

VI
SPM-11M-PLUS

REFERENCE MANUAL

OPERATING SYSTEM AND VERSION: RSX-11M-PLUS V1.0

SOFTWARE VERSION: SPM-11M-PLUS V1.0

Software Services,
digital equipment corporation,
maynard, massachusetts

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

DEC, DECnet, DECsystem-10, DECSYSTEM-20, DECUS,
DECwriter, DIBOL, Digital logo, EduSystem, IAS, MASSBUS,
PDP, PDT, RSTS, RSX, UNIBUS, VAX, VMS, VT

are trademarks of
Digital Equipment Corporation

Copyright © 1982 by Digital Equipment Corporation
Maynard, Massachusetts

CONTENTS

PREFACE		v
CHAPTER 1	OVERVIEW OF SPM-11M-PLUS.....	1-1
1.1	Resources	1-1
1.2	Users	1-2
1.3	Resource Model	1-2
1.4	Events	1-2
1.5	Task Metrics	1-3
1.6	System Metrics	1-4
CHAPTER 2	THE COMPONENTS OF SPM-11M-PLUS	2-1
2.1	Data Collection Component	2-2
2.1.1	Hooks	2-4
2.1.2	Hook Dispatcher	2-4
2.1.3	Hook Routine	2-5
2.1.4	Hook Record	2-5
2.1.5	KW11-P Clock and Handler	2-6
2.1.6	Hook Buffer	2-6
2.1.7	Logfile	2-7
2.1.8	Data Collection Task (...COL)	2-8
2.1.9	Termination Task (...TRM)	2-11
2.1.10	COL Message Task (COLMSG)	2-11
2.2	Data Reduction Component	2-11
2.2.1	MDR Pass 1	2-13
2.2.2	DRS Pass	2-13
2.2.3	MDR Pass 2	2-14
CHAPTER 3	GUIDE TO USING SPM-11M-PLUS.....	3-1
3.1	Installation	3-1
3.1.1	SPM-11M-PLUS Kit.....	3-1
3.1.2	Prerequisites to Installing SPM-11M-PLUS.	3-2
3.1.3	Installing SPM-11M-PLUS.....	3-3
3.2	Operation	3-5
3.2.1	Generating the MPB	3-5
3.2.2	Using COL	3-13
3.2.3	Using TRM	3-23
3.2.4	Using MDR	3-24
3.2.5	Using CLF	3-43
3.2.6	Using OVL	3-50
3.2.7	Writing User Information to the Logfile .	3-57

CHAPTER 4	DATA COLLECTION AND REDUCTION PROCEDURES	4-1
4.1	Data Collection Procedure	4-1
4.1.1	Measurement Period	4-1
4.1.2	Resource Types	4-3
4.1.3	Device Classes	4-3
4.1.4	Tasks	4-4
4.1.5	System Metrics	4-4
4.2	Data Reduction Procedure	4-4
CHAPTER 5	READING THE PERFORMANCE REPORT	5-1
5.1	Performance Report Data	5-1
5.1.1	Measurement Parameters	5-1
5.1.2	System Configuration	5-2
5.1.3	System Metrics	5-2
5.1.4	Task Metrics	5-4
CHAPTER 6	RESOURCES USED BY SPM-11M-PLUS.....	6-1
6.1	Resources Used by COL	6-1
6.2	Resources Used by MDR	6-2
CHAPTER 7	SPM-11M-PLUS USAGE NOTES.....	7-1
7.1	Taskbuilding COL	7-1
7.2	COL Buffers	7-1
7.3	State of System at COL Startup Time	7-2
7.4	Taskbuilding MDR.....	7-3
7.5	Maximum VT: Unit Supported by MDR	7-4
7.6	Performing Reduction Under Batch	7-4
7.7	Necessary UFD to Run MDR	7-5
7.8	Devices for TI: Filter	7-5
7.9	Manually Executing the DRS Pass	7-5
7.10	Taskbuilding CLF	7-5
7.11	Possible Hanging of CLF	7-6
7.12	Use with Layered Products	7-6
7.13	Reporting Problems and Errors	7-6
APPENDIX A	HOOK RECORD DEFINITIONS	A-1
APPENDIX B	MPB DEFINITIONS	B-1
APPENDIX C	INFORMATION RECORD DEFINITIONS	C-1
APPENDIX D	SYSTEM METRICS RECORD DEFINITIONS	D-1
APPENDIX E	LOGFILE FORMAT	E-1
APPENDIX F	EXAMPLE PERFORMANCE REPORT	F-1
APPENDIX G	EXAMPLE SMPGEN	G-1
APPENDIX H	EXAMPLE INSTALLATION	H-1

PREFACE

MANUAL OBJECTIVES

The SPM-11M-PLUS Reference Manual is a manual describing the structure and use of the SPM-11M-PLUS software supported by DIGITAL on RSX-11M-PLUS operating systems.

INTENDED AUDIENCE

This manual is for all users of SPM-11M-PLUS.

STRUCTURE OF THIS DOCUMENT

Chapter 1 is an overview of the architecture of SPM-11M-PLUS.

Chapter 2 describes the components of SPM-11M-PLUS.

Chapter 3 describes how to use SPM-11M-PLUS.

Chapter 4 describes data collection and reduction procedures.

Chapter 5 describes how to read the performance report.

Chapter 6 delineates the resources used by SPM-11M-PLUS.

Chapter 7 is the usage notes for SPM-11M-PLUS.

Appendix A is a specification of the Hook Records.

Appendix B is a specification of the Monitor Parameter Block (MPB).

Appendix C is a specification of the Information Records.

Appendix D is a specification of the System Metrics Record.

Appendix E is a specification of the logfile format.

Appendix F is a sample performance report.

Appendix G is a sample SMPGEN.CMD procedure.

Appendix H is a sample installation.

CHAPTER 1

OVERVIEW OF

SPM-11M-PLUS

The Software Performance Monitor (SPM-11M-PLUS) is a software instrument used to measure the performance of a computer system running the RSX-11M-PLUS operating system.

SPM-11M-PLUS is a natural counterpart to RSX-11M-PLUS: while RSX-11M-PLUS manages the resources of the computer system, SPM-11M-PLUS remembers how the resources were used over a period of time, similar to the way a tape recorder remembers a set of sound patterns. Moreover, SPM-11M-PLUS has the capability of "playing back" these recordings of resource usage, so an analyst can study how and when the resources were used.

SPM-11M-PLUS is an event driven performance monitor. As each event of interest occurs, the performance monitor is called.

1.1 RESOURCES

The key resources of interest are:

- o CPU
- o Main memory
- o I/O devices
- o File system
- o Task loader


Note that a resource can be a hardware resource -- such as the CPU, main memory or an I/O device; or a software resource -- such as the file system or the task loader.

1.2 USERS

A user is the entity which uses the resources managed by the operating system. We define a user to be an RSX-11M-PLUS task. Of primary concern is the workload the task places on the system, that is, how much and for how long the resources are used. To describe a task's workload we will define a set of metrics for resource usage.

1.3 RESOURCE MODEL

Because a number of tasks may concurrently request the use of a resource, and only one task may use the resource at a time, a queue will form for the resource. A simple queued resource is depicted as follows:



The vertical bars represent the queued resource requests and the square represents the resource. Three intervals are shown on the diagram. "W" denotes the interval during which a task is waiting for the resource. "U" denotes an interval when a task is using the resource. "S" is the sum of "W" and "U", called the service time.


1.4 EVENTS

The activity of the RSX-11M-PLUS operating system can be traced by defining a set of key events and recording a set of raw information associated with each event. The events of interest are those which capture the metrics of resource usage.

For example, we may want to know how long a task waits for the disk, how many times it uses the disk and how long it uses the disk. The following events are defined to capture the I/O activity:

1. Queue I/O packet to driver
2. Driver gets I/O packet
3. Driver posts I/O completion

Applied to the queued resource model for the disk, these I/O events appear as follows:



The complete set of events defined for RSX-11M-PLUS is specified in Appendix A.

1.5 TASK METRICS

The task metrics quantify the workload that individual tasks place on the system.

The task metrics are computed for all three types of intervals: U, W and S. The performance report will contain a line of task metrics for each resource used by the task. The task metrics are:

- o Count

A simple metric for resource usage is a count of the number of times the resource was used during the life of a task. For example, if a task does four QIOs to a disk, the I/O request count is four.

- o Total Time

Another simple metric is how long the task uses a resource. For example, if a task did four QIOs and the I/O requests took 20ms, 60ms, 10ms and 50ms to complete, the total usage time would be 140ms.

- o % Resident Time

Also of importance is what percentage the total usage time is of the time the task was in memory. For example, if the previous task was in memory for one second the percentage of resident time spent using the disk was $140\text{ms}/1000\text{ms} * 100\% = 14\%$.

- o Interval Statistics

The statistical nature of the usage time is of key importance.

The metrics defined for interval statistics are:

- Minimum interval
- Mean (average) interval
- Maximum interval
- Coefficient of Variation

The coefficient of variation is the standard deviation of the intervals divided by the mean interval.

The coefficient of variation is of particular importance because it tells whether the data is clustered about the mean (small c.v.) or widely scattered about the mean (large c.v.). In fact, if the c.v. is large, the mean may tell very little about the nature of the data, since very small values and very large values may combine to produce a mean which is not representative of any real value.

1.6 SYSTEM METRICS

The task workload metrics defined above are useful for determining who is using the system resources and how much each user is consuming.

Also of importance is the performance of the system as a whole. To this purpose we define a set of system metrics.

The system metrics vary as a function of time. To see how a given system metric varies with time, histograms are computed for the metric. Each histogram describes the per cent usage of a resource for one sampling interval. The sampling interval is a settable parameter described in Chapter 3, section 3.2.1.

CPU

The CPU can be in one of three general states:

- o User
- o Kernel (system)
- o Idle

A useful metric is the per cent of time spent in each state, that is:

- o % time in user state
- o % time in kernel state
- o % time in idle state

This gives us a general picture of how effectively the CPU is being utilized (% time in user state + % time in kernel state), as well as the overhead imposed by the operating system (% time in kernel state).

Memory

On RSX-11M-PLUS the memory space is divided into partitions. The user can select one system-controlled partition for monitoring by SPM-11M-PLUS. The performance monitor will compute the per cent of the partition that was in use at the end of each sampling interval and print a histogram.

The user will most commonly select the partition "GEN" for monitoring.

I/O Devices

The user can select one I/O device to be monitored by SPM-11M-PLUS for system-wide usage. SPM-11M-PLUS will compute the percent of time the device was utilized for each sampling interval and print a histogram. If the device is mounted Files-11, SPM-11M-PLUS will compute the percent disk space in use at the end of each sampling interval and print a histogram.

Primary and Secondary Pool

The RSX-11M-PLUS system pool contains the dynamically allocated data structures used by the executive. SPM-11M-PLUS will compute the per cent of primary pool and the per cent of secondary pool that was in use at the end of each sampling interval and print a histogram for each.

Checkpoint file(s)

SPM-11M-PLUS will monitor checkpoint files allocated by the ACS command. The performance monitor will compute the per cent of all active checkpoint file space in use at the end of each sampling interval and print a histogram.

CHAPTER 2

THE COMPONENTS OF SPM-11M-PLUS

The SPM-11M-PLUS performance monitor consists of two main components: the Data Collection Component and the Data Reduction Component. The Data Collection Component records the resource usage events of the operating system to a logfile; the Data Reduction Component reads the logfile and produces a performance report suitable for analysis.

This general architecture is depicted in Figure 1.




Figure 1

By decoupling the Data Collection Component from the Data Reduction Component in time, the performance monitor can collect detailed information about a huge number of events, albeit in a raw format. The Data Reduction Component can then be used to selectively analyze and filter the information in the logfile to produce a performance report. This architecture shifts the computational aspects of the performance monitor from the Data Collection Component to the Data Reduction Component. The benefit derived is that the performance monitor uses less CPU time, thus minimizing its interference with the CPU. The cost is more I/O activity, which in most systems can be directed to a secondary mass storage device not in the path of the major stream of I/O activity.

2.1 DATA COLLECTION COMPONENT

The Data Collection Component consists of both software and hardware distributed in various places within an RSX-11M-PLUS system.

The basic pieces of the Data Collection Component are:

- o Hooks
- o Hook Dispatcher
- o Hook Routines
- o Hook Records
- o KW11-P Clock and Handler
- o Hook Buffers
- o Logfile
- o Data Collection Task (...COL)
- o Termination Task (...TRM)
- o COL Message Task (COLMSG)

The relationship between these pieces is shown in Figure 2.

DATA COLLECTION COMPONENT




Figure 2

Heuristic Procedure for Data Collection

At this point a heuristic procedure for data collection will be given, although the terms used in the procedure have not been precisely defined.

When a resource usage event occurs in RSX-11M-PLUS, the corresponding hook calls the hook dispatcher, which uses the hook ID to call the appropriate hook routine. The hook routine collects information relevant to the hook and assembles it into a hook record, which resides in a hook buffer. When a hook buffer is filled with hook records, it is written to the logfile. The logfile is closed at the end of the measurement interval, and subsequently reopened by the Data Reduction Component. It reads the hook records and uses them to derive the resource usage metrics and print a performance report.

The details of these pieces will now be developed more fully.

2.1.1 Hooks

A hook is simply an implementation technique for capturing a resource usage event. For SPM-11M-PLUS, a hook is implemented as a subroutine call inserted at the point in the operating system where the event occurs. The subroutine call passes control and one argument, the hook ID, to the hook dispatcher. The form of an SPM-11M-PLUS hook is as follows:

```
MOV    @#PS,-(SP)
MOV    #H$XXXX,-(SP)
CALL   @$HKDSP
```

where XXXX is mnemonic identifying the particular hook. A complete specification of all hooks is given in Appendix A.

2.1.2 Hook Dispatcher

Hooks are resident within the operating system; their method of insertion is addressed in Chapter 3, section 3.1.3. However, most of the code and database for the Data Collection Component is not resident within the RSX-11M-PLUS executive.

The bulk of the data collection code is resident in the data collection task called COL (taskname is ...COL). The chief benefit derived from this design is that only a small amount of executive space is used by the performance monitor.

The purpose of the hook dispatcher is to map and call the hook routines which are resident in the COL task. Note that, though the hook routines are resident in the task partition for COL, they are executed in the context of a directive or interrupt service routine.

A one word dispatch vector labelled \$HKDSP is added to the module DRTBL. This word points to an entry point in the hook dispatcher. It is the port through which all hooks pass control from the RSX-11M-PLUS executive to the Data Collection Component of the performance monitor.

2.1.3 Hook Routine

A hook routine copies information which is generally scattered about the executive (typically in SYSCM or SYSTB) into a hook record which is resident in a static common or dynamic region. A timestamp is also recorded in the hook record.

There are two general types of hooks: task and system. The task hooks collect data relevant to resources used by tasks; the system hooks collect data which is system-wide (e.g., % CPU time in user, kernel and idle).

2.1.4 Hook Record

The information collected at each hook is formatted into a hook record. The hook record provides an identical representation of a hook's information content both in memory and in the logfile.

Each hook record contains both generic and hook-specific data. Generic data applies to all hooks. Generic data consists of:

- o Hook ID
- o Taskname
- o Timestamp
- o Task's TI (UCB address)

Hook specific data applies to only one or a small subset of the hooks. For example, the H\$CDRP (call directive processor) hook records the directive ID code as hook-specific data.

Appendix A specifies the information content of the hook records.

2.1.5 KW11-P Clock and Handler

For precise performance measurements, SPM-11M-PLUS needs a high speed clock running independent of the RSX-11M-PLUS system clock. It is therefore a requirement that a KW11-P clock be installed on a machine running SPM-11M-PLUS. This KW11-P clock must be in addition to the RSX-11M-PLUS system clock, be it a KW11-L or KW11-P.

The KW11-P is initialized by the COL task to run at 100 kHz (10 microsecond ticks) and to interrupt every 65536 ticks (approximately every 0.65 second).

The KW11-P clock time is maintained as a 2 word (32 bit) value which, at a rate of 100 kHz, will overflow in a little over 11 1/2 hours.

To preclude the possibility of overflow of the KW11-P clock time, a maximum measurement interval of 11 1/2 hours has been established. The task COL enforces this by setting an 11 1/2 hour mark-time when activated. If the timer expires, COL will stop data collection and exit.

The clock interrupt handler updates the high order clock time located in SYSCM at \$PCLKH. The low order clock time is not stored in memory; rather, each hook routine copies the current contents of the KW11-P clock's count register to the hook record directly to get the low order time word.

2.1.6 Hook Buffer

Hook records are stored in buffers which are resident in either a static common or a dynamic region. When a hook buffer is full, the hook routine sets an event flag which signals the COL task to write the buffer to the logfile.

The hook buffers are linked together into a circular list. There may be from 2 to 8 buffers in a list, and all the buffers must be of the same length. A buffer may be from 256 words to 4096 words long.

If the buffers are resident in a static common, the maximum size common which can be used is 4096 words. This limitation is imposed because the COL task uses one APR to map the static common.

If the buffers are resident in a dynamic region, the region can be from 512 words to 32K words long.

2.1.7 Logfile

The logfile contains four types of records:

- o Monitor Parameter Block (MPB)
- o Information Records
- o Hook Records
- o System Metric Records

See Appendix E for a specification of the logfile.

Monitor Parameter Block

The Monitor Parameter Block (MPB) is used by the COL task to initialize its measurement parameters. The MPB is formatted as an Information Record with record type I\$MPB. A detailed description of the MPB is given in Appendix B.

The MPB is generated by a command file called SMPGEN.CMD. See Chapter 3, section 3.2.1 for details.

Information Records

The Information Records are used to supplement the Hook Records in the logfile with system context useful to the Data Reduction Component.

When COL starts up it reads the MPB from the MPB file and writes it to the logfile. Following the MPB, COL writes a stream of Information Records which contain a number of RSX-11M-PLUS data structures defining the initial system state. These data structures are:

- o Task Control Block (TCB)
- o Partition Control Block (PCB)
- o Device Control Block (DCB)
- o Unit Control Block (UCB)

Each data structure is copied to a separate Information Record.

When COL begins a measurement interval (starts data collection), it records the RSX-11M-PLUS system time as well as the SPM-11M-PLUS time (KW11-P time) to an Information Record (I\$TMR).

When COL ends a measurement interval (stops data collection) it copies the RSX-11M-PLUS system time, SPM-11M-PLUS time, CPU time accrued in kernel state and CPU time accrued in idle state to an Information Record (I\$STM).

A complete specification of all Information Records is given in Appendix C.

Hook Records

As each hook point is encountered in the executive, a Hook Record is written to the logfile. See section 2.1.4 for details.

System Metrics Record

At the end of each sample interval, a System Metrics Record is written to the logfile. This record contains information which describes the system-wide usage of the selected resources.

The System Metrics Record is specified in Appendix D.

Logfile Media

The logfile can be resident on disk as a Files-11 file or on magtape as a sequence of tape records.

If the logfile is resident on disk, it consists of a sequence of virtual blocks (virtual block=256 words). The records within these virtual blocks are defined according to the record definitions for the MPB, Information Records, Hook Records and System Metric Records. Note that the logfile does not have any other record structure (such as FCS or RMS type records).

If the logfile is on a magtape, the MPB, Information Records and Hook Records are imbedded in a stream of tape records. Each tape record is equal in length to a hook buffer.

2.1.8 Data Collection Task (...COL)

This section gives a functional description of the COL task. A guide to using COL is provided in Chapter 3, section 3.2.2.

As noted before, this task contains both code which executes in executive context (as a directive or interrupt process) and code which executes in the context of a privileged task. The code which executes at executive level is the hook routines.

2.1.8.1 COL Executive Level Code -

The hook routines were made a part of the COL task for a number of reasons:

- o They must be resident in some partition outside of executive space.
- o They can easily be loaded as part of the COL task.
- o They implement routines which are common to hooks and to the COL task.

2.1.8.2 COL Task Level Code -

Most of the COL task consists of code which is executed at task level. The major functions of the COL task are:

- o Initialization
- o Starting data collection
- o Writing to logfile
- o Stopping data collection
- o Cleanup

2.1.8.2.1 Initialization -

When COL begins execution, there is very little SPM-11M-PLUS code elsewhere in the system. All that exists are the hooks in the executive which all call a dummy return subroutine in TDSCH. It is the job of COL to initialize all code and databases required for the SPM-11M-PLUS performance monitor. The specific sequence of initialization functions performed by COL are:

- o Read MPB file and store MPB in COL buffer.
- o Create dynamic region (if necessary) and set up ring of hook buffers. Mark dynamic region or static common as fixed and non-shufflable.
- o Write MPB and Information Records (RSX-11M-PLUS data structures) to logfile.
- o Mark COL task as fixed in memory and nonshufflable.
- o Initialize hooks specified in MPB.

- o Initialize KW11-P clock and turn it on.
- o Initialize measurement parameters as specified in MPB.

2.1.8.2.2 Starting Data Collection -

When data collection is disabled, the hook dispatch vector \$HKDSP points to the entry point \$HKEXT in the hook dispatcher, which simply returns to the instruction in the executive following the hook.

Just prior to starting data collection, COL writes a "start data collection" Information Record (I\$TMR) to the hook buffer.

To enable data collection the COL task simply points the hook dispatch vector \$HKDSP to the entry point \$HKENT in the hook dispatcher.

2.1.8.2.3 Writing Hook Records to Logfile -

The steady state function of the COL task is to write Hook Records to the logfile. When a hook routine fills a hook buffer, it sets an event flag for COL which wakes COL from its wait state. COL then writes the hook buffer to disk or magtape (waiting for completion of the write), and loops to see if another hook buffer is ready to write. When no more hook buffers are outstanding, COL returns to the wait state.

If the hook buffers fill up too rapidly for COL to write them to the logfile, COL will set the overflow status bit. If the overflow status is set when a measurement interval expires, COL will print a status message on the console device (CO:). In general, the data collected during a measurement interval should not be used for data reduction if the overflow condition occurred.

2.1.8.2.4 Stopping Data Collection -

To stop data collection, COL merely points the hook dispatch vector \$HKDSP to the dummy routine at \$HKEXT in the hook dispatcher. It then writes a stop data collection Information Record (I\$STM) to the hook buffer. Any outstanding records in the hook buffer will be written to the logfile.

2.1.8.2.5 Cleanup -

Before COL exits, it must release any system resources it has allocated and deactivate any hardware it has turned on. The following actions are thus performed:

- o Turn off KW11-P clock.
- o Unfix COL task.

2.1.9 Termination Task (...TRM)

The purpose of the termination task (...TRM) is to signal the COL task to stop data collection and exit. TRM provides a way to manually terminate the COL task, as opposed to automatic shutoff by timeout or buffers (see Chapter 3, section 3.2.1).

2.1.10 COL Message Task (COLMSG)

The COL message task contains the ASCII text for the error messages issued by the COL task. If the COLMSG task is installed and a COL error occurs, the full error message text will be printed. If a COL error occurs and the COLMSG task is not installed, the COL task will print an error number which can be correlated with the correct error message in Chapter 3, section 3.2.2.

2.2 DATA REDUCTION COMPONENT

The main purpose of the Data Reduction Component is to read the raw data from the logfile (hook records) and produce a performance report detailing the resources used by the various tasks in the system.

The Data Reduction Component accomplishes this in a multi-step process involving two tasks:

- o Monitor Data Reduction (MDR) Task
- o DRS Task

There are three major steps in the data reduction process:

1. Map Hook Records to Task Workload Records and generate system-wide metrics section of performance report (MDR Pass 1).
2. Sort Task Workload Records (DRS).
3. Summarize Task Workload Records and generate task metrics section of performance report (MDR Pass 2).

The Task Workload Records are intermediate data structures which contain the data collected for each task and each resource used by the task.

This general procedure is depicted in Figure 3.

DATA REDUCTION COMPONENT




Figure 3

From the figure we see there are four databases which are processed by the two tasks:

- o TRACE.LOG (logfile containing Hook Records).
- o RESOURCE.TMP (intermediate file containing unsorted Task Workload Records).
- o RESOURCE.SRT (intermediate file containing sorted Task Workload Records).
- o REDUCE.RPT (final performance report).

2.2.1 MDR Pass 1

MDR reads hook records from the logfile (default name is TRACE.LOG) and maps the information to the Task Workload Records in the intermediate file RESOURCE.TMP.

2.2.2 DRS Pass

DRS sorts the Task Workload Records using the following three keys:

- o Generic Taskname
- o Task's TI UCB address
- o Task start time.

NOTE: By "generic taskname" we mean the installed taskname of a multiuser task. For example, the generic taskname for the task PIPT5 is ...PIP. If the task is not multiuser, generic taskname is simply the run time taskname.

The three keys listed above are concatenated into one "big" key, with the generic taskname forming the highest order part of the key, the Task's TI UCB address forming the middle order part of the key, and the Task start time forming the low order part of the key. The DRS task sorts all the Task Workload Records using this one big key. With this technique the DRS task only needs to do one sort operation on the file RESOURCE.TMP.

When DRS is done sorting the Task Workload Records, the records will be arranged as follows:

- o All records with the same generic taskname will be clustered together.
- o Within a generic taskname cluster, all records with the same TI will be clustered together.
- o All records within a taskname-TI cluster will be sequenced by ascending task start time.

DRS writes the sorted Task Workload Records to the file RESOURCE.SRT.

2.2.3 MDR Pass 2

MDR reads the Task Workload Records from RESOURCE.SRT and generates a performance report. A description of the performance report is given in Chapter 5.

CHAPTER 3

GUIDE TO USING

SPM-11M-PLUS

3.1 INSTALLATION

3.1.1 SPM-11M-PLUS Kit

The SPM-11M-PLUS kit is distributed on 800 BPI magtape. It contains all the files necessary to install and use SPM-11M-PLUS. The files are in DOS tape format. When copied to disk, the files will occupy approximately 900 disk blocks.

