match. For example, a system is asked to create the following pseudo devices:

- FAKE:FILE1
- FAKE:FILE*
- FAKE:
- FAKE:FILE*.EXT

Normally, the following input will match the above pseudo devices:

INPUT NAME	SELECTED PSEUDO DEVICE
FAKE:	FAKE:
FAKE:FILE3	FAKE:FILE*
FAKE:FILE1	FAKE:FILE1
FAKE:FILE11	FAKE:
FAKE:FILEX.EXT	FAKE:FILE*.EXT
FAKE:FILEX.EX	FAKE:

When the code built by the ICREATE macro is issued specifying recipient nonexistent mode and the pseudo device FAKE:, the ICREATE function will not be performed because the pseudo device already exists. Consequently, when an ICREATE macro is used specifying recipient existent mode along with the pseudo device FAKE:FILE*, ICREATE will be executed because the pseudo device FAKE:FILE* already exists.

4.7 FUNCTIONAL SUMMARY OF SUPERVISOR CALL (SVC) INTERCEPTION

The following describes how interception works for each SVC and mode:

SVC1 caller	Any SVC1 issued by the specified task is intercepted.
SVC1 recipient existent	Any SVC1 directed to an lu assigned to the specified device or pseudo device is intercepted. (Note that disk volume interception is not supported for SVC1.)
SVC1 recipient nonexistent	The pseudo device is created and any SVC1 call specifying an lu assigned to this pseudo device is intercepted.
SVC2 code 7 caller	Any SVC2 code 7 issued by the specified task is intercepted.
SVC2 code 7 recipient existent	This call is invalid.

SVC2 code 7 recipient nonexistent	This call is invalid.
SVC3 caller	If the specified task goes to end of task for any reason, an SVC3 intercept will occur.
SVC3 recipient existent	This call is invalid.
SVC3 recipient nonexistent	This call is invalid.
SVC6 caller	Any SVC6 issued by the specified task is intercepted.
SVC6 recipient existent	Any SVC6 directed to the specified task or pseudo task is intercepted.
SVC6 recipient nonexistent	The pseudo task is created, and any SVC6 call directed to this pseudo task is intercepted.
SVC7 caller	Any SVC7 issued by the specified task is intercepted.
SVC7 recipient existent	Any SVC7 directed to the specified device, disk volume or pseudo device is intercepted.
SVC7 recipient nonexistent	The pseudo device is created and any SVC7 call specifying this pseudo device is intercepted.

4.8 FULL AND MONITOR CONTROL INTERCEPT PATHS

The ICREATE macro specifies the level of control that the intercept path allows an intercepting task to have over an application task.

A full control intercept path allows the intercepting task to exert full control over a task whose SVC has been intercepted. Specifically, the intercepting task can perform the following procedures:

- Make the task rollable via the IROLL macro. When an SVC is intercepted, the task that issued the SVC is placed in a wait state and made nonrollable. At the discretion of the intercepting task, the application task can be made rollable (assuming the application task can be rolled).

- Allow the application task to execute while it processes a proceed SVC via code built by the IPROCEED macro. When an SVC is intercepted, the application task that issued the SVC is placed in a wait state and made nonrollable. At the discretion of the intercepting task, the application task that issued the intercepted SVC can proceed with its execution while the intercepting task processes the SVC.
- Obtain data from the application task memory space via the IGET macro.
- Write data into the writable memory space of the application task via the IPUT macro.
- Send a task queue trap to the application task via the ITRAP macro. While processing the SVC, the intercepting task may find it necessary to send a task queue trap to the application task. The task queue item sent must have a valid OS/32 reason code in the high-order byte. In addition, the TSW of the application task must have the task queue entry (TQE) bit associated with the reason code set.

A monitor control intercept path allows the intercepting task to be notified whenever one of the designated SVCs is issued by an application task. Monitor control differs from full control in that once OS/32 has sent the task event trap to the intercepting task, the SVC is passed to the appropriate OS/32 executor and the task that issued the SVC proceeds with normal processing.

The following guidelines should be followed when assigning a level of control to the intercept path:

- Only one full control intercept path can be attached to a device or task (or pseudo device or task) for each type of SVC to be intercepted.
- A task or device (or pseudo task or device) can be attached to any number of monitor control intercept paths.

In the following example, a full control SVC7 intercept path is attached to device MAG:. A full control SVC1 intercept path is also attached to MAG:. No other SVC1 or SVC7 full control intercept paths can be attached. Of course, any number of SVC1 and SVC7 monitor control intercept paths can be attached to MAG;; here, one SVC7 and one SVC1 monitor control paths are attached.

Example:

```
ICREATE NAME=DEVNAME, MODE=RX, CONTROL=FC, SVC=(7)
ICREATE NAME=DEVNAME, MODE=RX, CONTROL=FC, SVC=(1)
ICREATE NAME=DEVNAME, MODE=RX, CONTROL=MC, SVC=(7)
ICREATE NAME=DEVNAME, MODE=RX, CONTROL=MC, SVC=(1)

DEVNAME DC C'          ' DEFINE 8 BLANK CHARACTERS

DC C'MAG '

DC C'          ' DEFINE 8 BLANK CHARACTERS (FD)

DC C'          ' DEFINE 4 BLANK CHARACTERS (EXTENSION)
```

4.9 HOW INTERCEPT PATHS HANDLE SUPERVISOR CALLS (SVCs) OCCURRING AT END OF TASK

SVC1 and SVC7 can be intercepted during end of task processing (including end of task processing after cancel), if intercept paths exist from these SVCs to devices assigned to the task's logical units. The intercepting task must be careful when writing into the operating system address space while executing these SVCs so as not to destroy the system's integrity.

If the application task is cancelled while the intercepting task is processing the SVC, SVC processing is aborted and the application task proceeds to end of task.

4.10 TERMINATING THE INTERCEPTED SUPERVISOR CALLS (SVCs)

When the intercepting task receives an SVC from a full control intercept path, the intercepting task has the option of returning the SVC to the operating system for processing. To do this, the intercepting task executes code built by an ICONT macro that allows the operating system to resume processing the intercepted SVC as if the intercept had never occurred. The ICONT macro cannot be used if an IPROCEED or IROLL macro has been used.

If the intercepting task chooses to process the SVC, the intercepting task executes code built by an ITERM macro after the SVC is processed. ITERM terminates the interception and, if no IPROCEED has been issued, allows the application task to resume execution with the instruction immediately following the intercepted SVC instruction.

Either ICONT or ITERM can be used to terminate interception from a monitor control intercept path. The system does not differentiate between the two calls in this case. Here the ICONT or ITERM macro replaces the RDB buffer address back on the circular list. It is very important that the ICONT or ITERM macro be used to replace the RDB.

Cancelling an application task under monitor or full control aborts the processing of the intercepted SVC in progress. The intercepting task must still issue an ICONT or ITERM to terminate the SVC interception.

4.11 HOW TO REMOVE INTERCEPT PATHS

An intercepting task can remove an intercept path by executing code built by an IREMOVE macro specifying the path to be removed. IREMOVE can be used for both immediate and delayed termination depending on whether the controlled shutdown or abort option is chosen.

The controlled shutdown option refuses all incoming requests and completes the servicing of all existing queued and executing SVCs. When processing of the last existing SVC intercepted by the path is completed, the path is removed from the system.

The abort option terminates all existing queued and executing SVCs before removing the intercept path from the system.

4.12 ERROR HANDLING

Run-time errors that result from executing intercept macro code are handled by user-written error routines within the intercepting task. When an error occurs, execution branches to the routine specified by either the IERRTST macro statement or the error parameter associated with each macro.

The IERRTST macro is issued immediately after a macro for which the error parameter has been omitted. If an error occurs, execution of the intercepting task will branch to a user-written error routine to handle the error. Error codes returned by the IERRTST macro are listed in Table 4-2. If no error occurs, execution continues at the instruction following the IERRTST macro.

