bmcsoftware

CMF® MONITOR

Batch User Guide and Reference

Supporting

CMF MONITOR 5.6

March 2007

ACTIVATE BUSINESS WITH THE POWER OF I.T.™

www.bmc.com

Contacting BMC Software

You can access the BMC Software website at http://www.bmc.com. From this website, you can obtain information about the company, its products, corporate offices, special events, and career opportunities.

United States and Canada

Address	BMC SOFTWARE INC 2101 CITYWEST BLVD HOUSTON TX 77042-2827 USA	**Telephone**	713 918 8800 or 800 841 2031	**Fax**	713 918 8000

Outside United States and Canada

Telephone	(01) 713 918 8800	**Fax**	(01) 713 918 8000

Customer support

You can obtain technical support by using the BMC Software Customer Support website or by contacting Customer Support by telephone or e-mail. To expedite your inquiry, see “Before contacting BMC.”

Support website

You can obtain technical support from BMC 24 hours a day, 7 days a week at http://www.bmc.com/support_home. From this website, you can

- read overviews about support services and programs that BMC offers
- find the most current information about BMC products
- search a database for issues similar to yours and possible solutions
- order or download product documentation
- report an issue or ask a question
- subscribe to receive proactive e-mail alerts when new product notices are released
- find worldwide BMC support center locations and contact information, including e-mail addresses, fax numbers, and telephone numbers

Support by telephone or e-mail

In the United States and Canada, if you need technical support and do not have access to the web, call 800 537 1813 or send an e-mail message to customer_support@bmc.com. (In the subject line, enter **SupID:*yourSupportContractID***, such as SupID:12345.) Outside the United States and Canada, contact your local support center for assistance.

Before contacting BMC

Have the following information available so that Customer Support can begin working on your issue immediately:

- product information
  - product name
  - product version (release number)
  - license number and password (trial or permanent)
- operating system and environment information
  - machine type
  - operating system type, version, and service pack or other maintenance level such as PUT or PTF
  - system hardware configuration
  - serial numbers
  - related software (database, application, and communication) including type, version, and service pack or maintenance level
- sequence of events leading to the issue
- commands and options that you used
- messages received (and the time and date that you received them)
  - product error messages
  - messages from the operating system, such as `file system full`
  - messages from related software

Contents

Chapter 8 Analyzer reports 327

Figures

Tables

About this book

This book contains detailed information about the CMF MONITOR product and is intended for data center operators, managers, and system programmers who use CMF MONITOR to gather and produce information to help perform system tuning tasks, and improve the performance and efficiency of their systems.

To use this book, you should be familiar with Multiple Virtual Storage (MVS) systems, job control language (JCL), the Interactive System Productivity Facility (ISPF), and how to respond to ISPF panels.

Related documentation

BMC products are supported by several types of documentation:

- online and printed books
- online Help
- release notes and other notices

Category	Document	Description
installation documents	*OS/390 and z/OS Installer Guide*	provides information about the installation of BMC products on OS/390 and z/OS systems
	Installation Checklist Generator (ICG)	creates a customized checklist of installation steps and information about installation requirements such as software, storage, and system requirements for your combination of products
core documents	*CMF MONITOR Customization Guide*	describes the installation procedures that are unique to CMF MONITOR
	CMF MONITOR Online Getting Started *CMF MONITOR Online User Guide*	explains how to use the MAINVIEW window interface and the CMF MONITOR Online views, as well as how to interpret the information that is presented
	CMF MONITOR CMFMON User Guide	explains how to use the CMFMON online facility and write facility, as well as how to generate CMFMON batch reports

Category	Document	Description
core documents (*continued*)	*DSO User Guide and Reference*	explains how to use the DATA SET OPTIMIZER (DSO) batch report control statements, and how to interpret the report information
	MAINVIEW Administration Guide	provides information about MAINVIEW operations, targets, single system image contexts, data sets, view customization, and diagnostic facilities
	MAINVIEW User Guide (formerly *Using MAINVIEW*)	provides information about working with MAINVIEW products in windows mode, in full-screen mode, and from MAINVIEW Explorer
	MAINVIEW Reference Summary (formerly *MAINVIEW Quick Reference*)	introduces the MAINVIEW family of products and lists the commands that are used to manage the MAINVIEW windows environment
	MAINVIEW Common Customization Guide	provides instructions for manually customizing the MAINVIEW environment for your product
	MAINVIEW Security Reference Manual (formerly *Implementing Security for MAINVIEW Products*)	provides complete information about the MAINVIEW security interfaces for windows mode, full-screen mode, and MAINVIEW Alternate Access
	MAINVIEW Security Guide	describes the basics of how to define security for MAINVIEW products with an external security manager (ESM)
	MAINVIEW Alternate Access Implementation and User Guide	explains how to configure, start, and stop VTAM and EXCP AutoLogon sessions to access MAINVIEW products without an active TSO communication support
	MAINVIEW Alarm Management Guide	explains how to create and install alarm definitions that indicate when exceptions occur in a sysplex
notices	*CMF MONITOR Release Notes*	provides information about what's new and what's changed in the current release of CMF MONITOR, and includes last-minute product information

Online and printed books

Like most BMC documentation, this book is available in printed and online formats. Visit the Customer Support page at http://www.bmc.com/support_home to request additional printed books or to view online books and notices (such as release notes and technical bulletins). Some product shipments also include the online books on a documentation CD.

NOTE

Online books are formatted as Portable Document Format (PDF) or HTML files. To view, print, or copy PDF books, use the free Acrobat Reader from Adobe Systems. If your product installation does not install the reader, you can obtain the reader at http://www.adobe.com.

Online Help

The CMF MONITOR product includes online Help. In the CMF MONITOR ISPF interface, access Help by pressing **PF1** from any ISPF panel.

To access the Messages & Codes application from any CMF MONITOR panel, type **MSG** on the **COMMAND** line.

Release notes and other notices

Printed release notes accompany each BMC product. Release notes provide current information such as

- updates to the installation instructions
- last-minute product information

In addition, BMC sometimes provides updated product information between releases (in the form of a flash or a technical bulletin, for example), and maintenance announcements accompany maintenance releases. The latest versions of the release notes and other notices are available on the Web at http://www.bmc.com/support_home.

Conventions

This section provides examples of the conventions used in this book.

General conventions

This book uses the following general conventions:

Item	Format	Example
information that you are instructed to type	bold	Type **CMFMON** and press **Enter**.
specific (standard) keyboard key names	bold	Press **Enter**.
field names, option names, directories, file names	bold	In the **COMMAND** field of the DEVV screen...
web addresses, e-mail addresses	underlined blue text	The BMC Software home page is at www.bmc.com.
view names, command names, nonspecific key names, keywords	uppercase	Use the HELP function key. On the **COMMAND** line of the DEVV screen...
commands that can be shortened	required letters capitalized	To enter delta mode, type **DElta**.
code examples, syntax statements, system messages, screen text	code typeface	`//STEPLIB DD`
emphasized words, new terms, variables	italics	The instructions that you give to the software are called *commands*. In this message, the variable *fileName* represents the file that caused the error.

This book uses the following types of special text:

NOTE

Notes contain important information that you should consider.

WARNING

Warnings alert you to situations that could cause problems, such as loss of data, if you do not follow instructions carefully.

Syntax conventions

Syntax statements appear in Courier. The following example shows a sample syntax statement:

```
COMMAND KEYWORD1 [KEYWORD2|KEYWORD3] KEYWORD4={YES|NO}
       fileName...
```

The following table explains conventions for syntax statements and provides examples:

Item	Example
Items in italic type represent variables that you must replace with a name or value.	`dtsbackup` *controlDirectory*
Brackets indicate a group of options. You can choose at least one of the items in the group, but none of them is required. Do not type the brackets when you enter the option. A comma means that you can choose one or more of the listed options. You must use a comma to separate the options if you choose more than one option.	`[tableName, columnName, field]`
Braces enclose a list of required items. You must enter at least one of the items. Do not type the braces when you enter the item.	{*DBDName* | *tableName*}
A vertical bar means that you can choose only one of the listed items. In the example, you would choose either *commit* or *cancel*.	`{commit	cancel}`
An ellipsis indicates that you can repeat the previous item or items as many times as necessary.	*columnName* . . .

Change bars

Change bars show where substantive technical changes have been made to the document since its previous publication. These changes include clarifications or corrections to existing information and changes that provide new information that corresponds to product changes. Editorial and formatting changes or typographical errors that were fixed are not noted unless these updates significantly affect your use of the product.

Part 1

CMF MONITOR user guide

This part presents the following topics:

Chapter

1

About CMF MONITOR

CMF MONITOR is a performance monitoring, statistics gathering, and reporting system developed and maintained by BMC Software. Statistics gathered by CMF MONITOR can be compiled and presented either in batch reports or through windowed online displays.

Other BMC Software products use the *Comprehensive Management Facility* (CMF) as a platform for collecting and reporting system performance information. BMC Software has developed a complete system of products based on CMF that manage and evaluate the demands on computer center resources and throughout an entire computing enterprise.

By using products that comply with the CMF architecture, data center managers can better control system performance, quickly respond to increased computer service needs, and plan for new capacity to meet service level objectives for projected corporate growth.

This chapter presents an overview of CMF MONITOR components, describes the integration with other BMC Software products and the MAINVIEW architecture, and compares CMF and RMF.

CMF MONITOR product components

CMF MONITOR has the following product components:

- Extractor
- Analyzer
- Online
- CMFMON
- Data Set Optimizer (DSO)

These components work together as a system to provide you with realtime and historical system performance information.

Relationship among the CMF MONITOR product components

The relationship among the CMF MONITOR components can be seen in Figure 1.

Figure 1 Relationship among CMF MONITOR product components

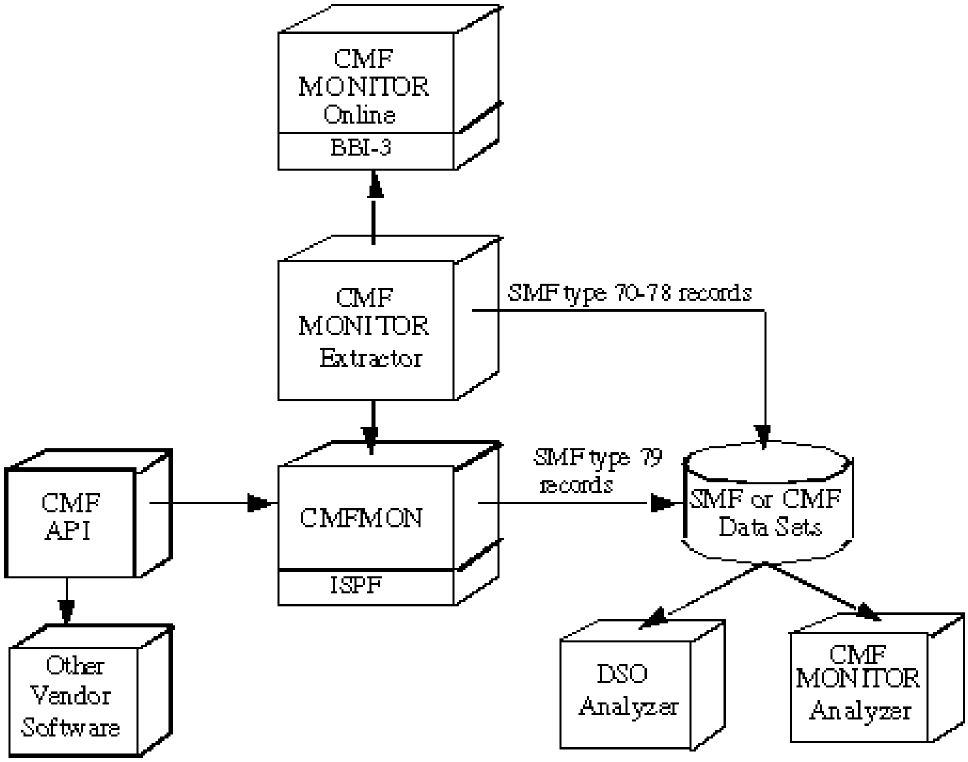

The CMF MONITOR Extractor collects information about system performance in common storage. From common storage, the data can be accessed by CMF MONITOR Online or written as records to System Management Facilities (SMF), CMF, or DSO data sets. The CMF and DSO Analyzers read the records from these data sets, and format and process them into reports that tell you about your computing enterprise.

CMF MONITOR Extractor

The CMF MONITOR Extractor collects information about configuration, CPU, software resource usage, and the system's workload, and stores records in one or more data sets. This data represents a statistical sample of system performance. Another function of the Extractor is that it dynamically prints summaries of system status on a periodic basis.

The Extractor data is used as follows:

- by the Analyzer to produce batch reports

- by CMF MONITOR Online (and some CMFMON screens) to provide realtime bottleneck detection and analysis

- by DSO to analyze the most efficient arrangement of data sets on your moveable head devices

The Extractor's centralized services are used by other measurement products to eliminate redundancy and reduce measurement overhead. The following BMC Software products can use data gathered by the CMF Extractor:

- DSO Analyzer, for detailed DASD analysis
- MAINVIEW for z/OS, for realtime and historical online performance analysis

You can specify the type of data to be gathered and the way that the data is to be used. You choose the activities to be monitored, the size of the sample to be collected, the storage medium to be used for output, and other data gathering characteristics. These tasks are accomplished by defining the CMF MONITOR Extractor control statements based on your informational requirements. Extractor control statements used to collect data and specify sampling rates are documented in Chapter 6, "Extractor control statements."

CMF MONITOR Analyzer

The CMF MONITOR Analyzer produces analytical reports from extracted data. You can use these reports for system tuning analysis.

When you submit a batch job, the Analyzer reads the records written by the CMF MONITOR Extractor and formats them into printed reports. Reports can contain data from the local system, or from one or more remote systems in your sysplex. These reports can be printed directly or downloaded to your PC to be formatted as Microsoft Excel spreadsheets. User-specified dates and times can be used to control the duration of the reporting period and the input records that are read to generate reports.

The Analyzer can produce a variety of graphics on almost 300 measurements. These graphs can be used to examine long-range data in many different ways.

The Analyzer provides both general and report control statements with parameters that you define to filter, order, and tailor the report data to your specifications. Analyzer control statements used to produce reports are documented in Chapter 7, “Analyzer control statements.” The Analyzer also provides an interface that uses ISPF panels to generate JCL and control statements for producing reports. This interface is described in “Generating JCL to produce Analyzer reports” on page 74.

The reports produced by each Analyzer control statement or combination of statements are documented in Chapter 8, “Analyzer reports.”

CMF MONITOR Online

CMF MONITOR Online monitors system activity, collecting information on all address spaces (TSO users, batch jobs, and Started Tasks), their use of various system resources, and the delays that each address space incurs while waiting for access to these resources.

Resources monitored are physical service entities, such as the processor, central storage, and DASD and tape devices; and logical entities, such as System Resource Manager (SRM), Hierarchical Storage Manager (HSM), and enqueue.

CMF MONITOR Online automatically detects resource use and contention, identifying delays that jobs encounter, resources that are contention bottlenecks, and jobs competing for those resources. CMF MONITOR Online provides this information through screen displays called *views*.

All CMF MONITOR Online views and commands are presented through the MAINVIEW cross-system architecture. This architecture provides concurrent multisystem access, windowing functions, and display customization.

The functions of CMF MONITOR Online and the use of the windowing and cross-system operations are discussed in *CMF MONITOR Online Getting Started* and *CMF MONITOR Online User Guide*.

CMFMON

The CMFMON component uses data-gathering application program interfaces (APIs) that create in-storage SMF type 79 record images. This information can then be displayed by the CMFMON online facility in one or more formatted screens—generated as batch reports, or written to DASD in the form of SMF type 79 records.

The functions of CMFMON are discussed in the *CMFMON User Guide*.

DSO Analyzer

The DSO component uses CMF Extractor data to report on the seek activity of devices with movable heads. The Extractor records seek activity by data set name. From these statistical records, the DSO Analyzer produces reports that specify an optimal ordering of data sets on your moveable head devices.

You might need to use the DSO Analyzer only when excessive seek time is caused by lengthy actuator travel between successive read / write operations. Devices with this problem can be made more efficient by being reorganized to minimize the distance between I / O operations. The DSO Analyzer can automatically generate control cards for FDR COMPAKTOR and DFDSS to do this reorganization.

DSO Analyzer operation and report formats for DSO are documented in the *DSO User Guide and Reference*.

CMF MONITOR and MAINVIEW

CMF MONITOR Online and many other BMC Software products run on the MAINVIEW architecture. For more information about the MAINVIEW platform and the products that run on it, see the *MAINVIEW Common Customization Guide*.

Required address spaces

All MAINVIEW products require at least three address spaces, which are described in this section. Figure 2 illustrates the communication between address spaces under the MAINVIEW architecture.