The command file PMONCPY.CMD can be used to copy the kit files from tape to tape, tape to disk, disk to tape or disk to disk. All files are contained in the directory [12,2]. The contents of the kit are:

<u>Component</u>	<u>File</u>
Kit copy command file	PMONCPY.CMD
SPM-11M-PLUS mini-reference guide	SPM.DOC
MPB GEN command file	SMPGEN.CMD
Macro definition prefix file	HKMAC.MAC
Executive source correction command file	PMONSLP.CMD
Executive source correction files	*.COR
F11ACP patch command file	PMONPAT.CMD
F11ACP patch source	DISPAT.PAT
COL/TRM build command file	PMONBLD.CMD
COL overlay descriptor file	COLBLD.ODL
COL/TRM object library	COL.OLB
MDR/CLF object library (EIS version)	MDREIS.OLB
MDR/CLF object library (FPP version)	MDRFPP.OLB
MDR/CLF/OVL build command file	MDRTKB.CMD
Data reduction sort task	DRS.TSK
DRS map file	DRS.MAP
DRS sort command file	DRS.CMD
CLF task	CLF.TSK
CLF map file	CLF.MAP
OVL object library	OVL.OLB
OVL task	OVL.TSK
OVL map file	OVL.MAP

3.1.2 Prerequisites to Installing SPM-11M-PLUS

COL Prerequisites

The following are the prerequisites for running the COL task:

- o Supervisor-mode library support
- o The supervisor-mode library FCSFSL must be installed
- o Memory Management (PLAS) Directives

MDR Prerequisites

All Prerequisites for MDR are included in a minimum functionality RSX-11M-PLUS system. To run MDR at least one system-wide checkpoint file must be allocated using the ACS command. For MDR to automatically spawn the task DRS, parent-offspring support must be included in the RSX-11M-PLUS executive.

3.1.3 Installing SPM-11M-PLUS

To install SPM-11M-PLUS, you must start with:

- o An RSX-11M-PLUS V1.0 distribution kit
- o An RSX-11M-PLUS autopatch kit at the current level
- o An SPM-11M-PLUS distribution kit

To install SPM-11M-PLUS you must perform an RSX-11M-PLUS SYSGEN. A SYSGEN is required because source code for SPM-11M-PLUS is inserted into the RSX-11M-PLUS sources.

SPM-11M-PLUS source code is added to the following RSX-11M-PLUS source modules:

<u>SPM-11M-PLUS Function</u>	<u>RSX-11M-PLUS Modules</u>
SPM-11M-PLUS Macros and Definitions.	RSXMC0,RSXMC
Hooks	DRDSP,DRQIO,DREIF,DREXP,DRGIN,DRTBL,DRCRV,REQSB,SYSXT,IOSUB,LOADR
Additional Stack Space	LOWCR
Common Database	SYSCM
Hook Dispatcher	TDSCH
KW11-P Clock Handler	TDSCH

For a complete example installation procedure, see Appendix H.

The SPM-11M-PLUS installation procedure consists of the following steps:

- o Copy RSX-11M-PLUS V1.0 kit to baseline disk.
- o Perform autopatch to current level.

Follow the instructions in the Autopatch notes for applying Autopatch.

Note: to run the Autopatch procedure, you must first invoke @[200,200] SYSGEN and answer "yes" to the following question:

```
>* DO YOU WANT TO APPLY AUTOPATCH [Y/N]:Y
```

When the Autopatch procedure finishes, control will be returned to SYSGEN.CMD. You must then answer "NO" to the following three questions:

```
>* DO YOU WANT TO DO SYSGEN PART 1 [Y/N] N
>* DO YOU WANT TO DO SYSGEN PART 2 [Y/N] N
>* DO YOU WANT TO DO SYSGEN PART 3 [Y/N] N
>ASN =
>@<EOF>
```

When the SYSGEN.CMD procedure exits, continue to the next step.

- o Assign SY: to the baseline disk.

- o Choose a directory to which you want the SPM-11M-PLUS kit files copied. This will be your SPM-11M-PLUS kit directory.

- o Assuming the SPM-11M-PLUS directory is [g,m], and SPM-11M-PLUS kit is physically mounted on device ddn:.

>SET /UIC=[g,m]

>UFD SY:[g,m]

>MOU ddn:/FOR

>FLX =ddn:[12,2]PMONCPY.CMD

>@PMONCPY

.
.
.

APPLY CORRECTIONS TO EXECUTIVE SOURCES? Y

.
.
.

APPLY PATCH TO F11ACP DISPATCHER MODULE (DISPAT)? Y

@<EOF>

>DMO ddn:

- o Invoke SYSGEN.CMD a second time

Invoke @[200,200]SYSGEN.CMD and answer "NO" to the following question:

>* DO YOU WANT TO APPLY AUTOPATCH [Y/N] N

Then perform a SYSGEN by answering "YES" to the following questions:

>* DO YOU WANT TO DO SYSGEN PART 1 [Y/N] Y

>* DO YOU WANT TO DO SYSGEN PART 2 [Y/N] Y

>* DO YOU WANT TO DO SYSGEN PART 3 [Y/N] Y

(Note: SYSGEN PART 3 is optional)

NOTE: DO NOT PATCH F11ACP if you want to use the FCPR03 version.

If the patch to F11ACP fails (checksum error), SPM-11M-PLUS will only calculate service times for F11ACP. If the patch succeeds, SPM-11M-PLUS will generate usage, wait and service times.

If you do not elect to correct the executive sources or patch F11ACP via PMONCPY.CMD you can do so later via the command files:

- 1) @PMONSLP for executive source corrections
- 2) @PMONPAT for F11ACP patch

On systems with an FP11 the file MDREIS.OLB is not needed and may be deleted. On systems without an FPP, the file MDRFPP.OLB is not needed and may be deleted. The file may be deleted before starting the SYSGEN.

- o When the SYSGEN completes, perform the following steps:

```
>SET /UIC=[g,m]
```

```
>@PMONBLD
```

NOTE: The command file will pause if you desire to edit the build files for CLF, OVL, or MDR. This will allow modifications to defaults for these programs. Instructions are included in the build files.

You can build the MDR, OVL, or CLF tasks separately from PMONBLD.CMD by invoking the command file MDRTKB.CMD.

- o You are now ready to run @SMPGEN to generate an MPB and run COL. This procedure is described in the next section.

3.2 OPERATION

3.2.1 Generating the Monitor Parameter Block (MPB)

The MPB is the structure which contains the measurement parameters.

To create an MPB you invoke the command file SMPGEN.CMD.

See Appendix G for an example SMPGEN.

The procedure is as follows:

```
>SET /UIC=[g,m]
```

```
>@SMPGEN
```

where [g,m] is your SPM-11M-PLUS kit directory.

The MPB is created in the following way: SMPGEN creates a source file (default filespec is TRACE.PAR), assembles and builds the code in the file to produce a binary MPB file (default filespec is TRACE.MPB). It is this file which is read by the data collection task COL to initialize the measurement parameters.

In addition to the files mentioned above, SMPGEN also will output, if requested, a saved answer file (default filespec is TRACE.CMD). This file can be input to another run of SMPGEN and will initialize the symbol table to the values of the variables saved in the file.

SMPGEN is driven from a menu of parameter categories. To display the menu, type an <ESC>. The menu appears as:

```
;
; LEGAL PARAMETER CATEGORIES ARE :
;   HOOK          TASK          MODE          LABEL
;   BUFFER        DEVICE        SYSTEM        DONE
;
```

The parameter categories can be specified in any order. If a parameter category is not specified, the parameters will receive default values from either:

- o The saved answer file, if one is specified.
- o The table of default parameters.

Table of Default Parameters

HOOKS	- all hooks enabled, excepting the LOAD OVERLAY, CALL DIRECTIVE PROCESSOR and USER INFORMATION hooks.
BUFFERS	- DYNAMIC region named TRCPAR created in partition GEN, 2 buffers with 8 blocks per buffer.
TASKS	- All tasks enabled for data collection.
DEVICES	- All devices enabled for data collection.
MODE	- MANUAL mode.
SYSTEM	- System Metrics enabled with 60 second sampling interval.
LABEL	- No label.

To select a parameter category, enter the name of the parameter category desired as shown in the following example:

```
ENTER PARAMETER CATEGORY? HOOK
```

Once a parameter category is selected, the detailed parameters are specified. Each parameter category has a menu of responses which are possible for that parameter category. To get the menu, type <ESC>. When you are finished entering parameters, type "DONE".

A detailed description of each parameter category follows.

o Label Field

To select the label parameter category, type: LABEL

The following query will be displayed:

```
ENTER MPB LABEL?
```

The response may be any string from 0 to 63 characters long. This user-definable string is stored in the MPB and will typically contain notes and comments applicable to a particular measurement run. It is printed as part of the heading on each page of the performance report.

o Hooks

To select the hooks parameter category, type: HOOK

The hooks determine which resource usage events are recorded to the logfile.

You do not select each hook individually. Rather, you select one or more groups of hooks. Each group of hooks contains the subset of hooks which will generate a consistent set of data for measuring the usage of a particular resource. Hooks may be chosen from the following groups:

- CPU USAGE HOOKS

This hook group measures task CPU usage.

It is required for all metrics for task CPU usage.

- MEMORY USAGE HOOKS

This hook group measures task memory usage.

It is required for all metrics for task memory usage and % resident time.

- I/O DEVICE USAGE HOOKS

This hook group measures task usage of I/O devices and the file system. The file system is regarded as any ACP which manages a device. Examples are F11ACP and MTAACP.

This hook group is required for all metrics relative to task usage of I/O devices or the file system.

- LOAD OVERLAY HOOK

This hook records the relative block number of a task's overlay to the logfile.

- CALL DIRECTIVE PROCESSOR HOOK

This hook records the directive ID of a directive to the logfile.

- USER INFORMATION HOOK

This hook is invoked when a task issues a directive of the form: GIN\$ 7, P1,P2,P3,P4,P5,P6,P7,P8. The eight parameters P1 - P8 are written to the logfile.

o Buffers

To select the buffers parameter category, type: BUFFER.

The hook buffers may be resident in either a static common or a dynamic region. The buffer type is specified by:

ENTER BUFFER TYPE?

The response may be either STATIC indicating a static common or DYNAMIC indicating a dynamic region.

NOTE: If "STATIC" is selected, the COL task must have been linked to a static common in the installation phase (see Chapter 3, section 3.1.3).

If "DYNAMIC" is selected, COL must not be linked to the static common.

If the buffer type is DYNAMIC, the name of the system-controlled partition in which to create the dynamic region is specified by:

ENTER NAME OF BUFFER MAIN (SYS) PARTITION?

The name of the buffer common/region is specified by:

ENTER NAME OF BUFFER COMMON/REGION?

The partition will be divided up into a number of hook buffers, determined by:

ENTER NUMBER OF BUFFERS [2.:8.:2.]?

Note that from 2 to 8 buffers may be specified.

Finally, the length of each buffer is specified by:

ENTER BUFFER SIZE IN 512 B BLOCKS?

A buffer can be from 1 to 16 blocks long; in other words, from 256 words to 4096 words long.

If a static common is used for buffering, the total space required for buffers must fit into the common. If a dynamic region is used, COL will create a region large enough to contain the buffers.

o Tasks

SPM-11M-PLUS can be set up to collect data for specific tasks. For example, there may be some tasks whose performance metrics are not desired at all. Typically, system "service" tasks, such as F11ACP, ...MCR and QMG... fall into this category.

To select the task parameter category, type: TASK.

The following query will appear:

DO YOU WANT TO TRACE ALL TASKS?

If your answer is N, you will be given two choices for selecting the tasks for data collection:

- Specify the tasks not to trace.
- Specify the tasks to trace.

The choice is made from the query:

DO YOU WANT TO SPECIFY THE TASKS NOT TO TRACE?

If the answer is Y, you are prompted with:

ENTER NAME OF TASK NOT TO TRACE?

If the answer is N, you are prompted with:

ENTER NAME OF TASK TO TRACE?

You then type, one per line, the names of the tasks desired or not desired. Each taskname is from 1 to 6 characters long. To terminate the list of tasks, type <CR>.

Up to 64 tasknames can be specified.

o Devices

To get the devices parameter category, type: DEVICE.

SPM-11M-PLUS can be set up to collect data for specific classes of devices. If you want to collect data only for selected classes of devices, answer N to the following query:

DO YOU WANT TO COLLECT DATA FOR ALL DEVICES?

A device in RSX-11M-PLUS may fall into one of four classes depending on the bit setting of the "Device Characteristics Word" U.CW1 in the UCB.

- Directory Device (DV.DIR)
- Sequential Device (DV.SQD)
- Terminal Device (DV.TTY)
- All Other Devices

To enable data collection for a device class, answer Y to the query for that device class. The queries are:

TRACE DISK (DIRECTORY DEVICE) I/O ACTIVITY?

TRACE TAPE (SEQUENTIAL DEVICE) I/O ACTIVITY?

TRACE TERMINAL I/O ACTIVITY?

TRACE ALL OTHER DEVICES I/O ACTIVITY?

o Mode

To select the mode parameter category, type: MODE.

The performance monitor can be set up to collect data in three different modes. You select the mode by answering:

ENTER MODE?

The possible modes are: MANUAL, AUTOSTOP and REPEAT. Each mode determines the manner in which a measurement interval begins and ends. A description of each mode follows:

- Manual

In manual mode, a measurement interval begins when the COL task is invoked and ends when the TRM (termination) task is run.

- Autostop

In autostop mode, a measurement interval begins when the COL task is invoked and ends when either of the following two events occur:

- . A mark-time expires.
- . A predetermined number of buffers have been filled.

You determine the type of autostop by responding to:

ENTER SAMPLING TYPE?

with either TIME or BUFFER.

If TIME is specified, you then have to specify the timeout interval:

ENTER SAMPLING INTERVAL TIME UNIT [D:MIN]

ENTER SAMPLING INTERVAL IN units ?

where "units" is the time unit specified in the previous query. Legal units are SEC, MIN or HR.

NOTE: Do not confuse this sampling interval with the system metrics sampling interval described in the next section.

If BUFFER is specified as the interval type, you must specify the number of buffers per sample:

ENTER NUMBER OF BUFFERS PER SAMPLE?

Note that each buffer is from 1 to 16 disk blocks (256 words) long.

- Repeat

Repeat mode is autostop mode repeated a number of times. The sampling interval is specified as for autostop mode. In addition, you must specify:

- . The number of intervals
- . The length of time between intervals.

You specify these parameters by answering the following queries:

ENTER MAX. NUMBER OF SAMPLES?

ENTER INTERVAL BETWEEN SAMPLES TIME UNIT [D:MIN]?

ENTER INTERVAL BETWEEN SAMPLES IN units ?

where "units" is the time unit specified in the previous query.

o System

To select the system-metrics parameter group, type "SYSTEM".

If you want the system metrics, answer Y to the next query:

DO YOU WANT THE SYSTEM METRICS? [Y/N]:

If you answered Y, you must then specify the sampling interval by:

ENTER SYSTEM METRICS SAMPLING INTERVAL (SEC) ?

o Done

When you are finished entering parameters, type "DONE". The SMPGEN procedure will then create:

- A monitor parameter block (MPB) file (filename.MPB).
- A parameter source file (filename.PAR).
- A saved answer file (filename.CMD, if specified).

Note: SMPGEN will prompt you for the parameter source file and saved answer file, and you must enter a valid file specification.

3.2.2 Using COL

The COL task is used to initialize the performance monitor, start data collection, collect data, stop data collection and clean up the performance monitor before exiting.

Installing COL

The COL task can be installed in any partition. If possible, create a separate 3.5K word partition for COL. If installed in a system-controlled partition, COL will automatically fix itself and mark itself as non-shufflable. It is therefore desirable to install and run COL when there is little or no system activity in the partition in which COL is installed so as to load COL in the lower end of the partition and reduce fragmentation of the partition.

COL is taskbuilt with a default priority of 230. Note that COL has a higher priority than F11ACP and ...MCR, but a lower priority than ...LDR. It is highly recommended that the priority of COL not be altered to a value lower than the default value of 230, otherwise excessive overflows may occur which will invalidate the measurement runs.

If full ASCII message text output is desired, install the COLMSG task.

If you are using a static common for COL buffers, be sure to install it before installing the COL task.

Install COL and COLMSG by:

```
>INS TRCPAR          ;ONLY FOR STATIC COMMON
>INS COL
>INS COLMSG
```

Running COL

To run COL, type a command line of the form:

```
>COL logfile/sw=MPBfile/sw
```

The default logfile specification is SY:[cur uic]TRACE.LOG.
The default MPBfile specification is SY:[cur uic]TRACE.MPB.

Stopping COL

There are two ways to terminate COL:

- o RUN TRM
- o Specify AUTOSTOP or REPEAT mode in SMPGEN.

RUN TRM will terminate COL regardless of the mode in which COL is running: MANUAL, AUTOSTOP or REPEAT.

COL will also terminate if any of the following abnormal conditions are detected:

- o Powerfail recovery

COL will issue a SPRA\$ directive. Upon powerfail recovery, COL will simply exit.

- o Write errors

If COL detects write errors, it will print a message and exit. Write errors will generally corrupt the logfile with invalid data.

- o Automatic 11 1/2 hour termination.

COL Switches

Switch defaults --

Output: SY:TRACE.LOG/DENS:800/EX:100./IF:100./WI:10.

Input: SY:TRACE.MPB/BF:2:8/DL:0/DV:LBO/PR:GEN

Output Switches	Description
/AP	Append records to end of magtape. COL will space over all existing files on the magtape and append the records for this run starting at the current end-of-volume point.
/DENS:density	Set tape density to 800 or 1600 bpi. Densities cannot be mixed on a single tape volume.
/EX:n	Specifies disk file extent size in 256W blocks.
/IF:n	Specifies disk file initial size in 256W blocks.
/WI:n	Specifies number of retrieval pointers in window block.
Input Switches	
/BF:m:n	Number of buffers = m ; size of buffer = n 256W blocks. This overrides the buffering description in the MPB.
/DL:n	Number of MINUTES COL delays before it starts data collection.
/DV:ddnn	Device selected for system metrics.
/PR:parnam	Partition selected for system metrics.
/TI:nn	Collect data only for tasks with TI of TTnn.

Example 1

>COL DB2:[3,300]CPUTEST/IF:1000/EX:200=DB1:[5,200]PARM1/DV:DB2:

logfile is DB2:[3,300]CPUTEST.LOG with initial size 1000 and extent size 200.

MPB file is DB1:[5,200]PARM1.MPB with system metrics specified for DB2.

Example 2

>COL MM1:/DENS:1600=DB1:[5,200]SAMPLE.MPB

logfile is on MM1: in 1600 BPI; MPB file is DB1:[5,200]SAMPLE.MPB.

Example 3

>COL =

logfile is SY:[current uic]TRACE.LOG; MPB file is SY:[current uic]TRACE.MPB.

NOTE: You cannot run COL with a null command line, i.e., >COL <CR>. The COL task expects to see explicit filespecs or a bare "=" indicating defaults on both input and output. Moreover, you cannot invoke COL with a command of the form:

>RUN COL

since no filespecs can be specified with this form. One final thought in this regard: only one copy of COL can be active in the system at a time. If multiple invocations of COL are attempted at different terminals, all but the first will be rejected with an error message "MONITORING ACTIVITY IN PROGRESS".

Once COL is activated, it will run in the mode specified in the MPB but always for a maximum of 11 1/2 hours. This time limit is imposed because the KW11-P clock time will overflow in a little over 11 1/2 hours.

COL Messages

COL outputs two types of messages: error messages and status messages.

Full ASCII error message text will be printed only if the COLMSG task is installed. If the COLMSG task is not installed, COL will print a message of the form "COL ERROR XXXXXX", where XXXXXX is the error number. The error messages are listed by error number in the following list. Error messages are issued if an error condition occurs while COL is starting up and initializing the performance monitor. Error messages are output to the terminal of issue (TI:) for COL.

Status messages are issued after COL has run its initialization phase. There are status messages for both normal and error conditions. Status messages are issued to the console terminal (CO:). The COLMSG task does not have to be installed to get the status messages; as they are printed directly by the COL task.

COL Error Messages

000000 COLMSG -- XXXXXX (DSW) GET COMMAND LINE ERROR

Explanation: A "Get MCR Command Line" directive failure occurred.

User Action: Reenter COL command line using correct form.

NOTE: Illegal command forms are:

1. >COL <CR>
2. >RUN \$COL <CR>

Legal command line forms are:

1. >COL filespec/sw = filespec/sw
2. >COL filespec/sw = /sw
3. >COL =filespec/sw
4. >COL =

000001 COLMSG -- COMMAND SYNTAX ERROR

Explanation: Entered incorrect syntax for command line.

User Action: Reenter command line with correct syntax.

000002 COLMSG -- MPB FILENAME SYNTAX ERROR

Explanation: Entered incorrect syntax for MPB filename.

User Action: Reenter MPB filename using correct syntax.

000003 COLMSG -- XXXXXX (FCS) MPB FILE ACCESS FAILURE

Explanation: COL failed to open MPB file or received a read error when attempting to read the MPB file. The code XXXXXX (FCS) indicates the FCS error code returned to COL.

User Action: Check that specified MPB file exists, the volume is mounted (if disk), or that the file protection does not prevent COL from reading the MPB.

000004 COLMSG -- XXXXXX (DSW) ASSIGN LUN FAILURE

Explanation: COL failed to assign a LUN to the logfile device.

User Action: Check that the device specified for logfile is in the system configuration.

000005 COLMSG -- ILLEGAL LOGFILE DEVICE TYPE

Explanation: The device specified for the logfile is not of the correct type. Valid logfile device types are: Directory structured (disk, DECTAPE), or Sequential (magtape).

User Action: Specify a device which is either directory structured or sequential.

000006 COLMSG -- LOG FILENAME SYNTAX ERROR

Explanation: Entered incorrect syntax for logfile.

User Action: Reenter logfile specification using correct syntax.

000007 COLMSG -- XXXXXX (FCS) LOGFILE ACCESS FAILURE

Explanation: COL failed to open the logfile. The code XXXXXX (FCS) is the FCS error returned to COL.

User Action: Check that the logfile volume is mounted native and the necessary directory exists (if disk), or the volume is mounted foreign and online (if magtape).

000010 COLMSG -- ILLEGAL BUFFER SPECIFIED

Explanation: The parameters specified by the /BF:n:m switch were not valid. Either the number of buffers (n) was out of range (<2 or >8) or the buffer size was out of range (>16).

User Action: Respecify a valid buffer count or size via the /BF switch.

000011 COLMSG -- XXXXXX (DSW) BUFFER PARTITION NOT IN SYSTEM

Explanation: The partition specified in the MPB is not in the system.

User Action: If the buffer partition is a static common, check that the static common partition exists and that the common is installed.

000012 COLMSG -- BUFFERS WILL NOT FIT INTO PARTITION

Explanation: The number of buffer blocks specified in the MPB (NUMBER OF BUFFERS * BUFFER SIZE IN 512 B BLOCKS) is greater than the number of 512 B blocks in the physical buffer partition.

User Action: Do a SMPGEN and specify the NUMBER OF BUFFERS and BUFFER SIZE IN 512 B BLOCKS such that the product of the two is less than the number of 512 B blocks in the physical buffer partition.

000013 COLMSG -- XXXXXX (DSW) CANNOT CREATE BUFFER PARTITION OR ADDRESS WINDOW

Explanation: COL received a DSW error when attempting to create the dynamic buffer partition specified in the MPB, or when attempting to create an address window to map the dynamic partition. The error XXXXXX (DSW) is the DSW error from the Create Region directive or from the Create Address Window directive.

User Action: Check that enough space exists in the partition to contain a dynamic region with the size specified in the MPB.

000014 COLMSG -- MONITORING ACTIVITY IN PROGRESS

Explanation: A multiple invocation of the COL task was attempted. Only one copy of COL can be active at a time.

User Action: Do not attempt to run more than one copy of COL at a time.

000015 COLMSG -- INVALID MPB FILE

Explanation: The MPB filespec specified in the COL command line is not the name of a valid MPB file.

User Action: Reenter a filespec for a valid MPB file.

000016 COLMSG -- INFORMATION RECORD ALLOCATION FAILURE

Explanation: Insufficient space exists in the buffer partition to contain all the Information Records.

User Action: Do a SMPGEN and specify enough blocks (NUMBER OF BUFFERS * BUFFER SIZE IN 512 B BLOCKS) to contain all the Information Records. The space required is roughly that needed to contain the largest of the following three system lists: 1)TCBs, 2)PCBs, 3) DCB-UCBs.

000017 COLMSG -- KW11-P VECTOR OR CSR UNAVAILABLE

Explanation: Either the KW11-P vector is in use, or the KW11-P CSR is non-existent on the UNIBUS.

User Action: Check that the correct KW11-P vector and CSR were specified in the procedure @PMONBLD.CMD, that the vector is not used by another RSX-11M-PLUS driver or task, and that the CSR is addressable on the UNIBUS.

000020 COLMSG -- CANNOT USE SYSTEM'S KW11-P AS SPM-11M-PLUS CLOCK

Explanation: The RSX-11M-PLUS system clock is a KW11-P and an attempt was made to use it as the SPM-11M-PLUS (performance monitor) clock.

User Action: Check that a separate KW11-P was installed for use by the SPM-11M-PLUS software, and that the vector and CSR specified in the procedure @PMONBLD.CMD was for this clock.

000021 COLMSG -- POOL NODE ALLOCATION FAILURE

Explanation: COL could not allocate a RSX-11M-PLUS pool node.

User Action: Re-run COL when sufficient pool is available.

000022 COLMSG -- BUFFER OVERFLOW

Explanation: COL failed to allocate space for a hook record because all buffers were full.

User Action: Re-run COL, using a larger buffer size specified either via SMPGEN.CMD or the /BF switch.

000023 COLMSG -- PARTITION SPECIFIED NOT IN SYSTEM

Explanation: The partition specified via the /PR switch is not configured in the system.

User Action: Respecify a partition configured in the system.

000024 COLMSG -- ILLEGAL DELAY VALUE SPECIFIED

Explanation: An illegal delay value has been specified via the /DL switch.

User Action: Re-specify a delay value in the correct range (0 to 28800).

COL Status Messages

COL -- PERFORMANCE MONITOR INITIALIZED

Explanation: COL has successfully initialized and is now ready to collect data.

User Action: None

COL -- PERFORMANCE MONITOR EXITING

Explanation: COL has stopped data collection and has cleaned up its databases. This status message is issued by COL immediately prior to exiting.