If the ERROR parameter is specified with an intercept macro and an error occurs, execution branches to the specified error routine within the intercepting task. If no error occurs, execution proceeds to the next executable statement. The error routine pointed to by the ERROR parameter can contain an IERRTST macro to identify what error has occurred.

TABLE 4-2 ERROR CODES RETURNED FOR INTERCEPT MACROS

ERROR CODE	MEANING	RELEVANT MACROS
MO	Invalid interception mode	ICREATE
AD	Invalid address in parameter control block (PCB)	ICREATE ITERM ICONT IREMOVE ITRAP IGET IPUT
EX	When MODE=RX, it indicates that no device or task of the given name exists. When MODE=RN, it indicates that the specified pseudo device or task could not be created because a device or task of the same name was found to exist.	ICREATE
SP	Insufficient system space to do request, or NINTC > 64 or PBSIZE > 998,	ICREATE ITERM ITRAP IGET IPUT
CT	Full control already selected.	ICREATE IROLL IPROCEED ITRAP IGET IPUT
HA	Invalid queue handler name.	ICREATE
FD	Invalid device name or task name.	ICREATE
ST	Invalid state for call; e.g., IROLL followed by ICONT or issuing INPUT with monitor control intercept path.	ICONT IREMOVE IROLL IPROCEED ITRAP IGET IPUT
TP	Task queue item not added.	ITRAP

TABLE 4-2 ERROR CODES RETURNED FOR INTERCEPT MACROS (Continued)

ERROR CODE	MEANING	RELEVANT MACROS
RD	Invalid RDB.	ITERM ICONT IROLL IPROCEED ITRAP IGET IPUT
ID	Intercept path corresponding to this path ID does not exist.	IREMOVE
WR	Attempt to copy SVC parameter block back into write-protected area.	ITERM
CD	Invalid subcode in SVC parameter block. SVC interception software not included at sysgen.	All
NT	Intercepted task has gone to end of task.	IROLL IPROCEED ITRAP IGET IPUT

4.13 MACROS USED WITH SUPERVISOR CALL (SVC) INTERCEPTION

Once configured for SVC interception, the operating system allows tasks to execute code built by macros for SVC interception provided the tasks were linked with the intercept option.

This section gives the syntax for the SVC macros described in the previous sections. See the OS/32 System Macro Library Reference Manual for a list of syntax rules.

4.13.1 ICREATE Macro

The ICREATE macro creates an intercept path for a particular SVC type. See Table 4-3 for valid combinations for the SVC, MODE and NAME parameters.

Format:

NAME	OPERATION	OPERAND
symbol	ICREATE	$\text{SVC} = \left\{ \begin{array}{l} (1) \\ (2, 7) \\ (3) \\ (6) \\ (7 [, X]) \end{array} \right\}$ $, \text{MODE} = \left\{ \begin{array}{l} \text{CL} \\ \text{RX} \\ \text{RN} \end{array} \right\}$, NAME=pointer , TID=pointer $, \text{CONTROL} = \left\{ \begin{array}{l} \text{FC} \\ \text{MC} \end{array} \right\}$, BUFFERL=pointer [, HANDLER=pointer] , PID=pointer , EXEC=pointer [, PBSIZE=n] [, SVAR=pointer] [, ERROR=pointer] [, PCB=pointer] [, FORM=L] [, NINTC=n] $\left[, \text{NOTIFY} = \left\{ \begin{array}{l} \text{EX} \\ \text{EN} \\ \text{NX} \end{array} \right\} \right]$ $\left[, \text{IOPT} = \left\{ \begin{array}{l} \text{YES} \\ \text{NO} \end{array} \right\} \right]$

Operands:

SVC= is an integer, enclosed by parentheses, that indicates the type of intercept path to be created:

- (1) indicates SVC1.
- (2,7) indicates SVC2 code 7.
- (3) indicates SVC3.
- (6) indicates SVC6.
- (7 ,X) indicates SVC7. X indicates an optional SVC7 intercept by lu when the lu is assigned to a file.

MODE= indicates one of the following interception modes:

- CL indicates caller mode
- RX indicates recipient existent mode
- RN indicates recipient nonexistent mode

When CL is specified, an intercept path is created for all SVCs (selected by the SVC parameter) issued from the task specified in the NAME or TID parameter.

When RX is specified, an intercept path is created for all SVCs (selected by the SVC parameter) directed to an existing task, device, pseudo task or pseudo device specified in the NAME parameter.

When RN is specified, a pseudo device is created for SVC1 or SVC7, or a pseudo task is created for SVC6. The pseudo device or task is attached to the intercept path created by the call.

A pseudo task or pseudo device is deleted when all intercept paths attached to it are removed. When a pseudo device is assigned without SVC7 interception, the requested access privileges are ignored and shared read/shared write privileges are granted. If an SVCl is attempted to a pseudo device without an interception in effect, an invalid function error (X'C0') is returned.

NAME=

indicates the address of the memory location specifying the name of a device task, pseudo device or pseudo task. This location must be fullword boundary-aligned and contain eight bytes of blanks followed by a standard file descriptor (fd) or task identifier (taskid). An fd must be packed, left-justified and padded with blanks within the fullword. A taskid must be left-justified and padded with blanks.

When RX or RN is specified by the MODE parameter, the standard fd or taskid given with the NAME parameter can include an asterisk (*) or a backward slash (\) to allow generic naming. See Section 4.6 for more information on the use of generic naming of pseudo devices and tasks.

TABLE 4-3 VALID COMBINATIONS FOR SVC, MODE AND NAME PARAMETERS

ICREATE PARAMETERS			FUNCTION
SVC=	MODE=	NAME=	
(1)	CL	taskid	Intercepts any SVC1 issued from the task.
	RX	fd	Intercepts any SVC1 directed to the existing device.
	RN	fd	Creates a pseudo device and intercepts any SVC1 directed to it.
(2,7)	CL	taskid	Intercepts any SVC2 code 7 issued from the task.
	RX	--	No function; specifying fd or taskid results in error.
	RN	--	Results in error.
(3)	CL	taskid	End of task interception; occurs no matter how a task terminates.
	RX	--	No function; specifying fd or taskid results in error.
	RN	--	Results in error.
(6)	CL	taskid	Intercepts any SVC6 issued from
	RX	taskid	Intercepts any SVC6 directed to the existing task.
	RN	taskid	Creates a pseudo task and intercepts any SVC6 directed to it.
(7)	CL	taskid	Intercepts any SVC7 issued from the task.
	RX	fd	Intercepts any SVC7 directed to the existing device.
	RN	fd	Creates a pseudo device and intercepts any SVC7 directed to it.

TID= indicates the address of a fullword location containing a taskid. This parameter, which is mutually exclusive with the NAME= parameter, can be used when MODE=CL or MODE=RX with SVC6, to identify the task to be intercepted. The TID can be obtained from field RDB.TID of an RDB from a previously intercepted SVC call.

CONTROL= contains a mnemonic indicating either full control (FC) or monitor control (MC) over intercepted SVCs.

When CONTROL=FC, an intercepting task can exert full control over an application task's intercepted SVCs.

When CONTROL=MC, an intercepting task acts as a monitor only; it has no control over an intercepted SVC.

BUFFERL= indicates the address of the standard circular list that contains the addresses of available RDB buffers.

The RDB used by the intercepting task to identify an intercepted SVC must not be moved to a new location after the interception takes place. The system ensures that the address of this RDB is the same as the address of the RDB that was passed to the intercepting task when the interception occurred.

HANDLER= indicates the address of a fullword location containing the name of a queue handler. This name, a maximum of eight characters, is left-justified and padded with blanks. If this parameter is omitted, the default queue handler is invoked.

NOTE

Currently, user-defined queue handlers are not supported.

PID= indicates the address of a halfword location that is used by the system to store the path identifier for the intercept path.