Figure 2 Communication between address spaces

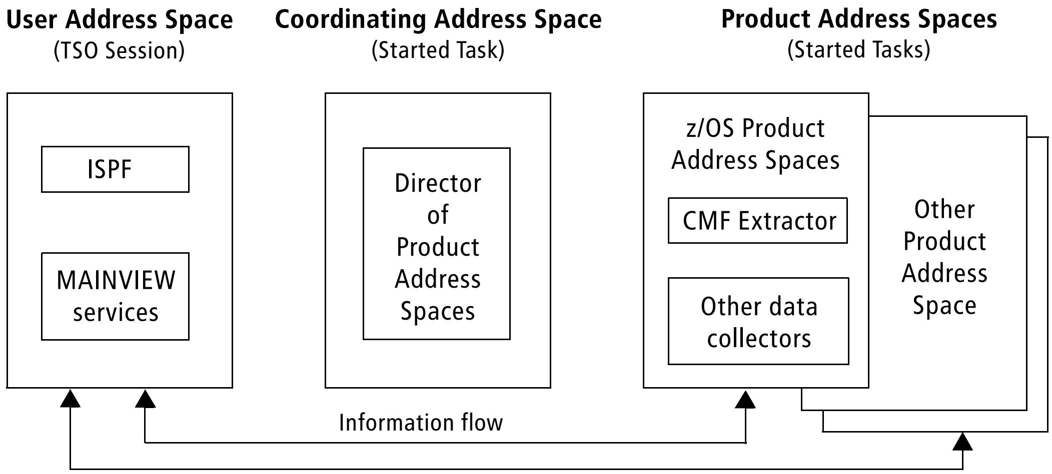

- User Address Space (UAS)

  The UAS is either a TSO session or a Started Task that provides VTAM or EXCP session support through the BMC Software MAINVIEW Alternate Access product.

- Coordinating Address Space (CAS)

  The CAS is a Started Task that runs as an MVS subsystem. There is one CAS per MVS image; each CAS provides various services to all MAINVIEW products running on that system.

- MAINVIEW Product Address Space (PAS)

  The MAINVIEW PAS runs as a Started Task and contains both the CMF Extractor and the online component of CMF MONITOR—called CMF MONITOR Online. The PAS also contains its own data collectors, which are in addition to the CMF Extractor samplers. And if you have the BMC Software MAINVIEW product, both MAINVIEW and CMF MONITOR run in the same MAINVIEW PAS on your system.

Although the CMF MONITOR Extractor is incorporated into the MAINVIEW PAS, the Extractor can be initialized separately from the data collectors belonging to the PAS. This arrangement allows you to run the CMF MONITOR Extractor without the overhead of the PAS data collectors. However, it is these data collectors in the PAS that provide support for both the MAINVIEW and the CMF MONITOR Online view displays.

NOTE

Initialization of the MAINVIEW PAS data collectors is controlled by the DC={START|STOP|CPM|IPM} parameter. This parameter is located in the PAS JCL and can be specified when starting the MVS PAS. See "Using the MODIFY command to change Extractor operation" on page 61 for more information about starting and stopping the PAS and using this parameter.

When initializing the MVS PAS with the following parameters, you can expect these results:

Parameter	Results
DC=STOP	You cannot access CMF MONITOR Online or MAINVIEW. However, you can later issue a MODIFY command with DC=START to access these products.
DC=START	Both the Extractor and the MVS PAS data collectors are initialized, and you can access CMF MONITOR Online and MAINVIEW.

CMF MONITOR and other BMC Software products

Other BMC Software products use data extracted by the CMF Extractor. These products are

- DSO Analyzer
- MAINVIEW for z/OS

The relationship among CMF MONITOR and other BMC Software products is shown in Figure 3.

Figure 3 Interaction among CMF MONITOR and other BMC Software products

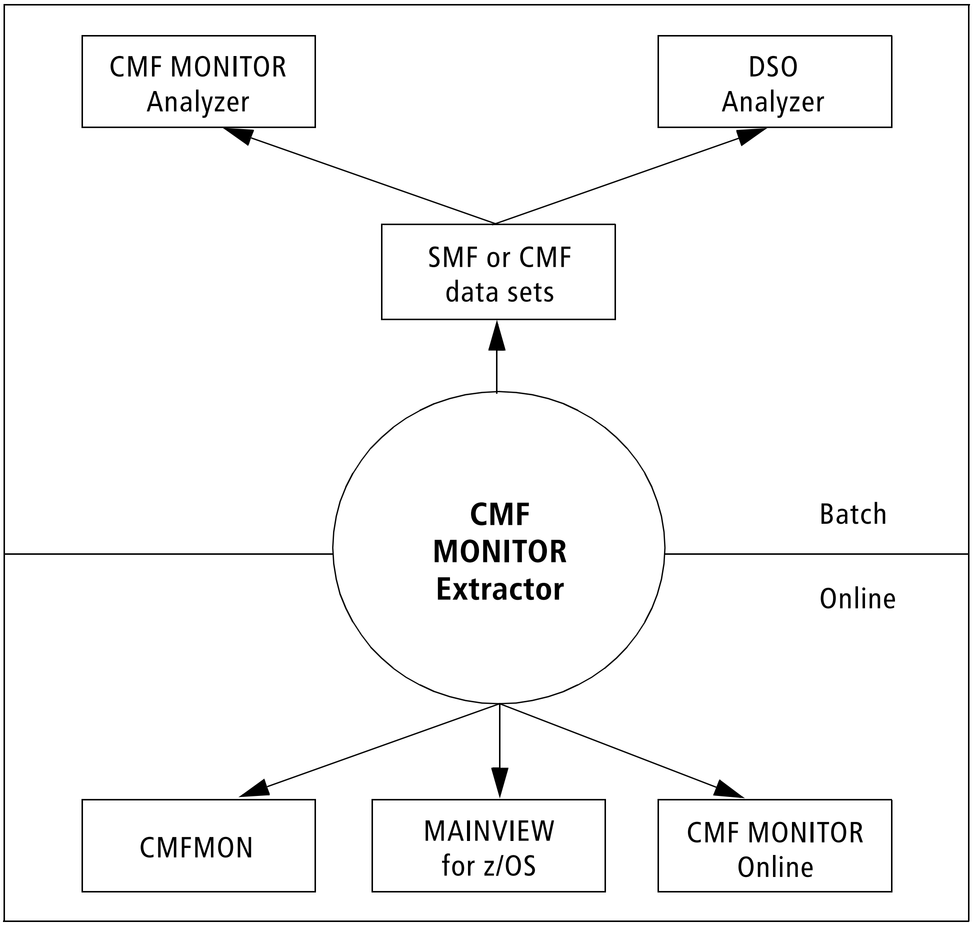

This section discusses products interaction with the CMF MONITOR Extractor, but it does not discuss CMFMON or the CMF and DSO Analyzers. (See “Relationship among the CMF MONITOR product components” on page 28 for more information about these CMF components.)

MAINVIEW for z/OS

MAINVIEW for z/OS is an MVS performance management monitoring product that is based on the MAINVIEW cross-system architecture. It provides online displays and performance management services for controlling and correcting your system's performance.

The MVS PAS collects data on over 5,000 discrete elements in your system and reports on their performance through MAINVIEW for z/OS views. The views present this information in both realtime and historical time frames, which allows for on-screen analysis of current and past resource performance.

The services offered through MAINVIEW for z/OS provide you with the ability to instantly respond to any performance circumstances that your system encounters. The set of services offered includes both System Programmer Services and Exception Monitoring Services.

MAINVIEW for z/OS uses CMF MONITOR Extractor to collect some of the data for its online views. To collect the correct data, CMF MONITOR Extractor control statements that execute specific samplers must be included in the Extractor JCL. (See "Extractor control statements used by BMC Software products" on page 54 for more information.)

MAINVIEW for z/OS operation and view data element descriptions and commands are documented in several books. See "Related documentation" on page 19 for more information.

CMF MONITOR compatibility with IBM RMF

CMF MONITOR has many compatibilities with RMF, and some differences. This section presents the compatibility issues of CMF MONITOR and RMF.

Similarities

- CMF MONITOR Extractor produces SMF record types that are compatible with z/OS 1.4 and later. You can use CMF MONITOR records with the RMF postprocessor, or you can use the CMF MONITOR Analyzer to process RMF data and produce reports.

  CMF MONITOR produces SMF type 7*x* series records that are compatible with the latest release of RMF running on the z/OS level where data is being extracted.

- CMF provides APIs (application programming interfaces) to integrate CMF data with other vendor products or user-written applications. CMF APIs allow you to retrieve the same data from CMF MONITOR as is available through RMF APIs.

Differences

CMF MONITOR can start and stop I/O monitoring of devices other than tape and DASD. When this feature is activated, CMF MONITOR makes sure that RMF control of the channel measurement blocks for nontape and non-DASD devices is maintained. RMF assumes that it has exclusive use of all nontape and non-DASD CMBs. CMF MONITOR does not perform start or stop I/O monitoring of this kind if RMF is active and sampling nontape and non-DASD device classes.

NOTE

CMF MONITOR device monitoring is controlled by the CLASS parameter of the Extractor DEVICE control statement. If SMF type 74 records that are compatible with those produced by RMF are desired, CLASS should be the only parameter used (see “DEVICE” on page 145).

CMF MONITOR functions that are unavailable in RMF

CMF MONITOR provides the following Extractor functions that RMF does not have:

- To reduce overhead, CMF MONITOR allows you to define individual sampling intervals for each sampler that is not event driven.

- Through its Extractor utilities, CMF allows you to browse SYS1.MAN*x* data sets online without switching SMF recording to another data set before gaining access to view this data.

- CMF MONITOR can run two Extractors from one monitor, a CPM and an IPM mode. (See "When to use CMF MONITOR Extractor CPM and IPM modes" on page 42.) RMF can run only one Extractor.

Chapter

2

Collecting data using the Extractor

The CMF MONITOR Extractor gathers and records system data for online displays and batch reports. This chapter discusses information about how the Extractor operates, the records it produces, and the BMC Software products that use it. Also explained in this chapter is how to define Extractor control statements to get the data you need.

Understanding the CMF MONITOR Extractor

The Extractor is a component of CMF and other BMC Software products that samples data in your system, stores the data in CSA, and then can record that data from CSA to either SMF or other data sets.

How the Extractor works

Figure 4 on page 40 shows how the Extractor collects and records data.

Figure 4 Extractor data collection and recording functions

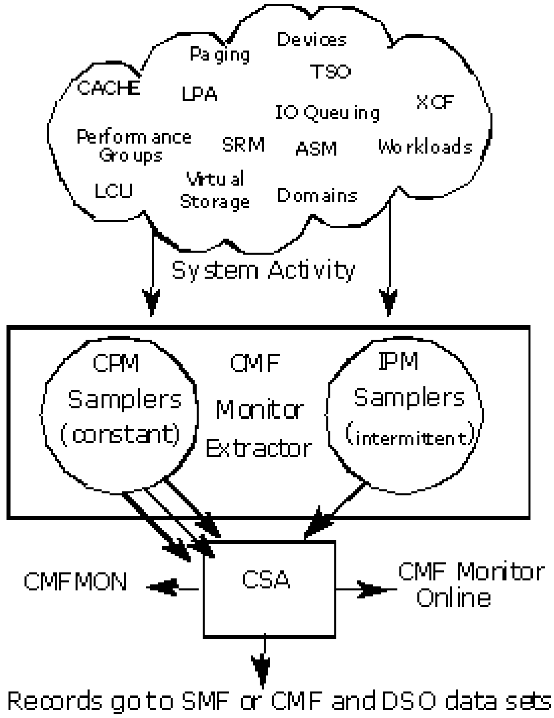

- Sampling data

  The Extractor uses samplers to measure hardware usage (CPU, channel paths, I/O queuing, and I/O devices) and software activity (the SRM, paging, enqueue contention, TSO usage, and common storage allocations). A sampler is a program module that uses DIE (disabled interrupt exit), SRB (service request block), or SRM (System Resources Manager) methods of measuring data; see “Sampling data and producing records” on page 43 for more information.

- Controlling data collection

  The function of each sampler that is running in your system is controlled by an Extractor control statement. (See Table 1 on page 47 for more information.) Most of the Extractor statements have parameters that control the manner in which a particular sampler functions. In addition, there is a global Extractor control statement that has an overall effect on the way the Extractor operates and samplers function.

- Two monitoring modes

  The Extractor is actually made up of two distinct submonitors that can be started, stopped, or modified independently, or they can both be running at the same time. The two submonitors in the CMF Extractor are

  — **CPM** (Continuous Performance Monitoring)—for low-resolution, long-duration monitoring

  — **IPM** (Intermittent Performance Monitoring)—for high-resolution, short-duration monitoring

  See “When to use CMF MONITOR Extractor CPM and IPM modes” on page 42 for more information about these monitoring modes.

- Writing records

  The sampled data is collected in CSA for an interval of time that you define. (This time interval can be synchronized with SMF recording.) When the interval time expires, another function of the Extractor writes the data in CSA as records to SMF, or CMF, or DSO data sets; see “Writing records” on page 49 for more information.

NOTE

BMC recommends the following actions for CMF MONITOR to adjust for Daylight Saving Time changes:

- For Spring Daylight Saving Time, stop the MVS PAS prior to the time change and start the PAS again after the time change.
- For Autumn Daylight Saving Time, stop the MVS PAS prior to the time change and wait one hour before restarting.

Although these changes are meant specifically to accommodate CMF MONITOR SMF records (which use local time), be aware that stopping the PAS affects not only CMF MONITOR but also other products that are running in the PAS, such as MAINVIEW for z/OS, MAINVIEW for UNIX System Services, or MAINVIEW SYSPROG Services.

How CMF MONITOR components use Extractor data

CMF MONITOR Online and CMFMON, the realtime components of CMF MONITOR, display the CMFMON, the realtime component of CMF MONITOR, displays the sampled data, as well as other data, directly from CSA, in windowed online views. The CMF MONITOR Analyzer reads records from SMF or CMF data sets, and DSO reads records from DSO data sets, to produce historical batch reports. See Chapter 4, “Producing and using Analyzer reports,” for more information about using the CMF Analyzer; see the *DSO User Guide and Reference* for more information about using the DSO Analyzer.

When to use CMF MONITOR Extractor CPM and IPM modes

CPM and IPM modes function separately from each other. Each mode can be started, stopped, or modified independently of the other mode, or both modes can be running at the same time. For information about how to start, stop, or modify CPM and IPM modes, see “Using the MODIFY command to change Extractor operation” on page 61.

The next two sections describe each monitoring mode.

CPM (Continuous Performance Monitoring)

The CPM mode is designed for low-resolution, consistent, long-duration monitoring. This mode is intended primarily for long-term system measurement, and BMC Software recommends that you run this submonitor 24 hours a day at low sampling rates. The data collected in CPM mode is best used for daily reporting because of its consistency, and for long-term trend analysis because of its duration.

By setting the samplers to run at low sampling rates in this mode, CPU consumption is minimized over the long term. In addition, CPM samplers that run 24 hours a day, but at low sampling rates, provide data that can be used best as a basis for identifying long-term performance trends in your system. However, this trend information can be developed only with continuous sampling.

A sample control statement for invoking the CPM monitor is shown in “Using the CMFCPM00 control statement set” on page 58.

IPM (Intermittent Performance Monitoring)

IPM mode provides you with a means for collecting two sets of data concurrently. By having two modes of the Extractor running at the same time, you can collect data in two different ways during the same time frame. For example, since the CPM mode is normally used for low-frequency, long-duration data collection, you might want to use the IPM mode for short-term, intermittent, high-resolution monitoring of various resources. By using two Extractors in this way, you can maintain the continuous sampling function while also obtaining the additional sampling data you need for another purpose. IPM is generally used to sample head movement, individual devices, and the like.

When you detect a trend that could be impacting performance, or when you have identified a distinct performance problem, you can use the IPM mode to run a set of samplers at high frequency rates. By concentrating your sampling of specific system resources at high intervals, the IPM data provides focused information about current performance problems for these resources. You can use the information from IPM mode to determine specific causes of system performance difficulties.

BMC Software recommends that you run this mode only when necessary to minimize system impact, or when you need to gather DSO data. Initiating IPM mode to investigate specific areas of activity or to monitor specific devices during peak periods of activity should be at the discretion of your system programmer.

Because IPM mode (with more frequent sampling rates) has a higher overhead associated with it and provides concentrated sampling, the default Started Task JCL shipped with the Extractor does not initialize IPM mode at startup.

To start the IPM submonitor, you can either use the MODIFY command (see "Using the MODIFY command to change Extractor operation" on page 61 for more information), or change the default setting on the MVS PAS PROC to DC=IPM (see "Defining Extractor control statements" on page 57 for more information).

Two Extractor control statements are invalid in this mode: CSMON and EXTSUM.

A sample control statement for invoking the IPM monitor is shown in "CMFIPM00 control statement set" on page 60.

Sampling data and producing records

The Extractor executes samplers to monitor hardware usage (CPU, channel paths, I/O queuing, and I/O devices) and software activity (the SRM, paging, enqueue contention, TSO usage, and common storage allocations), as well as other system resource information.