User Action: None

COL -- CANNOT MAP OUTPUT BUFFER

Explanation: COL received an error when attempting to map to the output buffer using the Map directive.

User Action: Check that the dynamic region used as the output buffer was successfully created.

COL -- LOGFILE WRITE ERROR

Explanation: COL received an FCS or QIO error when attempting to write to the logfile.

User Action: Check the integrity of the logfile. If repeated errors occur to the same logfile, rerun COL to create a new disk file or mount a fresh tape. Ensure the logfile volume is mounted via the MOU command.

COL -- LOGFILE CLOSE ERROR

Explanation: COL received an error when attempting to close the logfile. COL will exit shortly after printing this message with the logfile improperly closed.

User Action: Because the logfile was improperly closed, it may be difficult to read the logfile and produce a performance report.

COL -- BUFFER OVERFLOW

Explanation: Hook records were written to the hook buffer faster than COL could write them to the logfile.

User Action: In general, the data collected during a measurement interval which overflowed is not usable by the MDR task to produce a performance report. An attempt should be made to rerun the measurement to obtain data without overflow. Prior to attempting another measurement run, a SMPGEN should be run with larger buffers specified than for the measurement run which overflowed.

3.2.3 Using TRM

The TRM task is used to manually signal the COL task to stop data collection, cleanup and exit. TRM will cause COL to exit, regardless of the mode in which COL was running: manual, autostop or repeat.

TRM is run by the command >RUN TRM.

3.2.4 Using MDR

MDR is used to read the logfile and produce a performance report. MDR is used in conjunction with DRS (sorting) task to produce the performance report. The general procedure is as follows:

- o >INS DRS
>RUN MDR
MDR>LISTFILE/sw=LOGFILE/sw

MDR will read the command line and execute pass 1. MDR will then attempt to spawn the task DRS. If MDR detects an error in the spawn, or MDR has been built to prevent automatic spawn, the DRS pass must be executed manually. In this case, MDR will print the following message and suspend:

MDR -- YOU MUST NOW SORT THE INTERMEDIATE FILE BY
USING DRS.TSK

MDR --
>DRS RESOURCE.SRT=RESOURCE.TMP/FO:V:156/KE:1.12/PR:T

MDR -- AFTER THE SORT HAS COMPLETED, RESUME THIS TASK
(...MDR)

NOTE: Rather than type the command line to DRS as shown above, you can simply invoke the command file @DRS.

- o >@DRS sorts the intermediate file.

- o >RES

MDR executes pass 2 and prints the performance report.

- o MDR> Z

Certain MDR command switches restrict the input data which will be processed. If a logfile contains too much activity for MDR to reduce the data without running out of space, these switches (/AF, /BF, /-MD, /NT, /NU, /TS, /UC) may be used to reduce a logfile by successive runs of MDR, analyzing only parts of the workload in each run.

When using the switches /NT, /NU, /TS, or /UC, MDR will prompt for the task or device names after the command line is processed. The tasknames are specified as 1 to 6 RAD50 characters per entry. Only generic tasknames are used for /NT or /TS. The TI device names are entered as "DDnn:", with one entry per line.

MDR Switches

Switch defaults --

output: SY:REDUCE.RPT/SP

input: SY:TRACE.LOG/DM:2/LI:2/RW/SI:613/MD/SK:0/DE:800/-IC/-SC

Output Switches		Description
/SP	/-SP	SPool output file
Input Switches		
/AF:PChigh:PClow		Only process records from AftEr specified time
/BF:PChigh:PClow		Only process records from BeFore specified time
/DE:nnnn		Select input magtape DEnsity, 800 or 1600 (ignored unless /RW is also used)
/DM:nn		DuMp specified record types, nn is total of desired types: 01 - buffer headers 02 - unknown record types 04 - system configuration info records 10 - task event hook records 20 - system metric records 40 - user information hook records (H\$USER) 100 - load overlay hook records (H\$OVLY) 200 - call directive processor hook records (H\$CDRP)
/IC	/-IC	Replace the % resource column of the report with a count of incomplete intervals.
/LI:nn		Detail level of task workload metrics: 00 - no report produced 01 - system-wide summary only 02 - system and task summaries 03 - system, task, and TI: summaries 04 - all summaries and each complete task run 05 - all summaries and all task invocations

(continued on next page)

MDR Switches (continued)

Input Switches	Description
/MD /-MD	Report on tasks with missing data
/NT *	Do Not process hooks for listed Tasks
/NU *	Do Not process hooks for listed devices as TI:
/SC /-SC	Produce System Configuration section of report
/SI:nn	Selects the contents of Sampling Interval report by setting nn to the total for the desired items: 000 - no report produced 001 - data collection start/stop times 002 - % CPU usage 004 - % memory partition usage 010 - % device time busy 020 - % device space used 040 - % POOL used 100 - % CKPT space used 200 - clock time of metrics 400 - trailing blank line after the metrics
/SK:nnn	SKip nnn files when opening tape input file
/RW /-RW	ReWind tape input file before opening
/TS *	Process hook records only for listed TaSks
/UC *	Process hook records only for listed devices as TI:

- * - the switches /NT and /TS are mutually exclusive.
- the switches /NU and /UC are mutually exclusive.

MDR Messages

The MDR messages are issued to the TI device under which MDR runs. This section discusses the messages which MDR can issue.

```
MDR -- PLEASE ENTER TASKNAMES ONE AT A TIME FOLLOWING EACH
MDR -- WITH A <CR>. TYPE JUST A <CR> AFTER THE LAST ONE IS
      ENTERED
MDR -- TSKNAM>
```

Explanation: This is the prompt for the generic task-names for filtering by the /NT or /TS switch. The last line of this message is repeated for each line of input.

User Action: Enter generic task name(s).

```
MDR -- PLEASE ENTER TI DEVICES (DDNN:) ONE AT A TIME FOLLOWING
      EACH
MDR -- WITH A <CR>. TYPE JUST A <CR> AFTER THE LAST ONE IS
      ENTERED
MDR -- TI DEVICE>
```

Explanation: This is the prompt for device names to be used by the /UC or /NU filter. The last line of this message is repeated for each line of input.

User Action: Enter full device name(s) as "DDnn:".

```
MDR -- YOU MUST NOW SORT THE INTERMEDIATE FILE BY USING DRS.TSK
MDR -- (E.G.) MCR>DRS
      RESOURCE.SRT=RESOURCE.TMP/FO:V:156/KE:1.12/PR:T
MDR -- AFTER THE SORT HAS COMPLETED, RESUME THIS TASK (TSKNAM)
```

Explanation: This message indicates that the first pass of MDR has completed. Manual intervention is now required to run the DRS task with the given command. The task name under which MDR is running is given in parenthesis in the last line of the message.

User Action: Run the DRS task (@DRS.CMD), and when it has completed, resume the MDR task.

```
MDR -- DSW:##### F.ERR:##### IOSB:##### ##### PC:#####
      SP:#####
MDR -- STK:##### ##### ...
```

Explanation: This is a status message which gives detail of the internal state of MDR at the time when an error is detected. This message is always followed by

one of the messages discussed below, which describes the actual error condition.

User Action: See action specified for the message issued immediately following this status message.

MDR -- AN INCOMPLETE RESOURCE UTILIZATION REPORT WILL BE PRODUCED

Explanation: An error occurred which prevented the completion of pass II of MDR. Although a report will be produced, not all task-workload data has been summarized. A previous error message is issued giving the cause of this condition.

User Action: As specified for previous message.

MDR -- ATTEMPT TO PRINT NULL TWB

Explanation: An internal inconsistency occurred in the pass II processing of MDR.

User Action: * * * SPR !

MDR -- ATTEMPTING TO DEALLOCATE TWB NOT IN TSKHD

Explanation: An internal inconsistency occurred in the pass I processing of MDR.

User Action: * * * SPR !

MDR -- CHECKPOINT USAGE > 100%

Explanation: A Checkpoint metric indicates more than 100% usage of the checkpoint space. The metric is not reported.

User Action: Ensure logfile is not corrupted.

MDR -- CLOCK STOPPED OR HOOK RECORD OUT OF TIME SEQUENCE.

Explanation: The records in the logfile do not follow a strict time sequence. This condition invalidates the data in the logfile.

User Action: Ensure that the KW11-P clock is functioning. Re-collect the data.

MDR -- COMMAND READ ERROR

Explanation: An I/O error occurred while attempting to read the MDR command line.

User Action: Reenter the MDR command.

MDR -- COMMAND SYNTAX ERROR

Explanation: The command as entered has improper syntax.

User Action: Reenter the corrected MDR command.

MDR -- COMPLETION ERROR IN INPUT READ

Explanation: An I/O error occurred while reading the logfile. This condition will be treated as end-of-file and MDR will attempt to continue. This may be symptomatic of a corrupted logfile.

User Action: Check IOSB for error cause. Rerun MDR using another logfile if necessary.

MDR -- CONFLICTING SWITCHES /TS+/NT OR /UC+/NU

Explanation: The switches /TS and /NT were both specified, or the switches /UC and /NU were both specified. Only one of each pair of switches may be specified.

User Action: Reenter the command line without both of the conflicting switches.

MDR -- DEVICE USAGE > 100%

Explanation: MDR processed a system metric record which indicates greater than 100 percent of device utilization. No graphic output will be produced for the suspect record.

User Action: * * * SPR !

MDR -- DEVICE'S ACP NOT FOUND

Explanation: The system information records indicated that a device was mounted, yet no TCB for the ACP task

was found. The device may have been dismounted concurrently with COL's scan of the system configuration. A generic ACP name will be used for the device's ACP functions.

User Action: None.

MDR -- DYNAMIC SPACE ALLOCATION FAILURE

Explanation: MDR was unable to sufficiently extend itself to provide space for initialization of pass I. This indicates a severe shortage of necessary dynamic space.

User Action: Ensure MDR is checkpointable. Ensure the partition where MDR runs has enough space for MDR to extend. Use SET /MAXEXT to allow sufficient extension. Use of smaller buffers by COL will lower the MDR space requirement.

MDR -- ERROR ATTACHING, POSITIONING OR SELECTING DENSITY OF MAGTAPE

Explanation: An I/O error occurred when attempting to open the logfile on magtape. The drive may have been mounted to MTAACP which is incorrect since MDR does QIOs directly to the tape device driver. Ensure the magtape is mounted foreign.

User Action: Correct the condition and rerun MDR.

MDR -- ERROR CLOSING INPUT FILE

Explanation: An I/O error occurred when closing the input logfile.

User Action: None.

MDR -- ERROR CLOSING REPORT FILE

Explanation: An I/O error occurred when closing the report file. This file will not be properly closed or spooled, or MDR attempted to spool the output file but the task PRT... was not installed.

User Action: Correct the condition and manually close the file, or rerun MDR.

MDR -- ERROR CLOSING RESOURCE.TMP FILE

Explanation: An I/O error occurred at the end of pass I when closing the RESOURCE.TMP file. The DRS task will not be able to read this file until it is properly closed.

User Action: Correct the condition and manually close the file, or rerun MDR.

MDR -- ERROR DELETING RESOURCE.TMP FILE

Explanation: An I/O error occurred when MDR deleted its temporary file. The file is not properly closed, or deleted.

User Action: Manually unlock and delete this file.

MDR -- ERROR IN GTSK\$\$

Explanation: The directive failed preventing MDR from obtaining its taskname. The message requesting the sort to be run will not include MDR's taskname.

User Action: None.

MDR -- ERROR IN INPUT READ QIO

Explanation: An I/O error occurred while reading the logfile. This condition will be treated as end-of-file and MDR will attempt to continue. This may be symptomatic of a corrupted logfile.

User Action: Check DSW for error cause. Rerun MDR using another logfile if necessary.

MDR -- ERROR OPENING RESOURCE.SRT FILE

Explanation: An error was reported by FCS when MDR attempted to open the RESOURCE.SRT file for pass II. Pass II will not be able to run.

User Action: Correct the condition indicated by F.ERR and rerun MDR.

MDR -- ERROR OPENING RESOURCE.TMP FILE

Explanation: An error was reported by FCS when MDR attempted to open the RESOURCE.TMP file for write in pass I. Pass I will not be able to run.

User Action: Correct the condition indicated by F.ERR and rerun MDR.

MDR -- ERROR OUTPUTTING PRINT LINE

Explanation: An error was reported by FCS when MDR attempted to output a record to the report file. MDR will attempt to continue to produce the report. If this error is not caused by a transient condition the message will be repeated, and MDR will have to be aborted.

User Action: Correct the condition and rerun MDR.

MDR -- ERROR READING MPB

Explanation: An I/O error occurred when attempting to read the first block of the logfile.

User Action: Correct the condition and rerun MDR.

MDR -- ERROR READING RESOURCE.SRT FILE

Explanation: An I/O error occurred in pass II when attempting to read from the RESOURCE.SRT file. This will be treated as an end-of-file condition.

User Action: Correct the condition and rerun MDR.

MDR -- ERROR REOPENING RESOURCE.TMP FILE

Explanation: An I/O error occurred when attempting to re-open the RESOURCE.TMP file so it could be deleted. The RESOURCE.TMP file will not be deleted.

User Action: Delete the file manually.

MDR -- ERROR SKIPPING OVER MPB BUFFER

Explanation: An I/O error occurred when repositioning the tape logfile so the MPB could be reread once the buffer size is known.

User Action: Correct the condition and rerun MDR.

MDR -- ERROR WRITING RESOURCE.TMP FILE

Explanation: An I/O error was reported by FCS for a write operation during pass I.

User Action: Correct the condition and rerun MDR.

MDR -- EXPECTED RIB NOT FOUND

Explanation: An inconsistency occurred in MDR's internal dynamic data structures.

User Action: See Chapter 7, Section 7.3.

MDR -- EXPECTED RWB PREDECESSOR NOT FOUND IN OVERLAP

Explanation: An inconsistency occurred in MDR's internal dynamic data structures.

User Action: * * * SPR !

MDR -- EXPECTED RWB PREDECESSOR NOT FOUND IN MERGE

Explanation: An inconsistency occurred in MDR's internal dynamic data structures.

User Action: * * * SPR !

MDR -- IDLE TIME > 100%

Explanation: MDR processed a system metric record which indicates greater than 100 percent of CPU idle time. No graphic output will be produced for the suspect record.

User Action: See Chapter 7, Section 7.5.

MDR -- IMPROPER FORMAT OR NO SUCH DEVICE IN SYSTEM. PLEASE RETYPE.

Explanation: An improper response was entered to the prompt for TI filter devices.

User Action: Enter a corrected response.

MDR -- INPUT FILESPEC OR SWITCH SYNTAX ERROR

Explanation: An improper command line was used with MDR.

User Action: Reenter the command.

MDR -- INVALID DENSITY SELECTED, 800 OR 1600 ONLY

Explanation: A magtape density other than 800 or 1600 was specified.

User Action: Reenter the command with the correct density.

MDR -- INVALID DEVICE SPACE METRIC RECORDED

Explanation: A system metric record indicates more than 100% of device space used. The device space metric will not be reported. This metric is only available for devices mounted to F11ACP.

User Action: Ensure logfile not corrupted.

MDR -- INVALID INPUT DEVICE TYPE

Explanation: An attempt was made to reduce a logfile which is not on either a disk or tape volume. MDR can only use disk or tape for the input file.

User Action: Do not use MDR to read logfiles not residing on disk or tape.

MDR -- INVALID INPUT FILESPEC

Explanation: A syntax error occurred in the input filespec part of the MDR command.

User Action: Correct and reenter the command.

MDR -- INVALID LISTING LEVEL

Explanation: An invalid listing level was specified in the /LI switch. Acceptable values range from 0 to 5.

User Action: Correct and reenter the command.

MDR -- INVALID MPB FOUND BY ALC

Explanation: The logfile was found to contain improper data. This may be a result of using MDR to read a logfile produced by an incompatible COL version. This may also be the result of using MDR to read a file which was not produced by COL, or was corrupted.

User Action: Rerun MDR using a valid logfile.

MDR -- INVALID MPB FOUND BY INF

Explanation: The MPB did not pass validity testing when it was reread once the logfile buffer size was known. The logfile may have been corrupted while MDR was processing it.

User Action: * * * SPR !

MDR -- INVALID POOL METRIC RECORDED

Explanation: A Pool metric indicates more than 100% DSR usage. The metric is not reported.

User Action: Ensure logfile is not corrupted. Do not change DSR size while COL is running.

MDR -- INVALID RAD50 CHARACTER IN YOUR INPUT. PLEASE RETYPE.

Explanation: The response to the taskname prompt for filtered names was not a valid RSX-11M-PLUS taskname. All letters should be in uppercase.

User Action: Correct and retype the response.

MDR -- INVALID TIME VALUES; /AF MUST BE EARLIER THAN /BF

Explanation: An attempt to use time filter values which overlapped causes this message. This condition will not allow any records to pass the filter for processing.

User Action: Correct and reenter the MDR command.

MDR -- MEMORY USAGE > 100%

Explanation: MDR processed a system metric record which indicates greater than 100 percent of partition utilization. No graphic output will be produced for the suspect record.

User Action: * * * SPR !

MDR -- MISSING INFO FOR DEVICE METRIC

Explanation: The Info record describing the system metric device was not found. Device system metrics will not be reported.

User Action: Do not load or unload device drivers while COL is active. Ensure logfile is not corrupted.

MDR -- MISSING I\$SYS FOR POOL METRICS

Explanation: The DSR size was not found in the logfile. The Pool system metric is suppressed.

User Action: Ensure logfile is not corrupted.

MDR -- MISSING "=" AND INPUT FILESPEC

Explanation: The input command contains no "=" and no input filespec. Both input and output filespecs must be supplied or defaulted.

User Action: Enter a corrected MDR command including the "=".

MDR -- MORE THAN ONE INPUT FILESPEC NOT ALLOWED

Explanation: The MDR command contains more than one input filespec.

User Action: Enter a corrected MDR command with only one input filespec.

MDR -- MORE THAN ONE OUTPUT FILESPEC NOT ALLOWED

Explanation: The MDR command contains more than one output filespec.

User Action: Enter a corrected MDR command with only one output filespec.

MDR -- MPB FOUND AT LOCATION OTHER THAN FIRST RECORD OF FILE

Explanation: An MPB was read in the wrong position in the logfile. This is indicative of a corrupted logfile. MDR will bypass this MPB and attempt to continue to process the corrupted logfile. Reports produced from this logfile are suspect.

User Action: Use a different logfile for analysis.

MDR -- MPB NOT FOUND - NOT TRACE FILE

Explanation: The MDR input file did not begin with a MPB. Either the file was corrupted or it is not a logfile.

User Action: Use a different logfile for analysis.

MDR -- NO FOLLOWING PRIMARY RWB FOUND

Explanation: An internal inconsistency was detected in pass I.

User Action: * * * SPR !

MDR -- NO FPP OR BUILT /-FP

Explanation: MDR is unable to initialize for FPP ASTs. Either MDR was taskbuilt with the /-FP switch, or the processor where MDR is running has no FPP. This message cannot occur if the EIS version of MDR is used.

User Action: Rebuild MDR using the distributed .CMD file. Only run the FPP version of MDR on processors with an FP-11.

MDR -- OUTPUT FILESPEC OR SWITCH SYNTAX ERROR

Explanation: A syntax error was detected in the output filespec or an invalid output switch was specified.

User Action: Correct and reenter the MDR command.

MDR -- OVERFLOW BEFORE MPB - BAD TRACE FILE

Explanation: The logfile indicates buffer overflow occurred before the MPB was written. Either the logfile was corrupted, or the file was not produced by COL.

User Action: Use another file for MDR input.

MDR -- OVERFLOW IN FIRST PASS SUMMARIZATION

Explanation: A floating overflow occurred in pass I. Incorrect data will be reported for the particular task and summaries which this data represents.

User Action: Use the resulting report with caution. Use filtering to limit the amount of data processed to prevent this condition.

MDR -- OVERFLOW IN SECOND PASS SUMMARIZATION

Explanation: A floating overflow occurred in pass II. Incorrect data will be reported for the particular task and summaries which this data represents.

User Action: Use the resulting report with caution. Use filtering to limit the amount of data processed to prevent this condition.

MDR -- PARTITION RIB NOT FOUND FOR S\$MEM

Explanation: A system metric record was read for a partition which was not described by the system configuration information records. The graphic output of memory utilization will be suppressed.

User Action: None.

MDR -- PRINT LINE TOO LONG

Explanation: MDR generated a report file record longer than 132 bytes. The line is truncated by the logic which writes to the file.

User Action: * * * SPR !

MDR -- RIB LIST LINKAGE ERROR

Explanation: An inconsistency was detected in MDR's internal representation of the system configuration.

User Action: * * * SPR !

MDR -- RWB WITHOUT TWB FOUND IN RESOURCE.TMP

Explanation: A record in the RESOURCE.SRT file was encountered in an illogical position. The record is disgarded.

User Action: * * * SPR !

MDR -- SUMMARIZING FROM NULL TWB

Explanation: An inconsistency was detected in the pass II data structures.

User Action: * * * SPR !

MDR -- SUMMARIZING INTO NULL TWB

Explanation: An inconsistency was detected in the pass II data structures.

User Action: * * * SPR !

MDR -- SVTK FAILURE

Explanation: The RSX executive rejected MDR's attempt to specify an SST vector.

User Action: Correct the condition indicated by the DSW and rerun MDR.

MDR -- SYSTEM + IDLE TIME > 100%

Explanation: A system metric record was read which indicates the total of kernel and idle time is greater than the elapsed time. The graphic output for the record is suppressed.

User Action: See Chapter 7, Section 7.5.

MDR -- SYSTEM TIME > 100%

Explanation: A system metric record was read which indicates more kernel time than elapsed time. The graphic output for the record is suppressed.

User Action: See Chapter 7, Section 7.5.

MDR -- TI: READ ERROR

Explanation: An I/O error occurred when reading filter values from the user. The read with prompt which is used to obtain these inputs is not valid from a VT: (e.g. BATCH)

User Action: Rerun MDR.

MDR -- TRAP ###

Explanation: A floating point or conversion error occurred. The octal number in this message indicates the precise error cause.

User Action: * * * SPR !

MDR -- TRUNCATED INPUT FILE RECORD

Explanation: The end of a buffer was reached before all the records in that buffer were processed. This indicates a corrupted logfile.

User Action: Use a different logfile for analysis.

MDR -- TWB/RWB/RIB DYNAMIC SPACE ALLOCATION FAILURE

Explanation: MDR was unable to get needed virtual space for a dynamic structure in pass I. However sufficient dynamic storage was obtained to begin pass I.

User Action: Ensure MDR is checkpointable. Ensure the partition where MDR runs has enough space for MDR to extend. Use SET /MAXEXT to allow sufficient extension. Use of smaller buffers by COL will lower the MDR space requirement. Use of filtering of MDR input will also lower this space requirement. Use overlaid MDR.

MDR -- UNABLE TO ALLOCATE RESOURCE.SRT BUFFER

Explanation: MDR was unable to get needed virtual space for a dynamic structure in pass II.

User Action: Ensure MDR is checkpointable. Ensure the partition where MDR runs has enough space for MDR to extend. Use SET /MAXEXT to allow sufficient extension. Use of filtering of MDR input will lower this space requirement. Use overlaid MDR.

MDR -- UNABLE TO ALLOCATE RIB

Explanation: MDR was unable to sufficiently extend itself to provide space for initialization of pass I. This indicates a severe shortage of necessary dynamic space.

User Action: Ensure MDR is checkpointable. Ensure the partition where MDR runs has enough space for MDR to extend. Use SET /MAXEXT to allow sufficient extension. Use of smaller buffers by COL will lower the MDR space requirement.

MDR -- UNABLE TO OPEN INPUT FILE

Explanation: An I/O error occurred when MDR attempted to open the logfile.

User Action: Correct the condition and rerun MDR.

MDR -- UNABLE TO OPEN OUTPUT FILE

Explanation: An I/O error occurred when MDR attempted to open the report file.

User Action: Correct the condition indicated by F.ERR and rerun MDR.

MDR -- UNKNOWN INFO TYPE RECORD FOUND

Explanation: A corrupted logfile, or a file not produced by COL is being used as MDR input.

User Action: Use a different file for MDR input.

MDR -- UNKNOWN RECORD TYPE

Explanation: A corrupted logfile, or a file not produced by COL is being used as MDR input.

User Action: Use a different file for MDR input.

MDR -- VT UNIT NUMBER EXCEEDS PREDEFINED MAXVT

Explanation: An I\$CVT, I\$ELV, or I\$UCB record was encountered for a virtual terminal with a unit number which exceeds the predefined maximum MAXVT.

User Action: Rebuild MDR with a larger value defined for this maximum. See sections 7.4 and 7.5.

MDR -- VT0: DCB OR UCB NOT PROCESSED BY INF

Explanation: An I\$CVT or I\$ELV record was encountered without processing the I\$DCB and I\$UCB records for VT0: (prototype VT database). This indicates the database for the VT: devices was loaded after invoking the COL task.

User Action: See section 7.3.

MDR -- WILDCARD NOT ALLOWED IN FILESPEC

Explanation: The MDR command contains a wildcard.

User Action: Reenter the command without any wildcards.

MDR -- WRONG FOLLOWING PRIMARY RWB FOUND

Explanation: An inconsistency was detected in MDR's pass I dynamic data.

User Action: * * * SPR !

3.2.5 Using CLF

CLF is used to copy logfiles produced by COL. The three modes of CLF operation are:

- o Buffer Transfer Mode,

which copies a logfile from either disk or tape, to disk or tape, maintaining the logfile format which MDR uses for input.

- o FCS Record Transfer Mode,

which copies a logfile from either disk or tape, producing an output file with variable length FCS records.

- o Read-only Mode,

which reads the logfile, computes overflow and buffer counts, but does not create a new file.

The general procedure for using CLF is:

```
>RUN CLF
CLF> [OUTFILE/sw =] LOGFILE/sw
```

Note that the output filespec and "=" are optional. If these are not specified, CLF will operate in read-only mode, producing no output file.

When FCS mode is specified (by the /FCS switch), the records have the format specified in appendices A through D. The first word of each record is the record ID which uniquely identifies the record type. If overflows occur within the data stream, a special "overflow record" is placed in the output file at the point where records were lost due to overflow. The format of this (length = 4. bytes) record is:

600	0	600(8) = overflow record ID
###	2	count of records lost by overflow at this point

When the /FCS switch is used, the default output filespec is changed to TRACE.FCS.