NOTE

PID= is not to be confused with RDB.PID. PID= identifies the path issued by the IREMOVE statement.

EXEC= is the address of an SVC intercept executor routine within the intercepting task. This routine will process intercepted SVCs of the type specified with the SVC parameter. During SVC interception, the system removes an RDB specified by the list, fills it with information and queues a task event trap with the specified executor address to the intercepting task.

On entry to an executor routine, general register 0 contains the PID of the intercept path and general register 1 contains the address of the RDB buffer associated with the intercepted SVC. The executor routine executes as a task event service routine.

PBSIZE= specifies the number of bytes in the parameter block for the SVC indicated by the SVC parameter.

When this parameter is omitted, the parameter block size defaults to the standard sizes documented for each type of SVC in the OS/32 Supervisor Call (SVC) Reference Manual, except for SVC2 code 7 interception, which defaults to eight bytes.

The size of the RDB.PB field in the RDB for this interception path is the value of the PBSIZE parameter (or its default if PBSIZE is not specified).

SVAR= is the address of a fullword location containing user-defined data. This data is passed to the intercept logic. The queue handler named by the HANDLER parameter can later access the data. The SVAR parameter is for user-defined purposes when needed by a user-defined queue handler.

NOTE

Currently, user-defined queue handlers are not supported.

ERROR= is the address of an error routine within the intercepting task. If a run-time error occurs for this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following code built by the macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter.

When no PCB parameter is included, macro code automatically builds a new PCB and initializes it with values corresponding to the other specified parameters.

FORM= L requests a PCB to be built but not executed. Macro code constructs a PCB for this macro and initializes it with values. Subsequent macros can reference this PCB via the PCB parameter.

NINTC= n specifies the number of interceptions that can be handled concurrently for this intercept path. If there are more SVC interceptions outstanding than can be handled concurrently, the excess interceptions are queued. The default value for n is 1. The default value of 1 means that the intercepting task will receive notification of a single intercept, regardless of the number of RDBs. To allow an intercepting task to handle multiple concurrent intercepts (e.g., proceed I/O, halt I/O, etc.) this value should be greater than 1.

NOTIFY= designates the time at which a task is informed of an SVC7 intercept.

- EX intercepts on an exit from the SVC7 handler only
- EN intercepts on an entry from the SVC7 handler only
- NX intercepts on both entry and exit from the SVC7 handler

Where entry means the intercept occurs before the SVC7 is handled and exit means the intercept occurs after the SVC7 is handled.

IOPT= aids in preventing a system task deadlock between two system tasks that create intercept paths for SVC's issued to each other. If the YES option is specified, the system task creating the intercept path is temporarily removed from a wait state and rejects any intercept directed to it from another system task. Upon failing the intercept, the system task is placed back into intercept wait state. This option may be issued on a per path basis within a system task. The IOPT operand does not guarantee that a system deadlock will always be prevented; however, since the system task is temporarily removed from a wait state and dispatched to handle its TEQ, a user is given the capability to detect the possible deadlock situation and take corrective action. This option default is NO.

4.13.2 IREMOVE Macro

The IREMOVE macro allows an intercepting task to remove one or all previously created SVC intercept paths.

Format:

NAME	OPERATION	OPERAND
symbol	IREMOVE	PID=pointer ,TERM={ CS AB } [,ERROR=pointer] [,PCB=pointer] [,FORM=L]

Operands:

PID= is the address of the path identifier specifying the path being removed. A zero value in the PID halfword removes all existing intercept paths.

TERM= indicates either of two termination modes for intercepted SVCs already queued for the intercepting task:

- AB indicates abort. OS/32 aborts all currently queued requests before path removal.
- CS indicates controlled shutdown. OS/32 services only currently queued requests before path removal; requests made after TERM=CS is issued cannot be queued or processed.

If this parameter is omitted, AB is the default.

ERROR= is the address of an error routine within the intercepting task. If a run-time error occurs for this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter.

If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM=L requests that a PCB be built but not executed. A PCB is built by this macro and initialized with values. Subsequent macros can reference this PCB via the PCB parameter.

4.13.3 IGET Macro

The IGET macro allows an intercepting task to get data from the application task whose SVC is intercepted.

Format:

NAME	OPERATION	OPERAND
symbol	IGET	RDB=pointer ,ADST=pointer ,ADEND=pointer ,SDST=pointer ,SEND=pointer [,ERROR=pointer] [,PCB=pointer] [,FORM=L] [,DONE=addr]

Operands :

- RDB=** is the address of the RDB buffer built for the intercepted SVC.
- ADST=** is the start address of a data area within the application task whose SVC is intercepted. The contents of this area are transferred to an intercepting task data area.
- ADEND=** is the end address of the data area within the application task whose SVC is intercepted.
- SDST=** is the start address of a data area within the intercepting task. This area receives the data from the application task.
- SDEND=** is the end address of the data area within the intercepting task.
- ERROR=** is the address of an error routine within the intercepting task. If a run-time error occurs for this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the macro.
- PCB=** is the address of a PCB previously constructed and initialized by the FORM=L parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.
- FORM=** L requests a PCB be built but not executed. A PCB is built for this macro and initialized with values. Subsequent macros can reference this PCB via the PCB parameter.
- DONE=** is an address that specifies that the macro is to be a PROCEED call. When the call is completed, a task event interrupt occurs, using the routine specified by the address in the DONE parameter. This routine enters with R0 containing the error code for the call and R1 pointing to the macro's parameter block. Once this routine has finished processing, it exits using the code built by the TEXT macro.

The proceed form of the IGET macro must be used if an IROLL macro was issued to the application task whose SVC is intercepted. The system cannot guarantee that the application task is in memory or that it can be rolled into memory within a reasonable time.

4.13.4 IPUT Macro

The IPUT macro lets an intercepting task put data into a data area of the application task whose SVC is intercepted.

Format:

NAME	OPERATION	OPERAND
symbol	IPUT	RDB=pointer ,ADST=pointer ,ADEND=pointer ,SDST=pointer ,SDEND=pointer [,ERROR=pointer] [,PCB=pointer] [,FORM=L] [,DONE=addr]

Operands:

RDB=	is the address of the RDB buffer built for the intercepted SVC.
ADST=	is the start address of a data area within the application task. This area receives the contents of an intercepting task data area.
ADEND=	is the end address of the data area within the application task.
SDST=	is the start address of a data area within the intercepting task. The contents of this area are put into the application task data area.
SDEND=	is the end address of the data within the application task.
ERROR=	is the address of an error routine within the intercepting task. If a run-time error occurs for code built by this macro, execution branches to this error routine.

If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the macro.

PCB= is the address of a PCB previously constructed and initialized by the **FORM=L** parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM= L requests a PCB be built but not executed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the **PCB** parameter.

DONE= is an address that specifies that the macro is to be a proceed call. When the call is completed, a task event interrupt occurs, using the routine specified by the address in the **DONE** parameter. This routine enters with general register 0 containing the error code for the call and general register 1 pointing to the macro's parameter block. Once this routine has finished processing, it exits using the code built by the **TEXT** macro.

The proceed form of the **IPUT** macro must be used if an **IROLL** macro was issued to the application task. The system cannot guarantee that the application task is in memory or that it can be rolled into memory within a reasonable time.

4.13.5 ICONT Macro

The **ICONT** macro relinquishes control of an intercepted **SVC** by returning control to an **OS/32 SVC** executor.

Format:

NAME	OPERATION	OPERAND
symbol	ICONT	RDB=pointer [,ERROR=pointer] [,PCB=pointer] [,FORM=L]

Operands:

RDB= is the address of the RDB buffer built for the intercepted SVC.

ERROR= is the address of an error routine within the intercepting task. If a run-time error occurs for code built by this macro, execution branches to this error routine.

If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the code built by the macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter.