A sampler is a program module that creates data by using one of the following methods of measuring the performance of a specific area of your system:

- DIE (disabled interrupt exit)
- SRB (service request block)
- SRM (System Resources Manager)
- TCB (task control block)

Each sampler uses a combination of the four sampling methods to collect this data.

A sampler executes only when a corresponding Extractor control statement is defined in the Extractor JCL. (See "Defining Extractor control statements" on page 57 for more information.) The manner in which each sampler executes is controlled by parameters associated with each Extractor control statement. (See "How Extractor statements control sampler operations" on page 44 for more information.)

As a sampler gathers information, the data gets deposited in CSA. At the end of the interval time for the CPM or IPM monitoring mode, another function of the Extractor writes the data collected in CSA as records to SMF, or CMF, or DSO data sets. (See "When to use CMF MONITOR Extractor CPM and IPM modes" on page 42 for more information about monitoring modes, and "Writing records" on page 49 for more information about writing records.)

Four methods of sampling data

The Extractor uses four sampling methods to collect data:

DIE	Disabled Interrupt Exit—CMF keeps the time spent in the DIE samplers to a minimum to avoid degrading system performance
SRB	System Request Block—several samplers use the SRB sampling method Under SRB, higher resolution sampling is provided and system interrupts can still be honored.
SRM	System Resources Manager—the SRM event-counting method is used for sampling functions that are driven directly by SYSEVENTs (such as TSO or ENQUEUE) SRM sampling is driven by an event, not a time value.
TCB	Task Control Block—a separate subtask that can be dispatched on any processor, when required, by internal data-gathering mechanisms in the operating system

How Extractor statements control sampler operations

Most of the Extractor statements provide you with parameters that allow you to control the manner in which a particular sampler functions. Extractor statements that do not provide you with parameters are typically for samplers that are event-driven and, therefore, function only when the particular circumstance occurs in your system.

Parameters affect sampler operations such as the rate frequency at which a sampler gathers measured data and places it in CSA. System components can be measured at different user-selected sampling rates; for example, device activity can be sampled once a second, while CPU activity can be sampled 10 times a second.

Other parameters in Extractor control statements allow you to control the scope of resources being monitored, the type of information being collected, and many other data-gathering options.

All of this control through statement parameters is provided so that you can tune the Extractor to collect only the data your site requires. This procedure allows the Extractor to perform its monitoring functions on your system in the most efficient manner possible.

The sampling functions occur continuously. All samplers write out records based on the recording interval, with the exception of the samplers for these Extractor control statements:

- REPORT (GBLS sampler)
- HEADMOVE (HMOV sampler)

The GBLS and HMOV samplers write out records more frequently due to the volume of data that they are recording.

Record types

The Extractor samplers produce records that can be processed later by the CMF and DSO Analyzers or by the RMF postprocessor, or used by other BMC Software products. The SMF and CMF user record types produced by CMFMON and by the Extractor for the CMF MONITOR and DSO Analyzers are shown in Table 1 on page 47. User programs that are written to process SMF record types, such as MICS, MXG, or SLR, can process CMF MONITOR records.

NOTE

Refer to “DEVICE” on page 145 for more information about producing RMF-compatible type 74 records.

Changing the default CMF record type ID

A default CMF user record SMF ID of 240 is used. This default value can be changed, however, by defining a different ID value at the SMFRECID parameter of the REPORT control statement. (See “REPORT” on page 176 for more information.)

Producing your own reports using Extractor records

Records that are produced by Extractor samplers are in SMF type 70-79 format. In addition, CMF MONITOR produces SMF user records with a default type of 240. CMF MONITOR produces approximately 25 record subtypes. Refer to the *hilevel*.BBSAMP data set for data area maps for each of the CMF 240 user record subtypes. The member name for any record subtype is as follows:

Format	Member name
Assembler	CMFREC*xx*
C	CMFC*xx*
SAS	CMFSK*xx*
where *xx* is equivalent to the record subtype of 00 through 69	

Data area maps for SMF type 7*x* records are also available in BBSAMP members CMFSMF*xx*.

See Chapter 10, “Mapping CMF records created by CMF,” for more information about the BBSAMP members and writing your own programs using CMF records.

Numeric list of record types

Table 1 shows record types used by CMF MONITOR in ascending numeric order.

Table 1 Record types with corresponding Extractor statements and samplers (part 1 of 2)

Record type	Description	Sampler	Control Statement
SMF 70-1	CPU activity	CPUS	CPU
SMF 70-2	Cryptographic activity	CRYS	CRYPTO
SMF 71	paging activity	PAGS	PAGING
SMF 72-3	workload activity by service class	WLMS	WORKLOAD
SMF 72-4	resource usage and delay data by service class / period This sampler does not function unless the MVS PAS data collectors are active.	PGDS	PGDDLAY
SMF 73	channel activity	CHNS	CHANNEL
SMF 74-1	device activity	DEVS	DEVICE
SMF 74-2	Cross-System Coupling Facility (XCF) data	XCFS	XCFDATA
SMF 74-3	Open Edition MVS data	OMVS	OMVS
SMF 74-4	coupling facility data This sampler does not function unless the MVS PAS data collectors are active.	CFTS	CFDATA
SMF 74-5	cache data records	CA3H CA5H CA6H	CACHE
SMF 74-6	HFS statistics	HFSS	HFS
SMF 74-7	FICON Director statistics	FCSW	FICONSW
SMF 74-8	Enterprise Storage Server (ESS) statistics	CA3H CA5H CA6H	CACHE
SMF 75	page data set activity	ASMS	ASMDATA
SMF 76	system control block trace data	TRAS	TRACE76
SMF 77	enqueue activity This sampler does not run in IPM mode.	EQES	ENQUEUE
SMF 78-2	virtual storage data	VSMS	VSMDATA
SMF 78-3	I / O queuing data for 3090, ES / 9000 series, or later processors	IOQS	IOQ
Note: SMF records 79-1 through 79-12 are created by the CMFMON component of CMF MONITOR.			
SMF 79-1	address space state data	not applicable	ASD
SMF 79-2	address space resource data	not applicable	ARD

Table 1 Record types with corresponding Extractor statements and samplers (part 2 of 2)

Record type	Description	Sampler	Control Statement
SMF 79-3	central storage / processor / SRM activity	not applicable	SRCS
SMF 79-4	system paging activity data	not applicable	SPAG
SMF 79-5	address space SRM data	not applicable	ASRM
SMF 79-6	enqueue reserve data	not applicable	SENQR
SMF 79-7	enqueue contention data	not applicable	SENQ
SMF 79-9	device activity data	not applicable	DEV
SMF 79-11	page data set activity	not applicable	PGSPP
SMF 79-12	channel path activity	not applicable	CHANNEL
SMF 79-14	I / O queuing activity by logical control unit for the 3090, ES / 9000 series processors CMFMON	not applicable	IOQ
SMF103-1 SMF103-2	HTTP Server Report	not applicable	not applicable
SMF108-1 SMF108-3	LOTUS DOMINO Server Report	not applicable	not applicable
CMF 240-00	SRM constants, installation performance specifications, and Extractor control cards data	RECD	REPORT
CMF 240-01	CPU data	CPUS	CPU
CMF 240-02	ASM data	ASMS	ASMDATA
CMF 240-03	paging data	PAGS	PAGING
CMF 240-05	device data	DEVS	DEVICE
CMF 240-06	Extractor summary data This sampler does not run in IPM mode.	EXTS	EXTSUM
CMF 240-09	ASM data	ASMS	ASMDATA
CMF 240-11	global bit map	GBLS	REPORT
CMF 240-12	DASD head movement mount data	HMOV	HEADMOVE
CMF 240-13	DASD head movement seek data	HMOV	HEADMOVE
CMF 240-14	DASD head movement VTOC data	HMOV	HEADMOVE
CMF 240-16	LPA mapping data	LPAM	LINKMAP
CMF 240-18	CMF trace record data	TRCE	TRACE
CMF 240-20	TSO command summary record data	TSOS	TSODATA
CMF 240-21	TSO user summary record data	TSOS	TSODATA
CMF 240-24	disabled time sampling record data	DITS	DISTIM
CMF 240-29	COMMON STORAGE MONITOR records This sampler does not run in IPM mode.	CSMS	CSMON
CMF 240-50	output writer statistics data	not applicable	not applicable
CMF 240-98	used to identify invalid records	not applicable	not applicable
CMF 240-99	used to identify invalid records	not applicable	not applicable

Writing records

As the samplers defined to CPM and IPM modes gather data and deposit it in CSA, another function of the Extractor periodically takes the data collected in CSA and writes records for batch reports. Records can be written to SMF, CMF, or DSO data sets. For historical reporting purposes, or to run reports for long-term trend analysis, you will want to archive your data as these data sets become full (see “Archiving your data” on page 52 for more information).

The duration of time that the Extractor waits before writing the data collected in CSA to SMF or Extractor output data sets is called the *recording interval*. The recording interval is not variable; it is a set and established amount of time, but it can be customized. It is the systematic sampling and writing of records that provides integrity to CMF’s long-term trending data.

When writing records to an Extractor output data set, the Extractor issues a TCLOSE at the end of each interval, allowing the system to determine the correct end-of-file position even if the system fails.

Under certain conditions, the TCLOSE does not protect the data set. For example, if a blocked VBS record is being written and the system fails, there is no end-of-file marker. If a spanned VBS record is being written and the system fails, additional problems can result. Although data set damage rarely occurs in these circumstances, data sets can be recovered by copying the damaged data set to a new data set. BMC Software recommends using the CMF COPY VBS utility, discussed in Chapter 3, “Preprocessing Extractor data sets,”, to recover and copy a damaged data set.

Customizing the recording interval

You can synchronize the Extractor’s recording interval with your SMF recording interval. The recording interval is specified in the INTERVAL parameter of the REPORT control statement (see “REPORT” on page 176 for more information).

In most cases, the Extractor’s recording interval determines the rate at which data for CMF MONITOR Online is written into the historical data sets. However, if you need historical records for CMF MONITOR Online written at a faster rate than the Extractor writes to SMF, you can specify two different recording rates by using the CPM mode to write to SMF and the IPM mode to set the interval at which the data collectors write to the historical data sets for CMF MONITOR Online. See “When to use CMF MONITOR Extractor CPM and IPM modes” on page 42 for more information about setting different recording intervals.

Writing to SMF

To direct Extractor records to be written to the SMF data sets, you must specify SMF=YES and the SMFRECID keyword on the Extractor REPORT control statement. (See “REPORT” on page 176 for more information.)

The Extractor uses the IBM SMFEWTM macro to write to the SMF data set. If data is to be recorded to SMF, the SYS and / or SUBSYS parameter of member SMFPRM*xx* in SYS1.PARMLIB must be specified so that SMF type 70 through 79 records and the CMF user record type (as specified by the SMFRECID keyword on the Extractor REPORT statement) are written. For SMF purposes, the SUBSYS is CMF.

Writing to the Cross-System Data Server (XDS) buffers

Records are written automatically to XDS if both of the following conditions are present:

- The Extractor is writing records either to SMF or to an output data set.
- XDS is active with both of the following parameters:
  - TYPE CMF or TYPE ALL (or TYPE SMF, if SMF recording is active)
  - RECORDS 70:78 or RECORDS 7X or RECORDS ALL

A sample XDS member that contains the correct attributes is available in *hilevel.*UBBPARM with the member name CMFXDS01. See the *CMF MONITOR Customization Guide* for more information on starting XDS and using the sample members.

Writing to CMF or DSO data sets

If data is not written to SMF, you must specify that records be written to CMF or DSO output data sets.

CMF data sets should have been allocated during customization if you determined that you would not record Extractor data to SMF. Even if you did specify SMF recording but now want to record to CMF data sets, you first need to allocate these data sets. See the *CMF MONITOR Customization Guide* for more information about allocating these data sets.

When the current Extractor output data set becomes full or when the FLIP command is issued, the CMF Extractor writes records only to output data sets that are empty. If no empty data sets are available, recording is suspended. For information about how to empty CMF and DSO output data sets, see “Archiving your data” on page 52.

Specifying primary and alternate data sets to the Extractor

You can specify either same or different primary and alternate data sets for both CPM and IPM modes. If only one data set is specified, the Extractor cannot provide alternate data set support.

When you specify the same primary and alternate data sets for both modes, all records from both modes go to the same data sets.

When you specify different data sets, the records from each mode go to different data sets.

NOTE

In specifying data sets for both CPM and IPM modes

- If CPM and IPM data go to the same primary data set, they must also go to the same alternate data sets. You cannot specify the same primary data set and different alternate data sets.
- If CPM and IPM data go to different primary data sets, they must also go to different alternate data sets. You cannot specify different primary data sets and the same alternate data sets.

There are two ways to specify primary and alternate data sets to the Extractor. Use *one* of the following methods; do *not* use both.

- One method of identifying the primary and alternate data sets to the Extractor is the presence of DD statements in the Extractor JCL. Valid data set DD names for CMF and DSO are shown in Table 2.

  You can specify up to 101 data sets, with *xx* representing any one or two alphanumeric characters.

  NOTE

  The primary data set is the first one specified. The order in which the DD names are specified is the order in which they will be used.

Table 2 Primary and alternate data set DD names

Component	CPM	IPM
Extractor	//CMFCPM*xx* DD	//CMFIPM*xx* DD
DSO	//CMFCDS*xx* DD	//CMFIDS*xx* DD

The Extractor writes to these data sets automatically if the DD statements are present and SMF=YES is not specified on the Extractor report control statement. If only one statement is defined, alternate data set support is not provided.

For more information about changing the Extractor JCL, see the *MAINVIEW Common Customization Guide.*

- A second method of identifying primary and alternate data sets to the Extractor is through the DSNLIST parameter on the REPORT control statement. A DSNLIST parameter can be specified for dynamic allocation of up to 101 data sets in the REPORT control statement.

NOTE

The primary data set is the first one specified. The order in which the data set names are specified is the order in which they will be used.

See “REPORT” on page 176 for more information about the DSNLIST parameter and the REPORT Extractor control statement.

Archiving your data

For archiving Extractor output data sets, BMC Software recommends that you use either the IBM utility IFASMFDP or the CMF COPYVBS utility. JCL that executes the CMF COPYVBS utility is found in *hilevel*.UBBSAMP member CMFJCVBS. See Chapter 3, “Preprocessing Extractor data sets,” for information about using this utility.

NOTE

Using other copy utilities could result in lost data and data input errors when running the Analyzer.

After you have your CMF (or DSO) data sets archived, CMF MONITOR provides two members in *hilevel*.UBBSAMP that contain JCL to empty your CPM and IPM data sets.

- CMFJCLRS is a Started Task for clearing data sets.
- CMFJCLRB is a batch job for clearing data sets.

Archiving data sets simply copies the information; it does not empty the data sets to receive more information. When you use one of these members, the specified data sets are emptied. If you want to save your data, make sure that you have archived it before using one of these members.

Running CMF and RMF on the same system

You can run CMF on a system where you are also running RMF, but you should be aware of the following considerations:

- Both CMF and RMF Extractors produce identical type 70 through type 79 records, but CMF MONITOR should not write records to the SMF data set if CMF MONITOR and RMF are to be active at the same time.

  When RMF reads records containing both CMF and RMF data, RMF cannot distinguish between CMF-generated type 70-series records and RMF-generated type 70-series records; the RMF post processor produces reports that contain duplicate data.

  If you have inadvertently written CMF and RMF data together, you can use the CX10CVBS copy VBS utility to separate the records (see Chapter 3, "Preprocessing Extractor data sets").

- CMF MONITOR can start and stop I/O monitoring of devices other than tape and DASD. When this feature is activated, CMF MONITOR makes sure that RMF's control of the channel measurement blocks for nontape and non-DASD devices is maintained. RMF assumes that it has exclusive use of all nontape and non-DASD CMBs. If RMF is active and sampling nontape and non-DASD device classes, CMF MONITOR does not perform start or stop I/O monitoring of this kind.

NOTE

CMF MONITOR device monitoring is controlled by the CLASS parameter of the Extractor DEVICE control statement. If SMF type 74 records that are compatible with those produced by RMF are desired, CLASS should be the only parameter used. (See "DEVICE" on page 145.)

The similarities and differences between CMF MONITOR and RMF are discussed in "CMF MONITOR compatibility with IBM RMF" on page 36.

Using the Extractor trace facilities

CMF Extractor provides a trace facility that uses SRB and SRM sampling methods.

- **SRB**—allows you to code trace routines for specialized system sampling

  CMF MONITOR schedules global SRBs to perform many sampling functions. The trace facility permits you to interact with the SRB scheduling mechanism and introduce user-supplied trace routines.

  At intervals specified by you, the SRB routine receives control, and a data area of from 1 to 112 bytes is added as an entry to a CMF user type 240-18 trace record. Trace records vary in size, up to a maximum of 4 KB. They are composed of entries added to the trace record by the SRB routine.

- **SRM—**allows selected SYSEVENTs to be traced

  Using the SRM trace facility, you can trace selected SYSEVENTs as specified in the Extractor TRACE control statement. A trace for a SYSEVENT includes the name of the job for which the SYSEVENT was issued, along with the parameter registers zero and one. TSO SYSEVENTs, called TSEVENTs, contain a SYSEVENT of zero and include the TSO command name.