After CLF has completed reading the input file, the run statistics will be reported to CLF's TI:.

CLF Switches

Switch defaults --

output: SY:TRACE.LOG/-AP/DE:800/-FCS

input: SY:TRACE.LOG/DE:800/RW/SK:0

Output Switches	Description
/AP /-AP	Append file to the end of magtape. CLF will space over all existing files on the magtape and write the new file starting at the current end-of-volume point.
/DE:nnnn	Selects output magtape DEnsity, 800 or 1600
/FCS[:nn] /-FCS	Selects FCS sequential records in the output file, and changes the default output filespec to TRACE.FCS. If ":nn" is not specified, all record types will be copied. If ":nn" is specified, only selected record types will be written. "nn" is the total for desired record types: 01 - overflow records 04 - system configuration records 10 - task event hook records 20 - system metric records 40 - user inserted records 100 - load disk overlay records 200 - call directive processor records
Input Switches	
/DE:nnnn	Selects input magtape DEnsity, 800 or 1600 (ignored unless /RW is also used)
/RW /-RW	ReWind tape input file before opening
/SK:nnn	SKip nnn files when opening tape input file

CLF Messages

The CLF messages are issued to the TI device under which CLF is run. This section discusses the messages which CLF can issue.

CLF -- ##### BUFFERS READ
CLF -- ##### RECORDS READ
CLF -- ##### BUFFERS WITH OVERFLOW
CLF -- ##### RECORDS LOST BY OVERFLOW

Explanation: This statistics message is output at the completion of each run of CLF, indicating the statistics about the input file.

User Action: None.

CLF -- ##### FCS RECORDS WRITTEN

Explanation: This statistics message is output when CLF produces an FCS output file.

User Action: None.

CLF -- DSW:##### F.ERR:##### IOSB:##### ##### PC:#####
SP:#####
CLF -- STK:##### ##### ...

Explanation: This is a status message which gives detail of the internal state of CLF at the time when an error is detected. This message is always followed by one of the messages discussed below, which describes the actual error condition.

User Action: See action specified for the message issued immediately after this status message.

CLF -- COMMAND READ ERROR

Explanation: An I/O error occurred while reading the command line.

User Action: Reenter the command.

CLF -- COMMAND SYNTAX ERROR

Explanation: The command has improper syntax.

User Action: Reenter the corrected command.

CLF -- COMPLETION ERROR IN INPUT READ

Explanation: an I/O error occurred while reading the input file. This condition, which may be symptomatic of a corrupted logfile, will be treated as end-of-file.

User Action: Check F.ERR and IOSB for error cause, or use another logfile.

CLF -- ERROR ASSIGNING INPUT DEVICE

Explanation: CLF was unable to assign a LUN to the input file.

User Action: Check DSW and F.ERR for cause. Ensure the filespec syntax is correct.

CLF -- ERROR ASSIGNING OUTPUT DEVICE

Explanation: CLF was unable to assign a LUN to the output file.

User Action: Check DSW and F.ERR for cause. Ensure the filespec syntax is correct.

CLF -- ERROR ATTACHING, POSITIONING OR SELECTING DENSITY OF MAGTAPE

Explanation: An I/O error occurred when attempting to open a file on magtape.

User Action: Ensure the drive is mounted foreign if logfile format is used for that file. Ensure the output drive is selectable, and not write locked. Correct the error indicated by DSW or IOSB.

CLF -- ERROR CLOSING INPUT FILE

Explanation; An I/O error occurred when closing the input file.

User Action: None.

CLF -- ERROR CLOSING OUTPUT FILE

Explanation: An I/O error occurred when truncating or closing the output file.

User Action: Correct the condition and rerun CLF, or manually unlock the output file.

CLF -- ERROR IN INPUT READ QIO

Explanation: An I/O error occurred while reading the input file. This condition, which may be symptomatic of a corrupted logfile, will be treated as end-of-file.

User Action: Check DSW for cause. Rerun using another logfile if necessary.

CLF -- ERROR READING MPB

Explanation: An I/O error occurred when reading the first block of the logfile.

User Action: Correct the condition and rerun.

CLF -- ERROR WRITING TO OUTPUT FILE

Explanation: An I/O error occurred when writing to the output file.

User Action: Correct the condition and rerun.

CLF -- INPUT FILESPEC OR SWITCH SYNTAX ERROR

Explanation: A syntax error was detected in the input filespec.

User Action: Reenter the corrected command.

CLF -- INVALID DENSITY SELECTED, 800 OR 1600 ONLY

Explanation: A magtape density other than 800 or 1600 was specified.

User Action: Reenter the command with a correct density.

CLF -- INVALID INPUT DEVICE TYPE

Explanation: An attempt was made to read a file not on disk or tape. Disk and tape are the only acceptable media for logfiles.

User Action: Do not use CLF to read files not on disk or tape.

CLF -- INVALID MPB FOUND

Explanation: The first block of the input file does not pass certain validity tests which ensure the file is a logfile. This may be caused by a corrupted or obsolete logfile, or a logfile produced by an incompatible version of COL.

User Action: Use a different logfile.

CLF -- INVALID OUTPUT DEVICE TYPE

Explanation: An attempt was made to write in logfile format to a device other than disk or tape.

User Action: Only write logfile format to disk or tape.

CLF -- MORE THAN ONE INPUT FILESPEC NOT ALLOWED

Explanation: The command has improper syntax.

User Action: Reenter the corrected command.

CLF -- MORE THAN ONE OUTPUT FILESPEC NOT ALLOWED

Explanation: The command has improper syntax.

User Action: Reenter the corrected command.

CLF -- MPB NOT FOUND - NOT TRACE FILE

Explanation: The first block of the input file does not pass certain validity tests which ensure the file is a logfile. This may be caused by a corrupted or obsolete logfile.

User Action: Use a different logfile.

CLF -- OUTPUT FILESPEC OR SWITCH SYNTAX ERROR

Explanation: A syntax error was detected in the output filespec.

User Action: Reenter the corrected command.

CLF -- OVERFLOW BEFORE MPB - BAD TRACE FILE

Explanation: The first block of the input file does not pass certain validity tests which ensure the file is a logfile. This may be caused by a corrupted or obsolete logfile.

User Action: Use a different logfile.

CLF -- TRUNCATED INPUT FILE RECORD

Explanation: The end of a buffer was reached before all records in that buffer were processed. This indicates a corrupted logfile.

User Action: Use a different logfile.

CLF -- UNABLE TO OPEN INPUT FILE

Explanation: An I/O error occurred while opening the input file.

User Action: Correct the condition and rerun.

CLF -- UNABLE TO OPEN OUTPUT FILE

Explanation: An I/O error occurred while opening the output file.

User Action: Correct the condition and rerun.

CLF -- UNKNOWN RECORD TYPE

Explanation: An unknown record ID was found in the file.

User Action: Use a different logfile.

CLF -- WILDCARD NOT ALLOWED IN FILESPEC

Explanation: A wildcard was found in the command.

User Action: Reenter the command without wildcards.

3.2.6 Using OVL

OVL is used to read a logfile and produce a report showing the disk resident overlay activity for a task.

A memory-resident overlay segment is loaded from disk the first time that the segment is referenced. If a task has memory-resident overlays, OVL will count their initial loads in the overlay report. Subsequent memory-resident overlay "loads" will not be counted, since a disk load is not performed.

OVL has two input files:

- 1) A mapfile for the task whose disk resident overlays are to be analyzed.
- 2) A logfile.

In addition, the taskname switch /TS:tsknam must be appended to the logfile specification. If the taskname specified by /TS:tsknam is a generic taskname, e.g., ...XYZ, then all overlay hook records with the same generic taskname will be analyzed for the overlay report. If the taskname is not generic (not of the form ...XYZ), then only the hook records whose tasknames exactly match the /TS:tsknam taskname will be analyzed for the overlay report.

The overlay report lists all the segment names for a task and prints the number of times each segment was loaded.

In addition, you can trace each overlay load by using the /TR switch on the output file. As each overlay hook record is read, the segment name and the number of times that segment has been loaded up to that point will be printed.

The command syntax for OVL is:

```
>RUN OVL
OVL>  OUTFILE/SW = MAPFILE, LOGFILE/TS:tsknam
```


OVL Switches

Switch defaults --

output: SY:TRACE.RPT/-TR

input mapfile: SY:TRACE.MAP

input logfile: SY:TRACE.LOG/DE:800/RW/SK:0

Output Switches		Description
/TR	/-TR	Print a trace of each overlay segment, giving the segment name and the number of times the segment has been loaded up to that point.
Input Switches (to logfile)		
/TS:tsknam		This switch is required. If "tsknam" is of the form ...XYZ, all hook records with same generic task name will be included for overlay analysis; if "tsknam" is not of the form ...XYZ, then only hook records with tasknames exactly matching "tsknam" will be included for overlay analysis.
/DE:nnnn		Selects input magtape DEnsity, 800 or 1600 (ignored unless /RW is all used)
/RW	/-RW	Rewind tape input file before opening
/SK:nnn		Skip nnn files when opening tape input file

OVL Messages

The OVL messages are issued to the TI device under which OVL is run. This section discusses the messages which OVL can issue.

```
OVL --   DSW:##### F.ERR:##### IOSB:##### ##### PC:#####  
        SP:#####  
OVL --   STK:##### ##### ...
```

Explanation: This is a status message which gives detail of the internal state of OVL at the time when an error is detected. This message is always followed by one of the messages discussed below, which describes the actual error condition.

User Action: See action specified for the message issued immediately after this status message.

OVL -- COMMAND READ ERROR

Explanation: An I/O error occurred while reading the command line.

User Action: Reenter the command.

OVL -- COMMAND SYNTAX ERROR

Explanation: The command has improper syntax.

User Action: Reenter the corrected command.

OVL -- COMPLETION ERROR IN INPUT READ

Explanation: an I/O error occurred while reading the input file. This condition, which may be symptomatic of a corrupted logfile, will be treated as end-of-file.

User Action: Check F.ERR and IOSB for error cause, or use another logfile.

OVL -- ERROR ASSIGNING INPUT DEVICE

Explanation: OVL was unable to assign a LUN to the input file.

User Action: Check DSW and F.ERR for cause. Ensure the filespec syntax is correct.

OVL -- ERROR ASSIGNING OUTPUT DEVICE

Explanation: OVL was unable to assign a LUN to the output file.

User Action: Check DSW and F.ERR for cause. Ensure the filespec syntax is correct.

OVL -- ERROR ATTACHING, POSITIONING OR SELECTING DENSITY OF MAGTAPE

Explanation: An I/O error occurred when attempting to open a file on magtape.

User Action: Ensure the drive is mounted foreign. Correct the error indicated by DSW or IOSB.

OVL -- ERROR CLOSING INPUT FILE

Explanation: An I/O error occurred when closing the input file.

User Action: None.

OVL -- ERROR CLOSING OUTPUT FILE

Explanation: An I/O error occurred when closing the output file.

User Action: Correct the condition and rerun OVL or manually unlock the output file.

OVL -- ERROR IN INPUT READ QIO

Explanation: An I/O error occurred while reading the input file. This condition, which may be symptomatic of a corrupted logfile, will be treated as end-of-file.

User Action: Check DSW for cause. Rerun using another logfile if necessary.

OVL -- ERROR READING MPB

Explanation: An I/O error occurred when reading the first block of the logfile.

User Action: Correct the condition and rerun.

OVL -- ERROR PARSING MAP FILE

Explanation: OVL could not parse the specified file into overlay segment names and segment block numbers.

User action: Input a valid map file.

OVL -- ERROR PRINTING OVERLAY SUMMARY

Explanation: An I/O or formatting error occurred when attempting to print the summary section of the overlay report.

User action: Correct the condition and rerun.

OVL -- ERROR PRINTING OVERLAY TRACE

Explanation: An I/O or formatting error occurred when attempting to print the trace section of the overlay report.

User action: Correct the condition and rerun.

OVL -- INCONSISTENCY BETWEEN MAP FILE AND LOG FILE

Explanation: The overlay block numbers in the logfile for task specified by /TS:tsknam do not appear in the map file.

User Action: Be sure the map file is correct for all tasks in logfile with generic taskname specified by /TS:tsknam.

OVL -- INPUT FILESPEC OR SWITCH SYNTAX ERROR

Explanation: A syntax error was detected in the input filespec.

User Action: Reenter the corrected command.

OVL -- INVALID DENSITY SELECTED, 800 OR 1600 ONLY

Explanation: A magtape density other than 800 or 1600 was specified.

User Action: Reenter the command with a correct density.

OVL -- INVALID INPUT DEVICE TYPE

Explanation: An attempt was made to read a file not on disk or tape. Disk and tape are the only acceptable media for logfiles.

User Action: Do not use OVL to read files not on disk or tape.

OVL -- INVALID MPB FOUND

Explanation: The first block of the input file does not pass certain validity tests which ensure the file is a logfile. This may be caused by a corrupted or obsolete logfile.

User Action: Use a different logfile.

OVL -- MORE THAN TWO INPUT FILESPECS NOT ALLOWED

Explanation: The command has improper syntax.

User Action: Reenter the corrected command.

OVL -- MORE THAN ONE OUTPUT FILESPEC NOT ALLOWED

Explanation: The command has improper syntax.

User Action: Reenter the corrected command.

OVL -- MPB NOT FOUND - NOT TRACE FILE

Explanation: The first block of the input file does not pass certain validity tests which ensure the file is a logfile. This may be caused by a corrupted or obsolete logfile.

User Action: Use a different logfile.

OVL -- OUTPUT FILESPEC OR SWITCH SYNTAX ERROR

Explanation: A syntax error was detected in the output filespec.

User Action: Reenter the corrected command.

OVL -- OVERFLOW BEFORE MPB - BAD TRACE FILE

Explanation: The first block of the input file does not pass certain validity tests which ensure the file is a logfile. This may be caused by a corrupted or obsolete logfile.

User Action: Use a different logfile.

OVL -- TRUNCATED INPUT FILE RECORD

Explanation: The end of a buffer was reached before all records in that buffer were processed. This indicates a corrupted logfile.

User Action: Use a different logfile.

OVL -- UNABLE TO OPEN INPUT FILE

Explanation: An I/O error occurred while opening the input file.

User Action: Correct the condition and rerun.

OVL -- UNABLE TO OPEN OUTPUT FILE

Explanation: An I/O error occurred while opening the output file.

User Action: Correct the condition and rerun.

OVL -- UNKNOWN RECORD TYPE

Explanation: An unknown record ID was found in the file.

User Action: Use a different logfile.

OVL -- WILDCARD NOT ALLOWED IN FILESPEC

Explanation: A wildcard was found in the command.

User Action: Reenter the command without wildcards.

3.2.7 Writing User Information to the Logfile

The performance of an application task can be measured by instrumenting the application task with a special directive and sub-function code:

GIN\$ 7,P1,P2,P3,P4,P5,P6,P7,P8.

At each occurrence of this directive, the SPM-11M-PLUS performance monitor is invoked and a hook record of type H\$USER is written to the logfile. The eight parameters P1 to P8 are binary words whose values can be set to any value the user desires. The hook record written to the logfile is specified in appendix A, Page A-28.

The user can copy the logfile using CLF to obtain a logfile with H\$USER records in FCS variable length record format. Data reduction programs can then be written using the FCS record file as input.

Example:

GIN\$C 7,100,200,300,400,500,600,700,1000

will result in an H\$USER record with parameters 100, 200, 300, 400, 500, 600, 700, and 1000 written to the logfile.

To dump a stream of H\$USER records in a performance report, you run MDR and use the /DM:40 switch. If you only want the H\$USER records dumped, the command syntax is:

MDR> output = input/DM:40/LI:0/SI:0

To copy a stream of H\$USER records imbedded in a logfile with other record types, you can run CLF as follows:

>CLF output /FCS:40 = infile.

CHAPTER 4

DATA COLLECTION AND REDUCTION PROCEDURES

COL and MDR are complementary tools which can be used to collect large amounts of performance data and interpret the data to a great degree of precision. Both COL and MDR have a number of parameters and switches which provide flexibility in both collecting and interpreting the data. In the following sections we describe how to utilize the parameters of COL and MDR to collect the best possible set of data to analyze your computer system.

4.1 DATA COLLECTION PROCEDURE

There are a number of questions you should answer before running COL to collect data:

- o Length and type of measurement interval?
- o Resource types enabled for data collection?
- o Device classes enabled for data collection?
- o Tasks enabled for data collection?
- o System metrics enabled?

4.1.1 Measurement Period

The measurement period is defined as the time during which COL is collecting data. Typically this time is one continuous time interval; however, it may consist of a number of time intervals piecewise added together.

Under most conditions, you should select a continuous measurement period. This is done by running COL either in MANUAL mode or AUTOSTOP mode.

MANUAL mode should be used when:

- o There is an operator available to terminate COL.
- o You estimate the logfile space available is sufficient to contain all the data collected for the measurement period.

AUTOSTOP mode should be used when:

- o An operator is not available to terminate COL.
- o You want to run COL for a fixed amount of time (AUTOSTOP on TIME).
- o You want to set a limit on the amount of logfile space used (AUTOSTOP on BUFFER).

Under some conditions, particularly if the system is under heavy activity for long periods of time, COL will collect too much data given the available logfile space. To reduce the logfile space used, you can run COL in REPEAT mode. However, you should do so carefully. The chief concern when running COL in REPEAT mode is that "incomplete intervals" may result if a task or operation is outstanding exactly when COL starts or stops data collection. For example, if you start COL precisely when a task is active and has an I/O operation queued to the disk, an incomplete interval will be flagged by MDR for both the task invocation and for the I/O operation. This condition does not invalidate other information in the performance report, it just "leaves out" some information which potentially may be valuable. You can minimize the number of incomplete intervals by minimizing the number of times COL enables or disables data collection for a measurement period.

Let's say you want to collect data during a normal 8-hour daytime shift and you have only one 2400 foot magtape available for the logfile. However, you estimate that if COL runs all day, about two 2400 foot tapes will be needed. What you need to do is cut the amount of data collected in half. Consider the following alternative measurement schemes:

- o Sample time = 1 minute; Time between samples = 1 minute.

- o Sample time = 30 minutes; Time between samples = 30 minutes.
- o Sample time = 2 hours; Time between samples = 2 hours.

Scheme 1 will usually result in the highest number of incompletes, yet it probably will not "miss" any important system activity. Scheme 3 will probably result in the lowest number of incompletes, yet it is vulnerable to missing important intervals of system activity. Scheme 2 will probably result in a low number of incompletes, yet it has a good chance of capturing all periods of important system activity.

What we're observing here is an old problem in communications systems: How fast do you sample a signal to be able to reconstruct the signal from the sampled data? In our case, the problem is: how fast do you sample the workload in order to reconstruct it?

In communication systems, you sample the signal at twice the highest frequency in the signal. In computer systems, it would be reasonable to sample the workload at twice the highest frequency of variation in the workload. Thus if you observe the workload to fluctuate up and down about four times per shift, then 8 samples per shift (every 30 minutes) would reduce the amount of data and yet collect sufficient data to reconstruct the workload.

4.1.2 Resource Types

Normally you want to collect data for all resource types (CPU, memory, I/O). If the amount of data collected is too large, you can select one resource type per measurement period and run a number of measurement periods.

The hooks are grouped according to resource type. To select a particular resource type, you select the corresponding hook group.

4.1.3 Device Classes

The device classes are: Directory, Sequential, Terminal, and Other. If you want to measure the resource contention for a device, select it for data collection. Normally the disks are shared devices and should be selected. If the line printer is spooled and attached to the despooler task you may not need to collect usage data for it. Similarly, terminals are normally

singly used by one task at a time; the measurement results may not be useful from a resource contention point of view.

4.1.4 Tasks

The tasks selected for data collection are the definition of the workload that you want to measure. For example, if you want to analyze the application workload, select all the tasks in that workload and deselect the system service tasks such as F11ACP, MCR..., and QMG.... On the other hand, if you want to measure the system resources that a particular F11ACP is consuming, enable only F11ACP and disable all other tasks.

4.1.5 System Metrics

Normally system metrics should be enabled with a sample interval of 60 seconds. If you want to save logfile space and report paper, longer sample intervals may be used.

4.2 DATA REDUCTION PROCEDURE

Once you have a logfile on disk or tape, you are ready to read the logfile with MDR and reduce the data to a usable performance report. A guide to reading the performance report is given in Chapter 5.

The data reduction procedure should consist of a number of steps methodically performed. A general guideline is to first obtain a number of overviews of the system activity during the entire measurement period and then print detailed analyses of the subset measurement intervals of interest. For example, if during an 8 hour shift there were two intervals of peak activity, from 9 to 11 a.m. and from 2 to 4 p.m., you should first run MDR with /LI:0 to get a system metrics scan of the entire 8 hour shift, and then run MDR with /AF and /BF to get a report of only the peak morning activity and again with /AF and /BF to get only the peak afternoon activity. The general procedure is as follows:

- o Run MDR with /LI:0 for entire logfile.

Identify the intervals of interest and record their endpoints of time using the KW11-P clock timestamps printed above the system metrics.

- o Rerun MDR once for each subset measurement interval using /AF and /BF with the /LI:2 switch.

This will give both system and task summaries. Identify the largest consumers of resources and record their tasknames.

- o Rerun MDR with /LI:4 or /LI:5 and /TS for the tasks desired.

Observe a detailed analysis of the particular tasks which consumed the resources and the times when these resources were consumed.

CHAPTER 5

READING THE PERFORMANCE REPORT

5.1 PERFORMANCE REPORT DATA

When reading this section, refer to Appendix F, which contains a sample performance report.

The performance report consists of four sections:

- o Measurement Parameters
- o System Configuration
- o System Metrics
- o Task Metrics

5.1.1 Measurement Parameters

The measurement parameters are read from the Monitor Parameter Block (MPB) which is the first block of the logfile. The MDR task extracts the information from it and prints the measurement parameters on the first page of the performance report. The items printed are:

- o Label.
- o Measurement mode.
- o Buffer type, partition, number and size.
- o Resources selected for System Metrics.
- o TI filter device (if selected).
- o Tasks selected for data collection.

- o Devices selected for data collection.
- o Hooks enabled.
- o MDR command line

5.1.2 System Configuration

The system configuration information is contained in the Information Records and directly follows the MPB in the logfile. If the /SC switch is specified, the MDR task will extract the information from these records and print it to the logfile. The information printed is:

- o The primary pool (DSR) size in bytes, the executive feature mask words, and secondary pool size in 32 word blocks.
- o A list of the partitions in the system when COL was activated.
- o A list of the controllers and units in the system when COL was activated.
- o A list of the tasks installed in the system when COL was activated.
- o A list of the active checkpoint files and their sizes when COL was activated.

5.1.3 System Metrics

The system metrics are system-wide measurements of CPU usage, pool usage, the usage of a selected partition, checkpoint file usage, and the usage of a selected device. Measurements are made for each sampling interval and the results written to the logfile. The length of the sampling interval is selected in the SMPGEN.CMD procedure. The system metrics are:

- o CPU Usage

For each sampling interval, counters are kept for the total times spent in kernel and idle states. The remainder of the sampling interval time is assumed to be user state. To depict CPU usage, a histogram is printed on the performance report. Each 1% of time spent by the CPU in user mode is represented by a single letter "U". Each 1% of time spent by the CPU in kernel mode is represented by a single letter "K".

The user time is left justified and grows from left to right; the kernel time is right justified and grows from right to left. If any blank spaces exist in the line, they represent CPU idle time.

o Pool Usage

At the end of each sampling interval, COL scans the RSX-11M-PLUS primary and secondary pools and measures the percentage of each pool used. The percentage of pool used is represented by a histogram, where each "P" represents 1% of primary pool usage and each "S" represents 1% of secondary pool usage.

o Partition Usage

The user selects a partition to be monitored via the COL switch /PR:parnam. At the end of each sampling interval, COL scans the partition data structures for the selected partition to measure the percentage of the memory in use in the partition. The percentage of memory used in the partition is represented by a histogram, where each "M" represents 1% of partition usage.

o Checkpoint File Usage

When COL is activated, it scans the checkpoint file list and records a list of all active checkpoint files and their sizes.

At the end of each sampling interval, COL scans the checkpoint file list and measures the percentage of all active checkpoint file space in use. If there is more than one checkpoint file, COL sums the space in all checkpoint files to compute the percentage. The percentage of checkpoint file space used is represented by a histogram, where each "C" represents 1% of checkpoint file space usage.

o Device Usage

The user selects a device to be monitored via the COL switch /DV:ddnn. For each sampling interval, COL tallies all the I/O request times for the selected device to get the percentage of time that the I/O device was in use during the sampling interval. The percentage of time the device was used is represented by a histogram, where each "D" represents 1% of device usage time for the sampling interval.

addition, if the device is a Files-11 disk and is mounted native, the number of disk blocks used on the volume is computed at the end of the sampling interval. The percentage of disk space used is represented by a histogram where each "B" represents 1% of the total volume space used.

For each sampling interval, the histograms for CPU, pool, memory, checkpoint file space, and device are clustered together into six successive lines on the report.

The end of each sampling interval is timestamped and printed in two formats: KW11-P time and RSX-11M-PLUS system time.

The KW11-P time is printed as two octal numbers: the left number is the high order part and the right number is the low order part. The corresponding RSX-11M-PLUS system time is printed as HRS:MIN:SEC.TENTHS.

The KW11-P timestamp can be used to:

- o Correlate an individual task report page (which has the same form of timestamp) with a particular system metric line.
- o Select a particular measurement interval of interest by rerunning MDR using the /AF and /BF switches.

The totals for CPU kernel time, CPU idle time and device busy time are given for each measurement interval.

The last item in the system metrics section is the "Total Data Reduction Period". This is the total time for which data was reduced. If the report was generated using the /AF and /BF switches, only the data reduced between /AF and /BF is counted.