If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM=L requests a PCB be built but not accessed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the PCB parameter.

4.13.6 IPROCEED Macro

After an SVC has been intercepted, the intercepting task can execute code built by an IPROCEED macro to allow the application task that issued the SVC to proceed with its execution. Until the intercepting task executes code built by an IPROCEED macro, the application task is in a wait state.

Format:

NAME	OPERATION	OPERAND
symbol	IPROCEED	RDB=pointer [,ERROR=pointer] [,PCB=pointer] [,FORM=L] [,CC=n]

Operands:

RDB= is the address of the RDB buffer built for the intercepted SVC.

ERROR= is the address of an error routine within the intercepting task. If a run-time error occurs for code built by this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following code built by the macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM= L requests a PCB be built but not accessed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the PCB parameter.

CC= n is a decimal number specifying the setting of the application task program status word (PSW) condition code after the SVC instruction execution. If the CC parameter is omitted, the condition code of the application task PSW is set to zero.

4.13.7 IROLL Macro

After an SVC is intercepted, an IROLL macro lets an intercepting task change the status of the application task from nonrollable to rollable, provided that the task was established as rollable by Link. This allows OS/32 to roll out a task having an intercepted SVC that requires lengthy processing.

Format:

NAME	OPERATION	OPERAND
symbol	IROLL	RDB=pointer [,ERROR=pointer] [,PCB=pointer] [,FORM=L]

Operands:

RDB= is the address of the RDB buffer built for the intercepted SVC.

ERROR= is the address of an error routine within the intercepting task. If a run-time error occurs for this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM= L requests a PCB be built but not accessed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the PCB parameter.

4.13.8 ITERM Macro

The ITERM macro terminates SVC processing. It also allows an intercepting task to return the parameter block of the SVC it processed to the application task that issued the SVC. The returned parameter block can have updated information such as status, number of bytes transferred, etc.

Format:

NAME	OPERATION	OPERAND
symbol	ITERM	RDB=pointer ,TRAP=pointer ,COPY={ Y N } [,ERROR=pointer] [,PCB=pointer] [,FORM=L] [,CC=n]

Operands:

RDB=	is the address of the RDB buffer built for the intercepted SVC.
TRAP=	is the address of a fullword that contains an item to be added to the task queue of the application task whose SVC is intercepted.
COPY=	Y (yes) indicates that the SVC parameter block in the RDB is to be copied back into the parameter block of the intercepted SVC. N (no) indicates the copy operation is not performed. If this parameter is omitted, N is the default.
ERROR=	is the address of an error routine within the intercepting task. If a run-time error occurs for code built by this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the code built by the macro.
PCB=	is the address of a PCB previously constructed and initialized by the FORM=L parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.
FORM=	L requests that a PCB be built but not accessed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the PCB parameter.
CC=	n is a decimal number specifying the setting of the application task PSW condition code after the SVC instruction execution. If the CC parameter is omitted, the condition code of the application task PSW is set to zero.

4.13.9 ITRAP Macro

The ITRAP macro allows an intercepting task to send a task queue item to an application task whose SVC is intercepted. The task queue item can be any of the task queue items supported by OS/32.

Format:

NAME	OPERATION	OPERAND
symbol	ITRAP	{RDB=pointer} {TID=pointer} [,TRAP=pointer] [,ERROR=pointer] [,PCB=pointer] [,FORM=L] [,DONE=addr]

Operands:

RDB= is the address of the RDB buffer built for the intercepted SVC.

TID= is the address of a fullword containing the taskid for the task. Before issuing an ITRAP macro with the TID parameter, the intercepting task must have obtained the task identifier from an RDB and placed it into the fullword location.

NOTE

The TID form of this macro can be used to send a trap to a task that is not being intercepted.

TRAP= is the address of a fullword that contains an item to be added to the task queue of the application task having an SVC that is intercepted.

ERROR= is the address of an error routine within the intercepting task. If a run-time error occurs for code built by this macro, execution branches to this error routine. If this parameter is omitted and a run-time error occurs, execution resumes with the instruction following the code built by the macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM= L requests that a PCB be built but not accessed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the PCB parameter.

DONE= is an address that specifies that the macro is to be a PROCEED call. When the call is completed, a task event interrupt occurs, using the routine whose address is specified in the DONE parameter. This routine enters with general register 0 containing the error code for the call and general register 1 pointing to the macro's parameter block. Once this routine has finished processing, the intercepting task exits using code built by the TEXT macro.

The proceed form of the ITRAP macro must be used if an IROLL macro was specified in the application task having an SVC that is intercepted. The system cannot guarantee that the application task is in memory or that it can be rolled into memory within a reasonable time.

4.13.10 IERRTST Macro

The IERRTST macro allows an intercepting task to evaluate errors resulting from execution of code built by intercept macros in order to branch to appropriate error handling routines.

Format:

NAME	OPERATION	OPERAND
symbol	IERRTST	xx=pointer . . [xx=pointer] [ELSE=pointer] [PCB=pointer] [FORM=L]

Operands:

xx= is a 2-character alphabetic string specifying one of the error codes for the intercept macros. See Table 4-2.

pointer specifies the name of an intercepting task error routine that handles errors having a returned error code identical to the one specified by the xx parameter. For instance, an IERRTST macro might include these parameters for evaluating an IPUT macro:

```
IERRTST AD=pointer,NT=pointer,RD=pointer
```

These parameters specify the addresses of the error routines to which execution will branch whenever the returned error code equals AD, NT or RD.

ELSE= is the name of an error routine to be executed for errors other than those specified in the xx parameter. If this parameter is omitted, one of the following actions occurs for returned errors:

- If the returned error code corresponds to the one specified by the xx parameter, execution branches to a specific error routine.
- If the returned error code does not correspond to the one specified by the xx parameter, execution branches to the instruction immediately following the code built by the IERRTST macro.

PCB= is the address of a PCB previously constructed and initialized by the FORM=L parameter. If this parameter is omitted, a new PCB is automatically built and initialized with values corresponding to the other specified parameters.

FORM= L requests that a PCB be built but not accessed. A PCB is built for this macro and initialized with values. Subsequent macros can refer to this PCB via the PCB parameter.

4.13.11 \$RDB Macro

The \$RDB macro is used to define a structure containing the symbolic names for all of the RDB fields. It is recommended that symbolic names be used to refer to the RDB fields instead of coding the hexadecimal offsets to the fields.

Format:

NAME	OPERATION	OPERAND
symbol	\$RDB	

4.14 SAMPLE SUPERVISOR CALL (SVC) INTERCEPTION PROGRAMS

The following program uses SVC interception software to intercept SVC1 to the existing real device MAG1. Each time an SVC1 is issued to MAG1, the program prints the following message:

SVC 1 CALL INTERCEPTED

SVC1 is terminated with a device unavailable error code, X'A0'.

```

    $RDB                                DEFINES AN RDB STRUCTURE
*   ADD AN RDB BUFFER ADDRESS TO THE RDB BUFFER ADDRESS LIST.
    LA      0,RDB                        LOAD THE ADDRESS OF THE RDB
*                                           INTO REGISTER 0
    ABL     0,BUFLIST                    ADD THE ADDRESS OF THE RDB
*                                           TO THE CIRCULAR LIST
*   CREATE THE INTERCEPT PATH
    ICREATE NAME=INTNAME,                FD FOR DEVICE NAME                X
        MODE=RX,                          RECIPIENT-EXISTENT MODE                X
        CONTROL=FC,                       GIVES INTERCEPTING TASK FULL CONTROL X
        SVC=(1),                          ALL SVC 1 ARE TO BE INTERCEPTED    X
        EXEC=INTRTN,                      POINTS TO THE SVC EXECUTOR ROUTINE    X
        BUFFERL=BUFLIST,                  ASSIGNS POINTER TO FREE BUFFER LIST   X
        PID=PATHID,                      DATA AREA FOR INTERCEPT PATH ID    X
        ERROR=BOMBOUT                     ERROR ROUTINE FOR ICREATE MACRO