For more information on invoking the trace facility, see “TRACE” on page 183.

Extractor control statements used by BMC Software products

This section lists the CMF Extractor control statements and corresponding samplers used by CMF MONITOR and other BMC Software products.

Chapter 6, “Extractor control statements” provides detailed information about all Extractor control statements used by BMC Software products.

CMF MONITOR Extractor control statements

The following chart lists specific control statements with their appropriate samplers.

Control statement	Sampler
ASMDATA	ASMS
CACHE	CA3H, CA5H, CA6H
CFDATA	CFTS
CHANNEL	CHNS
CPU	CPUS
CRYPTO	CRYS
CSMON	CSMS
DEVICE	DEVS
DISTIM	DITS
ENQUEUE	EQES
EXTSUM	EXTS
FICONSW	FCSW
HEADMOVE	HMOV
HFS	HFSS
IOQ	IOQZ
LINKMAP	LPAM
OMVS	OMVS
PAGING	PAGS
PGDDLAY	PGDS
REPORT	GBLS, RECD
TRACE	TRCE
TRACE76	TRAS
TSODATA	TSOS
USER	USER
VSMDATA	VSMS
WORKLOAD	WLMS
XCFDATA	XCFS

For a list of samplers that provide SMF type 79 records, see the *CMF MONITOR CMFMON User Guide*, Appendix B.

DSO control statements

The following chart lists required control statements with their appropriate samplers.

Control statement	Sampler
HEADMOVE	HMOV
REPORT	GBLS,RECD

NOTE

BMC Software recommends that the DSO samplers run under IPM mode, not CPM mode.

See the *DSO User Guide and Reference* and "HEADMOVE" on page 159 and "REPORT" on page 176 for more factors to consider when setting up the CMF Extractor JCL and control statements to collect measurement data for the DSO reports.

MAINVIEW for z/OS control statements

The following chart lists required control statements with their appropriate samplers.

Control statement	Sampler
ASMDATA	ASMS
CACHE	CA3H, CA5H, CA6H
CPU	CPUS
DEVICE	DEVS
PAGING	PAGS
REPORT	GBLS, RECD

The DEVICE statement must be defined twice: once with the CLASS=DASD parameter, and again with the parameters CLASS=TAPE and OFFLINE=NO.

Defining Extractor JCL

The CMF MONITOR Extractor is wholly incorporated in the MAINVIEW architecture as part of the MVS product address space (PAS). When the PAS is started, the Extractor is started because the Extractor's program execution and DD statements (JCL) have been incorporated into the PAS Started Task procedure (PROC). In addition, when you start the PAS, you have the option of starting CMF MONITOR Online as well as MAINVIEW for z/OS, if your site uses this product. (See information about the DC parameter in "Using the MODIFY command to change Extractor operation" on page 61.)

During either AutoCustomization or manual customization, the MVS PAS PROC statement and Extractor JCL are modified to accommodate your site requirements. The MVS PAS PROC is described fully in the *MAINVIEW Common Customization Guide*. The *MAINVIEW Administration Guide* contains information about starting and stopping the address spaces that are required for the MAINVIEW architecture.

Defining Extractor control statements

This section discusses the default Extractor control statement sets for both CPM and IPM modes that are shipped with CMF MONITOR.

During AutoCustomization or manual customization, Extractor control statement members for both CPM and IPM were customized to your site requirements. If you want to change or create additional control statement members, you can use the default members in *hilevel*.UBBPARM as a starter set of statements. However, any members that you create must follow this specific naming convention:

- for CPM mode, members must be named CMFCPM*xx*
- for IPM mode, members must be named CMFIPM*xx*

where *xx* is a unique two-character identifier.

Default CPM and IPM control statement sets

The *hilevel*.UBBPARM data set contains two sample Extractor control statement members that comprise a starter set for initial execution of the CMF MONITOR Extractor.

CMFCPM00	invokes a CPM monitor that runs continuously, sampling most functions
CMFIPM00	invokes an IPM monitor that runs for 60 minutes, sampling most functions valid in the IPM mode

Using the CMFCPM00 control statement set

CMFCPM00 invokes the CPM monitor, which runs continuously. (See “When to use CMF MONITOR Extractor CPM and IPM modes” on page 42 for more information.) Records are written to the Extractor data set every 15 minutes. The CMF user record ID is 240 (X'F0'). Records are written to SMF, because SMF=YES is specified on the REPORT control statement.

Sample control statements in CMFCPM00 are shown in Table 3. For an explanation of the control statements, see Chapter 6, “Extractor control statements.”

Table 3 Extractor control statements in starter set CMFCPM00 (part 1 of 2)

```
***********************************************************************
*                                                                     *
*   STARTER CONTROL STATEMENTS FOR THE EXTRACTOR IN CPM MODE          *
*                                                                     *
*   USE THIS MEMBER WHEN THE FOLLOWING PRODUCT IS RUNNING BY ITSELF   *
*   IN CPM MODE:                                                      *
*                                                                     *
*      - CMF MONITOR                                                  *
*                                                                     *
*   WHEN YOU ARE READY TO PLACE YOUR PRODUCTS INTO A PRODUCTION       *
*   ENVIRONMENT AND RMF IS NOT PRESENT, YOU MAY WANT TO MAKE SOME     *
*   OF THE FOLLOWING CHANGES TO THE CONTROL STATEMENTS:               *
*                                                                     *
*   1. TO WRITE RECORDS TO THE SMF DATA SET (RATHER THAN TO CMF OUTPUT*
*      DATA SETS), CHANGE THE SMF= PARAMETER ON THE REPORT CONTROL    *
*      STATEMENT FROM SMF=NO TO SMF=YES.                              *
*                                                                     *
*   2. MODIFY THE EXTSUM CONTROL STATEMENT TO INCLUDE YOUR            *
*      INSTALLATION JOB CLASSES AND PERFORMANCE GROUPS.               *
*                                                                     *
***********************************************************************
* CHANGE LOG:                                                         *
*    CREATED BY ?USER ON ?DATE AT ?TIME.                              *
*                                                                     *
***********************************************************************
```

(*continued on next page*)

Table 3 Extractor control statements in starter set CMFCPM00 (part 2 of 2)

```
 REPORT  CPM,INTERVAL=15,SYNCH=00,CSA=512,SMFRECID=240,
         RUNTIME=1440,SMF=NO
 ASMDATA SAMPLE=2000
***********************************************************************
*                                                                     *
*     THE CACHE CONTROL STATEMENT NEEDS TO BE INVOKED FROM ONLY       *
*     ONE SYSTEM IF ALL CACHE SUBSYSTEMS ARE SHARED.                  *
*                                                                     *
***********************************************************************
 CACHE
 CFDATA
 CHANNEL
 CPU     SAMPLE=2000
*CRYPTO
 DEVICE  SAMPLE=2000,CLASS=DASD
 DEVICE  SAMPLE=2000,CLASS=TAPE,OFFLINE=YES
***********************************************************************
*                                                                     *
*    THE MAJOR=SYSDSN PARAMETER SPECIFIES THAT ONLY RESOURCES HAVING  *
*    THE MAJOR NAME OF SYSDSN ARE MONITORED.  YOU CAN CHANGE THE      *
*    VALUE OF THIS PARAMETER TO THE MAJOR NAME OF YOUR CHOICE OR YOU  *
*    CAN REMOVE THE MAJOR= PARAMETER TO MONITOR ENQUEUE CONTENTION    *
*    FOR ALL RESOURCES.                                               *
*                                                                     *
***********************************************************************
 ENQUEUE MAJOR=SYSDSN
***********************************************************************
*                                                                     *
*    THE EXTSUM STATEMENT SHOULD BE MODIFIED TO MONITOR SPECIFIC      *
*    INSTALLATION JOB CLASSES.                                        *
*                                                                     *
***********************************************************************
 EXTSUM  SPINOFF=NO,JES=NO,SAMPLE=2000,
        JOBCLASS=(JC=A,JD=CLASSA,
                  JC=B,JD=CLASSB)
*FICONSW
***********************************************************************
*                                                                     *
*    THE HEADMOVE SAMPLER CAN HAVE HIGH CPU OVERHEAD, DEPENDING ON    *
*    THE SAMPLE RATE AND THE NUMBER OF DEVICES MONITORED.             *
*                                                                     *
***********************************************************************
*HEADMOVE ALL,SAMPLE=250,VTOC=YES
 IOQ
*LINKMAP
 PAGING  SAMPLE=6000
*PGDDLAY
*TSODATA LIMIT=50,USER=YES,SAMPLE=2000
***********************************************************************
*                                                                     *
*    THE VSMDATA SAMPLER CAN HAVE HIGH CPU OVERHEAD.                  *
*                                                                     *
***********************************************************************
 VSMDATA  SAMPLE=6000
 WORKLOAD
 XCFDATA
***********************************************************************
*                                                                     *
*    END OF CPM CONTROL STATEMENTS                                    *
*                                                                     *
***********************************************************************
```

NOTE

This set is a sample. It does not include control statements for all possible options.

CMFIPM00 control statement set

CMFIPM00 invokes an IPM monitor that runs for 60 minutes. (See “When to use CMF MONITOR Extractor CPM and IPM modes” on page 42 for more information.) Records are written to the Extractor data set every 15 minutes. The CMF user record ID is 240 (X'F0'). Records are written to the data set defined by the //CMFIPM1 DD statement.

- To start the IPM monitor, either issue the MODIFY command, or specify **DC=IPM** on the MVS PAS PROC.
- To stop the IPM monitor, issue the MODIFY command **IPM=STOP**.

Table 4 on page 60 shows a sample set. For an explanation of the control statements, see Chapter 6, “Extractor control statements.”

Table 4 Extractor control statements in starter set CMFIPM00

```
***********************************************************************
*                                                                     *
*   CONTROL STATEMENTS FOR THE EXTRACTOR IN IPM MODE                  *
*                                                                     *
*   USE THIS MEMBER WHEN THE FOLLOWING PRODUCT IS RUNNING BY ITSELF   *
*   IN IPM MODE:                                                      *
*                                                                     *
*      - CMF MONITOR                                                  *
*                                                                     *
***********************************************************************
* CHANGE LOG:                                                         *
*    CREATED BY ?USER ON ?DATE AT ?TIME.                              *
*                                                                     *
***********************************************************************
 REPORT IPM,INTERVAL=15,CSA=512,SMFRECID=240,RUNTIME=60,SMF=NO
 CACHE
 CHANNEL
 CPU      SAMPLE=500
 DEVICE   SAMPLE=500,CLASS=DASD
 DEVICE   SAMPLE=500,CLASS=TAPE,OFFLINE=YES
 HEADMOVE ALL,SAMPLE=33
 IOQ
***********************************************************************
*                                                                     *
*    END OF IPM CONTROL STATEMENTS                                    *
*                                                                     *
***********************************************************************
```

NOTE

This is a sample set. It does not include control statements for all possible options.

Using the MODIFY command to change Extractor operation

At some point while running the Extractor in your system, you might need to modify the configuration of your Extractor control statement set, start and stop the CMF Extractor data samplers, or start and stop IPM mode.

Descriptions and examples of valid MODIFY commands that can be used to control CMF MONITOR Online and the Extractor are as follows:

```
F  jobname[,MSGFREE] [,CPM={xx|STOP}][,IPM={xx|STOP}][,STATUS]
[,FLIP={IPM|CPM}][,PROFILE][,DC={START|STOP|STATUS|CPM|IPM}]
[,HMOVRESCAN=CPM|IPM|BOTH][,XDS={xx|STOP}]
```

Table 4 on page 60 explains the MODIFY commands. Table 6 on page 64 provides examples of how to issue these commands.

Table 5 CMF MONITOR MODIFY commands (part 1 of 2)

Command	Explanation
jobname	name of the Extractor job to receive the MODIFY command; this is the name of MVS PAS PROC
MSGFREE	causes the Extractor message file to be dynamically deallocated and spun off for printing The message file, defined at the //CMFMSG DD statement, is reallocated immediately after deallocation with no loss of data.
CPM=	specifies that CPM mode is to be started, stopped, or executed under a different control statement set; *xx* specifies a new control statement set The current control statement set, if one is running, is terminated before the new one is executed. If CPM=STOP is specified, the Extractor terminates the CPM monitoring mode. If the IPM mode is not active when CPM=STOP is specified, the address space is cancelled.
IPM=	specifies that IPM mode is to be started, stopped, or executed under a different control statement set; *xx* specifies a new control statement set The current control statement set, if one is running, is terminated before the new one is executed. If IPM=STOP is specified, the Extractor terminates the IPM monitoring mode.
STATUS	Displays information on the console regarding the status of the active submonitors, CPM and IPM; this is the same information you can view through the STATUS Extractor utility in CMF MONITOR Online (See the *CMF MONITOR Online User Guide* for more information.)
FLIP=	causes the Extractor to begin writing (flip) to the next available alternate data set for either IPM or CPM mode If an alternate data set was not defined, the MODIFY command is rejected.
PROFILE	displays system configuration information on the console This information is also available by using the CONFIG utility option available through CMF MONITOR Online. (See the *CMF MONITOR Online User Guide* for more information.)

Table 5 CMF MONITOR MODIFY commands (part 2 of 2)

Command	Explanation
DC=	specifies that the data collectors are to be stopped or started, or that status information about the data collectors should be displayed on the console If both the CAS and MVS PAS are initialized, you can control whether CMF MONITOR Online is executing by issuing this command. When a MODIFY command is issued with ■ DC=STOP, the Extractor continues to function, but CMF MONITOR Online becomes unavailable and the PGDDLAY and CFDATA samplers stop performing their sampling functions. ■ DC=START, CMF MONITOR Online is initialized and the PGDDLAY and CFDATA samplers, if defined, begin or resume their sampling functions. ■ DC=CPM, the data collectors initiate in CPM mode. ■ DC=IPM, the data collectors initiate in IPM mode. To view status information about the data collectors on the console, specify the DC=STATUS attribute.
HMOVRESCAN=	causes the HEADMOVE samplers to initiate a VTOC scan, resulting in a new set of CMF 240-14 records written to the extractor output dataset ■ HMOVRESCAN=CPM—the PAS initiates a new VTOC scan in the CPM HEADMOVE sampler ■ HMOVRESCAN=IPM—the PAS initiates a new VTOC scan in the IPM HEADMOVE sampler ■ HMOVRESCAN=BOTH—the PAS initiates a new VTOC scan in both the CPM and IPM HEADMOVE samplers
XDS=	specifies that the CMF Cross System Data Server (XDS) is to be started, stopped, or executed under a different control statement set ■ XDS=xx—causes the PAS to activate XDS using *xx* as the suffix of a *hilevel*.BBPARM member CMFXDS*xx* ■ XDS=STOP—causes the PAS to disable XDS

MODIFY command examples

Examples of how to issue the attributes for the MODIFY command are provided in Table 6.

Table 6 MODIFY command examples for CMF MONITOR

You type	System response
F MVSPAS,IPM=03,CPM=04	invokes the IPM and CPM sampling modes using the control statement packets CMFIPM03 and CMFCPM04, respectively
F MVSPAS,MSGFREE	frees the //CMFMSG DD data set for printing, and reallocates it
F MVSPAS,IPM=XY	invokes the IPM mode with control statement packet CMFIPMXY
F MVSPAS,STATUS	produces the CMF Extractor Status Display on the console
F MVSPAS,FLIP=CPM	causes the CPM mode to start writing to an alternate output data set
F MVSPAS,PROFILE	produces the CMF system configuration display on the console
F MVSPAS,DC=STOP	causes the MVS PAS data collectors to stop functioning (rendering CMF MONITOR Online unavailable), and suspends the sampler for the PGDDLAY and CFDATA control statements, if defined This command also renders MAINVIEW for z/OS running in the same PAS unavailable.
F MVSPAS,DC=START	invokes the data collectors in the MVS PAS; CMF MONITOR Online is initialized, and the sampler for the PGDDLAY and CFDATA control statements, if defined, begin or resume their sampling functions
F MVSPAS,DC=STATUS	produces a status display of the MVS PAS data collectors on the console
F MVSPAS, HMOVRESCAN=BOTH	causes the MVS PAS HEADMOVE samplers for both IPM and CPM modes to initiate a VTOC scan and write CMF 240-14 records

Chapter

3

Preprocessing Extractor data sets

There are several reasons why you might want to preprocess CMF Extractor data sets before producing reports:

- **Extractor data set damage**

  System outages can damage the integrity of Extractor output data sets and cause a variety of QSAM-related errors when using the Analyzer. The CX10CVBS utility that is distributed with CMF MONITOR repairs damaged Extractor data sets and copies them to a new data set. CX10CVBS can process sequential as well as VSAM files.

- **specific record type creation**

  You might want to decrease the size of a data set used as input to CMF, RMF, or a user-written report program, so that the data set contains only specific record types. A smaller data set can decrease processing time for your reports, or a user-written program might accept only certain types of SMF records.