5.1.4 Task Metrics

Levels of Task Summarization

The task metrics describe the usage of the resources in the system on a task-by-task basis. There are a number of levels of summarization available in the performance report, which enables the user to choose between a very detailed report or a shorter summarized report. The levels of summarization and their corresponding switches are as follows:

- o Individual Task Invocations (/LI:4 or /LI:5)

An individual task invocation is defined as a single run of a task, that is, the time between an executive Run Task event and an executive Exit Task event (the exact placement of the hooks for these events is specified in Appendix A). A complete set of task metrics will be printed for each task invocation. If most or all tasks in the system are enabled for data collection, a report containing individual task invocations may be quite long.

- o Task by TI Summary (/LI:3)

In this level of summary, the task metrics for tasks with the same generic taskname (see Chapter 2, section 2.2.2) and the same TI are grouped together when the statistics are computed. The task metrics then reveal the measurements summed across all tasks with the same generic taskname from the same TI.

- o Task Summary (/LI:2)

In this level of summary, the task metrics for tasks with the same generic taskname (see Chapter 2, section 2.2.2) are grouped together when the statistics are computed. The task metrics then reveal the measurements summed across all tasks with the same generic taskname.

- o System-Wide Summary (/LI:1)

In this level of summary, the task metrics for all tasks in the system that were enabled for data collection are grouped together when the statistics are computed. The task metrics then reveal the measurements summed across all tasks in the system.

Task Metric Data

The following data is presented in the performance report for all levels of task summarization.

- o Generic TASK

The generic taskname (see Chapter 2, Section 2.2.2) is printed on the top of the page to identify the task for which the data applies.

o Elapsed Times

The elapsed time is the time from a Run Task event to an Exit Task event.

o Number of Task Invocations

The Number of Task Invocations is the number of times that a Run Task event occurred for the task. If the COL task is started while a task is in progress, the Run Task event will not be recorded in the logfile for that task. When the task exits, an Exit Task event will be recorded to the logfile. When MDR reads an Exit Task event without a corresponding Run Task event, it reports an "incomplete" invocation in the Number of Task Invocations line.

o Memory Times

There are three memory times:

- In (memory)
- Out (memory) - includes time waiting for memory on initial task load and all time checkpointed (waiting and stopped).
- Wait (memory) - includes time waiting for memory on initial task load and time checkpointed (waiting only).

o Memory Size in K Words

Size of the task region.

o Space-Time Product

The space-time product is the amount of memory that a task uses integrated over the time that the memory was used. Each increment of memory used (in K words) is multiplied by the time (in SECs) that the increment was used, and all products are summed to get the space-time product. If a task changes its size via an Extend Task directive, the corresponding change in memory usage is computed.

o ACP Counts

The ACP Counts are tallied across all ACPs in the system. They are not broken down by separate ACP.

There are seven categories of ACP functions printed on the performance report. They are:

- Access/Deaccess

These functions are: Access File for Read, Access File for Read/Write, Access File for Read/Write/Extend, Deaccess File, Connect Logical Link, and Disconnect Logical Link, and Close Out LUN.

- Create/Delete

These functions are: Create File, Delete File and Truncate File.

- Read/Write

These functions are: Read Virtual Block, Write Virtual Block, Receive Message and Transmit Message.

- Directory

These functions are: Find File Name in Directory, Enter File Name in Directory and Remove File Name from Directory.

- File Attribute

These functions are: Read File Attributes and Write File Attributes.

- Extend

This is the Extend File function.

- Other

These functions are all ACP functions other than those listed above, such as: Unlock Block, User Magtape Control Function and Network Control Function.

o Memory Residency

The Memory Residency count is incremented each time a task successfully allocates memory for an initial task load or for a checkpoint read.

o CPU Timeslices

A CPU Timeslice is the time between a Load Task Context event and a Save Task Context event, minus the idle time accrued during that period. Note that if a task issues a directive and the processor switches to kernel mode, the corresponding time in kernel mode is charged to the active task's CPU timeslice. Also, if an interrupt occurs while a task's context is active, the corresponding kernel mode time is charged to the active task's CPU timeslice.

o Initial Loads

Initial Loads is the number of Initial Task Load events for the task. The USAGE times indicate the time it took the loader task to load the task; the WAIT times indicate how long a task's load request was in the loader task's queue.

o Checkpoint R+W

Checkpoint R+W is the number of Checkpoint Read and Checkpoint Write events for the task. The USAGE times indicate the time it took the loader task to checkpoint (read or write) the task. (Note: the USAGE times do not indicate how long a task was checkpointed. This information appears under "Memory Times" on the task metric page.)

o DDnn Ovly Loads

DDnn Ovly Loads is the number of overlay read (IO.LOV) QIOS the task issued to device DDnn.

o DDnn QIOs

DDnn QIOs is the number of logical QIOs a task issued to device DDnn. Please note the following caveats:

- DDnn Ovly Loads is included in this count, since IO.LOV is a Logical block QIO.

- Task IO.RVB and IO.WVB (virtual) QIOs which are successfully and entirely mapped by the executive or F11ACP to IO.RLB and IO.WLB (Logical) QIOs are included in this count.
- Task IO.RVB and IO.WVB (virtual) QIOs which are sent to F11ACP, and which are performed by F11ACP on behalf of the task (partially mapped, multiple transfers), are not included in this count. This logical block I/O is charged to F11ACP.
- The IO.ATT, IO.DET and IO.KIL QIOs are not included in this count.
- Drivers for devices with the characteristic UC.QUE may directly process an I/O request without the IRP being queued. This will cause an incomplete measurement of Usage and Wait times. (e.g., VT:, TT:)

USAGE Times

The USAGE times indicate how long a task used a resource.

WAIT Times

The WAIT times indicate how long a task was waiting for the use of a resource.

SERVICE Times

The SERVICE time is the USAGE time plus WAIT time which occurred for each resource request. Note that the minimum SERVICE time is not necessarily equal to the minimum USAGE time plus minimum WAIT time, since the minimum USAGE time may have occurred for a different resource request than the resource request which had the minimum WAIT time. The same applies to the maximum.

ARRIVALS PER SECOND

The mean task rate is the number of resource requests issued by a task divided by the task's elapsed time. It tells, from a task's point of view, the rate of resource requests issued by a task. The mean system rate is the number of resource requests issued by a task divided by the data reduction period. It tells, from a resource's point of view, the rate of resource requests issued by a task.

Resource Metrics

For each resource used by a task, the following metrics are printed on the report:

- o Count
The number of times a task used a resource.
- o Total Time
The total time during which a task used the resource (equal to the sum of all resource usage intervals).
- o % Resident Time
% Resident Time is $(\text{Total Time} / \text{Resident Time}) * 100\%$.
- o Minimum
The smallest interval during which a task used a resource.
- o Mean
The arithmetic average of all resource usage intervals for a task.
- o Maximum
The largest interval during which a task used a resource.
- o Coefficient of Variation
The coefficient of variation is the standard deviation of the intervals divided by the mean interval.
- o % Resource
% resource is the total time a task used the resource divided by the total time of the Data Reduction Period (total resource time available).
- o Incomplete
If the logfile does not contain a "complete" set of events needed to define an interval, MDR will tally the interval as an "incomplete". For example, if MDR reads a I/O Done event from the logfile for a certain

task and device and there is no corresponding Get I/O Packet event for the task and device, MDR will regard the I/O interval as an incomplete. Incomplete intervals can occur if COL starts data collection while an operation (interval) is already in progress.

The only data recorded for incomplete intervals is the count; no timing information is computed since the information is not precise.

Incomplete intervals are printed in lieu of the % resource metric if the /IC switch is specified.

CHAPTER 6

RESOURCES USED

BY SPM-11M-PLUS

The purpose of this section is to give an estimate of the resources used by the COL and MDR tasks.

6.1. RESOURCES USED BY COL

- o CPU

When a task metric hook is encountered, from 50 to 75 instructions are executed at system level (PR=0, \$STKDP=0) before returning to the main flow of the executive.

- o Executive I-space

Following a SYSGEN, SPM-11M-PLUS occupies approximately 290 words of executive I-space. This is mainly accounted for by the hooks, hook dispatcher and the KW11-P handler.

- o Primary Pool

When the COL task is activated approximately 260 words of primary pool are used.

- o Physical Memory

The COL task occupies about 3.5K words of physical memory. The COL buffers (either in a static common or dynamic region) can vary from 0.5K words to 32K words.

- o I/O Requests

The COL task issues a FCS WRITE\$ to write a buffer to disk; it issues a QIO\$ to write a buffer to magtape. The rate of I/O requests generated by COL depends on the system activity that COL is tracing.

6.2 RESOURCES USED BY MDR

- o Memory

The code and static data of MDR requires varying amounts of memory depending upon which version is built. See the TASK IMAGE SIZE in MDR.MAP. Additional space is required for I/O buffers and dynamic structures which are manipulated within MDR's virtual address space. MDR will automatically extend its virtual space as needed to accommodate these structures. Typically MDR will require an additional 8KW for all the dynamic storage.

- o Pool

All MDR I/O is singly buffered, and synchronous, minimizing the use of pool.

- o I/O

Each pass of MDR reads a single file and writes to the report file. During pass 1 the RESOURCE.TMP file is also being written. Terminal I/O is used for command input, and error reporting. The DRS task will have up to 7 files concurrently opened for I/O.

CHAPTER 7

SPM-11M-PLUS

USAGE NOTES

7.1 TASKBUILDING COL

COL must be linked to the supervisor-mode library [1,1]FCSFSL. This means that supervisor-mode support must be included in the system at SYSGEN.

7.2 COL BUFFERS

When COL starts up it copies three system lists to the log file:

- 1) TCBs
- 2) PCBs
- 3) DCBs-UCBs

There must be enough buffer space to store the largest of these three system lists, otherwise COL will print the error message COLMSG -- INFORMATION RECORD ALLOCATION FAILURE.

Before COL is left to run unattended for long periods of time, COL should be run for 5 to 10 minutes under a typical heavy or peak system workload and the console terminal should be monitored for the occurrence of overflow errors, i.e., COLMSG -- BUFFER OVERFLOW. If buffer overflows occur, the buffer size for COL can be increased by using the /BF switch (see Chapter 3, section 3.2.2).

Another cause of COL overflows may be small disk file extent sizes. When COL attempts to write a buffer to the logfile, F11ACP may have to do an extend file operation before the actual data write can be done. The time it takes F11ACP to extend a file may be long, and COL may overflow as a result of this. To remedy this situation, two switches have been built into COL:

- o Initial File Size /IF
- o Extend File Size /EX

If an estimate can be made of the maximum logfile size before COL is run, it is best to preallocate all the logfile space via the /IF switch. In this case COL will not need F11ACP to do any extend file operations during its run and this may prevent overflows from occurring.

7.3 STATE OF SYSTEM AT COL STARTUP TIME

When COL starts up, it copies the TCBs, PCBs, DCBs and UCBs in RSX-11M-PLUS to the logfile. The MDR task later refers to these data structures to derive the necessary context for printing a performance report.

To assure that MDR will have all the context it needs to print the report, you should insure that the following conditions are satisfied before starting up COL:

- o All tasks installed

If you want to disable a task for data collection, that task must be installed when COL starts up. If the system supports multi-user protection, then the following is true: if you disable a task of the form ...XYZ, then all multi-user copies of that task will also be disabled for data collection.

- o All partitions set up

If you want memory metrics kept for a partition, the partition must be set up
(via SET /MAIN=XYZPAR:m:n:type) before starting COL.

- o All drivers loaded

If you want the I/O metrics kept for a device-unit, the driver managing that device-unit must be loaded before starting COL.

The MDR program is not capable of analyzing a logfile which contains task event records for an I/O device, but no system configuration information for that device. To ensure that unusable logfiles are not created, no device driver database may be loaded while the task ...COL is active.

- o All devices mounted

If you want the correct ACP taskname for a mounted volume to be printed on the performance report, the volume must be mounted before starting COL. Otherwise, MDR will print a "pseudo" taskname for the ACP.

Example: If the device DK1: is mounted to the F11ACP task after data collection is in progress for COL, MDR will print the pseudo name DK1***ACP on the report.

- o All checkpoint files allocated

If you want the system metrics for checkpoint file usage, be sure to allocate all checkpoint files with the ACS command before starting COL.

7.4 TASKBUILDING MDR

MDR will run on processors with or without hardware FPP (floating point processor) support. If your processor does not have hardware FPP, it must minimally have hardware EIS instructions.

You must taskbuild MDR to reflect the type of hardware (FPP or EIS) on your processor. You can build MDR either overlaid or non-overlaid with both FPP and EIS versions. To build MDR, invoke the build command file:

```
>@MDRTKB
```

This command file will prompt you to build any of the four available versions:

1. FPP support, overlaid
2. FPP support, non-overlaid
3. EIS support, overlaid
4. EIS support, non-overlaid

You can change certain MDR defaults by editing the taskbuild command file produced by MDRTKB.CMD. The comments in the MDR build file will prompt you for changing the defaults. The defaults which can be changed are:

- o Number of lines per page in REDUCE.RPT file
- o Default value for /LI switch
- o Default value for /SI switch
- o Default magtape logfile density (/DE)
- o Default value for /DM switch
- o Whether MDR automatically spawns DRS
- o The maximum unit number of a VT in the measured system

7.5 MAXIMUM VT: UNIT NUMBER SUPPORTED BY MDR

MDR has the capability to process the dynamic creation and deletion of virtual terminals during a period of data collection. One aspect of the implementation of this feature is that MDR must pre-allocate the data structures which describe all the VTs while processing the I\$UCB for VT0:.. Thus on a system with no VT: support these structures will not be allocated. Although device unit numbers may range from 0 to 377 (8), RSX-11M-PLUS always assigns the lowest possible unit number to a newly created VT:.. Since each such data structure used by MDR requires 6 words, allowing for 377(8)VTs would require more than 1.5 K Words of MDRs virtual space. To reduce this excessive use of space to a much smaller value, 40(8), is used as the maximum VT: unit number which MDR will recognize. Should this prove insufficient, MDR may be rebuilt with the parameter MAXVT increased.

7.6 PERFORMING REDUCTION UNDER BATCH

MDR may not be run in BATCH if any of the switches /NT, /NU, /TS, or /UC is used. This occurs because MDR uses "read with prompt" to obtain the filter parameters, but the VT: driver does not support this function. Since I/O error at terminals are often correctable, or due to operator error (break key), MDR will repeatedly retry I/O which the VT: driver does not support. This causes an endless loop.

Manual intervention to run the DRS pass is impractical when running MDR under a batch processor.

The tasks CLF, DRS, and MDR terminate their execution via the EXIT\$ directive. This implies that the resulting status will always be EX\$SUC even when an error occurs.

7.7 NECESSARY UFD TO RUN MDR

The intermediate files used by MDR will always be created in the directory SY:[current UIC]. Therefore to run MDR, this directory must exist, on a mounted Files-11 disk.

7.8 DEVICES FOR TI: FILTER

All devices which are specified in the MDR TI: filter (/NU or /UC) must be known to MDR via the I\$UCB records. Thus to use a VT: in this filter, the VT: (or BATCH Processor) must be created before starting data collection.

7.9 MANUALLY EXECUTING THE DRS PASS

MDR cannot determine whether DRS has completed successfully or with errors. If errors are likely in this pass (e.g., available disk space is minimal) manual intervention to execute the DRS pass may be desirable. MDR will issue instructions for this manual operation if either of the following two conditions are true:

- o MDR is built with the GBLPAT which prevents automatic spawn of DRS
- o The task ...DRS is not installed when MDR is running.

7.10 TASKBUILDING CLF

Generally CLF need not be taskbuilt since it is supplied as a task image. However to change defaults or apply patches, the build command file is invoked:

```
>@MDRTEKB
```

The defaults which may be changed are:

- o default magtape density
- o default value of the /FCS:nn switch (if specified as /FCS)

7.11 POSSIBLE HANGING OF CLF

If CLF is writing an FCS output file to the null device (NL:), and no records are actually written, CLF will hang in WAITFOR state when attempting to close the output file. To recover from this error, it is necessary to abort CLF.

7.12 USE WITH LAYERED PRODUCTS

SPM-11M-PLUS inserts hooks into the RSX-11M-PLUS executive at standard places where the executive does allocations and deallocations of resources.

SPM-11M-PLUS measures kernel CPU time by hooking the standard interrupt save (\$INTSV, \$INTSE, \$INTSC) and system exit (\$DIRXT) points.

If any layered product or user privileged component uses nonstandard techniques for resource allocation or interrupt entry, the SPM-11M-PLUS measurements will be inaccurate.

For example, DECNET uses a nonstandard interrupt save routine, which means the kernel CPU time measurement will be inaccurate.

7.13 REPORTING PROBLEMS AND ERRORS

A Software Performance Report (SPR) is used to report problems, errors, or suggested enhancements concerning SPM-11M-PLUS.

Whenever possible, the following information should be submitted with the SPR:

- o A map of the failing component
- o The MPB source file (TRACE.PAR)
- o The logfile on Magtape (TRACE.LOG)
- o The faulty performance report (REDUCE.RPT)
- o The faulty CLF output file on Magtape
- o The faulty OVL output file.

All media will be returned.

APPENDIX A

HOOK RECORD

DEFINITIONS

NOTE: All record offsets in this appendix
are expressed in octal bytes.

All record lengths are expressed in
decimal bytes.

1.0 H\$CDRP - CALL DIRECTIVE PROCESSOR

1.1 LOCATION

DRDSP.MAC

```
After  $EMTRP::DIRSV$  
After  50$:  MOV      R5,R2  
Follow      MOV      $HEADR,R4
```

1.2 DATA COLLECTION

1.2.1 Storage Dependencies -

\$TKTCB (Current Task TCB)
\$DICSV (Directive ID Code)

1.2.2 Record Format -

Hook I.D	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
DIC	14

1.2.3 Record Length -

14. bytes.

2.0 H\$RTSK - REQUEST TASK TO RUN

2.1 LOCATION

```
      REQS.B.MAC
      After   $TSKRP::
      After   1$:
      Follow          MOV      R2,T.UCB(R0)
```

2.2 DATA COLLECTION

2.2.1 Storage Dependencies -

R0=TCB address

2.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
Request UIC	14

2.2.3 Record Size -

14. bytes.

3.0 H\$XTSK - TASK EXIT

3.1 LOCATION

DREIF.MAC

After \$CEXIT::

Follow 130\$: MOV \$TKTCB,R0

3.2 DATA COLLECTION

3.2.1 Storage Dependencies -

\$TKTCB=TCB address

3.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
Idletime	14

3.2.3 Record Length

16. bytes.

4.0 H\$SCTX - SAVE CONTEXT OF TASK

4.1 LOCATION

SYSXT.MAC

After RESCH: CLR

Follow

73\$:

MOV

@\$RQSCH

H.GARD(R3),R2

4.2 DATA COLLECTION

4.2.1 Storage Dependencies -

\$TKTCB (Current Task TCB)

4.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
Idle time	14

4.2.3 Record Length

16. bytes

5.0 H\$LCTX - LOAD CONTEXT OF TASK

5.1 LOCATION

SYSXT.MAC

After RESCH: CLR @\$RQSCH
Follow 120\$: MOV H.GARD(R2),R0

5.2 DATA COLLECTION

5.2.1 Storage Dependencies -

\$TKTCB (Current Task TCB)

5.2.2 Record Format -

Hook I.D	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
Idle time	14

5.2.3 Record Length

16. bytes

6.0 H\$USTP - UNSTOP TASK

6.1 LOCATION

REQSB.MAC
After \$EXRQN::BIT #T2.STP*2!T2.STP,T.ST2(R0)
Follow BIC #T2.STP*2!T2.STP,T.ST2(R0)

6.2 DATA COLLECTION

6.2.1 Storage Dependencies -

R0=TCB address

6.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12

6.2.3 Record Length -

12. bytes.

7.0 H\$QAST - QUEUE AST TO TASK

7.1 LOCATION

REQSB.MAC
After \$QASTT::ADD #T.ASTL,R0
Follow BEQ 5\$

REQSB.MAC
After \$QUEXT::
Follow Mov R2, R0

7.2 DATA COLLECTION

7.2.1 Storage Dependencies -

R0=TCB address

7.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12

7.2.3 Record Length -

12. bytes.

8.0 H\$EXTK - EXTEND TASK

8.1 LOCATION

DREXP.MAC

After \$DREXP::

Follow 71\$: MOV R1,P.SIZE(R0)

8.2 DATA COLLECTION

8.2.1 Storage Dependencies -

\$TKTCB=TCB address

R0=PCB address

8.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
Partition Size (P.SWSZ)	14

8.2.3 Record Length -

14. bytes.

9.0 H\$QPAR - QUEUE REQUEST FOR PARTITION

9.1 LOCATION

REQSB.MAC
After \$TSKRP::
After 4\$:
Follow MOV R3,R1

9.2 DATA COLLECTION

9.2.1 Storage Dependencies -

R1=PCB address

9.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12

9.2.3 Record Length -

12. bytes.

10.0 H\$QLDR - QUEUE TASK LOAD REQUEST TO LOADER QUEUE

10.1 LOCATION

REQSB.MAC

After \$LOADT::

Follow 50\$: MOV \$LDRPT,R0

10.2 DATA COLLECTION

10.2.1 Storage Dependencies -

R1=PCB address

10.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
P.SIZE	14
P.STAT	16

10.2.3 Record Length -

16. bytes.

11.0 H\$GLDR - GET REQUEST FROM LOADER QUEUE

11.1 LOCATION

LOADR.MAC

After \$LOADR::CALL \$SWSTK,40\$
Follow 5\$: MOV R1, R3

11.2 DATA COLLECTION

11.2.1 Storage Dependencies -

R1=PCB address

11.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12

11.2.3 Record Format -

12. bytes.

12.0 H\$FLDR - FINISH LOADER REQUEST

12.1 LOCATION

LOADR.MAC

After \$LOADR::CALL \$SWSTK,40\$
Follow 81\$: CALL \$SWSTK,\$LOADR

12.2 DATA COLLECTION

12.2.1 Storage Dependencies -

R4=PCB address

12.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12

12.2.3 Record Length -

12. bytes.

13.0 H\$RLPR - RELEASE PARTITION

13.1 LOCATION

REQSB.MAC

After \$RLPAR::
Follow BNE \$NXTSK

13.2 DATA COLLECTION

13.2.1 Storage Dependencies -

R1 = PCB address

13.2.2 Record Format -

Hook ID	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
T.ST2	14

13.2.3 Record Length -

14. bytes.

14.0 H\$QDRV - QUEUE I/O PACKET TO DRIVER

14.1 LOCATION

DRQIO.MAC

After \$DRQIO::

Follow \$DRQRQ::MOV U.SCB(R5),R4

14.2 DATA COLLECTION

14.2.1 Storage Dependencies -

R1=IRP address

R5=UCB address

14.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
UCB Address	14
Function code	16
LUN	20
IRP Address	22

14.2.3 Record Length -

20. bytes.

15.0 H\$GPKT - DRIVER OR ACP GETS AN I/O PACKET

15.1 LOCATION

```
IOSUB.MAC
    After  $GTPKT::
    After  127$:
    Follow                                MOVB U.UNIT(R5),R2

DISPAT.MAC (F11ACP)
    After  .START::
    Follow                                MOV $HEADR,R5
```

15.2 DATA COLLECTION

15.2.1 Storage Dependencies -

R1=IRP address

15.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
IRP Address	14

15.2.3 Record Length -

14. bytes.

16.0 H\$IODN - I/O DONE PROCESSING

16.1 LOCATION

IOSUB.MAC
Follow \$IODON::MOV U.SCB(R5),R4

16.2 DATA COLLECTION

16.2.1 Storage Dependencies -

R4= SCB Address

16.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
IRP Address	14

16.2.3 Record Length -

14. bytes.

17.0 H\$IOFN - COMPLETE I/O POST PROCESSING

17.1 LOCATION

IOSUB.MAC

Follow \$IOFIN::

17.2 DATA COLLECTION

17.2.1 Storage Dependencies -

R3=IRP address

17.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
IRP Address	14

17.2.3 Record Length -

14. bytes.

18.0 H\$QAC1 - QUEUE I/O PACKET TO ACP

18.1 LOCATION

```
DRQIO.MAC
    After $DRQIO::
    Follow      MOV  U.ACP(R5),R0
```

18.2 DATA COLLECTION

18.2.1 Storage Dependencies -

R1=IRP address

18.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
UCB Address	14
Function Code	16
LUN	20
IRP Address	22

18.2.3 Record Length -

20. bytes.

19.0 H\$QAC2 - QUEUE I/O PACKET TO ACP

19.1 LOCATION

IOSUB.MAC

After \$GTPKT::

Follow 95\$:

19.2 DATA COLLECTION

19.2.1 Storage Dependencies -

R1=IRP address

19.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
UCB Address	14
Function Code	16
LUN	20
IRP Address	22

19.2.3 Record Length -

20. bytes.

20.0 H\$CRVT - CREATE VIRTUAL TERMINAL

20.1 LOCATION

DRCRV.MAC

After \$DRCRV::

After 70\$:

Follow

MOVB D.UNIT(R1),2(SP)

20.2 DATA COLLECTION

20.2.1 Storage Dependencies -

R1=DCB address

20.2.2 Record Format -

I\$CVT	0
UCB Address	2
Unit Number	4

20.2.3 Record Length -

6. bytes.

21.0 H\$ELVT - ELIMINATE VIRTUAL TERMINAL

21.1 LOCATION

DRCRV.MAC

After \$DRCRV::

Follow \$DEAVT:: MOV (R0),R0

21.2 DATA COLLECTION

21.2.1 Storage Dependencies

R0 = DCB Address

21.2.2 Record Format

I\$ELV	0
Unit Number	2

21.2.3 Record Length -

4. bytes.

22.0 H\$OVLY - LOAD OVERLAY

22.1 LOCATION

```
DROIO.MAC
      After $DRQIO::
      Follow      BNE IE0VR
```

22.2 DATA COLLECTION

22.2.1 Storage Dependencies -

R4 = <Address +2> of Overlay Relative Block Number

22.2.2 Record format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
Relative Block Number	14

22.2.3 Record Length -

14. bytes.

23.0 H\$SYEN - SYSTEM ENTRANCE

23.1 LOCATION

SYSXT.MAC
After \$DIRSV::
Follow MOV SP, @\$HEADR

SYSXT.MAC
After \$INTSI::
Follow MOV SP, @\$HEADR

SYSXT.MAC
After \$INTSC::
Follow MOV SP, @\$HEADR

23.2 DATA COLLECTION

23.2.1 Storage Dependencies -

None

23.2.2 Routine -

Timestamps system entrance time \$BSTIM.