```

* IF ERROR OCCURS IN ICREATE MACRO ENABLE TASK EVENT TRAP SO TASK
 * CAN GO INTO TRAP WAIT FOR INTERCEPTS TO OCCUR

* LOAD TSW WITH WAIT STATE SET AND TASK EVENT TRAPS ENABLED

LTSW TETS,WT

* COME HERE IF ERROR OCCURS IN ICREATE MACRO

BOMBOUT SVC 3,1 FAIL TASK ON ERROR

* ALLOCATE DATA AREA FOR ICREATE

	ALIGN 4		
INTNAME	DC	C' '	NODE NAME
	DC	C' '	RESERVED
	DC	C'MAG1'	DEVICE NAME
	DC	C' '	FILE NAME PART 1
	DC	C' '	FILE NAME PART 2
	DC	C' '	EXTENSION

BUFLIST DLIST 1 DESIGNATE 1 RDB IN CIRCULAR LIST

RDB DS RDB.PB+20 ALLOCATES SIZE OF RDB + SVC 1

PATHID DS 2 DESIGNATE AREA FOR PATH ID

* TRAP EVENT SERVICE ROUTINE

* THE FOLLOWING ROUTINE IS EXECUTED WHEN AN SVC IS INTERCEPTED

INTRTN	SVC 2,NOTIFY	LOG MESSAGE THAT SVC 1 WAS INTER-
*		CEPTED
	LHI 0,X'A000'	RETURN DEVICE UNAVAILABLE STATUS
*		FOR INTERCEPTED SVC 1
	STH 0,RDB.PB+2(1)	SAVE SVC 1 STATUS IN STATUS FIELD
*		OF RDB

*

* TERMINATE THE INTERCEPTED CALL, COPYING THE MODIFIED SVC
 * PARAMETER BLOCK IN THE RDB BACK OVER THE USER'S SVC PARAMETER
 * BLOCK.

ITERM RDB=(1),COPY=Y

TEXIT EXIT THE TASK EVENT ROUTINE

ALLOCATE DATA AREA FOR TRAP EVENT SERVICE ROUTINE

	ALIGN 4
NOTIFY	DB 0,7,0,22
	DC C'SVC 1 CALL INTERCEPTED'
	END

The following program creates a pseudo device to which a user task (u-task) can assign and write. The user's data buffer is passed to the OS/32 command processor via SVC2 code 14 to be executed as a command line.


```

IRDR    EQU    *
        SVC    2,PEEK01          GET NAME OF SYSTEM CONSOLE
        L      R00,CON
        ST     R00,SVC7.VOL+SVC7CON
        LHI   R00,SV7.ASGN!SV7.SRW
        SLL   R00,16            ASSIGN LU 0 SRW
        ST     R00,SVC7.OPT+SVC7CON
        SVC    7,SVC7CON        ASSIGN TO SYSTEM CONSOLE
        LB     R00,SVC7.STA+SVC7CON
        LR     R00,R00          WAS THE ASSIGN OK?
        BNZ   BADCON           NO
        LIS   R00,0            CHANGE SVC 7 TO FETCH ATTR
        STH   R00,SVC7.OPT+SVC7CON
        SVC    7,SVC7CON        FETCH ATTRIBUTES ON CON:
        LB     R00,SVC7.STA+SVC7CON
        LR     R00,R00          WAS THE FETCH OK?
        BNZ   BADCON           NO
        LHI   R00,SV7.CLOS      CHANGE SVC 7 TO CLOSE
        SRLS  R00,8            DO NOT DESTROY DEVICE CODE
        STB   R00,SVC7.OPT+SVC7CON
        SVC    7,SVC7CON        CLOSE THE SYSTEM CONSOLE
        LB     R00,SVC7.STA+SVC7CON
        LR     R00,R00          WAS THE CLOSE OK?
        BNZ   BADCON           NO
        LHI   R00,X'7FFF'       BAD LENGTH FOR SVC 2,14 TO GET
        STH   R00,COMMAND+4     MAX LENGTH ALLOWED BY SYSTEM
        SVC    2,COMMAND        WILL GET ERROR STATUS 3
        LH    R00,COMMAND+6     USE AS IRDR LENGTH
        STH   R00,SVC7.LRC+SVC7CON
        SPACE 1
        LHI   R00,RDBNUM        NUMBER OF RDB'S
        LA    R01,RDBPOOL       ADDRESS OF RDB POOL
INTRDB  EQU    *
        ATL   R01,RDBP          ADD RDB TO QUEUE
        AHI   R01,RDBSIZE       ADDRESS OF NEXT RDB
        SIS   R00,1            ALL RDB'S ADDED TO QUEUE?
        BNZ   INTRDB           NO
        SPACE 1
        ICREATE SVC=(7),MODE=RN,NAME=NAME,                X
                CONTROL=FC,BUFFERL=RDBP,PID=PID,EXEC=INT7
        IERRTST FD=BADFD,EX=BADEX,ELSE=BADALL
        ICREATE SVC=(1),MODE=RX,NAME=NAME,PBSIZE=SVCLX,   X
                CONTROL=FC,BUFFERL=RDBP,PID=PID,EXEC=INT1
        IERRTST FD=BADFD,EX=BADEX,ELSE=BADALL
        SPACE 1
        LTSW  WT,TETS          ENTER TRAP WAIT
        SPACE 3
BADFD   SVC    2,LOGFD
        SVC    3,1
BADEX   SVC    2,LOGEX
        SVC    3,1
BADALL  SVC    2,STRANGE
        SVC    3,1
BADCON  SVC    2,LOGCON
        SVC    3,1
        SPACE 1

```

```

        ALIGN 4
LOGFD   DC      H'7',H'8'
        DC      C'FD ERROR'
LOGEX   DC      H'7',H'8'
        DC      C'EX ERROR'
STRANGE DC      H'7',H'8'
        DC      C'!! ERROR'
LOGCON  DC      H'7',H'12'
        DC      C'!!CON ERROR '
        SPACE 1
NAME    DC      C'          IRDR          '
PID     DSF     1
        SPACE 1
RDBNUM  EQU     3                      NUMBER OF RDB'S IN POOL
RDBSIZE EQU     RDB.+SVC7.             MAXIMUM SIZE OF RDB
RDBP    DLIST   RDBNUM                 RDB POOL
RDBPOOL DS      RDBSIZE*RDBNUM        RDB BUFFERS
        TITLE  SVC 7 TEQ HANDLER
*****
*
*          SVC 7 INTERCEPT EXECUTOR
*
*****
INT7     EQU     *
        LR      R10,R01                SAVE RDB POINTER
        LR      R11,R10
        AH      R11,RDB.OFF(R10)       ADDRESS OF SVC 7 PBLK
        LB      R00,SVC7.OPT(R11)     GET SVC 7 OPTIONS
        LR      R00,R00                FETCH ATTRIBUTES?
        BZ      DOFETCH                YES
        CLHI   R00,X'FF'              EXTENDED SVC 7 FUNCTIONS?
        BE      INT7.NS                YES - NOT SUPPORTED
        THI    R00,X'40'              ASSIGN?
        BNZ    DOOPEN                  YES
        THI    R00,X'04'              CLOSE?
        BNZ    DOCLOSE                 YES
        THI    R00,X'21'              CHAP OR CHECKPOINT?
        BNZ    INT7.IG                 YES - IGNORE
        SPACE 1
INT7.NS EQU     *
        SVC     2,UNPACK7              PUT SVC 7 OPTION IN ERROR MESSAGE
        SVC     2,LOG7ERRC             AND LOG ERROR MESSAGE
        LIS    R00,1                   RETURN ILLEGAL FUNCTION TO USER
        STB    R00,SVC7.STA(R11)       AS AN ERROR STATUS
        ITERM  PCB=TERM,RDB=(R10)     TERMINATE THIS SVC 7
        TEXTIT PCB=EXIT                EXIT FROM TEQ HANDLER
        SPACE 3
*
*          IGNORE SVC 7 COMMAND PROCESSOR
*
INT7.IG EQU     *
        ITERM  PCB=TERM,RDB=(R10)     IGNORE THIS SVC 7
        TEXTIT PCB=EXIT                EXIT FROM TEQ HANDLER
        SPACE 3
*
*          OPEN PROCESSOR