- **RMF identical record creation**

  By processing the CMF data using the CX10CVBS utility, CMF flags that are set in a reserved field in the CMF records are reset, making the CMF records identical to RMF records.

How to use the CX10CVBS utility

This utility uses the BMC Software READVBS subroutine to copy VBS records from the data set defined in the //SYSUT1 DD statement to the data set defined in the //SYSUT2 DD statement. The valid records are written RECFM=VBS, LRECL=32762. Invalid records are written to the output location defined by either an optional //SNAPVBS DD or a //SNAPREC DD statement. See "How data is copied using the CX10CVBS utility" on page 71 for more information about how records are written. Do not specify any DCB characteristics on the //SYSUT2 DD statement.

NOTE

The SYSUT1 input and SYSUT2 output data sets can be either tape or disk.

Use *hilevel*.UBBSAMP member CMFJCVBS, shown in Figure 5 on page 67, to execute the CX10CVBS utility.

After you have used the CX10CVBS utility to copy the data, S001 and S002 abends are eliminated when using the CMF MONITOR Analyzer because of the presence of an end-of-file (EOF) mark.

Figure 5 Sample execution JCL for CX10CVBS

```
//JOBCARD JOB
//*
//*--------------------------------------------------------------------
//*
//* SAMPLE JCL FOR EXECUTING THE COPYVBS UTILITY.
//*
//*   REVIEW THE JCL FOR APPLICABILITY TO YOUR INSTALLATION
//*   STANDARDS
//*
//*   CHANGE ?BBCHILV TO THE HIGH-LEVEL QUALIFIER YOU CHOSE FOR
//*   THE CMF MONITOR LIBRARIES.
//*
//*   CHANGE ?BBASMFID TO THE SMF ID (SYSTEM ID) OF THE TARGET
//*   SYSTEM.
//*
//*   SPECIFY COPY OPTIONS IN THE PARM= FIELD OF THE EXEC STATEMENT.
//*
//*   CHANGE THE SYSUT1 DD STATEMENT TO POINT TO THE DATA SET YOU
//*   WISH TO COPY RECORDS FROM.
//*
//*   CHANGE THE SYSUT2 DD STATEMENT TO POINT TO THE DATA SET YOU
//*   WISH TO COPY RECORDS INTO.
//*
//*--------------------------------------------------------------------
//*
//CMFCVBS  EXEC PGM=CX10CVBS,REGION=4096K,PARM='TYPE=CPM'
//*
//STEPLIB  DD   DISP=SHR,                 - BBLINK LOAD LIBRARY
//         DSN=?BBCHILV.BBLINK
//SYSUT1   DD   DISP=SHR,                 - "COPY FROM" DATA SET
//         DSN=SYS1.MAN1
//SYSUT2   DD   DISP=SHR,                 - "COPY INTO" DATA SET
//         DSN=?BBCHILV.SYS?BBASMFID.CPMOUT1
//
```

Defining the PARM parameter

By entering the selection criteria through the PARM parameter, you can make the CX10CVBS program selectively copy records. If no PARM parameter is entered, the entire data set is copied. The PARM parameter format for the program is

PARM='*TYPE*=*xxx*,*SMFrecID*[,SUB=*nn*][/[STARTAFT=*n*][,STOPAFT=*n*]]

Table 7 defines the CX10CVBS parameters. Examples of how to define PARM parameter values are shown in Table 8 on page 70.

Table 7 CX10CVBS parameters (part 1 of 2)

Parameter	Definition	
TYPE=*xxx*	defines the type of records to be processed from the input data set defined in the //SYSUT1 DD statement and copied to the output data set defined in the //SYSUT2 DD statement One of the following values can be defined for *xxx*:	
	RMF	Only RMF data recorded to the input data set by the RMF data gatherer is copied to the output data set. This value can verify that the RMF data is not corrupted by a system outage.
	SMF	Only SMF data recorded to the input data set is copied to the output data set. This value can verify that the SMF data is not corrupted by a system outage.
	CPM	Only CPM data recorded to the input data set by the CMF Extractor or type 79 records produced by CMFMON's Write Facility are copied to the output data set.
	IPM	Only IPM data recorded to the input data set by the CMF Extractor is copied to the output data set.
	CPR	Only CPM data recorded to the input data set by the CMF Extractor is copied, and CMF flags in the reserved field SMF7*x*RV2 are reset, thereby making the CMF data look exactly like RMF data.
	IPR	Only IPM data recorded to the input data set by the CMF Extractor is copied, and CMF flags in the reserved field SMF7*x*RV2 are reset, thereby making the CMF data look exactly like RMF data.
SMF*rec*ID	defines the ID of the specific SMF or CMF user record type(s) to be selected from the input data set You can define multiple values to the *SMFrecID* subparameter, if several record types are to be copied at once. In this case, the values must be enclosed in parentheses. Up to 16 record IDs can be specified. The default ID for CMF MONITOR records is 240. The following values show default ranges of record types for the *SMFrecID* subparameter, when a specific TYPE value is defined. Specific record types can also be defined if they are within the valid range of default values for any of the TYPE subparameters.	

Table 7 CX10CVBS parameters (part 2 of 2)

<table>
<tr><th>Parameter</th><th colspan="2">Definition</th></tr>
<tr><td></td><td>RMF</td><td>70 through 79

If the SMFrecID value is omitted for TYPE=RMF, all RMF records (70 through 79) are copied.</td></tr>
<tr><td></td><td>SMF</td><td>1 through 255

If the SMFrecID value is omitted for TYPE=SMF, all SMF records (1 through 69 and 80 through 127) are copied.</td></tr>
<tr><td></td><td>CPM</td><td>70 through 79, and 128 through 255

If the SMFrecID value is omitted for TYPE=CPM, all CPM records are copied.</td></tr>
<tr><td></td><td>IPM</td><td>70 through 79, and 128 through 255

If the SMFrecID value is omitted for TYPE=IPM, all IPM records are copied.</td></tr>
<tr><td></td><td>CPR</td><td>70 through 79 CPM records, formatted to be like RMF records</td></tr>
<tr><td></td><td>IPR</td><td>70 through 79 IPM records, formatted to be like RMF records</td></tr>
<tr><td>SUB=nn</td><td colspan="2">valid only with TYPE=CPM or TYPE=IPM; the following values are valid:

00 01 02 03 04 05 06 07 09 11 12 13 14 16 18 19 20 21 23 24 27 29 50 69

If the SUB=nn value is omitted, all CPM or IPM records are copied, and the SMFrecID value determines the SMF record ID for CMF Extractor user records.

If the SMFrecID value is 70 through 79, then only the specific SMF record type is copied, and the SUB=nn value is ignored.

Note: You can define multiple values to the SUB=nn subparameter, if several record types are to be copied at once. In this case, the values must be enclosed in parentheses. Up to 16 subrecord IDs can be specified.</td></tr>
<tr><td>STARTAFT=n</td><td colspan="2">defines a starting point in the input data set for copy records

This subparameter causes CX10CVBS to skip n records before beginning the copy operation. This subparameter must be preceded by a slash.</td></tr>
<tr><td>STOPAFT=n</td><td colspan="2">defines an ending point in the input data set for copy records

This subparameter causes CX10CVBS to stop the copy operation after n records are skipped or copied. This subparameter must be preceded by a slash if no STARTAFT parameter is defined. If a STARTAFT parameter is defined, the STOPAFT parameter must be preceded with a comma.</td></tr>
</table>

Because of the many different record types and ID combinations that can be specified, examples are shown in Table 8.

Table 8 Examples of PARM values used to copy records selectively

Types of records to be copied	PARM=Value	Notes
CMF CPM device activity user records (240-05)	TYPE=CPM,,SUB=05	SMF RECID 240 is defaulted
CMF CPM device activity user records (222-05)	TYPE=CPM,222,SUB=5	SMF RECID 222 is used
CMF IPM RMF type enqueue records (77)	TYPE=IPM,77	none
all RMF records (70 through 79)	TYPE=RMF	none
SMF user records (128)	TYPE=SMF,128	none
RMF CPU records (70)	TYPE=RMF,70	none
all CMF CPM records (240)	TYPE=CPM	SMF RECID 240 is defaulted
all CMF CPM records (230)	TYPE=CPM,230	SMF RECID 230 is used
CMF CPM global and LPA user records (240-11) and (240-16)	TYPE=CPM,,SUB=(11,16)	SMF RECID 240 is defaulted
CMF CPM global and LPA user records (241-11) and (241-16)	TYPE=CPM,241,SUB=(11,16)	SMF RECID 241 is used
RMF CPU and channel records	TYPE=RMF,(70,73)	none
CMF CPM type 70 and 74 records and associated user records (240s)	TYPE=CPM,(240,70,74),SUB=(01,5)	none
5th through 8th RMF CPU records	TYPE=RMF,70/STARTAFT=4,STOPAFT=8	none
copy CPM data removing CMF flags to make records identical to RMF-generated records.	TYPE=CPR	none
copy 20 CPM records	Type=CPM/STOPAFT=20	none

NOTE

The MVS operating system version number is stored in the product section of all 70 series records in packed format field SMF*xx*MFV.

CX10CVBS return codes

Table 9 describes the return codes that are issued by the I/O routine and control statement parsing routines of the CX10CVBS utility.

Table 9 CX10CVBS return codes

Return code	Description
00	normal
04	no records copied
06	VSAM problem encountered
08	SYSUT1 open failed
12	obtain failed
16	invalid keyword
20	invalid TYPE= option
24	invalid delimiter
28	non-numeric record ID
32	invalid record ID for IPM or CPM
36	invalid record ID for RMF or SMF
40	invalid subrecord ID for IPM or CPM
44	reserved
48	list exceeds 16 elements

How data is copied using the CX10CVBS utility

The CX10CVBS utility copies only valid blocks of data. Blocks with an invalid BDW, RDW, or SDW are skipped. The bad block is snapped to ddname //SNAPVBS (and //SNAPREC, if present). Bad block descriptor words are defined as follows:

- **BDW**—The second halfword is not zeroes. The first halfword is not greater than 8.
- **RDW**—The first halfword is not greater than 4.
- **SDW**—Spanned records are out of sequence.

A SYNAD exit is also used, so blocks that encounter an I/O error are treated as bad blocks; that is, copying is suppressed and the block is snapped to ddname SNAPVBS, if it is present.

RMF (70 through 79) and CMF MONITOR (subtypes 00 through 99) records are copied only if the length and count of the triplets agree with the RDW. Short records are snapped to ddname SNAPREC, if it is present.

CX10CVBS copies records only up to and including the track pointed to by the DS1LSTAR field of the DSCB. This procedure eliminates the possibility of a missing EOF when old or bad data, or both, might be copied.

Concatenated data sets are copied in the order of concatenation. If any data set in a concatenation has a missing EOF with no intervening TCLOSEs, the entire data set is bypassed, and the next one in the concatenation is processed.

Some valid blocks of data that follow a bad block can be skipped, which can happen with concatenated data sets and with numerous spanned records (caused by small block sizes). Eliminating these conditions reduces the possibility of skipping valid blocks due to the presence of bad blocks.

CX10CVBS supports recovery of VSAM-created files.

NOTE

To select records by date or by system ID, use the SMF utility IFASMFDP.

Chapter

4

Producing and using Analyzer reports

Once the CMF MONITOR Extractor or RMF has gathered data, the CMF MONITOR Analyzer can be used to produce reports and graphs. This chapter discusses information about how reports are generated, using and interpreting reports, and generating and defining Analyzer JCL. Also explained in this chapter is how to write your own programs to process Extractor data.

How reports are generated

Reports are generated when Analyzer JCL is submitted as a batch job. The Analyzer reads records produced by the CMF Extractor or RMF from either SMF or CMF data sets and filters, calculates, and formats the data into reports or graphs.

NOTE

A report is not created if a record type required for the report is missing.

SMF data is also available in the XDS data buffer. For information on activating XDS, see the *CMF MONITOR Customization Guide*.

Analyzer JCL contains DD statements and two types of Analyzer control statements:

- **general**—used to establish global characteristics for reports

  Some statements can also be used to affect specific reports. (See "Using general control statements" on page 96 for more information.)

- **report**—used to define specific reports or graphs to be produced

  (See "Using report control statements" on page 97 for more information.)

By defining control statements with specific parameters, you can generate reports that are customized to your needs.

How you can use Analyzer JCL and control statements is described in this chapter. Specific information about the purpose and function of each control statement and its parameters is given in Chapter 7, "Analyzer control statements."

Generating JCL to produce Analyzer reports

Analyzer JCL statements are used to

- define how a report batch job should run in your system
- point to the data set containing Extractor records for reports
- direct diagnostic and error messages should your job encounter problems
- specify where report output should be directed, as well as other batch reporting variables

You can generate JCL that produces CMF Analyzer reports by using the ISPF interface. To access this interface, invoke your MAINVIEW CLIST. The first panel you see looks similar to the panel in Figure 6.

Figure 6 MAINVIEW Selection Menu

```
--------------------------- MAINVIEW Selection Menu ---------------------------
OPTION  ===>                                                  DATE   -- YY/MM/DD
                                                              TIME   -- 14:20:55
      0    Parameters and Options                             USERID -- BCVAXT1
      E    Alerts and Alarms                                  MODE   -- ISPF 4.8
      P    PLEX Management (PLEXMGR)
      U    Utilities, Tools, and Messages

   Solutions for:
      A    Automated Operations
      C    CICS
      D    DB2
      I    IMS
      L    Linux
      N    Network Management
      S    Storage Management
      T    Application Management and Performance Tuning
      W    WebSphere and MQSeries
      Z    z/OS and USS
```

All BMC Software MAINVIEW products are accessed from this panel. To display a menu that includes choices for the CMF Analyzer, select option Z (z/OS and USS) from the MAINVIEW Selection Menu. The z/OS and USS Solutions panel is displayed, as shown in Figure 7.

Figure 7 z/OS and USS Solutions panel

```
-------------------------- z/OS and USS Solutions --------------------------
OPTION  ===>                                                  DATE   -- YY/MM/DD
                                                              TIME   -- 14:22
   Performance                                                USERID -- BCVAXT1
      1  MVzOS       MAINVIEW for z/OS                        MODE   -- ISPF 4.8
      2  MVUSS       MAINVIEW for Unix System Services
      3  CMF         CMF MONITOR
      4  SYSPROG     MAINVIEW SYSPROG Services

   Operations
      5  CSMON       Common Storage Monitor
      6  CMFMON      CMFMON realtime analysis
      7  CMFUTIL     CMF Extractor Online Utilities
      8  ANALYZER    Generate CMF Analyzer batch reports
      E  ALERTS      Alert Management

   General Services
      M  MESSAGES    Messages and Codes
      P  PARMS       Parameters and Options
```

From the panel as shown in Figure 7, select the **ANALYZER** option (option **8**), to display the main menu for generating JCL and control statements for CMF Analyzer batch reports. This menu is shown in Figure 8.

Figure 8 Produce CMF Analyzer Batch Reports main menu

```
---------------------- Produce CMF Analyzer Batch Reports --------------------
Option  ===>

    0  Setup          Set up CMF Analyzer JCL
    1  Input/Output   Specify input and output of CMF Analyzer
    2  Reports        Select CMF Analyzer reports
    3  Filter         Filter input data for Reports
    4  Generate       Generate CMF Analyzer JCL
    5  Edit/Submit    Edit/Submit existing CMF Analyzer JCL

    X  Exit           Terminate
```

Using the report generation panels

You can select any of the options on the main menu at any time.

The first time that you use these panels, you cannot select option 4 (Generate) until you have first saved the information in panels 0 (Setup), 1 (Input/Output), and 3 (Filter). These panels require information that is specific to your environment. Saving information in these panels helps prevent potential JCL errors.

Each of the panels provides the JCL generator with information needed to produce the reports you request in the format that you specify. Each panel is described in "Panels for generating CMF Analyzer JCL" on page 76.

Panels for generating CMF Analyzer JCL

The following panel is displayed when you select option **0** (**Setup**) from the main menu for the interface.

Figure 9 Panel for setting up CMF Analyzer JCL

```
 ----------------------- Set Up CMF Analyzer JCL ------------------------------
 Command ===>

 Job Statement Information:
  ===> //ANALYZER JOB (NNNN),'CMF ANALYZER',CLASS=F,
  ===> //          NOTIFY=XXXX,MSGCLASS=Z
  ===> //*
  ===> //*JOBPARM R=XX,S=SYSX
  ===> //*

  Company Name     ===>
                       (Specify a report header of up to 52 characters)

  Company Address  ===>
                       (Specify a subheader of up to 40 characters)

  JCL Generator Work Library ===> 'BMVSLC.BBSAMP'
     Temporary Workfile Unit ===> VIO
             Include STEPLIB ===> YES

Press END to save changes and return to the previous panel
Type CANCEL to return to the previous panel without saving changes
```

This panel is used for creating a jobcard for the Analyzer job and for specifying the header and subheader to be used on all reports. This panel also contains a field for specifying a JCL Generator work library. The default value for this field is *hilevel*.BBSAMP, but you can change it to conform to your site's naming conventions.

Values for the fields on this panel need to be specified once, and then it is unlikely that they will have to be changed.