23.2.3 Record Format -

None

24.0 H\$SYXT - SYSTEM EXITS TO USER TASK

24.1 LOCATION

SYSXT.MAC

```
After $DIRXT::  
After 9$:  MOV (SP)+,R0  
Follow      MOV (SP)+,R3
```

24.1.1 Storage Dependencies -

None

24.1.2 Routine -

Copies current clock time, subtracts system entrance time (\$BSTIM), and adds the resulting delta time to the total system time \$TSTIM.

24.1.3 Record Format -

None

25.0 H\$EIDL - ENTER EXECUTIVE IDLE LOOP

25.1 LOCATION

SYSXT.MAC

After RESCH:

After 30\$:

Follow

CLRB @\$CURPR

25.2 DATA COLLECTION

25.2.1 Storage Dependencies -

None

25.2.2 Routine -

Copies current clock time, subtracts the system entrance time (\$BSTIM), and adds the resulting delta time to the total system time \$TSTIM. Timestamps idle entrance time \$BIDLE.

25.2.3 Record Format -

None

26.0 H\$XIDL - EXIT EXECUTIVE IDLE LOOP

26.1 LOCATION

SYSXT.MAC

After \$QFORK::
Follow CLRBL \$IDLFL

26.2 DATA COLLECTION

26.2.1 Storage Dependencies -

None

26.2.2 Routine -

Copies current clock time, subtracts the idle loop entrance time (\$BIDLE), and adds the resulting delta time to the total idle time \$IDLE. Timestamps system entrance time \$BSTIM.

26.2.3 Record Format -

None

27.0 H\$USER - USER INFORMATION HOOK

27.1 LOCATION

DRGIN.MAC
Follow \$SFPMN::

27.2 DATA COLLECTION

27.2.1 Storage Dependencies -

\$TKTCB (Current Task TCB)
R3 = Address of DPB

27.2.2 Record Format -

Hook I.D.	0
Timestamp	2
Taskname (RAD50)	6
TI UCB Address	12
GIN\$ Parameter 1	14
GIN\$ Parameter 2	16
GIN\$ Parameter 3	20
GIN\$ Parameter 4	22
GIN\$ Parameter 5	24
GIN\$ Parameter 6	26
GIN\$ Parameter 7	30
GIN\$ Parameter 8	32

27.2.3 Record Length -

28. bytes.

APPENDIX B

MPB DEFINITION

Monitor Parameter Block

The Monitor Parameter Block (MPB) is the data structure which contains the measurement parameters. These measurement parameters are defined by invoking the procedure @SMPGEN.CMD, which produces:

- o The binary MPB file (filename.MPB)
- o The source MPB file (filename.PAR)
- o The MPB saved answer file (filename.CMD)

The COL task reads the MPB and initializes its databases from the MPB.

The length of an MPB in bytes is M.LGTH (defined by MPBDF\$ in HKMAC.MAC)

The MPB has the following format and contents:

<u>Parameter</u>	<u>Offset</u>	<u>Size(wds)</u>	<u>Values</u>
MPB ID word (I\$MPB)	M.PID	1	I\$MPB
Measurement label	M.LABL	32.	
Buffer Type	M.BTYP	1	DYNAMIC=0 STATIC=1
Buffer Count	M.BCNT	1	
Buffer Size (bytes)	M.BSIZ	1	
Buffer Size (64b blk)	M.BS64	1	

<u>Parameter</u>	<u>Offset</u>	<u>Size(wds)</u>	<u>Values</u>
Buffer Partition Name	M.BPAR	2	
SYS Partition Name	M.BPNM	2	
Buffer Partition Addr.	M.BPAD	1	
Buffer Partition Size	M.BTSZ	1	
Buffer Partition Flags	M.BFLG	1	
Measurement Mode	M.MODE	1	MANUAL=0 AUTOSTOP=1 REPEAT=2
Sample Count (Buffers)	M.SCNT	1	
Sample Interval Mag.	M.SINT	1	
Sample Interval Units	M.SUNT	1	
Repeat Count	M.RCNT	1	
Repeat Interval Mag.	M.RINT	1	
Repeat Interval Units	M.RUNT	1	
Hook Bit Table	M.HTBL	4	
I/O Selection Word	M.IOS	1	
Other Devices Sel. Word	M.OTHR	1	
Device Class Mask	M.CW1M	1	
Device Selection Mask	M.CW1S	1	
Sys. Metrics Samp. Int.	M.SAMP	1	
Sys. Metrics Partition	M.PART	2	
Sys. Metrics Device	M.DEVC	2	
TI Filter Device	M.TIDV	1	
TI Filter Device Num.	M.TINM	1	
MPB Validation Word	M.VALD	1	P\$\$VER
Task Selection Mask	M.TKS	1	
Task Selection Table	M.TTBL	128.	

APPENDIX C

INFORMATION RECORD

DEFINITIONS

Information Record

The Information Record provides system context data which can be correlated with the Hook Records to identify tasks, partitions, and I/O device-units.

In addition, there are Information Records which capture essential data when COL starts and stops data collection.

For lengths of the following data structures, see the macro IBDEF\$ in HKMAC.MAC.

I\$SYS - System Common Information

I\$SYS	0
Initial Pool Size	2
\$FMASK	4
\$FMASK+2	6
Sys. Met. Dev. UCB Address	10
SEC Pool Partition Size	12

I\$TCB - Task Control Block

I\$TCB	0
I\$TCB Record Size	2
TCB Address	4
TCB Contents . . .	6

I\$PCB - Partition Control Block

I\$PCB	0
I\$PCB Record Size	2
PCB Address	4
PCB Contents . . .	6

I\$DCB - Device Control Block

I\$DCB	0
I\$DCB Record Size	2
DCB Address	4
DCB Contents . . .	6

I\$UCB - Unit Control Block

I\$UCB	0
I\$UCB Record Size	2
UCB Address	4
UCB Contents . . .	6

I\$CKP - Checkpoint File Information

I\$CKP	0
CKP. File Device UCB Addr.	2
CKP. File Size	4

I\$TMR - Start Data Collection

I\$TMR	0
System Time (Year)	2
System Time (Month)	4
System Time (Day)	6
System Time (Hour)	10
System Time (Minute)	12
System Time (Second)	14
System Time (Tick)	16
Ticks Per Second	20
High KW11-P Time	22
Low KW11-P Time	24

I\$STM - Stop Data Collection

I\$STM	0
System Time (Year)	2
System Time (Month)	4
System Time (Day)	6
System Time (Hour)	10
System Time (Minute)	12
System Time (Second)	14
System Time (Tick)	16
Ticks Per Second	20
High KW11-P Time	22
Low KW11-P Time	24
CPU Kernel Time (High)	26
CPU Kernel Time (Low)	30
CPU Idle Time (High)	32
CPU Idle Time (Low)	34
Device Usage Time (High)	36
Device Usage Time (Low)	40

APPENDIX D

SYSTEM METRICS

RECORD DEFINITION

System Metrics Record

The System Metrics Record contains data which represents the system-wide usage of the CPU, pool, a selected memory partition, checkpoint file space, and a selected device.


S\$MET	0
High KW11-P Time	2
Low KW11-P Time	4
CPU Kernel Time (High)	6
CPU Kernel Time (Low)	10
CPU Idle Time (High)	12
CPU Idle Time (Low)	14
# Free Pool Nodes	16
Total Free Pool Space (bytes)	20
Smallest Pool Node (bytes)	22
Largest Pool Node (bytes)	24
Main Partition PCB Address	26
Number of Subpartitions	30
Total Space Used (32W blks)	32
Total CKP.File Space	34
Total CKP.File Space Used	36
UCB Address of Device	40
Device Usage Time (High)	42
Device Usage Time (Low)	44
Free Disk Blocks (High)	46
Free Disk Blocks (Low)	50
SEC Pool Free (32 W blks)	52

APPENDIX E

LOGFILE FORMAT

1.0 DISK LOGFILE FORMAT

A logfile on disk is a Files-11 file with the following format:



Each BUFFER has the following general format (byte offsets):

Record Count	0
Overflow Count	2
x\$xxx	4
.	
.	
.	
x\$xxx	
.	
.	
.	
Unused Space	

x\$xxx is I\$, H\$ or
S\$ record ID.
See appendices
A, B, C and D.

All BUFFERS are the same length. The BUFFER size in 512B blocks is derived from the BUFFER size in bytes which is stored in the MPB in BUFFER #1:

BUFFER #1 Format (byte offsets)


Record Count	0
Overflow Count	2
I\$MPB	4
⋮	
BUFFER Size (bytes)	4+M.BSIZ
⋮	
I\$SYS	
⋮	
x\$xxx	
⋮	
Unused Space	

Note:
I\$MPB is the 1st record;
I\$SYS is the 2nd record;
x\$xxx is I\$, H\$ or S\$
record ID.
See appendices
A, B, C and D.

$$\text{BUFFER size (blocks)} = \frac{\text{BUFFER size (bytes)}}{512.}$$

2.0 MAGTAPE LOGFILE FORMAT

A magtape can contain one or more logfiles. Each BUFFER is terminated by an IRG. Each logfile is terminated by a tape mark (EOF). The magtape is terminated by a double tape mark (EOF EOF). Each BUFFER has the same format as for disk.



3.0 FCS LOGFILE FORMAT

The CLF utility task can be used to convert a logfile from the formats specified in section 1.0 and 2.0 to a file containing a stream of variable length FCS records. Each FCS record contains one record of type H\$xxxx, I\$xxxx or S\$xxxx. The first word of each record contains the record ID, H\$xxxx, I\$xxxx or S\$xxxx. The record formats are specified in Appendices A, B, C and D.

See Chapter 3, Section 3.2.5 for using CLF.

APPENDIX F
EXAMPLE PERFORMANCE
REPORT

Another test ... No Noise Tasks!

PERFORMANCE MONITOR DATA COLLECTION PARAMETERS

Mode: MANUAL

Buffering in a DYNAMIC region called TRCPAR in Partition GEN , of 6 Buffers with 4 Blocks per Buffer

System-wide metrics recorded every 120 seconds for the CPU, Partition GEN , and Device LBO:

Task Workload data was limited to the 11: of 115:

Task Workload DEVICES Selected: DISKS TERMINALS

Task Workload Data Collection Enabled for All EXCEPT the following Task(s):

TKIN	...TRM	...RMD	FLMSG	MIAACP	SHE...	...DCL	DCL...	...DMO	...MCR	MCR...	...MOU	...CA.	COT...	ERRLOG	HRC...
SYSLOG	FXR...	...INS	...UFD	QMG...	TTP11	TTP12	PRT...	...AT.	...SHA	...ERF	...ACS	...BRO	...BYE	...CON	...HEL
...	LOA	...	...	...	...	...	...	...	...	...	...	...	...	...	...

the following HOOK POINTS are ENABLED:

(Directives)	RTSK	XTSK														
(Run/Exit)	SCTX	LCTX														
(CPU Usage)	USTP	QAST	EXTK	QPAR	QLDR	GLDR	FLDR	RLPR								
(Memory/LDR)	QDRV	GPST	IODN	IOFN	QAC1	QAC2	CRVT	ELVT								
(I/O, ACF's)	OVLV															
(Overlay)	SYEN	SYXT	EIDL	XIDL												
(Kernel/Idle)																
(User Define)																

MDR command line>SAMPLE/-SP=SAMPLE/SI:777

SPM-11M-PLUS V1.0 (for RSX-11M-PLUS V1.0) 4-DEC-81 16:41:37 Page 2 Sampling Intervals

Another test ... No Noise Tasks!

Data Collection Started 4-DEC-81 16:31:53.9 00003 161170

	0	20	40	60	80	100	(% utilization)
ending at 000272:175031 (16:33:53.9):							
CPU Usage	UUUUUUUUUUUUUUUUUUUU						
POOL Usage	PPPPPPPPPPPPPPPPPPPP						
SECPOL Usage	SSSSSSSSSSSSSSSSSSSS						
GEN	MMMMMMMMMMMMMMMMMMMM						
CKPT Space	CCC						
DRO: Busy	DDDDDDDD						
DRO: Bks	RRRRRRRRRRRRRRRRRRRR						
ending at 000562:013215 (16:35:53.9):							
CPU Usage	UUUUUUUUUUUUUUUUUUUU						
POOL Usage	PPPPPPPPPPPPPPPPPPPP						
SECPOL Usage	SSSSSSSSSSSSSSSSSSSS						
GEN	MMMMMMMMMMMMMMMMMMMM						
CKPT Space	CCC						
DRO: Busy	DDDDDDDDDD						
DRO: Bks	RRRRRRRRRRRRRRRRRRRR						
ending at 001051:027670 (16:37:53.9):							
CPU Usage	UUUUUUUUUUUUUUUUUUUU						
POOL Usage	PPPPPPPPPPPPPPPPPPPP						
SECPOL Usage	SSSSSSSSSSSSSSSSSSSS						
GEN	MMMMMMMMMMMMMMMMMMMM						
CKPT Space	CCC						
DRO: Busy	DDDDDDDD						
DRO: Bks	RRRRRRRRRRRRRRRRRRRR						
ending at 001340:043643 (16:39:53.9):							
CPU Usage	UUUUUUUUUUUUUUUUUUUU						
POOL Usage	PPPPPPPPPPPPPPPPPPPP						
SECPOL Usage	SSSSSSSSSSSSSSSSSSSS						
GEN	MMMMMMMMMMMMMMMMMMMM						
CKPT Space	CCC						
DRO: Busy	DDDDDDDDDDDDDDDDDDDD						
DRO: Bks	RRRRRRRRRRRRRRRRRRRR						
Data Collection Ended 4-DEC-81 16:40:32.7 001433 065013							
SECONDS (System-wide) Elapsed=518.73683 Kernel=42.36473 Idle=409.92571 Device=49.53258							
Data Reduction Period (Seconds) = 518.73683							

SPM-11M-PLUS V1.0 (for RSX-11M-PLUS V1.0)

4-DEC-81 16:41:37

Page 3

Task Summary

Another test ... No Noise Tasks!

Generic TASK = ...KED

Elapsed Times (seconds):

6.920	18.17	29.42	0.62	36.3377
minimum	mean	maximum	coef(var)	total

Number of Task Invocations = 2

(2 complete) (0 incomplete)

Memory Times seconds:

36.34	0.1560E-02	0.8400E-03	X elapsed:	100.0	0.0	0.0
in	out	wait		in	out	wait

Memory Size in KWords:

14.75	14.75
min	max

Space-Time Product:

2	535.959	267.979	0.62	2
increments	total KWS	mean KWS	coef(var) KWS	incomplete

ACP Counts: 12

access/deaccess create/delete read/write

8	1	2	0	0
create/delete	read/write	file attribute	extend	other

RESOURCE

	count	total time	% resident time	minimum	mean	maximum	coef(var)	% resource
Memory Residency	2	36.3362	N/A	6.9187	18.1681	29.4175	0.62	N/A
CPU Timeslices	164	2.48250	6.8	0.0018	0.0150	0.5149	3.33	0.5
Initial Loads	2	0.125100	0.3	0.0470	0.0626	0.0781	0.25	0.0
DR0: Dvly Loads	34	0.749450	2.1	0.0042	0.0220	0.0632	0.65	0.1
DR0: QIOs	34	0.749450	2.1	0.0042	0.0220	0.0632	0.65	0.1
DR1: QIOs	21	0.534100	1.5	0.0086	0.0254	0.0643	0.71	0.1
TT5: QIOs	198	30.7584	84.6	0.0004	0.1553	7.6509	4.48	5.9

RESOURCE

	count	total time	% resident time	minimum	mean	maximum	coef(var)
Memory Residency	2	0.840000E-03	N/A	0.0004	0.0004	0.0004	0.00
Initial Loads	2	0.294000E-02	0.0	0.0015	0.0015	0.0015	0.01
DR0: Dvly Loads	34	0.114240	0.3	0.0004	0.0034	0.0422	2.88
DR0: QIOs	34	0.114240	0.3	0.0004	0.0034	0.0422	2.88
DR1: QIOs	21	0.348600E-01	0.1	0.0017	0.0017	0.0262	3.30
TT5: QIOs	198	0.983300E-01	0.3	0.0004	0.0005	0.0007	0.08

RESOURCE

	count	total time	% resident time	minimum	mean	maximum	coef(var)
Memory Residency	2	36.3370	N/A	6.9192	18.1685	29.4179	0.62
Initial Loads	2	0.128040	0.4	0.0485	0.0640	0.0796	0.24
DR0: Dvly Loads	34	0.863690	2.4	0.0066	0.0254	0.1054	0.83
DR0: QIOs	34	0.863690	2.4	0.0066	0.0254	0.1054	0.83
DR1: QIOs	21	0.568760	1.6	0.0090	0.0271	0.0648	0.69
TT5: QIOs	251	30.8769	85.0	0.0004	0.1230	7.6514	5.05
DR01F1 QIOs	34	2.38404	6.6	0.0024	0.0701	0.2925	0.95

RESOURCE

	# arrivals	mean task rate	ARRIVALS PER SECOND	mean system rate
Memory Residency	2	0.06	0.00	0.32
CPU Timeslices	164	4.51	0.32	

Another test ... No Noise Tasks!

Initial Loads	2	0.06	0.00
DRO: Dvly Loads	34	0.94	0.07
DRO: QIOs	34	0.94	0.07
DRI: QIOs	21	0.58	0.04
TTS: QIOs	251	6.91	0.48
DRO1F1 QIOs	34	0.94	0.07

SPM-11M-PLUS V1.0 (for RSX-11M-PLUS V1.0)

4-DEC-81 16:41:37

Page 5

Task Summary

Another test ... No Noise Tasks!

Generic TASK = ...MAC

Elapsed Times (seconds):

68.12	68.12	68.12	0.00	68.1162
minimum	mean	maximum	coef(var)	total

Number of Task Invocations = 1

(1 complete)

(0 incomplete)

Memory Times seconds: 68.12

0.1140E-02	0.7800E-03	% elapsed:	100.0	0.0	0.0
in	out	wait	in	out	wait

Memory Size in KWords:

13.34	26.79	28.81
min	mean	max

Space-Time Product:

17	1825.01	107.354	2.63
increments	total KWS	mean KWS	coef(var) KWS

ACP Counts: 46

create/delete	3	1	24	0
access/deaccess	create/write	directory	file attribute	extend

18

0

other

RESOURCE

	count	total time	% resident time	minimum	mean	maximum	coef(var)	% resource
Memory Residency	2	68.1151	N/A	5.9924	34.0575	62.1226	0.82	N/A
CPU Timeslices	327	48.6094	71.4	0.0003	0.1487	6.1600	3.30	9.4
Checkpoint R + W	2	0.143860	0.2	0.0503	0.0719	0.0935	0.30	0.0
Initial Loads	1	0.799800E-01	0.1	0.0800	0.0800	0.0800	0.00	0.0
DR0: Only Loads	359	8.41797	12.4	0.0067	0.0234	0.0756	0.43	1.6
DR0: QIOs	426	9.83298	14.4	0.0051	0.0231	0.0756	0.46	1.9
DR1: QIOs	263	5.02539	7.4	0.0042	0.0191	0.0684	0.80	1.0

RESOURCE

	count	total time	% resident time	minimum	mean	maximum	coef(var)
Memory Residency	2	0.780000E-03	N/A	0.0004	0.0004	0.0004	0.05
Checkpoint R + W	2	0.180000E-02	0.0	0.0005	0.0009	0.0014	0.47
Initial Loads	1	0.185000E-02	0.0	0.0019	0.0019	0.0019	0.00
DR0: Only Loads	359	0.432650	0.6	0.0004	0.0012	0.0460	4.08
DR0: QIOs	426	0.557930	0.8	0.0004	0.0013	0.0483	4.22
DR1: QIOs	263	0.263510	0.4	0.0004	0.0010	0.0361	3.67

RESOURCE

	count	total time	% resident time	minimum	mean	maximum	coef(var)
Memory Residency	2	62.1232	N/A	0.0002	31.0616	62.1230	1.00
Checkpoint R + W	2	0.145740	0.2	0.0508	0.0729	0.0949	0.30
Initial Loads	1	0.818300E-01	0.1	0.0818	0.0818	0.0818	0.00
DR0: Only Loads	359	8.85062	13.0	0.0071	0.0247	0.0906	0.50
DR0: QIOs	426	10.3909	15.3	0.0056	0.0244	0.1115	0.55
DR1: QIOs	263	5.28890	7.8	0.0046	0.0201	0.0926	0.81
DR00F1 QIOs	34	0.963570	1.4	0.0032	0.0283	0.1038	1.10
DR01F1 QIOs	58	2.79098	4.1	0.0024	0.0481	0.1750	0.96

RESOURCE

	# arrivals	mean task rate	ARRIVALS PER SECOND	mean system rate
Memory Residency	3	0.04	0.01	0.01

Another test ... No Noise Tasks!

CPU Timeslices	327	4.80	0.63
Checkpoint R + W	2	0.03	0.00
Initial Loads	1	0.01	0.00
DR0: Ovlv Loads	359	5.27	0.69
DR0: QIOs	426	6.25	0.82
DR1: QIOs	263	3.86	0.51
DR00F1 QIOs	34	0.50	0.07
DR01F1 QIOs	58	0.85	0.11

SPM-11M-PLUS V1.0 (for RSX-11M-PLUS V1.0)

4-DEC-81 16:41:37

Page 7

Task Summary

Another test ... No Noise Tasks!

Generic TASK = ...PIP

Elapsed Times (seconds): 1.303 minimum 1.634 mean 2.228 maximum 0.26 coef(var) 4.90299 total

Number of Task Invocations = 3 (3 complete) (0 incomplete)

Memory Times seconds: 4.901 in 0.2360E-02 out 0.1250E-02 wait 100.0 in 0.0 out 0.0 wait

Memory Size in KWords: 11.34 min 11.34 mean 11.34 max

Space-Time Product: 3 55.5915 18.5305 mean KWS 0.26 coef(var) KWS incomplete 3

ACP Counts: 6 access/delete create/delete read/write 2 directory 12 file attribute 3 extend 0 other 0

RESOURCE

RESOURCE	count	total time		% resident time		USAGE TIMES IN SECONDS		coef(var)		% resource
		min	max	min	max	mean	maximum	mean	maximum	
Memory Residency	3	4.90063	N/A	1.3019	1.6335	0.0005	0.2131	0.0005	2.2249	0.26
CPU Timeslices	49	0.553460	11.3	0.0005	0.1113	0.0631	0.0750	0.0005	0.2131	2.76
Initial Loads	3	0.211490	4.3	0.0005	0.0631	0.0631	0.0750	0.0005	0.0750	0.07
DRO: Only Loads	35	0.836200	17.1	0.0063	0.0239	0.0063	0.0661	0.0063	0.0661	0.61
DRO: QIOs	35	0.836200	17.1	0.0063	0.0239	0.0063	0.0661	0.0063	0.0661	0.61
DRI: QIOs	3	0.892200E-01	1.8	0.0053	0.0297	0.0053	0.0662	0.0053	0.0662	0.88
TTS: QIOs	25	1.35006	27.5	0.0420	0.0540	0.0420	0.0909	0.0420	0.0909	0.23

RESOURCE

RESOURCE	count	total time		% resident time		WAIT TIMES IN SECONDS		coef(var)	
		min	max	min	max	mean	maximum	mean	maximum
Memory Residency	3	0.125000E-02	N/A	0.0004	0.0004	0.0004	0.0004	0.0004	0.01
Initial Loads	3	0.493000E-02	0.1	0.0015	0.0016	0.0015	0.0018	0.0015	0.08
DRO: Only Loads	35	0.151100E-01	0.3	0.0004	0.0004	0.0004	0.0005	0.0004	0.07
DRO: QIOs	35	0.151100E-01	0.3	0.0004	0.0004	0.0004	0.0005	0.0004	0.07
DRI: QIOs	3	0.138000E-02	0.0	0.0004	0.0005	0.0004	0.0006	0.0004	0.14
TTS: QIOs	25	0.114600E-01	0.2	0.0005	0.0005	0.0005	0.0005	0.0005	0.01

RESOURCE

RESOURCE	count	total time		% resident time		SERVICE TIMES IN SECONDS		coef(var)	
		min	max	min	max	mean	maximum	mean	maximum
Memory Residency	3	4.90188	N/A	1.3023	1.6340	0.0721	0.0768	0.0721	0.07
Initial Loads	3	0.216420	4.4	0.0646	0.0721	0.0646	0.0768	0.0646	0.59
DRO: Only Loads	35	0.851310	17.4	0.0068	0.0243	0.0068	0.0666	0.0068	0.59
DRO: QIOs	35	0.851310	17.4	0.0068	0.0243	0.0068	0.0666	0.0068	0.59
DRI: QIOs	3	0.906000E-01	1.8	0.0058	0.0302	0.0058	0.0666	0.0058	0.87
TTS: QIOs	25	1.36152	27.8	0.0425	0.0545	0.0425	0.0914	0.0425	0.23
DRO1F1 QIOs	29	1.99097	40.6	0.0026	0.0687	0.0026	0.2080	0.0026	0.82

RESOURCE

RESOURCE	# arrivals	ARRIVALS PER SECOND		mean system rate	
		mean task rate	mean system rate	mean task rate	mean system rate
Memory Residency	3	0.61	0.01	0.01	0.09
CPU Timeslices	49	9.99	0.09	0.01	0.09

SPM-11M-PLUS V1.0 (for RSX-11M-PLUS V1.0)	4-DEC-81 16:41:37	Page 8	Task Summary
Another test ... No Noise Tasks!			
Initial Loads			
DR0: Only Loads	3		0.01
DR0: QIOs	35		0.07
DR1: QIOs	35		0.07
TT5: QIOs	3		0.01
DR0F1 QIOs	25		0.05
	29		0.06

Another test ... No Noise Tasks!