```

```

*
DOOPEN EQU *
LB R15,SVC7.OPT+1(R11) GET ACCESS PRIVILEGES
SRLS R15,5 SRO = 0 & ERO = 1
CLHI R15,2 REQUESTING READ ONLY ACCESS?
BL OPEN.ERR YES - ERROR
B OPEN.OK SKIP SECURITY CHECK
SPACE 2
*-----USER DEFINED SECURITY CHECK FOLLOWS-----
L R15,RDB.TID(R10) MOVE TID FOR PEEK03
ST R15,TID
SVC 2,PEEK03 INFO ON USER TASK
LM R14,MONITOR GET NAME OF USERS MONITOR
CLI R14,C'.MTM' TASK A SUB-TASK OF MTM?
BNE OPEN.OK NO
CLI R15,C' ' BE SURE
BNE OPEN.OK NO?
LM R14,TASKNAME GET NAME OF USER
CLI R14,C'LEE ' IS IT ME?
BNE OPEN.ERR NO
CLI R15,C' ' BE SURE
BNE OPEN.ERR NO?
L R15,LEGACY GET NAME OF USERS TERMINAL
CLI R15,C'CT42' IS IT MINE?
BNE OPEN.ERR NO
L R15,ACCT.P GET USERS PRIVATE ACCOUNT NUMBER
CLHI R15,29 AM I IN MY ACCOUNT?
BNE OPEN.ERR NO
L R15,ACCT.G GET USERS GROUP ACCOUNT NUMBER
CLHI R15,18 DO I HAVE MY CORRECT GROUP ACCOUNT?
BNE OPEN.ERR NO
*-----
OPEN.OK EQU *
ICONT PCB=CONT,RDB=(R10) RETURN TO OS SVC 7 EXECUTOR
TEXTIT PCB=EXIT EXIT FROM TEQ HANDLER
SPACE 2
OPEN.ERR EQU *
LIS R15,9 RETURN ASSIGNMENT ERROR TO USER
STB R15,SVC7.STA(R11)
ITERM PCB=TERM,RDB=(R10) RETURN BAD STATUS TO USER
TEXTIT PCB=EXIT EXIT FROM TEQ HANDLER
SPACE 3
*
* CLOSE PROCESSOR
*
DOCLOSE EQU *
ICONT PCB=CONT,RDB=(R10) RETURN OS OS SVC 7 EXECUTOR
TEXTIT PCB=EXIT EXIT FROM TEQ HANDLER
SPACE 3
*
* FETCH ATTRIBUTES PROCESSOR
*
DOFETCH EQU *
LA R09,SVC7CON GET ADDRESS OF FETCH ATTR OF CON
LB R15,SVC7.OPT+1(R09) MOVE DEVICE CODE

```

```

        STB   R15,SVC7.OPT+1(R11)
        LIS   R15,0          GOOD STATUS
        STB   R15,SVC7.STA(R11)
        L     R15,SVC7.KEY(R09)  DEVICE ATTR & RECORD LENGTH
        ST    R15,SVC7.KEY(R11)
        L     R15,NAME+8      IRDR DEVICE NAME
        ST    R15,SVC7.VOL(R11)
        LM    R12,SVC7.FNM(R09)
        STM   R12,SVC7.FNM(R11)
        ITERM PCB=TERM,RDB=(R10) RETURN SVC 7 FETCH PBLK TO USER
        TEXT  PCB=EXIT       EXIT FROM TEQ HANDLER
        TITLE SVC 1 TEQ HANDLER
*****
*
*       SVC 1 INTERCEPT EXECUTOR
*
*****
INT1    EQU   *
        LR    R10,R01        SAVE RDB ADDRESS
        LR    R07,R10
        AH    R07,RDB.OFF(R10) ADDRESS OF SVC 1 PBLK
        LIS   R14,0          NO ERROR ON COMMAND FUNCTION
        LIS   R15,0          LENGTH OF TRANSFER
        LB    R13,SVC1.FC(R07) GET FUNCTION CODE
        THI   R13,SV1.CMDF   COMMAND FUNCTION?
        BNZ   ECHODONE       YES - TREAT AS A NOP
        THI   R13,SV1.WRIT   IS USER DOING A WRITE?
        BNZ   INT1.WRT       YES
*
*       A read from the internal reader will give the
*       user an illegal function status.
*
        LHI   R14,X'C000'    ILLEGAL FUNCTION ON READ
        B     ECHODONE       FINISH UP
*
*       Queue the user's command line to the internal reader
*
INT1.WRT EQU *
        L     R11,SVC1.SAD(R07) GET START ADDRESS
        L     R12,SVC1.EAD(R07) AND END ADDRESS
        IGET  RDB=(R10),SDST=BUFFER,SDEND=BUFEND,
              ADST=(R11),ADEND=(R12)
        SR    R12,R11        GET LENGTH-1 OF STRING
        LR    R15,R12
        AIS   R15,1          LENGTH OF USER COMMAND LINE
        STH   R15,COMMAND+4
        SVC   2,COMMAND      PASS COMMAND TO IREADER
        LH    R14,COMMAND+2  COMMAND QUEUED TO IREADER?
        BZ    ECHODONE       YES
        LHI   R14,X'A000'    NO - GIVE DEVICE UNAVAILABLE
        SPACE 1
ECHODONE EQU *
        STH   R14,SVC1.STA(R07) RETURN STATUS
        ST    R15,SVC1.LXF(R07) RETURN LENGTH
        THI   R13,SV1.WAIT   IS USER REQUEST A WAIT?
        BNZ   ECHOWAIT       YES - NO NEED FOR A TRAP

```

	SPACE	1	
	L	R15,RDB.PAD(R10)	GET ADDRESS OF USER SVC 1 PBLK
	OI	R15,Y'08000000'	I/O PROCEED COMPLETION PARAMETER
	ST	R15,TRAP	
	SPACE	1	
	ITERM	RDB=(R10),TRAP=TRAP,COPY=Y	TERMINATE WITH TRAP
	TEXTIT	PCB=EXIT	EXIT FROM TEQ HANDLER
	SPACE	1	
ECHOWAIT	EQU	*	
	ITERM	PCB=TERM,RDB=(R10)	TERMINATE THIS SVC 1
	TEXTIT	PCB=EXIT	EXIT FROM TEQ HANDLER
	EJECT		
	ALIGN	4	
UNPACK7	DB	2,6,0,0	PUT SVC 7 ERROR CODE IN
	DAC	SVC7ERRC	
LOG7ERRC	DB	0,7	
	DC	Z(LOG7ERRX-*)	
	DB	C'UNSUPPORTED SVC 7 FUNCTION '	
SVC7ERRC	DB	C'.. INTERCEPTED '	
LOG7ERRX	EQU	*-1	
	SPACE	2	
	ALIGN	4	
TRAP	DS	4	I/O PROCEED COMPLETION TRAP
	SPACE	2	
	ALIGN	4	
CONT	ICONT	FORM=L	CONTINUE SVC
	SPACE	2	
	ALIGN	4	
TERM	ITERM	FORM=L,COPY=Y	TERMINATE SVC
	SPACE	2	
	ALIGN	4	
EXIT	TEXTIT	FORM=L	EXIT SVC
	SPACE	2	
	ALIGN	4	
PEEK01	DB	1,19	
	DS	22	
CON	DS	4	SYSTEM CONSOLE NAME
	SPACE	2	
	ALIGN	4	
SVC7CON	DS	SVC7.	
	SPACE	2	
	ALIGN	4	
PEEK03	DB	3,19,0,0	GET INFO AN USER TASK
TID	DSF	1	USER TASK ID
TASKNAME	DSF	2	NAME OF USER TASK
CTSW	DSF	1	CURRENT TASK STATUS WORD
TOPT	DSF	1	TASK OPTIONS
WAITS	DSF	1	TASK WAITS
ACCT.P	DSF	1	USER'S PRIVATE ACCOUNT NUMBER
ACCT.G	DSF	1	USER'S GROUP ACCOUNT NUMBER
L.VOL	DSF	1	LOAD VOLUME NAME
L.FD	DSF	2	LOAD FILE NAME
L.EXT	DSF	1	LOAD EXTENSION & FILE CLASS
MONITOR	DSF	2	NAME OF MONITOR TASK