Specifying Analyzer input and output

The following panel is displayed when you select option **1** (**Input/Output**) from the main menu for the interface.

Table 10 Panel for specifying source of CMF Analyzer data

```
-------------------- Specify Source of CMF Analyzer Data --------------------
Command ===>

CMF Record Type              ===> 240     REGION Size (in K) ===> 6000
Data Type                    ===> CPM     (C)PM, (I)PM, (R)MF

Reports on Multiple Systems ===> SEPARATE   (S)EPARATE, (C)OMBINED

Source of Input Data ===> DATASET   (XDS     - Cross-system data server buffer)
                                    (DATASET - Extractor output data set)

Input Data Set Names (not used if XDS is the source of input data)
                     ===> 'SYS1.MAN1'
                     ===>
                     ===>
                     ===>
                     ===>
                     ===>
                     ===>
                     ===>
                                   Validate Data Sets ===> YES  (YES,NO)

Press END to save changes and continue with the next panel
Type CANCEL to return to the previous panel without saving changes
```

Use this panel to specify the CMF record type used by your site, whether the data is CPM, IPM, or RMF, and whether reports from multiple systems should be separate or combined. This panel also enables you to specify which of the three data sources (XDS data, data from Extractor output data sets, or SMF data sets) is used as the source of Analyzer reports.

- If XDS is specified, the EXTDATA DD statement is not used in the JCL, and report data is obtained from the XDS data buffer. For more information on using XDS data, see Chapter 9, “Using the CMF MONITOR APIs.”
- When DATASET is specified, an EXTDATA DD statement is generated for each data set that is used as input for the Analyzer reports.

NOTE

The **Validate Data Sets** field provides a check of the input data sets. If you specify YES, you will not be able to exit this panel unless all of the input data sets exist.

When you press **End**, the panel for specifying your report output conditions is displayed, as shown in Figure 10.

Figure 10 Panel for specifying CMF Analyzer output destination

```
------------------- Specify CMF Analyzer Output Destination ------------------
Command ===>


Sysout Class        ===> *
                        (Specify a SYSOUT class for report output)

Data Set Name       ===>
                        (Specify a data set name for report output - optional)

Volume Serial       ===>
                        (If not cataloged)

Press END to save changes and return to the Primary Menu
Type CANCEL to return to the Primary Menu without saving changes
```

Use this panel to specify a SYSOUT class for Analyzer messages, the CMF Log, and Analyzer report output. In addition, you can specify a data set name for just your report output. This arrangement is useful for generating reports to be converted to spreadsheets, as described in Chapter 5, “Using the Analyzer Spreadsheet Converter.”

NOTE

The Analyzer has been enhanced so that you can send data to both SYSOUT and a data set.

If you only want to keep your report output, specify a data set name and a SYSOUT class (such as Z), which is purged automatically.

Selecting Analyzer reports

Figure 11 is displayed when you select option **2** (**Reports**) from the main menu for the interface.

Figure 11 Panel for selecting CMF Analyzer reports

```
--------------------------- CMF Analyzer Report List ------------- Row 1 of 36
Command ===>                                                    Scroll  ===> CSR
 Primary Commands: ALL , DEFAULT , (L)OCATE , NONE, SORT NAME or INCLUDE
 Line Commands: B - Browse report parameters   H - Help   I - Include report
               S - List/Update report parameters   X - Exclude report

 Press END to save changes and return to main menu
 Type CANCEL to return to the main menu without saving changes
                                                                          Update
LC Name      Include   Report Title                                       Parms
-- --------  --------  -------------------------------------------------  -----
   AUXSTOR     YES     Auxiliary Storage Report                            NONE
   CACHEACT    YES     Cache Subsystem Overview Report                     NO
   CFACT       NO      Coupling Facility Activity Report                   NONE
   CHANNEL     YES     Channel Path Activity Report                        NO
   CMFSTAT     YES     CMF Records Statistics Report                       NONE
   CMFSUM      YES     CMF Summary Report                                  YES
   COMMSTOR    YES     Common Storage Usage Summary Report                 NO
   CPU         YES     CPU Utilization Report                              NONE
   CPUCON      YES     Processor Concurrency Report                        NONE
   CRYPTO      YES     Cryptographic Hardware Activity Report              NONE
   DASD        YES     Direct Access Report                                NONE
   DEVACT      YES     Device Activity Report                              NO
   DOMINO      NO      Lotus Domino Server Summary and Detail Report       NO
   ENQUEUE     YES     Enqueue Conflict Report                             NO
   ESS         NO      ESS Statistics Report                               NONE
   EXCEPTS     NO      Exceptions Subreport                                NO
   EXCEPTS     NO      Exceptions Trace Detail Report                      NO
   FICONSW     NO      FICON Director Activity Report                      NONE
   GRAPH       NO      Various Graphical Reports                           NO
   GRAPH       NO      Graphics Trace Detail Report                        NO
   HFS         NO      HFS Statistics Report                               NONE
   HTTP        YES     HTTP Server Summary and Detail Report               NO
   IOQ         YES     I/O Queuing Activity Report                         NO
   LINKPACK    NO      Link Pack Area Report                               NO
   OMVS        NO      OMVS Kernel Activity Report                         NONE
   PERFSUM     YES     Performance Summary Report                          NO
   PROTKEY     YES     CPU Utilization by Protect Key Report               NO
   PRSM        YES     Logical Partition Report                            NO
   SHARDEV     YES     Shared Device Activity Report                       NO
   SRM         YES     System Resources Manager Report                     NONE
   STORAGE     YES     Storage Management Report                           NONE
   TRACE       YES     Trace Report                                        NO
   TSOPERF     YES     TSO Command Summary Report                          NO
   TSOUSER     YES     TSO User Summary Report                             NO
   VIRTSTOR    YES     Virtual Storage Activity Report                     NO
   VOLSER      NO      Direct Access Report Plot of Volume                 NO
   WLMGL       YES     Workload Manager Goal Mode Report                   NO
   XCF         YES     Cross-System Coupling Facility Report               NO
```

This panel lists all possible reports. You can include or exclude a report by using the I or X line commands. After a report is included, and if it has parameters, you can use the S line command to select that report for further modification. If an included report does not have parameters, or if its parameters are not modified, a default version of that report is produced.

NOTE

Currently, only the CMF Summary Report can be modified by using this panel. To modify other reports, you must edit the appropriate control statement after the JCL has been generated.

Example of a report parameter panel

The following panel is displayed when you type **S** next to CMFSUM on the panel for selecting Analyzer reports.

Figure 12 Panel for specifying CMFSUM parameters

```
------------------------ CMFSUM MEASURES Selection List ---- Row 1 to 29 of 45
Command ===>                                                Scroll  ===> PAGE

 Interval ===> E          (E)XTRACTOR, hh:mm:ss,
                          (H)OURLY, (D)AILY, (W)EEKLY,
                          (M)ONTHLY, (Q)TRLY, (S)EMIANNL, (F)OREVER
 Measures ===>            (A)LL, (R)MF, or blank to include 1 or more from list

 Primary Command: ALL , DEFAULT , (L)OCATE , NONE , SORT MEASURE or INCLUDE
 Line Commands: I - Include Measure    X - Exclude Measure

 Press END to save changes and return to the previous panel
 Type CANCEL to return to the previous panel without saving changes
 Type DEFAULT to include the original set of measures

LC Measure   Include  Corresponding Report Fields
-- --------  -------  -------------------------------------------------------
   AFQUEUE     NO     Average Available Frames Queue
   APPCAVG     YES    APPC Average
   APPCMAX     YES    APPC Maximum
   AVGREADY    NO     Average Ready Queue
   BATCHAVG    YES    Batch Average
   BATCHMAX    YES    Batch Maximum
   CAPRATIO    NO     Average Capture Ratio
   CHPUTIL     NO     Channel Path Utilization Rate
   CHPBUSY     NO     Channel Path Busy
   CPUBUSY     YES    CPU Busy
   CPUBZMVS    NO     MVS CPU Busy
   CPUSERV     NO     CPU Service Rate
   CSALLOC     NO     Average CSA Allocated
   DASDRATE    YES    DASD Rate
   DASDRESP    YES    DASD Response Time
   DPAGING     YES    Demand Paging
   ECSALLOC    NO     Average ECSA Allocated
   EPGRATE     NO     Expanded Storage Page Rate
   ESFRAME     NO     Expanded Storage Frames
   ESQALLOC    NO     Average ESQA Allocated
   EXCPRATE    NO     Average EXCPs Rate
   FIXFRAME    NO     Average Fixed Frames
   HIGHUIC     NO     High Unreferenced Interval Count
   HSFRAME     NO     Average Hiperspace Frames
   INTERVAL    YES    Interval HH.MM.SS
   IOSERV      NO     I/O Service Rate
   LPARDISP    YES    LPAR Dispatch Percentage
   MIGRAGE     NO     Migration Age
   MIGRATE     NO     Migration Rate
```

Use this panel to select parameters that modify the CMF Summary Report. After you select the interval and the measures you want, press **End** to return to the Report List, where you can continue selecting reports to be included in the JCL report list.

Filtering input data

The following panel is displayed when you select option **3** (**Filter**) from the main menu for the interface.

Figure 13 Panel for specifying filters on input data

```
------------------------ Filter Input Data for Reports -----------------------
Command ===>

Start Date (dd mmm yyyy) ===>                   End Date ===>
Start Time (hh:mm:ss)    ===>                   End Time ===>

Report Cycle    ===> ALL
    (All,DAILY,WEEKLY,BIWEEKLY,MONTHLY,WORKWKLY,WEEKENDS)

                   Start shift (hh:mm:ss)           End Shift (hh:mm:ss)
Report Shift 1  ===>                            ===>
Report Shift 2  ===>                            ===>
Report Shift 3  ===>                            ===>

System Identification ===> ALL       (ALL, SYSNAME, SYSID)
===>            ===>            ===>            ===>            ===>
===>            ===>            ===>            ===>            ===>
===>            ===>            ===>            ===>            ===>
===>            ===>            ===>            ===>            ===>
===>            ===>            ===>            ===>            ===>
===>            ===>            ===>            ===>            ===>

Press END to save changes and return to the previous panel
Type CANCEL to return to the previous panel without saving changes
```

This panel allows you to filter your input data so that reports include only the specified dates, times, shifts, and cycles. You can also specify particular SYSNAMEs or SYSIDs to be included in your reports.

Generating the JCL

The panel in Figure 14 on page 83 is displayed when you select option **4** (**Generate**) from the main menu for the interface. You will be able to use this interface only if you have saved information from panels 0, 1, and 3.

Figure 14 Panel for submitting Analyzer JCL

```
----------------------- Generate CMF Analyzer JCL ----------------------------
Command ===>

JCL Data Set           ===> 'CXA40.CAENG.UBBSAMP'

JCL Member Name        ===>           (1-8 character member name)

Replace JCL Member?    ===>           (YES,NO)

JCL Member Description ===>

Edit generated JCL ===> YES           (YES, NO)
            (NO submits batch job when you press ENTER)
            (YES displays edit panel when you press ENTER)

 Press ENTER to generate JCL
 Press END to save changes and return to the previous panel
 Type CANCEL to return to the previous panel without saving changes
```

This panel allows you to create a new JCL member that includes CMF Analyzer parameters, based on the information you provided in the previous panels.

To create a new member, specify a name for that member in the **JCL Member Name** field and a description in the **JCL Member Description** field, and then press **Enter**. If you specified NO in the **Edit generated JCL** field, the batch job is submitted and you are returned to the main menu. If you specified YES in the **Edit generated JCL** field, you are placed in an edit session for the data set member that you specified.

Listing previously specified members

The following panel is displayed when you select option **5** (**Edit/Submit**) from the main menu of this interface.

Figure 15 JCL Member List (left half)

```
------------------------------- JCL Member List ------------------- Row 1 of 1
Command ===>                                                Scroll  ===> CSR


 Line Commands: B - Browse JCL   DEL - Delete JCL   E - Edit JCL
                SUB - Submit JCL


Press END to return to main menu
                                                                          >>>
LC Member   Description                              Date       Time
-- -------- ---------------------------------------- ---------- --------
   SAMP1    Default Reports                          1996/03/29 10:29:21
******************************* Bottom of data********************************
```

If you scroll to the right, you will see additional information about the member, as shown in Figure 16.

Figure 16 JCL Member List (right half)

```
------------------------------- JCL Member List -------------------------------
Command ===>                                                   Scroll  ===> PAGE

 Line Commands: B - Browse JCL   DEL - Delete JCL   E - Edit JCL
                SUB - Submit JCL

 Press END to return to main menu
                                                                       >>>
LC  Member   Description                             Date       Time
--- -------- ---------------------------------------- ---------- --------
    SAMP1    'SLC1.BBSAMP'                                                SLC1
 ******************************* Bottom of data********************************
```

This panel allows you to select a report from a list of those previously set up. If you have a set of reports that you want to run multiple times, you can use this panel to submit your job directly, without having to use other panels in this interface.

Defining Analyzer JCL manually

If you decide not to use the JCL generator described in “Generating JCL to produce Analyzer reports” on page 74, you can create your own JCL.

A sample JCL member is shipped with CMF MONITOR and is discussed in “Using the default Analyzer JCL member” on page 85. This sample member contains all but two of the JCL statements that are used by the Analyzer.

One of the JCL statements that is not included in the sample member is needed only when producing graphics reports on a JES2 system; see “Producing graphics reports on a JES2 system” on page 91 for more information about this statement.

The other JCL statement that is not included in the sample member is needed only when printing reports or graphs on a laser printer; see “Defining your report output to a laser printer” on page 91 for more information about this statement.

Using the default Analyzer JCL member

The *hilevel*.UBBSAMP data set contains a sample Analyzer JCL member, called CMFJANL, that comprises a starter set for initial execution of the CMF MONITOR Analyzer. The CMFJANL member is shown in Table 11.

The JCL statements are described in Table 12 on page 86.

Table 11 Example of sample JCL member CMFJANL

```
//JOBCARD JOB                                                         00010000
//*                                                                   00020000
//*------------------------------------------------------------------ 00030000
//*                                                                   00040000
//*   SAMPLE JCL FOR EXECUTING THE CMF ANALYZER.                      00050000
//*                                                                   00060000
//*   REVIEW THE JCL FOR APPLICABILITY TO YOUR INSTALLATION           00070000
//*   STANDARDS                                                       00080000
//*                                                                   00090000
//*   CHANGE ?BBCHILV TO THE HIGH-LEVEL QUALIFIER YOU CHOSE FOR       00100000
//*   THE CMF MONITOR LIBRARIES.                                      00110000
//*                                                                   00120000
//*   CHANGE ?BBUNIT ON THE UNIT= KEYWORD TO MATCH YOUR               00130000
//*   SITE'S STANDARD.  NOTE: BMC SOFTWARE RECOMMENDS THAT            00140000
//*   YOU OMIT THE DMSSMAIN DD STATEMENT IN ORDER TO                  00150000
//*   IMPROVE JOB RUN TIME PERFORMANCE THROUGH THE USE OF HIPERSPACE. 00160000
//*                                                                   00170000
//*   CHANGE ?BBASMFID TO THE SMF ID (SYSTEM ID) OF THE TARGET        00180000
//*   SYSTEM.                                                         00190000
//*                                                                   00200000
//*------------------------------------------------------------------ 00210000
//*                                                                   00220000
//CMFRPTS  EXEC PGM=CMFANLYZ,REGION=6M                                00230000
//*                                                                   00240000
//STEPLIB  DD   DISP=SHR,                - ANALYZER LOAD LIBRARY      00250000
//         DSN=?BBCHILV.BBLINK                                        00260000
//*DMSSMAIN DD   UNIT=?BBUNIT,            - DMSS WORK FILE            00270000
//*         SPACE=(CYL,(10)),                                         00280000
//*         DISP=NEW                                                  00290000
//EXTDATA  DD   DISP=SHR,                - EXTRACTOR INPUT DATA       00300000
//         DSN=?BBCHILV.SYS?BBASMFID.CPMOUT1                          00310000
//*IPSLIB  DD   DISP=SHR,                - INPUT FOR GRAPH TYPE=DOMAIN 00320000
//*        DSN=SYS1.PARMLIB                                           00330000
//SYSIN    DD   DISP=SHR,                - ANALYZER CONTROL STATEMENTS 00340000
//         DSN=?BBCHILV.SYS?BBASMFID.UBBPARM(ANLYSAMP)                00350000
//*CMXREC  DD   DISP=SHR,                - COPY FILE, IF USING        00380000
//*        DSN=?BBCHILV.SAVEDATA         - CMFREC STATEMENT           00390000
//RPTCONTS DD   SYSOUT=*                 - REPORT TABLE OF CONTENTS   00400000
//CMFLOG   DD   SYSOUT=*                 - COLLECTION PHASE LOG       00410000
//SYSPRINT DD   SYSOUT=*                 - REPORTS                    00420000
//CMFPRINT DD   DUMMY                    - REPORTS                    00430000
//CMXTRACE DD   SYSOUT=*                 - TRACE DATA                 00440000
//SNAPS    DD   SYSOUT=*                 - ANALYZER SNAPS             00450000
//SNAPVBS  DD   SYSOUT=*                 - INVALID RECORD SNAPS       00460000
//SYSUDUMP DD   SYSOUT=*                 - ANALYZER ABEND             00470000
//                                                                    00480000
```

Descriptions of Analyzer JCL statements

All of the JCL statements shown in Table 11 on page 85 are described in Table 12.