Generic TASK = ...TKB

Elapsed Times (seconds):

	40.17 minimum	40.17 mean	40.17 maximum	0.00 coef(var)	40.1703 total
Number of Task Invocations = 1					

Memory Times seconds: 40.17 0.7800E-03 0.4200E-03 % elapsed: 100.0 out 0.0 wait 0.0

Memory Size in KWords: 18.84 18.84 18.84 max

Space-Time Product: 2 756.945 378.472 mean KWS 0.99 coef(var) KWS incomplete 1

ACP Counts: 34 access/delete create/delete read/write 1 8 0 7 0 other

RESOURCE

Memory Residency

CPU Timeslices

Initial Loads

DRO: Ovlly Loads

DRO: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

DRI: QIOs

RESOURCE

Memory Residency

CPU Timeslices

Initial Loads

DRO: Ovlly Loads

DRO: QIOs

DRO: QIOs

DRO: QIOs

DRO: QIOs

DRO: QIOs

DRO: QIOs

DRO: QIOs

DRO: QIOs

arrivals

1

112

1

49

mean task rate

0.02

2.79

0.02

1.22

mean system rate

0.00

0.22

0.00

0.09

Another test ... No Noise Tasks!

DRO: QIOs	423	10.53	0.82
DRI: QIOs	659	16.41	1.27
DROOF1 QIOs	10	0.25	0.02
DROIF1 QIOs	44	1.10	0.08

Another test ... No Noise Tasks!

Elapsed Times (seconds): 1.303 minimum 21.36 mean 68.12 maximum 1.12 coef(var) 149.527 total

Number of Task Invocations = 7 (7 complete) (0 incomplete)

Memory Times seconds: 149.5 in 0.5840E-02 out 0.3290E-02 wait 100.0 in 0.0 out 0.0 wait

Memory Size in KWords: 11.34 min 21.22 mean 28.81 max

Space-Time Product: 24 3173.51 132.229 mean KWS 2.13 coef(var) KWS incomplete 6

ACP Counts: 98 access/deaccess create/delete read/write 5 directory 55 file attribute 5 extend 25 other 0

RESOURCE

	count	total time	% resident time	N/A	minimum	maximum	coef(var)	% resource
Memory Residency	8	149.521	N/A	1.3019	18.6902	62.1226	1.14	N/A
CPU Timeslices	652	69.0182	46.2	0.0003	0.1059	6.1600	3.48	13.3
Checkpoint R + W	2	0.143840	0.1	0.0503	0.0719	0.0935	0.30	0.0
Initial Loads	7	0.519490	0.3	0.0470	0.0742	0.1029	0.21	0.1
DR0: Gvly Loads	477	10.9495	7.3	0.0062	0.0230	0.0756	0.49	2.1
DR0: QIOs	918	18.5823	12.4	0.0043	0.0202	0.0756	0.57	3.6
DR1: QIOs	946	20.8997	14.0	0.0041	0.0221	0.0696	0.84	4.0
TTS: QIOs	223	32.1085	21.5	0.0004	0.1440	7.6509	4.56	6.2

RESOURCE

	count	total time	% resident time	N/A	minimum	maximum	coef(var)
Memory Residency	8	0.329000E-02	N/A	0.0004	0.0004	0.0004	0.04
Checkpoint R + W	2	0.188000E-02	0.0	0.0005	0.0009	0.0014	0.47
Initial Loads	7	0.115000E-01	0.0	0.0015	0.0016	0.0019	0.10
DR0: Gvly Loads	477	0.654970	0.4	0.0004	0.0014	0.0460	4.02
DR0: QIOs	918	1.50539	1.0	0.0004	0.0016	0.0583	4.10
DR1: QIOs	946	0.577430	0.4	0.0004	0.0006	0.0361	3.48
TTS: QIOs	223	0.109790	0.1	0.0004	0.0005	0.0007	0.08

RESOURCE

	count	total time	% resident time	N/A	minimum	maximum	coef(var)
Memory Residency	8	143.532	N/A	0.0002	17.9415	62.1230	1.22
Checkpoint R + W	2	0.145740	0.1	0.0508	0.0729	0.0949	0.30
Initial Loads	7	0.530990	0.4	0.0485	0.0759	0.1047	0.21
DR0: Gvly Loads	477	11.6045	7.8	0.0066	0.0243	0.1054	0.57
DR0: QIOs	918	20.0877	13.4	0.0047	0.0219	0.1137	0.72
DR1: QIOs	946	21.4771	14.4	0.0045	0.0227	0.0926	0.83
TTS: QIOs	276	32.2384	21.6	0.0004	0.1168	7.6514	5.08
DR00F1 QIOs	44	1.34208	0.9	0.0032	0.0305	0.1051	1.10
DR01F1 QIOs	165	8.67549	5.8	0.0024	0.0526	0.2925	1.01

RESOURCE

	# arrivals	mean task rate	ARRIVALS PER SECOND	mean system rate
	+	+	+	+

SPM-11M-PLUS V1.0 (for RSX-11M-PLUS V1.0)	4-DEC-81 16:41:37	Page 12	System-Wide Summary
Memory Residency	9	0.06	0.02
CPU Timeslices	652	4.36	1.26
Checkpoint R + W	2	0.01	0.00
Initial Loads	7	0.05	0.01
DR0: Only Loads	477	3.19	0.92
DR0: QIOs	918	6.14	1.77
DR1: QIOs	946	6.33	1.82
TT5: QIOs	276	1.85	0.53
DR00F1 QIOs	44	0.29	0.08
DR01F1 QIOs	165	1.10	0.32
Another test ... No Noise Tasks!			

Another test ... No Noise Tasks!

76. Buffers Read
8849. Records Read
0. Buffers With Overflow
0. Records Lost by Overflow

APPENDIX G
EXAMPLE SMPGEN

```

@SMPGEN
>
>*** Software Monitor Parameter GENERation command file ***
>
>*** Version 401 ***
>
>Copyright (C) 1981,1982 Digital Equipment Corporation, Maynard, Mass.
>
>
> Initialize default parameters
>
>* READ SAVED ANSWER FILE FOR DEFAULT ANSWERS? [Y/N]:N
>
> Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]:
>
> LEGAL PARAMETER CATEGORIES ARE :
>      HOOK          TASK          MODE          LABEL
>      BUFFER        DEVICE        SYSTEM        DONE
>
>* ENTER PARAMETER CATEGORY [S]: HOOK
>
> Select/deselect hook points
>
>* CPU USAGE HOOKS? [Y/N]:Y
>* MEMORY USAGE HOOKS? [Y/N]:N
>* I/O DEVICE USAGE HOOKS? [Y/N]:Y
>* USER INFORMATION HOOK? [Y/N]:N
>* LOAD OVERLAY HOOK? [Y/N]:Y
>* CALL DIRECTIVE PROCESSOR HOOK? [Y/N]:N
>
> Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]: BUFFER
>
> Setup buffer management parameter, type <esc> for help
>
>* ENTER BUFFER TYPE [D:DYNAMIC] [S R:0-7]:
>* ENTER NAME OF BUFFER MAIN (SYS) PARTITION [D:GEN] [S R:0-6]: BUFPAR
>* ENTER NAME OF BUFFER COMMON/REGION [D:TRCPAR] [S R:0-6]:
>* ENTER NUMBER OF BUFFERS [D R:2.-8. D:2.] : 4
>* ENTER BUFFER SIZE IN 512B BLOCKS [D R:1.-16. D:8.] :
>
> Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]: TASK
>
> Select/deselect tasks to be traced
>
>* DO YOU WANT TO TRACE ALL TASKS? [Y/N]:
>* DO YOU WANT TO SPECIFY THE TASKS NOT TO TRACE? [Y/N]:Y
>
> Enter one taskname per line. Terminate list by <CR>.
>
>* ENTER NAME OF TASK NOT TO TRACE [S]: F11ACP
>* ENTER NAME OF TASK NOT TO TRACE [S]: MCR...
>* ENTER NAME OF TASK NOT TO TRACE [S]: QMG...
>* ENTER NAME OF TASK NOT TO TRACE [S]: LPP0
>* ENTER NAME OF TASK NOT TO TRACE [S]: TKTN
>* ENTER NAME OF TASK NOT TO TRACE [S]: F11MSG
>* ENTER NAME OF TASK NOT TO TRACE [S]: MTAACP

```

```

>* ENTER NAME OF TASK NOT TO TRACE [S]: ...MCR
>* ENTER NAME OF TASK NOT TO TRACE [S]: DCL...
>* ENTER NAME OF TASK NOT TO TRACE [S]: ...DCL
>* ENTER NAME OF TASK NOT TO TRACE [S]: ...HEL
>* ENTER NAME OF TASK NOT TO TRACE [S]: ...BYE
>* ENTER NAME OF TASK NOT TO TRACE [S]: ...MOU
>* ENTER NAME OF TASK NOT TO TRACE [S]: ...DMO
>* ENTER NAME OF TASK NOT TO TRACE [S]: ...AT.
>* ENTER NAME OF TASK NOT TO TRACE [S]:
>
>; Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]: DEVICE
>* DO YOU WANT TO COLLECT DATA FOR ALL DEVICES? [Y/N]:N
>
>; Select device types to be traced
>
>* TRACE DISK (DIRECTORY DEVICE) I/O ACTIVITY? [Y/N]:Y
>* TRACE TAPE (SEQUENTIAL DEVICE) I/O ACTIVITY? [Y/N]:N
>* TRACE TERMINAL I/O ACTIVITY? [Y/N]:N
>* TRACE ALL OTHER DEVICES I/O ACTIVITY? [Y/N]:Y
>
>; Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]: MODE
>
>; Define data collection mode, type <esc> for help
>
>* ENTER MODE [D:MANUAL] [S R:0.-8.]:
>
>;
>;
>; LEGAL MODES ARE:
>;     MANUAL          AUTOSTOP          REPEAT
>;
>* ENTER MODE [D:MANUAL] [S R:0.-8.]: REPEAT
>* ENTER SAMPLE TYPE [S R:0-6]:
>
>; LEGAL SAMPLE TYPES ARE:
>;
>;     BUFFER          TIME
>;
>* ENTER SAMPLE TYPE [S R:0-6]: BUFFER
>* ENTER NUMBER OF BUFFERS PER SAMPLE [D R:1.-32000. D:16.]: 400
>* ENTER MAX. NUMBER OF SAMPLES [D R:1.-32000. D:1.]: 25
>* ENTER INTERVAL BETWEEN SAMPLES TIME UNIT [D:MIN] [S R:0-3]:
>* ENTER INTERVAL BETWEEN SAMPLES IN MIN [D R:1.-10000. D:1.]: 45
>
>; Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]: SYSTEM
>
>; Select system metrics
>
>* DO YOU WANT THE SYSTEM-WIDE METRICS? [Y/N]:Y
>* ENTER SYSTEM METRICS SAMPLING INTERVAL (SECS) [D R:15.-28800. D:60.]: 300
>
>; Select parameter category, type <esc> for help
>
>* ENTER PARAMETER CATEGORY [S]: LABEL
>
>; Define 1 to 63 character label strings, type <esc> for help
>
>* ENTER MPB LABEL [S R:0-63]: This is a sample SMPGEN.
>
>; Select parameter category, type <esc> for help
>
>;

```

```
>* ENTER PARAMETER CATEGORY [S]: DONE
>
> Write saved answer file and parameter source file
> Enter <esc> if no output file is desired
>
>* ENTER NAME OF SAVED ANSWER FILE [D:TRACE.CMD] [S]: SAMPLE.CMD
>* ENTER NAME OF PARAMETER SOURCE FILE [D:TRACE.PAR] [S]: SAMPLE.PAR
>MAC SAMPLE=HKMAC,SAMPLE.PAR
>TKB @PARBLD
>PIP SAMPLE.OBJ;*/DE,PARBLD.CMD;*
>@ <EOF>
>
```


APPENDIX H
EXAMPLE INSTALLATION

BACKUP RESTORE UTILITY V01

```
>tim 08:19 12/9/81
>run cnf
>
ENTER INPUT DEVICE [/CSR=NNN/VEC=NNN]: mm:
ENTER OUTPUT DEVICE [/CSR=NNN/VEC=NNN]: dr:/vec=254

>run fmt
>
FMT>dr:
```

** WARNING - DATA WILL BE LOST ON DRO: **

CONTINUE? [Y OR N]: y

START FORMATTING

OPERATION COMPLETE

```
FMT>^Z
run bad
>
BAD>dr:/li
BAD -- DRO: BAD BLOCK FOUND - LBN= 24863.
BAD -- DRO: TOTAL BAD BLOCKS= 1.
BAD>^Z
run bru
>
BRU>/bac:drsys/ve
FROM: mm:
TO: dr:
BRU - STARTING TAPE 1 ON MM0:
```

BRU - END OF TAPE 1 ON MM0:

BRU - STARTING VERIFY PASS TAPE 1 ON MM0:

BRU - END OF TAPE 1 ON MM0:

BRU - COMPLETED

BRU>^Z

```
RSX-11M-PLUS V01 BL6 124K
>RED DR:=SY:
>RED DR:=LB:
>RED DR:=SP:
>MOU DR:RSX11MPBL6
>@12,541BASTART.CMD
>;
>; RSX-11M-PLUS Distribution kit installation Phase-2
>;
>; The system disk now contains only a bootable baseline system. The
>; remainder of the distribution kit must be loaded from the
>; distribution tape by BRU (the BACKUP and RESTORE utility). These
>; instructions will guide you through the remainder of the
>; installation process.
>;
>; 1) It is important to specify the correct date and time. Use the
>; format of "DD-MMM-YY HH:MM".
>;
```

```

>* Please enter the date and time [S]: 9-dec-81 8:35
>TIME 9-dec-81 8:35
>
> 2) Now allocate checkpoint space for use by system utilities.
>
>ACS SY:/BLKS=1024.
>
> 3) We will now bring online all devices which were SYSGEN-ed into
> this baseline image and which exist in your hardware system.
>
> If there is hardware in your system occupying a control
> register address which conflicts with a standard DEC register
> assignments, it is possible that the "CON ONLINE ALL" command
> might cause a system crash. This is because a DIGITAL device
> driver is attempting to manipulate foreign hardware or
> the CSR address assignments in your system are different
> from those assumed in this baseline configuration. On
> the first pass through this file please answer "NO<CR>"
> to the following question. If such a mismatch exists and
> a crash does occur, reboot and answer "YES<CR>" to this
> question. Instructions will be provided to help
> circumvent the problem.
>
>* Do you need to do manual configuration? [Y/N]:n
>
>CON ONLINE ALL
>
> 4) We will now create directory entries for the files to be loaded
> from the common backup set. The UFD commands will not be displayed.
>
>
> 5) Please supply the logical name of the magnetic tape drive
> which should still contain the distribution tape. The
> name is of the form:
>
> MM0: or
> MT2: etc...
>
>* What is the device name of the distribution tape [S]: ^Z
>@ <EOF>
>ins $brv
>mov mm:/for/dens=800
>brv /noinit/sup/dens:800 mm: dr:
BRU - STARTING TAPE 1 ON MM0:
BRU - END OF TAPE 1 ON MM0:
BRU - COMPLETED
>dmo mm:
DMO -- TIO: DISMOUNTED FROM MM0: *** FINAL DISMOUNT ***
08:39:23 *** MM0: -- DISMOUNT COMPLETE

```

```

>@12,54]bastart
>
>      RSX-11M-PLUS Distribution Kit installation Phase-2
>
> The system disk now contains only a bootable baseline system. The
> remainder of the distribution kit must be loaded from the
> distribution tape by BRU (the BACKUP and RESTORE utility). These
> instructions will guide you through the remainder of the
> installation process.
>
> 1) It is important to specify the correct date and time. Use the
>     format of "DD-MMM-YY HH:MM".
>
> * Please enter the date and time [S]: 09-dec-81 08:42
>TIME 09-dec-81 08:42
>
> 2) Now allocate checkpoint space for use by system utilities.
>
>ACS SY:/BLKS=1024.
ACS -- CHECKPOINT FILE ALREADY IN USE ON DEVICE
>
> 3) We will now bring online all devices which were SYSGEN-ed into
>     this baseline image and which exist in your hardware system.
>
>     If there is hardware in your system occupying a control
>     register address which conflicts with a standard DEC register
>     assignments, it is possible that the "CON ONLINE ALL" command
>     might cause a system crash. This is because a DIGITAL device
>     driver is attempting to manipulate foreign hardware or
>     the CSR address assignments in your system are different
>     from those assumed in this baseline configuration. On
>     the first pass through this file please answer "NO<CR>"
>     to the following question. If such a mismatch exists and
>     a crash does occur, reboot and answer "YES<CR>" to this
>     question. Instructions will be provided to help
>     circumvent the problem.
>
> * Do you need to do manual configuration? [Y/N]:n
>
>CON ONLINE ALL
>
> 5) Please supply the logical name of the magnetic tape drive
>     which should still contain the distribution tape. The
>     name is of the form:
>
>         MM0: or
>         MT2: etc...
>
> * What is the device name of the distribution tape [S]: mm:
> * Is mm: written at 1600 BPI? [Y/N]:n
>
> 6) We will now mount that tape for access by BRU.
>
>MOU mm:/FOR/DENS=800
>
> 7) Ensure that the BRU task is installed.
>
>
> 8) We will now load the remainder of the distribution kit.
>
>TIME
08:43:22 09-DEC-81
>BRU /REW/DENS:800/BAC:COMMON/NOI mm: SY:
BRU - STARTING TAPE 1 ON MM0:

```

BRU - END OF TAPE 1 ON MMO:

BRU - MOUNT TAPE 2 ON MMO:

BRU - MOUNT TAPE 2 ON MMO:

BRU - STARTING TAPE 2 ON MMO:

BRU - END OF TAPE 2 ON MMO:

BRU - COMPLETED

>TIME

09:04:52 09-DEC-81

>DMD mm:

DMD -- TTD: DISMOUNTED FROM MMO: *** FINAL DISMOUNT ***

09:04:53 *** MMO: -- DISMOUNT COMPLETE

>:

>: 9) The installation of the RSX-11M-PLUS kit is now complete. The next
>: time this disk is booted, this command dialog will not be repeated.

>: We will now proceed through the normal startup command sequence.

>:

>SET /UIC=[2,54]

>INS \$PIP/INC=16000

>PIP [200,200]SYSGEN.CMD/RE/NV=SYSGEN.CMD

>PIP LB:[1,1]KITIDENT.DAT/AP=LB:[2,54]INSTALOG.DAT

>:

>:

>:

>SET /UIC=[200,200]

>:

PLEASE NOTE

>:

>: If you have not yet read the system release notes do so now before
>: attempting to perform a SYSGEN or utilize the new features of
>: this system.

>:

>CON ONLINE ALL

>RUN ERRLOG

ERL -- ERROR LOG INITIALIZED

>START/QUEUE/MANAGER

>:

>INS \$GMGPRT/TASK=PRT.../SLV=NO

09:05:33 GMP -- UNEXPECTED EXIT PROCESSOR LPP0

>INITIALIZE/QUEUE LPP0

09:05:35 GMP -- UNEXPECTED EXIT PROCESSOR BAPO

INI -- QUEUE ALREADY EXISTS

>INITIALIZE/PROCESSOR LPP0/FLAG:2/FORM:0/LOWER

INI -- DEVICE OFFLINE

>INITIALIZE/PROCESSOR BAPO

INI --*WARNING*-- ACCOUNTING NOT ACTIVE OR NOT INSTALLED

>ASSIGN/QUEUE PRINT LPP0

ASS -- PROCESSOR DOES NOT EXIST

>ASSIGN/QUEUE BATCH BAPO

>@ <EOF>

>@[200,200]sysgen

>: RSX-11M-PLUS SYSGEN BLB

>:

>: COPYRIGHT (C) 1979, 1980

>: DIGITAL EQUIPMENT CORP., MAYNARD MA., 01754

>:

>SET /UIC=[200,200]

>TIM

09:08:39 09-DEC-81

>:

>: EXPANDED COMMENTS PROVIDE A DESCRIPTION OF EVERY STEP IN THIS

```

>; SYSGEN COMMAND FILE.  ON THE OTHER HAND, SHORT COMMENTS
>; PROVIDE VIRTUALLY NO EXPLANATORY TEXT.
>;
>INS LB:[3,54]MACBIG/INC=20000
>INS LB:[3,54]LBR/INC=20000
>INS LB:[3,54]TKBBIG/INC=20000
>INS LB:[3,54]EDI
>INS LB:[3,54]VMR
>INS LB:[3,54]FLX/INC=20000
>* DO YOU WANT EXPANDED COMMENTS? [Y/N]:n
>* ENTER DEVICE_UNIT OF RELEASE DISK [DDU:] [S R:2-5]: dr:
>ASN dr:=IN:
>ASN dr:=OU:
>ASN dr:=SY:
>ASN dr:=LB:
>ASN dr:=WK:
>ASN dr:=TK:
>ASN dr:=BC:
>ASN dr:=LI:
>ASN dr:=OB:
>ASN dr:=EX:
>ASN dr:=MP:
>* DO YOU WANT TO APPLY AUTOPATCH ( <ESC> FOR DETAILS )? [Y/N]:y
>;
>; ENTER THE DEVICE_UNIT OF THE DRIVE ON WHICH THE AUTOPATCH TAPE
>; KIT IS LOADED OR "SY:" IF THE AUTOPATCH KIT HAS ALREADY
>; BEEN READ FROM THE DISTRIBUTION MEDIA.
>;
>* AUTOPATCH KIT DEVICE [S]: mm:
>SET /UIC=[1,11]
>;
>; THE PREVIOUS CONTENTS OF THE AUTOPATCH WORK DIRECTORY ARE NOW PURGED
>;
>PIP LB:[1,11]*.*/*DE/NM
>;
>; READ THE AUTOPATCH COMMAND FILE FROM THE AUTOPATCH TAPE
>;
>* IS mm: WRITTEN AT 1600 BPI? [Y/N]:n
>MOU mm://FOR/DENS=800
>BRU/REW/NOINIT/BAC:AUTOPATCH/DENS:800 mm: SY:
BRU - STARTING TAPE 1 ON MMO:

BRU - END OF TAPE 1 ON MMO:

BRU - COMPLETED

>SET /UIC=[200,200]
>PIP [200,200]/NV=[1,11]AUTOPATCH.CMD
>;
>; Autopatch 1E OCTOBER, 1981
>;
>; Create the FORTRAN interface Routine
>; object file UFD
>;
>UFD dr:[52,20]
>;
>; Create the SYSLIB object file UFD
>;
>UFD dr:[53,20]
>;
>; Copy the updated Files to dr:
>;
>BRU /REWIND/NOINIT/SUPERSEDE/BAC:UPDATES mm: dr:
BRU - STARTING TAPE 1 ON MMO:

```

BRU - END OF TAPE 1 ON MMO:

BRU - COMPLETED

```
>SET /UIC=[1,1]
>;
>; Insert the updated file system object files into SYSLIB
>;
>PIP FCS.OBS/NV=[50,20]*.OBJ
>LBR SYSLIB/RP=FCS.OBS
MODULE "CONTRL" REPLACED

MODULE "PNTMRK" REPLACED

>PIP FCS.OBS;O/DE
>;
>; Insert the updated FORTRAN Interface Routine object files into SYSLIB
>;
>PIP FIR.OBS/NV=[52,20]*.OBJ
>LBR SYSLIB/RP=FIR.OBS
MODULE ".CNMRK" REPLACED

MODULE ".DSCKP" REPLACED

MODULE ".ENAST" REPLACED

MODULE ".ENCKP" REPLACED

MODULE ".INAST" REPLACED

MODULE "SPSUB " REPLACED

MODULE ".WFLOR" REPLACED

>PIP FIR.OBS;O/DE
>;
>; Insert the updated SYSLIB object files into SYSLIB
>;
>PIP SLB.OBS/NV=[55,20]*.OBJ
>LBR SYSLIB/RP=SLB.OBS
MODULE "OVCTR " REPLACED

MODULE "$SUPL " REPLACED

MODULE "$SUPL " REPLACED

>PIP SLB.OBS;O/DE
>LBR SYSLIB/RP=QIOSYM.OBJ/SS
MODULE "QIOSYM" REPLACED

>;
>; Insert the updated macro files in RSXMAC.SML
>;
>LBR RSXMAC.SML/RP=MACRO.MCS
MODULE "SCAL$S" REPLACED

MODULE "VSDA$ " REPLACED

>PIP MACRO.MCS;*/DE
>* DO YOU WISH TO COPY THE CORRECTION AND PATCH FILES TO dr:? [Y/N]:y
>;
>; Create any necessary UFDs
>;
>UFD dr:[24,60]
```

```

>UFD dr:[32,60]
>UFD dr:[33,60]
>UFD dr:[34,60]
>UFD dr:[36,60]
>UFD dr:[45,60]
>UFD dr:[52,60]
>UFD dr:[62,60]
>UFD dr:[70,60]
>UFD dr:[106,60]
>UFD dr:[117,60]
>UFD dr:[200,60]
>;
>; Copy the correction and patch files to dr:
>;
>BRU /REW/NOINIT/NEW_VERSION/BAC:CORRECTIONS mm: dr:
BRU - STARTING TAPE 1 ON MMO:

BRU - END OF TAPE 1 ON MMO:

BRU - COMPLETED

>* DO YOU WISH TO PAUSE TO COPY OR EDIT ANY CORRECTION FILES? [Y/N]:n
>;
>; APPLICATION OF THE AUTOPATCH IS COMPLETE. A RECORD OF THIS
>; WILL BE APPENDED TO THE KITIDENT FILE.
>;
>SET /UIC=[1,1]
>DMO mm:
DMO -- TTO:    DISMOUNTED FROM MMO:    *** FINAL DISMOUNT ***
09:21:50 *** MMO: -- DISMOUNT COMPLETE
>SET /UIC=[200,200]
>* DO YOU WANT TO DO SYSGEN PART 1 (ASSEMBLE THE EXEC)? [Y/N]:n
>* DO YOU WANT TO DO SYSGEN PART 2 (BUILD THE EXEC AND PRIV TASKS)? [Y/N]:n
>* DO YOU WANT TO DO SYSGEN PART 3 (BUILD UTILITY TASKS)? [Y/N]:n
>ASN =
>@ <EOF>
>

>ufd dr:[12,2]
>set /uic=[12,2]

>mou mm:/for
>flx /rs=mm:pmocpy.cmd/do
>
>@pmocpy
>* IS SOURCE DEVICE A MAGTAPE? [Y/N]:y
>* IS DESTINATION DEVICE A MAGTAPE? [Y/N]:
>* ENTER SOURCE DEVICE (DDN:) [S]: mm:
>* ENTER DESTINATION DEVICE (DDN:) [S]: dr:
>* IS dr: MOUNTED? [Y/N]:y
>;
>; PLEASE SPECIFY THE GROUP AND MEMBER NUMBERS OF THE SOURCE KIT DIRECTORY.
>; THE ORIGINAL KIT DISTRIBUTION IS IN DIRECTORY [12,2].
>;
>* ENTER GROUP NUMBER OF SOURCE KIT DIRECTORY [O R:1-377 D:12]:
>* ENTER MEMBER NUMBER OF SOURCE KIT DIRECTORY [O R:1-377 D:2]:
>;
>; PLEASE SPECIFY THE GROUP AND MEMBER NUMBERS OF THE DIRECTORY TO WHICH
>; THE SPM-11M-PLUS KIT FILES WILL BE COPIED.
>;