LEGACY DSF 1
 PRIO DS 1
 DS 3
 SPACE 3
 ALIGN 4
 COMMAND DB 1,14,0,0
 DCX 0
 DCX 0
 DC A(BUFFER)
 SPACE 1
 ALIGN 4
 BUFFER DS 128
 BUFEND EQU *-1
 END IRDR

NAME MTM USERS TERMINAL
 TASK PRIORITY
 (RESERVED)

QUEUE COMMAND TO IREADER
 STATUS

ADDRESS OF BUFFER

APPENDIX A
SUPPORTED VERTICAL FORMS CONTROL (VFC) CHARACTER SET

HEX	CHAR	OPERATIONS AFFECTING LINE SPACING
09	HT	Horizontal tab
0B	VT	Set vertical tabs (EVFU, no print)
20	b	1 line b/print
2B	+	No line advance
2D	-	3 lines b/print
30	0	2 lines b/print
31	1	Top of form b/print
32	2	Select VFU-2 b/print
33	3	Select VFU-3 b/print
34	4	Select VFU-4 b/print
35	5	Select VFU-5 b/print
36	6	Select VFU-6 b/print
37	7	Select VFU-7 b/print
38	8	Select VFU-8 b/print
39	9	Select VFU-9 b/print
41	A	Select VFU-10 b/print
42	B	Select VFU-11 b/print
43	C	Select VFU-12 b/print
45	E	1 line a/print
46	F	No line advance
47	G	3 lines a/print
48	H	2 lines a/print
49	I	Top of form a/print
4A	J	Select VFU-2 a/print
4B	K	Select VFU-3 a/print
4C	L	Select VFU-4 a/print
4D	M	Select VFU-5 a/print
4E	N	Select VFU-6 a/print
4F	O	Select VFU-7 a/print
50	P	Select VFU-8 a/print
51	P	Select VFU-9 a/print
52	R	Select VFU-10 a/print
53	S	Select VFU-11 a/print
54	T	Select VFU-12 a/print
60	'	No line advance
61	a	1 line b/print
62	b	2 lines b/print
63	c	3 lines b/print
64	d	4 lines b/print
65	e	5 lines b/print
66	f	6 lines b/print
67	g	7 lines b/print

HEX	CHAR	OPERATIONS AFFECTING LINE SPACING
68	h	8 lines b/print
69	i	9 lines b/print
6A	j	10 lines b/print
6B	k	11 lines b/print
6C	l	12 lines b/print
6D	m	13 lines b/print
6E	n	14 lines b/print
6F	o	15 lines b/print
70	p	16 lines b/print
71	q	17 lines b/print
72	r	18 lines b/print
73	s	19 lines b/print
74	t	20 lines b/print
75	u	21 lines b/print
76	v	22 lines b/print
77	w	23 lines b/print
78	x	24 lines b/print
79	y	25 lines b/print
7A	z	26 lines b/print
7B	{	27 lines b/print
7C	!	28 lines b/print
7D	}	29 lines b/print
7E	~	30 lines b/print
7F	DEL	31 lines b/print
80		32 lines b/print
81		33 lines b/print
82		34 lines b/print
83		35 lines b/print
84		36 lines b/print
85		37 lines b/print
86		38 lines b/print
87		39 lines b/print
88		40 lines b/print
89		41 lines b/print
8A		42 lines b/print
8B		43 lines b/print
8C		44 lines b/print
8D		45 lines b/print
8E		46 lines b/print
8F		47 lines b/print
91		48 lines b/print
92		50 lines b/print
93		51 lines b/print
94		52 lines b/print
95		53 lines b/print
96		54 lines b/print
97		55 lines b/print
98		56 lines b/print
99		57 lines b/print
9A		58 lines b/print
9B		59 lines b/print

HEX	CHAR	OPERATIONS AFFECTING LINE SPACING
9C		60 lines b/print
9D		61 lines b/print
9E		62 lines b/print
9F		63 lines b/print
A0		64 lines b/print
A1		65 lines b/print
A2		66 lines b/print
A3		67 lines b/print
A4		68 lines b/print
A5		69 lines b/print
A6		70 lines b/print
A7		71 lines b/print
A8		72 lines b/print
A9		73 lines b/print
AA		74 lines b/print
AB		75 lines b/print
AC		76 lines b/print
AD		77 lines b/print
AE		78 lines b/print
AF		79 lines b/print
B0		No line space
B1		1 line a/print
B2		2 lines a/print
B3		3 lines a/print
B4		4 lines a/print
B5		5 lines a/print
B6		6 lines a/print
B7		7 lines a/print
B8		8 lines a/print
B9		9 lines a/print
BA		10 lines a/print
BB		11 lines a/print
BC		12 lines a/print
BD		13 lines a/print
BE		14 lines a/print
BF		15 lines a/print
C0		16 lines a/print
C1		17 lines a/print
C2		18 lines a/print
C3		19 lines a/print
C4		20 lines a/print
C5		21 lines a/print
C6		22 lines a/print
C7		23 lines a/print
C8		24 lines a/print
C9		25 lines a/print
CA		26 lines a/print
CB		27 lines a/print
CC		28 lines a/print
CD		29 lines a/print
CE		30 lines a/print
CF		31 lines a/print

HEX	CHAR	OPERATIONS AFFECTING LINE SPACING
D0		32 lines a/print
D1		33 lines a/print
D2		34 lines a/print
D2		35 lines a/print
D3		36 lines a/print
D4		36 lines a/print
D5		37 lines a/print
D6		38 lines a/print
D7		39 lines a/print
D8		40 lines a/print
D9		41 lines a/print
DA		42 lines a/print
DB		43 lines a/print
DC		44 lines a/print
DD		45 lines a/print
DE		46 lines a/print
DF		47 lines a/print
E0		48 lines a/print
E1		49 lines a/print
E2		50 lines a/print
E3		51 lines a/print
E4		52 lines a/print
E5		53 lines a/print
E6		54 lines a/print
E7		55 lines a/print
E8		56 lines a/print
E9		57 lines a/print
EA		58 lines a/print
EB		59 lines a/print
EC		60 lines a/print
ED		61 lines a/print
EE		62 lines a/print
EF		63 lines a/print
F0		64 lines a/print
F1		65 lines a/print
F2		66 lines a/print
F3		67 lines a/print
F4		68 lines a/print
F5		69 lines a/print
F6		70 lines a/print
F7		71 lines a/print
F8		72 lines a/print
F9		73 lines a/print
FA		74 lines a/print
FB		75 lines a/print
FC		76 lines a/print
FD		77 lines a/print
FE		78 lines a/print
FF		79 lines a/print