Table 12 JCL control statements for the CMF MONITOR Analyzer (part 1 of 5)

JCL control statement	Description
//CMFRPTS EXEC	specifies the program name (CMFANLYZ) for the Analyzer, the region size, and other processing parameters BMC Software recommends a region size of 6 MB. The **PARM** field defines either CTRLSIZE or NLOG, or both. ■ PARM='CTRLSIZE=*xxx*K defines the amount of dynamic work area that the Analyzer is to use. BMC Software recommends omitting this parameter in most situations. For information about changing this value, see "Setting values of region, DMSS reserve, and CTRLSIZE" on page 91. ■ PARM='NLOG' eliminates printing of Extractor characteristics, the IPS, and the SRM Constants Report of the Collection Phase Log. See "Preliminary reporting information" on page 330 for more information about the Collection Phase Log reports. ■ PARM='CTRLSIZE=*xxx*K,NLOG' specifies the amount of dynamic work area and eliminates printing of Extractor characteristics.
//STEPLIB DD	required if *hilevel*.BBLINK is not in a LINKLIST data set; specifies a partitioned data set that contains the Analyzer load modules
//DMSSMAIN DD	(*optional*) when specified, defines a BDAM work file Allocate a minimum of 5 cylinders, and add 1 cylinder for every 10,000 records of input from EXTDATA. Secondary extents are ignored. The Analyzer issues messages that specify the number of spaces used. For reduced EXCP and improved performance, BMC Software recommends omitting this statement. If omitted or specified as DD DUMMY, a hiperspace is used as the work file.

Table 12 JCL control statements for the CMF MONITOR Analyzer (part 2 of 5)

JCL control statement	Description
//EXTDATA DD	defines the SMF or CMF data set containing Extractor records from which reports are to be produced If you want to use data in the XDS data buffer, this statement must be omitted. If you want to use data from an Extractor output data set, this statement is required. SMF data sets residing on DASD are VSAM data sets and cannot be concatenated. Records for the same interval must remain in the original order in which they were written. If, for some reason, the records become disordered, their original order can be restored by specifying the following statement, which will re-order the records by SYSID, Date, and Time: `SORT FIELDS=(15,4,CH,A,11,4,BI,A,7,4,BI,A),EQUALS` This statement might not work if the records have been sorted without the EQUALS parameter or if record types have been separated and are now being merged. If this SORT statement fails to re-order the records, you must return to the original data if it is available.
//SYSIN DD	defines input for the Analyzer control statements
//CMXREC DD	(*optional*) defines a sequential output data set where records accepted for analysis are to be written If you need this statement defined, you must remove the comment (*) character from the CMFJANL member.
//RPTCONTS DD	(*optional*) defines a print file for the Report Table of Contents If you use this statement, it should be inserted in front of the //CMFLOG DD and //SYSPRINT DD statements. You must define the optional operand OPTCD=J to cause laser printer control characters to be generated for the Table of Contents output. This statement allows the output to be printed on a laser printer. The CHARS= operand also needs to be defined for laser printing because it specifies the fonts to be used. This operand can be defined either in this statement or in a //*name*. OUTPUT statement.

Table 12 JCL control statements for the CMF MONITOR Analyzer (part 3 of 5)

JCL control statement	Description
//CMFLOG DD	defines an optional print file to direct the Collection Phase Log reports to an alternate data set These reports are automatically produced by the Analyzer, unless the SHIFT statement is defined with RPTS=INTERVAL or RPTS=DAILY. If the SHIFT statement is defined with RPTS=INTERVAL or RPTS=DAILY, the Collection Phase Log reports are automatically suppressed. If a //CMFLOG DD statement is not defined, the reports are written to a CMFLOG print file dynamically allocated by the system. By specifying the //CMFLOG DD statement, you can direct the Collection Phase Log reports to a different print file. You can define this print file to a valid data set name, or as DUMMY or NULLFILE. It can be useful to define a data set on DASD as the //CMFLOG DD print file destination, in case of an Analyzer error. If problems occur while using NLOG, rerun the Analyzer without NLOG defined to obtain the Collection Phase Log. The Extractor characteristics are sometimes helpful in diagnosing problems. Each time the Analyzer is run, this data set is overwritten. If used, the //CMFLOG DD statement should be inserted after the //RPTCONTS DD statement and before the //SYSPRINT DD statement. In addition, a subset of the Collection Phase Log reports can be suppressed by defining PARM='NLOG' on the EXEC statement. To send the Collection Phase Log report output to a laser printer, you must define an OPTCD=J parameter to the //CMFLOG DD statement. The OPTCD=J parameter causes laser printer control characters to be generated with the output, so it can be printed on a laser printer. The CHARS operand also needs to be defined for laser printing because it specifies the fonts to be used. This operand can be defined either in this statement or in a //*name* OUTPUT statement.

Table 12 JCL control statements for the CMF MONITOR Analyzer (part 4 of 5)

JCL control statement	Description

//SYSPRINT DD	defines a print file or an Analyzer output data set for the requested CMF MONITOR reports and graphs
	
	If you define an output data set, it must be allocated with the following characteristics:
	
	`RECFM=FBM`
	`LRECL=133`
	`DSORG=PS`
	
	The //SYSPRINT DD statement must be inserted after the optional //RPTCONTS DD and //CMFLOG DD statements.
	
	You must define the optional operand OPTCD=J to cause laser printer control characters to be generated for the reports and graphs output. This statement allows the output to be printed on a laser printer.
	
	The CHARS operand also needs to be defined for laser printing because it specifies the fonts to be used. This operand can be defined either in this statement or in a //*name* OUTPUT statement.

//CMFPRINT DD	(*optional*) defines an Analyzer output data set for the requested CMF MONITOR reports and graphs
	
	This statement can be used in conjunction with the SYSPRINT DD statement, so that report output can be directed to both SYSOUT and an output data set.
	
	The output data set must be allocated with the following characteristics:
	
	`RECFM=FBM`
	`LRECL=133`
	`DSORG=PS`
	
	The //CMFPRINT DD statement, if used, should be inserted after the optional //RPTCONTS DD and //CMFLOG DD statements.

//CMXTRACE DD	(*optional*) defines a print file for the CMF MONITOR Trace Report output
//SNAPS DD	(*optional*) defines a print file for snap dumps issued by the Analyzer
//SNAPVBS DD	(*optional*) defines a print file for snap dumps issued by the Analyzer
//SYSUDUMP DD	provides for a dump if a program fails

Table 12 JCL control statements for the CMF MONITOR Analyzer (part 5 of 5)

JCL control statement	Description

//CMFSTAGE DD
 	
//CMFSTAGO DD
	(*optional, unless dynamic allocation fails*) defines a temporary staging data set when RPTS=INTERVAL or RPTS=DAILY is defined on the SHIFT general control statement; (see "SHIFT" on page 299 for more information about the UNIT parameter relating to SHIFT)
	
	BMC Software recommends that you first try defining the UNIT= parameter on the SHIFT statement before defining these DD statements. These DD statements should be defined only if dynamic allocation of the temporary staging data sets fails when UNIT= is defined. There are three methods to define the statements:
	
	1. Define BOTH DD statements this way:
	
	`UNIT=VIO,DSN=&&CMFSTAGE,SPACE=(CYL,nn)` and no other parameters.
	
	The space should be the same as that on the DMSSMAIN DD statement.
	
	2. Add an IEFBR14 Step prior to the CMFANLYZ Step, which allocates a data set with the following characteristics:
	
	`RECFM=VBS`
	`LRECL=32760`
	`DSORG=PS`
	`BLKSIZE=8192`
	
	Point to this data set in the CMFSTAGE DD statement with DISP=SHR, and in the CMFSTAGO DD statement with DISP=MOD.
	
	3. Create a permanent data set with the following characteristics:
	
	`RECFM=VBS`
	`LRECL=32760`
	`DSORG=PS`
	`BLKSIZE=8192`
	
	Point to this data set in the CMFSTAGE DD statement with DISP=SHR, and in the CMFSTAGO DD statement with DISP=MOD.
	
	Occasionally, you might experience problems while trying to allocate the staging data sets on UNITS, either temporarily or dynamically, with error messages such as IEC141I RC013-34. These messages are IEC data set open/close DFP messages and can be dependent on
	
	■ how the units are managed (for example, are the volumes SMS managed?)
	■ are they mounted PRIVATE or STORAGE?
	■ other vendor software hooks relating to Open/Close processing
	■ the DFP maintenance release level in place, and so on
	
	When these problems happen, the best solution is to preallocate the staging data set with the vendor software hooks method (third bullet), and point CMFSTAGE and CMFSTAGO to it in the Analyzer JCL.

Producing graphics reports on a JES2 system

When you are producing graphics reports such as Kiviat or pie charts on a system that runs JES2, you need to include the following JES JOBPARM statement in the JCL:

```
/*JOBPARM  LINECT=62
```

This statement causes an override of the default lines-per-page value defined to your printer device. The override is necessary for the maximum size full-page graphs, which contain 62 lines of output, because it allows the entire graph to print on a single page. Otherwise, a portion of the graph, as much as is defined to your printer device's default setting, would print, a page eject would occur, and then the remaining portion of the graph would print on the next page.

Defining your report output to a laser printer

If you want to send your report output to be printed on a laser printer, you must define a //*name* OUTPUT statement containing a CHARS= parameter in the Analyzer JCL. This statement defines an output name and the fonts to be used by a laser printer. If this statement is defined, additional parameters must be defined to the //SYSPRINT DD statement as well.

Setting values of region, DMSS reserve, and CTRLSIZE

The amount of private area storage below 16 MB that is available to the CMF Analyzer is controlled by the MVS region size. The Analyzer storage management routines divide this available region size into three storage areas. The size of these storage areas is determined by three factors:

- REGION value specified on the JCL EXEC statement
- CTRLSIZE value specified in the **PARM** field of the JCL EXEC statement
- RESERVE value specified in the DMSS Analyzer control statement

Figure 17 illustrates how these storage areas are configured.

Figure 17 Storage area configuration

```
0K    +------------------------------------------------------------+
      |                                                            |
      |                CTRLSIZE=600K                               |
      |                                                            |
600K  +------------------------------------------------------------+
      |                                                            |
      |                                                            |
      |                                                            |
      |                                                            |
      |                DMSS INIT,RESERVE=2048                      |
      |                                                            |
      |                                                            |
      |                                                            |
      |                                                            |
      |                                                            |
2648K +------------------------------------------------------------+
      |                                                            |
      |                                                            |
      |                                                            |
      |                                                            |
      |                                                            |
      |                                                            |
      |                REGION=6144K or REGION=6M                   |
      |                                                            |
      |            The RESERVE= and CTRLSIZE= values               |
      |            are subtracted from the total region            |
      |            size and the remaining area is                  |
      |            allocated for DMSS buffers.                     |
      |                                                            |
      |            (6144K-600K-2048K = 3396K)                      |
      |                                                            |
6144K +------------------------------------------------------------+
```

The default values for RESERVE and CTRLSIZE are dynamically computed by the Analyzer as follows:

- CTRLSIZE is set to about 20% of the OS Region size, or to the amount of storage that is available if the Analyzer is dynamically invoked by some other program.

- RESERVE is set to about 40% of the rest of available storage. This figure will usually end up being about 33–38% of the OS Region size.

BMC Software recommends that you omit these parameters and let the Analyzer calculate the amounts. This action will almost always provide adequate resources for running your Analyzer jobs. Should you experience problems like those described in this section, simply increase the OS Region size.

BMC Software recommends a minimum region size of 1 MB, with at least 3 MB being preferable.

If you are producing a large number of reports or processing a large amount of data, or both, the values for REGION, RESERVE and CTRLSIZE might not be big enough. This situation can result in user abends U008, U100, or U0999, or system abends S80A, S878, or S106.

If you experience any of these user abends, follow these guidelines:

- First, set the REGION parameter on the EXEC statement in the Analyzer JCL to the maximum allowed by your installation, and rerun the Analyzer job. BMC Software recommends specifying REGION=6M, which is the value of the sample Analyzer JCL member CMFJANL, of *hilevel*.UBBSAMP.

NOTE

If REGION=0K or REGION=0M is specified on either the JOB card or the EXEC card for the Analyzer, CMF will not run. The LSQA will be insufficient to load programs.

- The RESERVE parameter of the DMSS Analyzer control statement determines the amount of storage available for DMSS control blocks and index areas. The storage that is required increases with the amount of input records that are to be processed. The RESERVE value should be increased if the Analyzer experiences S80A, S878, or S106 system abends, and if you have set your REGION size to its maximum. The default values computed by the Analyzer for various REGION sizes are

  — RESERVE=1000K, CTRLSIZE=184,320 bytes, and RESERVE=339,968 bytes

  — RESERVE=3000K, CTRLSIZE=593,920 bytes, and RESERVE=1,158,544 bytes

  — RESERVE=6000K, CTRLSIZE=1,208,320 bytes, and RESERVE=2,387,968 bytes

  Thus, if you were using a 6000K region when you had a problem, you need to specify a RESERVE greater than 2,400,000 bytes.

- The CTRLSIZE value that is specified in the **PARM** field of the JCL EXEC statement determines the amount of storage available for control blocks associated with each report request. A good guideline is to specify 1 K per report that is produced on the Report Table of Contents. This value should be increased if the Analyzer experiences U008 or U999 user abends, and if you have set your REGION size to its maximum.

NOTE

If the COMMSTOR control statement is specified, you should add an additional 100 K.

The SHIFT and SYSPLEX report control statements can have a dramatic effect on the number of report requests. For example, if you specify SHIFT `DINTV=(080000,96,001500), RPTS=SEPARATE`, and you specified 30 report control statements, 96 reports are produced for each report control statement specified. Therefore, 2880 (96 * 20) reports are produced, which means that you will need a CTRLSIZE value of at least 2.9 MB.

NOTE

Abends can occur when the RESERVE and CTRLSIZE values are too large relative to the region size.

Defining Analyzer control statements

All control statements appear either after the //SYSIN DD * statement or after the data set pointed to by the //SYSIN DD statement in the Analyzer JCL. Control statements are used to define global or specific report and graph characteristics to Analyzer batch jobs.

Using the default Analyzer control statement member

The *hilevel*.UBBPARM data set contains a sample Analyzer control statement member called ANLYSAMP, which comprises a starter set for initial execution of the CMF MONITOR Analyzer.

ANLYSAMP can be used to check the installation of the Analyzer, read Extractor data, and print a sampling of CMF MONITOR reports and graphs. It also provides you with a starting control statement set that you can modify for your particular site requirements.

Table 13 on page 95 contains an example of the ANLYSAMP member.

The reports produced by and function of each statement defined in ANLYSAMP, shown in Table 13 on page 95, are discussed in Chapter 8, “Analyzer reports.”

Table 13 ANLYSAMP control statement

```
***********************************************************************
*    GENERAL CONTROL STATEMENTS                                       *
***********************************************************************
*
     RECTYPE 240              -  set user record type(default)
     SYSPLEX TYPE=SYSNAME,RPTS=SEPARATE
*
***********************************************************************
*    REPORT CONTROL STATEMENTS                                        *
***********************************************************************
*
*******> GENERAL
*
     CMFSTAT                  -  CMF records statistics report
     CMFSUM                   -  CMF summary report
     PERFSUM                  -  performance summary report
*
*******> STORAGE
*
     AUXSTOR                  -  auxiliary storage report
     COMMSTOR REPORT=SUMMARY  -  common storage usage detail report
     SRM                      -  systems resources manager report
     STORAGE                  -  storage management report
     VIRTSTOR DETAIL=YES      -  virtual storage activity report
*
*******> CPU
*
     CPU                      -  CPU utilization report
     CHANNEL                  -  Channel Path Activity Report
     CPUCON                   -  processor concurrency report
     ENQUEUE THRESHLD=100     -  enqueue conflict report
     PROTKEY                  -  CPU utilization by protect key report
     PRSM                     -  logical partition report
     TRACE                    -  trace report
     XCF                      -  cross-system coupling facility report
*
*******> PERIPHERALS
*
     CACHEACT REPORT=SUBSYS   -  cache subsystem reports
     DASD                     -  direct access report
*     VOLSER vvvvvv,wwwwww  <=== enter volume(s) for detail DASD report
     DEVACT                   -  device activity report
     ESS                      -  ESS Statistics Report
     FICONSW                  -  FICON Director Activity Report
     IOQ                      -  I/O queuing activity report
*
*******> GRAPHS
*
     EXCEPTS INTERVAL=00:30:00,MIN=8,MAX=20,
        MEASURE=PAGESEC,
        ASSOC=(PAGEINS,PAGEOUTS),CPU=ALL
     GRAPH TYPE=PLOT,INTERVAL=00:10:00,
        MEASURE=(PAGESEC,PAGETIME,CPU),
        LIMIT=(10000,100000),CPU=ALL
*
*******> TSO
*
     TSOPERF                  -  TSO command/interval summary reports
     TSOUSER                  -  TSO user summary report
*
***********************************************************************
*    The following reports must always be                             *
*    combined for all systems in a sysplex.                           *
***********************************************************************
*
     SYSPLEX TYPE=SYSNAME,RPTS=COMBINED
*
     CFACT                    -  coupling facility activity reports
     SHARDEV                  -  shared devices actvity report
     WLMGL                    -  goal-mode workload reports
*
```

Specifying the CMF user record type

The ANLYSAMP control statement member causes CMF Analyzer to read user records with an ID of 240 (X'F0'). The value of 240 is the default user type defined to the Extractor under which user records are written to data sets; however, this value can be changed. (See “REPORTS” on page 290 for more information.)