```



```

>* ENTER GROUP NUMBER OF TARGET DIRECTORY [O R:1-377]: 12
>* ENTER MEMBER NUMBER OF TARGET DIRECTORY [O R:1-377]: 2
>SET /UIC=[12,2]
>FLX dr:/RS=mm:[12,2]SMPGEN.CMD,SPM.DOC/DO
>FLX dr:/RS/-RW=mm:[12,2]HKMAC.MAC,PMONSLP.CMD/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]SGNPER.COR,DRCRV.COR/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]DRDSP.COR,DREIF.COR,DREXP.COR/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]DRGIN.COR,DRGIO.COR,DRTBL.COR/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]LOWCR.COR,SYSCM.COR/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]REQSB.COR,LOADR.COR/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]IOSUB.COR,SYSXT.COR,TDSCH.COR/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]PMONPAT.CMD,DISPAT.PAT/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]COL.OLB,COLBLD.ODL,PMONBLD.CMD/DO/-RW
>FLX dr:/CO/RS/BL:60./-RW=mm:[12,2]CLF.TSK/DO/-RW
>FLX dr:/CO/RS/BL:75./-RW=mm:[12,2]OVL.TSK/DO/-RW
>FLX dr:/CO/RS/BL:175./-RW=mm:[12,2]DRS.TSK/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]OVL.MAP,OVL.OLB/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]CLF.MAP,DRS.CMD,MDRTKB.CMD/DO/-RW
>FLX dr:/RS/-RW=mm:[12,2]MDREIS.OLB,MDRFPP.OLB,DRS.MAP/DO/-RW
>* APPLY CORRECTIONS TO EXECUTIVE SOURCES? [Y/N]:y
>INS $SLP
>SET /UIC=[200,200]
>PIP RSXMC0.MAC/AP=[12,2]HKMAC.MAC
>SLP @ [12,2]SGNPER.COR
>SET /UIC=[11,10]
>SLP @ [12,2]DRCRV.COR
>SLP @ [12,2]DRDSP.COR
>SLP @ [12,2]REQSB.COR
>SLP @ [12,2]SYSXT.COR
>SLP @ [12,2]DREIF.COR
>SLP @ [12,2]DREXP.COR
>SLP @ [12,2]LOADR.COR
>SLP @ [12,2]DRGIO.COR
>SLP @ [12,2]IOSUB.COR
>SLP @ [12,2]SYSCM.COR
>SLP @ [12,2]LOWCR.COR
>SLP @ [12,2]DRTBL.COR
>SLP @ [12,2]DRGIN.COR
>SLP @ [12,2]TDSCH.COR
>REM ...SLP
>;
>;          *** NOTE ***
>;
>;          DO NOT APPLY THE FOLLOWING PATCH TO F11ACP IF YOU ARE GOING
>;          TO SELECT THE FCPR03 (DEFAULT) F11ACP AT SYSGEN.
>;
>* APPLY PATCH TO F11ACP DISPATCHER MODULE (DISPAT)? [Y/N]:y
>INS $PAT
>SET /UIC=[12,2]
>MAC DISPAT.POB=HKMAC.MAC,DISPAT.PAT
>LBR DISPAT.OBJ;1=[1,24]FCP/EX:DISPAT
>PAT DISPAT.OBJ;2=DISPAT.OBJ;1/CS:044214,DISPAT.POB/CS:017664
>* WAS THE CHECKSUM SUCCESSFUL? [Y/N]:y
>SET /UIC=[1,24]
>LBR FCP/RP=[12,2]DISPAT.OBJ;2
MODULE "DISPAT" REPLACED

>REM ...PAT
>SET /UIC=[12,2]
>@ <EOF>
>dmo mm:
DMO -- T10:      DISMOUNTED FROM MMO:      *** FINAL DISMOUNT ***
09:41:19 *** MMO: -- DISMOUNT COMPLETE
>

```

```

>
>set /uic=[200,200]
>@sysgen
>; RSX-11M-PLUS  SYSGEN  BL6
>;
>; COPYRIGHT (C)  1979, 1980
>; DIGITAL EQUIPMENT CORP., MAYNARD MA., 01754
>;
>SET /UIC=[200,200]
>TIM
09:42:23 08-DEC-81
>;
>; EXPANDED COMMENTS PROVIDE A DESCRIPTION OF EVERY STEP IN THIS
>; SYSGEN COMMAND FILE.  ON THE OTHER HAND, SHORT COMMENTS
>; PROVIDE VIRTUALLY NO EXPLANATORY TEXT.
>;
>* DO YOU WANT EXPANDED COMMENTS? [Y/N]:n
>* ENTER DEVICE UNIT OF RELEASE DISK [DDU:] [S R:2-5]: dr:
>ASN dr:=IN:
>ASN dr:=OU:
>ASN dr:=SY:
>ASN dr:=LB:
>ASN dr:=WK:
>ASN dr:=TK:
>ASN dr:=BC:
>ASN dr:=LI:
>ASN dr:=OB:
>ASN dr:=EX:
>ASN dr:=MP:
>* DO YOU WANT TO APPLY AUTOPATCH ( <ESC> FOR DETAILS )? [Y/N]:n
>* DO YOU WANT TO DO SYSGEN PART 1 (ASSEMBLE THE EXEC)? [Y/N]:y
>;
>; RSX-11M-PLUS  SYSGEN PART 1  <VERSION 1.0>
>;
>;      DETERMINE SYSTEM FEATURES AND ASSEMBLE THE EXECUTIVE
>;
>* HAVE YOU ALREADY RUN SGN? [Y/N]:n
>SET /UIC=[11,24]
>PIP IN:*.OBJ;*/DE/NM,*.TTY;*,*.OBS;*
>SET /UIC=[1,24]
>PIP IN:[11,10]RSXMC.MAC;*/DE/NM,SYSTB;*
>PIP IN:[11,24]RSXASM.CMD;*/DE/NM
>PIP IN:[1,24]RSXBLD.CMD;*/DE/NM,DSP11M.CMD;*,LDR11M.CMD;*,RSX11M.CMD;*
>PIP IN:[1,24]DIR11M.CMD;*/DE/NM,DIRCOM.CMD;*,DIR.CMD;*
>SET /UIC=[200,200]
>PIP RSXMC.SAV/PU/NM,SGNSYM.CMD
>;
>;      START SGN --  VERSION 06.00B
>;
>;      BEGIN PART A OF SYSGEN PHASE 1
>;
>;
>; EXPANDED COMMENTS RESULT IN THE DISPLAY OF EXPLANATORY TEXT
>; PRECEDING EACH QUERY.  SHORT COMMENTS OMIT THE TEXT.  EXPANDED
>; COMMENTS ARE AVAILABLE ON A PER QUERY BASIS BY ENTERING ESCAPE
>; OR ALTMODE FOLLOWING THE DISPLAY OF THE QUERY.  FOR QUESTIONS
>; THAT CONTAIN A RANGE SPECIFICATION, AN ANSWER THAT FALLS
>; WITHIN THE RANGE MUST PRECEDE THE ESCAPE.  THIS ANSWER WILL
>; HAVE NO EFFECT ON THE QUESTION WHEN IT IS ASKED AGAIN.
>;
>; TO EXIT FROM THE SGN INDIRECT FILES AT ANY TIME, TYPE CNTRL/Z.
>;
>; IF YOU ARE UNSURE OF THE ANSWER TO A QUESTION FOR WHICH A DE-
>; FAULT ANSWER EXISTS, USE THE DEFAULT ANSWER.

```

```

>;
>* LONG DIALOGUE DESIRED FOR EXECUTIVE/PROCESSOR OPTIONS? [Y/N]:n
>* DO YOU WANT THE STANDARD FULL-FUNCTIONALITY EXECUTIVE? [Y/N]:y
>* ASSEMBLY LISTING FILES DESIRED? [Y/N]:y
>* WHAT IS TO BE THE LISTING DEVICE? [DDU:] [S R:2-5]:.dr:
>* SPOOL LISTING FILES? [Y/N]:n
>;
>; THE TASK BUILDER MAP OF THE EXECUTIVE WILL BE SAVED ON THE SYSTEM DISK.
>;
>; BEGIN EXECUTIVE SERVICE OPTIONS.
>;
>* DO YOU WANT DECNET SUPPORT? [Y/N]:n
>* ARE YOU PLANNING TO INCLUDE A USER WRITTEN DRIVER? [Y/N]:n
>* DO YOU WANT TO INCLUDE THE EXECUTIVE DEBUGGING TOOL (XDT)? [Y/N]:y
>* ENTER CRASH NOTIFICATION DEVICE CSR ADDRESS [D R:160000-177700 D:177564]:
>* ON WHAT DEVICE DO YOU WISH CRASH DUMPS TO BE WRITTEN [DD:] [S R:2-3]: mm:
>* WHAT IS THE PHYSICAL UNIT NUMBER OF THE CRASH UNIT [D R:0-7 D:0]: 0
>;
>; END EXECUTIVE SERVICE OPTIONS.
>;
>; BEGIN PROCESSOR OPTIONS.
>;
>* ENTER MEMORY SIZE (1024 WORD BLOCKS) [D R:128.-1920. D:256.]: 320
>* DO YOU WANT FLOATING POINT PROCESSOR SUPPORT? [Y/N]:y
>* IS YOUR SYSTEM CLOCK PROGRAMMABLE (KW11-P)? [Y/N]:n
>* IS YOUR LINE FREQUENCY 50 HERTZ? [Y/N]:n
>;
>; END PROCESSOR OPTIONS.
>;
>; END PART A OF SYSGEN PHASE 1
>;
>* DO YOU WANT TO CONTINUE WITH PART B OF THIS SGN? [Y/N]:y
>;
>; WE WILL NOW CONTINUE WITH PART B OF SYSGEN PHASE 1
>;
>; BEGIN PERIPHERAL OPTIONS.
>;
>* EXPANDED COMMENTS DESIRED FOR PERIPHERAL OPTIONS? [Y/N]:n
>* HOW MANY RH CONTROLLERS DO YOU HAVE [D R:0-17 D:4]: 2
>;
>* HOW MANY RP04/05/06 DISK PACK DRIVE UNITS DO YOU HAVE? [D R:0-77]: 0
>;
>* HOW MANY RM02/03/05 DISK DRIVE UNITS DO YOU HAVE? [D R:0-77]: 2
>* ARE ANY OF THE UNITS DUAL-ACCESS? [Y/N]:n
>;
>* WHAT IS THE PHYSICAL UNIT NUMBER OF DR0: [D R:0-7 D:0]: 0
>* TO WHICH RH CONTROLLER IS DR0: CONNECTED [S]: A
>* ENTER THE DISK TYPE FOR DR0: (RM02,RM03,RM05) [S]: rm03
>;
>* WHAT IS THE PHYSICAL UNIT NUMBER OF DR1: [D R:0-7 D:1]:
>* TO WHICH RH CONTROLLER IS DR1: CONNECTED [S]: A
>* ENTER THE DISK TYPE FOR DR1: (RM02,RM03,RM05) [S]: rmo\o\03
>;
>* HOW MANY RS03/04 DISK DRIVE UNITS DO YOU HAVE? [D R:0-77]: 0
>;
>* HOW MANY TU16/45/77/TE16 TAPE UNITS DO YOU HAVE? [D R:0-100]: 1
>* HOW MANY MASTER DRIVE UNITS ARE THERE [D]: 1
>;
>* WHAT IS THE PHYSICAL UNIT NUMBER OF MM0: [D R:0-7 D:0]:
>* TO WHICH RH CONTROLLER IS MM0: CONNECTED [S]: B
>;
>; THE AVAILABLE SLAVE UNITS HAVE ALREADY BEEN ALLOCATED TO OTHER
>; MASTER DRIVES. MASTER DRIVE MM0: IS THEREFORE ASSUMED TO HAVE
>; NO SLAVE DRIVE UNITS ATTACHED TO IT.
>;

```

H-11

```

>PIP [200,200]RSXMC.MAC/PU/NM
>LBR IN:RSX11M/CR:256.:1800.:256.=OU:[11,24]RSX11M.OBJ
>LBR IN:TTDRV/CR:40.:392.:128.=OU:[11,24]TTDRV.OBS
>;
>; END SYSGEN PHASE 1
>;
>* DO YOU WANT TO DO SYSGEN PART 2 (BUILD THE EXEC AND PRIV TASKS)? [Y/N]:y
>;
>; IF YOU HAVE ALREADY DONE A SYSGEN ON THE TARGET DISK, A
>; SYSTEM MAY RESIDE IN [1,54]. YOU MAY WANT TO SAVE THAT
>; SYSTEM.
>;
>* HAVE YOU DONE A SYSGEN ON THE TARGET DISK BEFORE? [Y/N]:n
>;
>; RSX-11M-PLUS SYSGEN PART 2 <VERSION 1.0>
>;
>; BUILD THE EXECUTIVE AND ALL PRIVILEGED TASKS
>;
>* HAVE YOU ALREADY BUILT THE EXEC? [Y/N]:n
>SET /UIC=[1,24]
>TIM
10:33:20 09-DEC-81
>TKB @RSXBLD
>TIM
10:35:04 09-DEC-81
>* DO YOU WANT TO CONTINUE AND BUILD THE TASKS NOW? [Y/N]:y
>SET /UIC=[1,24]
>* WHAT FCP DO YOU WANT (DEFAULT = FCPR03)? [S]: FCPR02
>* DO YOU NEED TO EDIT ANY OF THE BUILD COMMAND FILES? [Y/N]:n
>SET /UIC=[1,54]
>PIP SYSVMR.CMD=VMRTTY.CMD/AP
>* DO YOU WANT TO EDIT SYSVMR.CMD? [Y/N]:n
>;
>; NOW WE BUILD THE PRIVILEGED TASKS AND LOADABLE DRIVERS
>;
>ASN NL:=MP:
>* DO YOU WANT THE MAPS OF THE PRIVILEGED TASKS? [Y/N]:n
>SET /UIC=[200,200]
>;
>; MAPS OF THE DRIVERS WILL BE SAVED ON THE SYSTEM DISK
>;
>;
>; IF ANY OF THE DRIVERS YOU ARE BUILDING HAVE A LOADABLE DATA BASE,
>; YOU NEED TO EDIT THE FILE DRIVERS.BLD TO ADD THE NAME OF
>; YOUR DATA BASE MODULE AFTER THE NAME OF YOUR DRIVER MODULE.
>;
>; FOR EXAMPLE:
>;
>; *PL LB:XXDRV
>; SY:[1,24]RSX11M/LB:XXDRV
>; *AP :XXTAB
>; SY:[1,24]RSX11M/LB:XXDRV:XXTAB
>; *ED
>;
>* DO YOU NEED TO EDIT DRIVERS.BLD? [Y/N]:n
>SET /UIC=[1,24]
>TKB @[200,200]DRIVERS.BLD
>PIP SY:[1,34]TTCOM.TMP;*/DE/NM
>SET /UIC=[200,200]
>;
>; BLDDRV SAVES THE COMMAND FILE USED TO BUILD THE DRIVER(S)
>; IN [200,200], AND NAMES IT DRIVERS.BLD
>;
>SET /UIC=[1,24]
>TIM
10:39:52 09-DEC-81

```

```

>TKB @ACCB LD
>TKB @ACNBLD
>TKB @ACSBLD
>TKB @BOOBLD
>TKB @BPRBLD
>TKB @BROBLD
>TKB @BYEBLD
>TKB @COTBLD
>TKB @DCLBLD
>TKB @DCLDSPBLD
>TKB @DMOBLD
>TKB @ERFBLD
>TKB @ERLFSLBLD
>TKB @FCPRO2BLD
>TKB @FXRBLD
>TKB @F11MSGBLD
>TKB @HELFSLBLD
>TKB @HRCBLD
>TKB @INDFSLBLD
>TKB @INIBLD
>TKB @INSFSLBLD
>TKB @LOABLD
>TKB @LPPBLD
>TKB @MCRBLD
>TKB @MCRDSPBLD
>TKB @MOUBLD
>TKB @MTABLD
>TKB @PMDFSLBLD
>TKB @QMGFSLBLD
>TKB @RMDBLD
>TKB @SAVBLD
>TKB @SHABLD
>TKB @SHFBLD
>TKB @SHTBLD
>TKB @SYLFSLBLD
>TKB @TKNBLD
>TKB @UFDBLD
>TKB @UNLBLD
>TIM
11:16:51 09-DEC-81
>ASN LB:=OU:
>SET /UIC=[1,54]
>ASN LB:=SY:
>PIP OU:RSX11M.SYS/CO/NV/BL:498.=IN:RSX11M.TSK
>VMR @SYSVMR

```

LOADING EXECUTIVE DATA SPACE
DATA SPACE LOADING COMPLETED

VMR -- INSTALLED TASKS OR COMMONS MAY NO LONGER FIT IN PARTITION

SET /TOP=FCPPAR:--*

VMR -- INSTALLED TASKS OR COMMONS MAY NO LONGER FIT IN PARTITION

SET /TOP=DRVPAR:--*

```

SYSVAR 026320 00202500 00012000 MAIN
LDRPAR 026254 00214500 00003000 MAIN
        024644 00214500 00003000 TASK [...LDR]
FCPPAR 026210 00217500 00051500 MAIN
SECPOL 026040 00271200 00010000 SEC POOL
DRVPAR 025774 00301200 00074200 MAIN
        025730 00301200 00001400 DRIVER (RD:)
        025664 00302600 00001300 DRIVER (UT:)
        025620 00304100 00000100 DRIVER (NL:)
        025554 00304200 00003100 DRIVER (CD:)
        025510 00307300 00004600 DRIVER (MM:)
        025444 00314100 00004700 DRIVER (DR:)
        025400 00321000 00035400 RW COM !TTCOM !
        025314 00356400 00017000 DRIVER (TT:)

```

```

GEN      025250 00375400 02002400 MAIN
...LDR 09.01 025060 LDRPAR 248. 00003000 LB0:-00156234 FIXED
TKTN   03.16 113160 SYSPAR 248. 00010000 LB0:-00161016
...RMD 00.01 114220 GEN      225. 00025300 LB0:-00161554
F11MSG 00011 116110 GEN      200. 00040000 LB0:-00157424
MTAACP 0010 114534 GEN      200. 00013500 LB0:-00161445
SHE... 01.02 113600 GEN      200. 00013500 LB0:-00160607
...DCL 0113A 116740 GEN      160. 00040000 LB0:-00157636
DCL... 0113A 116634 SYSPAR 160. 00010000 LB0:-00157451
...DMO 23.02F 116530 GEN      160. 00040000 LB0:-00157006
...MCR 01.01 115050 GEN      160. 00040000 LB0:-00161127
MCR... 3.3 114744 SYSPAR 160. 00010000 LB0:-00161327
...MOU 23.03 114640 GEN      160. 00040000 LB0:-00160500
F11ACP M0320 026104 FCPPAR 149. 00051500 LB0:-00157347
COT... 00.02 117044 GEN      145. 00015700 LB0:-00156756
ERRLOG 05 116320 GEN      145. 00016200 LB0:-00157126
HRC... 02 115700 GEN      140. 00053400 LB0:-00160422
PMD... 3.1 114430 GEN      140. 00023700 LB0:-00160530
SYSLOG 01.01 113264 GEN      130. 00020400 LB0:-00162070
SHF... 03.01 113474 SYSPAR 105. 00010000 LB0:-00160635
FXR... 01 116214 GEN      100. 00003100 LB0:-00157147
...INI 20.01 115470 GEN      100. 00040000 LB0:-00160644
...INC 4 115364 GEN      100. 00030000 LB0:-00157576
...SAV 04.38 114114 GEN      100. 00040000 LB0:-00161704
...UFD 00410 113054 GEN      100. 00040000 LB0:-00160722
BAPO 01.01 117464 GEN      80. 00043000 LB0:-00157157
GMG... 01.04 114324 GEN      75. 00027000 LB0:-00160556
LPPO 1.06A 115154 GEN      70. 00014200 LB0:-00161037
...ACC 01.01 024524 GEN      65. 00036700 LB0:-00156527
...AT. 04.27 115574 GEN      65. 00057700 LB0:-00157525
...SHA 01.01 114010 GEN      65. 00025500 LB0:-00161762
...ERF 04 116424 GEN      61. 00040000 LB0:-00157024
SHC... 01.01 113704 GEN      60. 00051700 LB0:-00161363
ACNT 3 024420 GEN      50. 00055700 LB0:-00157033
...ACS 01.1 117674 GEN      50. 00040000 LB0:-00156634
...BOO 04.12 117570 GEN      50. 00040000 LB0:-00156654
...BRD 03A 117360 GEN      50. 00031200 LB0:-00156703
...BYE 01.93 117254 GEN      50. 00014400 LB0:-00156737
...CON 01.01 117150 GEN      50. 00123200 LB0:-00040355
...HEL 01.25 116004 GEN      50. 00030200 LB0:-00157472
...LOA 03.3 115260 GEN      50. 00040000 LB0:-00160740
SHUTUP 02.01 113370 GEN      50. 00010700 LB0:-00160707
...UNL 03.4 112750 GEN      50. 00040000 LB0:-00162113
RHA OFL CPA CSR=176700 VEC=254 PRI=5
RHB OFL CPA CSR=172440 VEC=224 PRI=5
YLA OFL CPA CSR=177560 VEC=60 PRI=5
YZA OFL CPA CSR=160070 VEC=300 PRI=5
YZB OFL CPA CSR=160100 VEC=310 PRI=5
DR0: RHA0: OFL DRIVER
DR1: RHA1: OFL DRIVER
MM0: RHB0: OFL DRIVER
CO0: OFL DRIVER
TT0: YLA0: OFL DRIVER
TT1: YZA0: OFL DRIVER
TT2: YZA1: OFL DRIVER
TT3: YZA2: OFL DRIVER
TT4: YZA3: OFL DRIVER
TT5: YZA4: OFL DRIVER
TT6: YZA5: OFL DRIVER
TT7: YZA6: OFL DRIVER
TT10: YZA7: OFL DRIVER
TT11: YZB0: OFL DRIVER
TT12: YZB1: OFL DRIVER
TT13: YZB2: OFL DRIVER
TT14: YZB3: OFL DRIVER
TT15: YZB4: OFL DRIVER

```

```

TT16:  YZB5:      OFL    DRIVER
TT17:  YZB6:      OFL    DRIVER
TT20:  YZB7:      OFL    DRIVER
NLO:   OFL    DRIVER
VTO:   OFL    DRIVER
RDO:   ONL    DRIVER
POOL=1200:13428.:13444.:2025
>* DO YOU WANT TO DO SYSGEN PART 3 (BUILD UTILITY TASKS)? [Y/N]:n
>ASN =
>@ <EOF>
>
>hdd [1.54]rsx11m
XDT: 6

XDT>
RSX-11M-PLUS V01 BL6

>
TIM 11:30 9-DEC-81
>TIM
11:30:02 09-DEC-81
>SAV /WB/NOCON

RSX-11M-PLUS V01 BL6 320K
>RED DR:=SY:
>RED DR:=LB:
>RED DR:=SP:
>MOD DR:RSX11MPBL6
>@DR:[1.2]STARTUP
>;
>;
>; IF you have not yet read the system release notes do so now before
>; attempting to perform a SYSGEN or utilize the new features of
>; this system.
>;
>* Please enter time and date (HH:MM MM/DD/YY) [S]: 11:31
>* Please enter time and date (HH:MM MM/DD/YY) [S]: 11:31 12/9/81
>TIME 11:31 12/9/81
>ACS SY:/BLKS=1024.
>CON ONLINE ALL
>RUN ERRLOG
ERL -- ERROR LOG INITIALIZED
>START/QUEUE/MANAGER
>
>INS $QMGPR/TASK=PRT.../SLV=NO
11:31:07 QMG -- UNEXPECTED EXIT PROCESSOR BAPO
>INITIALIZE/QUEUE LPPO
INI -- QUEUE ALREADY EXISTS
>START/ACCOUNTING
>INITIALIZE/PROCESSOR LPPO/FLAG:2/FORM:0/LOWER
INI -- DEVICE DOES NOT EXIST
>INITIALIZE/PROCESSOR BAPO
>ASSIGN/QUEUE PRINT LPPO
ASS -- PROCESSOR DOES NOT EXIST
>ASSIGN/QUEUE BATCH BAPO
>@ <EOF>

```



```

>SET /UIC
UIC=[12,2]
>@PMONBLD
>* ENTER KW11-P VECTOR ADDRESS [O R:100-774 D:104]:
>* ENTER KW11-P CSR ADDRESS [O R:160000-177778 D:172540]:
>* WILL THE BUFFERS BE IN A STATIC PARTITION? [Y/N]:n
>TKB @COLBLD
>TKB @TRMBLD
>TKB @COLMSGBLD
>;
>; Taskbuild the MDR, OVL & CLF Programs.
>;
>* EIS or FPP <FPP> [S]: eis
>; you must specify either "EIS" or "FPP" without the quotes.
>* EIS or FPP <FPP> [S]: EIS
>;
>* Do you wish to build the task OVL? [Y/N]:n
>;
>* Do you wish to build the task CLF? [Y/N]:n
>;
>* Do you wish to build the task MDR? [Y/N]:y
>* Build Overlaid Version of MDR.TSK? [Y/N]:n
>* Do you want to edit the MDR taskbuild command file? [Y/N]:n
>TKB @MDREISNOV
>@ <EOF>
>

```