INDEX

A			
Account Reporting Utility	1-14	CPU	
ACPRIVILEGE	2-9	task transferring	3-16
ADCCP	1-19	CPU-override status	1-12
Advanced data communications control procedure. See ADCCP.		CSS	1-20
APU	1-2	D	
customizing fault	3-19	D-tasks	1-7
disabled	3-3	ACPRIVILEGE option	2-9
enabled	3-3	DISC option	2-9
internal task control execution	3-17	Data Collection facility	1-14
monitor task	3-11	Data structures	2-5
monitoring task execution only, queue operating states	3-4	DCB	1-19
operating states	3-3	Device control block. See DCB.	
preempting task execution	3-10	Diagnostic tasks. See d-tasks.	
queue pointer	3-16	DISPLAY ACCOUNTING command	1-14
status signal	3-10	DISPLAY ERRORS command	1-17
task execution	3-2	DLIST instruction	4-6
task queue ordering	3-8	DT state	1-7
task transferring	3-16	Dynamic	
verifying task transfer	3-18	scheduling	1-9
Asynchronous events	1-11	time-slice	1-10
ATTN command	1-20	E	
Auxiliary processing unit. See APU.		E-tasks	1-7
B		data structures used by	2-4
Basic data communications subsystem	1-19	relocatable	2-2
C		RELOCATE option	2-2
Caller mode. See CL.		writing	2-2
CDS	1-17	writing relocatable	2-3
Central processing unit. See CPU.		Error recording subsystem	1-17
CL	4-18	Error Reporting Utility	1-17
Command processor subsystem	1-20	ES state	1-8
Command substitution system. See CSS.		ET state	1-7
Commands		Executive tasks. See e-tasks.	
ATTN	1-20	F,G	
DISPLAY ACCOUNTING	1-14	File management	
DISPLAY ERRORS	1-17	subsystem	1-16
MEMORY	1-18	support services	1-16
OPTION	2-1	Floating point	
QUEUE	2-4	instructions	1-21
SET PRIORITY	1-10	subsystem	1-21
Computational-intensive task	3-2	H	
Console monitor subsystem	1-19	HASP	1-19
Control/Diagnostic System. See CDS.		HDLC	1-19
		High level data link controller. See HDLC.	

I, J, K	
I/O	
intensive processors	1-2
subsystem	1-17
ICONT macro	4-28
ICREATE macro	4-8
	4-16
IERRTST macro	4-34
IGET macro	4-25
Input/output processor. See IOP.	
Input/output. See I/O.	
Intercept paths	4-13
creating	4-8
full and monitor control	4-11
how to remove	4-14
Interception program	
sample SVC	4-36
Internal interrupt subsystem	1-21
IOP	3-1
IPROCEED macro	4-29
IPUT macro	4-27
IREMOVE macro	4-24
IROLL macro	4-30
IS state	1-8
ISU state	1-8
ITERM macro	4-31
ITRAP macro	4-32
L	
LFC	1-10
LIB	1-18
Line frequency clock. See LFC.	
LLE	1-12
Load-leveling executive. See LLE.	
Load power fail monitor. See LPFM.	
Loader and segmentation subsystem	1-18
Loader information block. See LIB.	
Logical processing unit. See LPU.	
Logical processor mapping table. See LPMT.	
LPFM	3-4
LPMT	1-2
LPU	1-2
mapping	3-6
M	
Machine architecture	1-6
Macro libraries	2-5
MTMSTRUC.MLB	2-4
SYSSTRUC.MLB	2-4
Macros	
ICONT	4-28
ICREATE	4-16

Macros (Continued)	
IERRTST	4-34
IGET	4-25
IPROCEED	4-29
IPUT	4-27
IREMOVE	4-24
IROLL	4-30
ITERM	4-31
ITRAP	4-32
\$RDR	4-36
Memory	1-14
local	1-15
map	1-18
preventing access	
conflicts	3-20
shared	1-15
system	1-15
MEMORY command	1-18
Memory diagnostics subsystem	1-18
MTM	1-2
data structures macro	
libraries	2-8
Multi-terminal monitor. See MTM.	
Multiprocessing	3-1
support	1-2
N	
Nested calls	1-20
NS state	1-7
NSU state	1-8
O	
OPTION command	2-1
	2-4
Optional user SVC subsystems	1-21
OS/32	
basic data communications	1-19
command processor	1-20
console monitor	1-19
data structures macro	
libraries	2-5
error recording and reporting	1-17
file management	1-16
floating point	1-21
I/O management	1-17
I/O subsystem	1-17
internal interrupt	1-21
job accounting	1-14
loader and segmentation	1-18
memory diagnostics	1-18
memory management	1-14
optional user SVC	1-21
software support	1-4
subsystems	1-3
system initialization	1-20
task management	1-6
timer management	1-15

P		SVC handling	3-19
PIC	1-16	SVC interception	4-15
Precision interval clock.		error codes	4-14
See PIC.		error handling	4-10
Priority scheduling	1-9	functional summary	4-2
Privileged tasks		how it works	4-16
e-tasks	2-1	macros used with	4-13
Processor queue		occurring at the end of	4-3
assigning tasks to	3-7	task	4-1
Program status word. See		preparing a task for	4-13
PSW.		system macros for	3-5
Pseudo device	4-8	terminating intercept	3-9
generic naming for	4-9	calls	3-14
PSW	1-7	SVC13	
Q		Synchronous data logic	
QPB	3-20	control. See SDL.	
Queue		Synchronous events	1-11
priority assignments	3-9	Sysgen	1-1
timer	1-15	System	
QUEUE command	3-9	configuration	1-3
Queue parameter block. See		deadlock avoidance	3-21
QPB.		initialization subsystem	1-21
R		time-slice	1-10
RDB	4-4	System generation. See	
circular list for	4-6	sysgen.	
\$RDR macro	4-36	System pointer table. See	
Read real time clock. See		SPT.	
RRTC.		T	
Real-time support module.		Task	
See RTSM.		computational-intensive	3-2
Recipient existent mode.		controlling order of	
See RX mode.		execution	3-8
Recipient nonexistent mode.		event trap	4-8
See RN mode.		mode	1-17
Relocation/protection	1-18	rollable	1-13
Request descriptor block.		Task common. See TCOM.	
See RDB.		Task control block. See TCB.	
RN mode	4-18	Task event trap	4-8
RRTC	1-16	Task images	
RS state		impure segments	1-18
alternate	1-7	pure segments	1-18
RSA state	1-7	Task management	
RTSM	1-16	3200MPS Family of	
RX mode	4-18	Processors	1-11
S		roll function	1-1
Sample SVC interception		scheduling	1-11
programs	4-36	subsystem	1-6
SCL	1-15	system performance	1-1
SDL	1-19	Task management subsystem	
Segment control list. See		machine state control	1-7
SCL.		self-initialization	1-6
SET PRIORITY command	1-10	task dispatcher	1-6
SPT	2-2	Task priority levels	
Subsystems. See OS/32.		dispatch	1-9
SVC, mode and name parameters		maximum	1-8
valid combinations	4-20	run	1-9
		task	1-8
		Task queues	
		APU	1-7
		CPU ready	1-7
		roll-in	1-7



PUBLICATION COMMENT FORM

We try to make our publications easy to understand and free of errors. Our users are an integral source of information for improving future revisions. Please use this postage paid form to send us comments, corrections, suggestions, etc.

1. Publication number _____
2. Title of publication _____
3. Describe, providing page numbers, any technical errors you found. Attach additional sheet if necessary. _____

4. Was the publication easy to understand? If no, why not? _____

5. Were illustrations adequate? _____

6. What additions or deletions would you suggest? _____

7. Other comments: _____

From _____ Date _____

Position/Title _____

Company _____

Address _____

FOLD

FOLD




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 22 OCEANPORT, N.J.

POSTAGE WILL BE PAID BY ADDRESSEE

Concurrent Computer Corporation
2 Crescent Place
Oceanport, NJ 07757



**ATTN:
TECHNICAL SYSTEMS PUBLICATIONS DEPT.**

FOLD

FOLD

STAPLE

STAPLE