For the Analyzer to read CMF user records, the same type defined to the SMFRECID parameter of the Extractor REPORT control statement must be defined to the Analyzer RECTYPE control statement. (See “RECTYPE” on page 288 for more information.)

NOTE

If you have another product that generates type 240 records, you must specify a record type other than 240 on both the RECTYPE Analyzer control statement and the SMFRECID parameter of the REPORT Extractor control statement. The XDS RECTYPE parameter might also need to be changed.

Specifying records from CPM or IPM monitoring modes

During a single batch job, the Analyzer can read records collected by either CPM or IPM mode, but cannot read records from both modes simultaneously. By default, the Analyzer reads type 70 series records and CMF type 240 user records that were collected in CPM mode. If you need the Analyzer to read IPM mode records or CMF user records gathered under an SMF record ID other than 240, define the RECTYPE Analyzer general control statement to your Analyzer JCL.

Using general control statements

The Analyzer’s general control statements establish global characteristics for reports, and some statements can be used to affect specific reports, as well.

Most general control statements appear at the beginning of the control statement set, directly after the //SYSIN DD * statement or in the data set pointed to by the //SYSIN DD statement. However, some statements can be used only within the report control statement set. They are fully described in the General Control Statements section in Chapter 7, “Analyzer control statements.”

The following list shows all Analyzer general control statements:

CMFREC	PERIOD
CYCLE	RECTYPE
DATA	SEVERITY
DATETIME	SHIFT
DMSS	SYSPLEX
HEADERS	

Using report control statements

Report control statements appear after the general control statement set. In general, the Analyzer's report control statements define specific reports to be produced and provide parameters for organizing or filtering report contents.

NOTE

The PERFORM statement does not cause a report to be produced but modifies other report statements.

The report control statements are fully described in "Report control statements" on page 210.

The following list shows all report control statements:

AUXSTOR	DEVACT	LINKPACK	TSOPERF
CACHEACT	DOMINO	OMVS	TSOUSER
CFACT	ENQUEUE	PERFORM (modifier)	VIRTSTOR
CHANNEL	ESS	PERFUM	VOLSER
CMFSTAT	EXCEPTS	PROTKEY	WLMGL
CMFSUM	FICONSW	PRSM	XCF
COMMSTOR	GRAPH	SHARDEV	
CPU	HFS	SRM	
CPUCON	HTTP	STORAGE	
DASD	IOQ	TRACE	

Using and interpreting reports

The information in the CMF MONITOR reports can be used to define thresholds for key system resources. The reports summarize performance data by job class or service class periods; they also report on DASD head movement activity, CPU usage, link pack area activity, and TSO usage data.

Some reports are produced automatically, such as the System Resources Manager (SRM) Constants Report. Automatically produced reports are found in the Collection Phase Log of the Analyzer output. See “Preliminary reporting information” on page 330 for more information.

Knowing what reports you need

The reports that Analyzer statements produce fall into different categories. Some reports belong to more than one category.

The following sections show the categories of reports and the Analyzer control statements that produce reports belonging to each category. See “Report control statements” on page 210 for more information about each control statement and the reports it produces.

WORKLOAD reports

The following reports are included:

OMVS
TSOPERF
TSOUSER
WLMGL

CPU reports

The following reports are included:

CPU
CPUCON
PROTKEY
PRSM

SYSTEM RESOURCE reports

The following reports are included:

AUXSTOR	STORAGE
COMMSTOR	TRACE
ENQUEUE	VIRTSTOR
LINKPACK	XCF
SRM	

DEVICE reports

The following reports are included:

AUXSTOR	ESS
CACHEACT	FICONSW
DEVACT	IOQ
DASD	SHARDEV

Web-related reports

The following reports are included:

DOMINO
HTTP

Miscellaneous reports

The following reports are included:

CMFSTAT	GRAPH
CMFSUM	HFS
EXCEPTS	TRACE

An example of each CMF MONITOR report is provided in Chapter 8, "Analyzer reports," as well as field descriptions and calculations. For instructions on capturing these reports, go to "Capturing CMF Analyzer and CMFMON reports" on page 107.

Interpreting report field information

Data field results can vary due to environmental factors such as the version of MVS running in your environment or the configuration of your system resources. An explanation of these variations, together with the field descriptions and calculations for all report data, are provided in Chapter 8, “Analyzer reports.”

In some cases, report fields contain data presented using one of the following conventions:

- **dashes (---)**

  If a report requires input from more than one Extractor record and one of them is not present, dashes (---) appear in the fields that need data from the missing record.

- scientific notation

  If a number is too large to be displayed in the space provided in a report field, the number is displayed in scientific notation. Scientific notation is provided only for nine-character or longer field values. The format for numeric values displayed in scientific notation is

  sn.n{.nnn}Eyxx

s	is the sign, either + or -
n.nnn	is a real number greater than zero (0) and less than ten (10)
y	is the sign of the exponent, either + or -
xx	is the exponent

  The precision of the real number varies depending on the size of the field; for example, the number 9.37E+07 is

$$9.37 \times 10^{7} = 93{,}700{,}000$$

- E’s (EEEEE...)

  If a value is too large to be displayed, even in scientific notation, it is represented as a string of Es.

Understanding report headings

A report heading is printed automatically at the top of each report page. Each heading includes the product name and version number, a report title, page number, and the current report date and time. All requested CMF MONITOR reports receive the heading illustrated in Figure 18.

Figure 18 Example of a standard report heading

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    WORKLOAD MANAGER MAP REPORT                  RPTSEQ      7 PAGE    181
BMC SOFTWARE, INC.                                           XYZ COMPANY                          REPORT DATE: DD MMM YY 13.26
ACTL 10 JUN YY 17.00.00  11 JUN YY 17.00.00            WORLDWIDE HEADQUARTERS                     SYSTEM ID: **ALL**  COMB-MVS

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  72-3/7,896/0/46.64
```

Descriptions of report headings fields

A description of each field in a report heading is included in Table 14.

Table 14 Field descriptions for a report heading (part 1 of 2)

Field name	Description
PRODUCED BY	name and version number of the product
REQD	requested beginning and ending date-time range
SHFT	day and time based on SHIFT control statement
ACTL	actual beginning and ending date-time range encountered ■ The first date-time pair under ACTL is the date and time of the first record encountered in the input data set that contained information for the report (see Figure 18). ■ The second or end date-time pair is the date and time of the last record encountered in the input data set that contained information for that report; the end date-time also includes the interval time for the last record. ■ For the Extractor Characteristics Subreport of the Collection Phase Log, the date-time of the first record is in effect from begin date-time and the end date-time is the date-time of the last record encountered; however, the end date-time does not include the interval time for the last record. ■ When using these date-times for DATETIME or CYCLE selection, selection criteria are based on the record start date-time.

Table 14 Field descriptions for a report heading (part 2 of 2)

Field name	Description
BASED ON	Records used for this report, in the format REC TYPE/# RECS/# SAMPLES/REC HOURS, where REC TYPE record type and subtype # RECS number of records # SAMPLES number of samples REC HOURS duration of recording period (to nearest hundredth of an hour)
(report title)	title of report, followed by user-generated data from title and location fields, as specified on the optional HEADERS control statement
RPTSEQ	sequence number of the report and page number
REPORT DATE	date and time control statements were processed
SYSTEM ID	system identifier ■ If the report contains data from a single MVS image, the sysname or the sysid appears in the **SYSTEM ID** field of the report heading. ■ If the report contains data from multiple MVS images, *MULTI* or **ALL** is printed in the **SYSTEM ID** field of each report heading.
(MVS system release number)	appears to the right of the **SYSTEM ID** field ■ If the report contains data from a single MVS release, the release number appears in this area. ■ If the report contains data from multiple MVS releases, COMB-MVS appears in this area.
REPORT CYCLE	report cycle based on CYCLE control statement

Writing your own programs to process Extractor data

You can write your own programs by using the SMF record format members presented in CMF MONITOR’s *hilevel*.BBSAMP data set.

Three groups of members in BBSAMP provide

- SAS
- C structures
- Assembler MACROs

For more information about CMF MONITOR’s SMF record formats, see Chapter 10, “Mapping CMF records created by CMF.”

Chapter

5

Using the Analyzer Spreadsheet Converter

The CMF Analyzer Spreadsheet Converter automatically changes CMF Analyzer report data into Microsoft Excel spreadsheets. The spreadsheets can be used for interactive, detailed analysis, creating graphs, or producing specialized reports. The Spreadsheet Converter can be installed on an unlimited number of PCs and used by any number of people in your organization.

The Spreadsheet Converter detects if you are using a language other than English, and displays the instruction screen in the same language as your copy of Excel, if possible.

NOTE

The Spreadsheet Converter is compatible only with Microsoft Excel 2000 and later, and does not need any additional hardware that is not part of the minimum requirements for running Excel.

The Spreadsheet Converter is a conversion tool only, designed to make viewing or manipulating the data in the CMF Analyzer or CMFMON reports easier. It *only* converts CMF Analyzer reports. BMC Software does not sell or support Microsoft Excel. Any questions about using Excel should be directed to Microsoft Corporation.

Installing the Spreadsheet Converter on your PC

The Spreadsheet Converter is distributed with CMF MONITOR as member CX98SSCX of the BBSAMP and UBBSAMP libraries.

To uninstall an old version

1 If an old version of the Spreadsheet Converter is installed on your system, you must uninstall it before installing the new version.

A Open the old Spreadsheet Converter (.xla) file.

B Under **Tools** on the Excel Menu bar, select **Customize** to open the Customize pop-up window.

C Under **Toolbars** in the Customize window, select the **BMC Software** option and click the **Delete** button on the right side of the window.

D Close the window and the Excel application.

To install the new version

2 Transfer the file CX98SSCX to your PC using any file transfer method, such as IND$FILE, TCP / IP file transfer protocol, or e-mail.

NOTE

Make sure that you specify a binary file transfer, since the the file is already in PC format in BBSAMP.

3 Rename the file CX98SSCX.XLA

4 You might want to save the file in its own directory or folder for easy access.

NOTE

Information about downloading the Spreadsheet Converter program is also contained in the *CMF MONITOR Customization Guide.*

Capturing CMF Analyzer and CMFMON reports

The Spreadsheet Converter accepts any report generated by the CMF Analyzer and exported CMFMON reports; however, only selected reports receive special formatting during the conversion. You can verify the exact list of reports by reviewing the Conversion Log for your converted report, or by reading a note attached to the Spreadsheet Converter. To display the note, start the Spreadsheet Converter and select the Insert / Note menu item. Any reports not on this list will be stored as a series of records in a spreadsheet. Such reports can be parsed into columns by way of the Excel *text-to-columns* feature.

Table 15 lists the CMF Analyzer reports (and their respective Analyzer control statements) that receive special formatting by the Spreadsheet Converter. Reports that are *not* listed in this table are loaded into a worksheet without reformatting or special processing.

Table 15 CMF Analyzer reports that receive special formatting by the Spreadsheet Converter

Report name	Analyzer control statement
CMF Summary Report	CMFSUM
CPU Utilization Report	CPU
Device Activity Report	DEVACT
I/O Queuing Activity Report	IOQ
Performance Summary Report	PERFSUM
Storage Management Report	STORAGE
Workload Manager Goal Mode Report	WLMGL

Capturing CMF Analyzer reports

In order to successfully convert a report, the SYSPRINT output from the CMF Analyzer run must be captured into a data set. This can be accomplished in one of the following ways:

- Use the SYSPRINT DD statement in the CMF Analyzer JCL to point to the data set that is to contain the reports.
- Use the CMFPRINT DD statement in your CMF Analyzer JCL, and have it point to the data set that is to contain your reports.
- Use the IBM Sysout Display and Search Facility (or equivalent) to copy the data from the SYSOUT queue to a data set.

Whichever method you use to capture the reports, be sure to include the carriage control information in column 1 of the file (RECFM=FBA or VBA). The Spreadsheet Converter requires this information in order to properly identify the pages of the reports.

Capturing CMFMON reports

The Spreadsheet Converter also loads reports captured from CMFMON.

The CMFMON EXPORT command has a *comma separated values* (CSV) option, which is a standard format for spreadsheet input. The Spreadsheet Converter loads these types of reports and applies standard Excel formatting. For additional information on CMFMON reports or the EXPORT command, please refer to the *CMF MONITOR CMFMON User Guide.*

Downloading reports to the PC

Reports captured on the mainframe must be transferred to the PC in order to be processed by the Spreadsheet Converter. You can transfer reports with IND$FILE or any other file transfer program. When transferring your report to the PC, you must make sure that

- the file is transferred in ASCII format
- CRLF marks are inserted at the end of each line
- the suffix .txt is used to designate CMF Analyzer report files
- the suffix .csv is used to designate CMFMON report files

If a new directory was created to contain the Spreadsheet Converter, it can also be used to receive these output files.

Running the Spreadsheet Converter

The Spreadsheet Converter (file CX98SSCX.XLA) runs as an Excel Add-in. The Spreadsheet Converter can be started by any of the following methods:

- Start Microsoft Excel, and then use File / Open to load the Spreadsheet Converter.

- Start Microsoft Excel, and then use Tools / Add-ins to add the Spreadsheet Converter.

- Double-click the CX98SSCX.XLA file in Explorer or File Manager.

Next, follow the procedures beginning on the next page.

1 After the Spreadsheet Converter is loaded, begin execution by selecting the BMC Software toolbar from the View menu item, as shown in Figure 19.

Figure 19 Select BMC Software toolbar from the Excel Menu

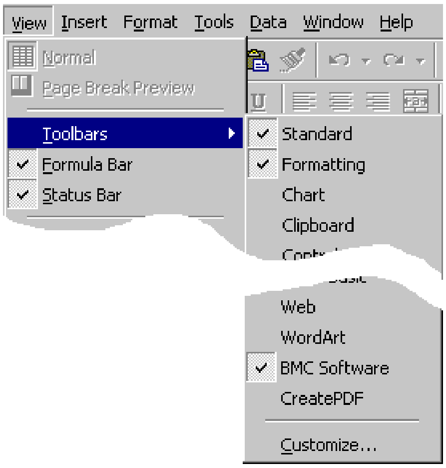

2 When you select **Instructions**, the BMC Software Spreadsheet Converter instruction screen is displayed, as shown in Figure 20.

Figure 20 Spreadsheet Converter instruction screen

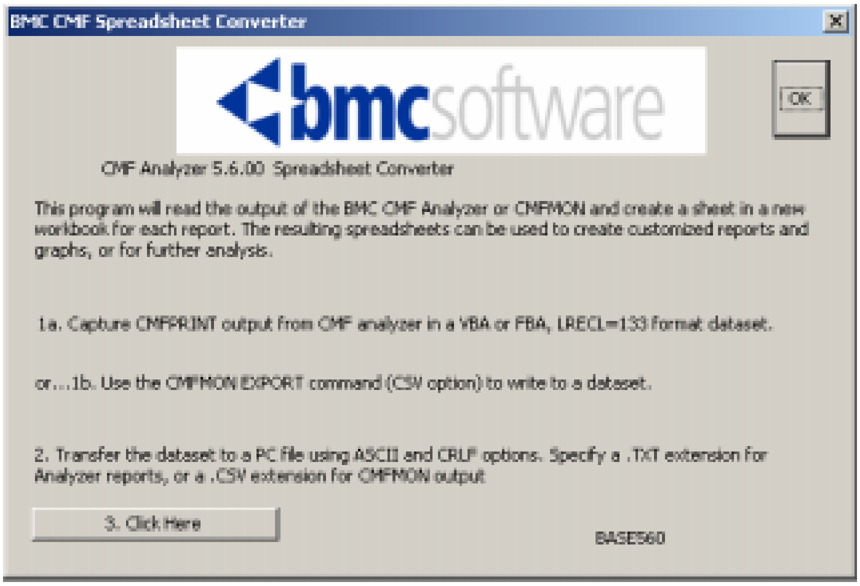

3 When you select **Click Here**, the Spreadsheet Converter displays a list of the files that you downloaded to your PC, as shown in Figure 21 on page 110.

Figure 21 Selecting the file containing reports to convert to Excel

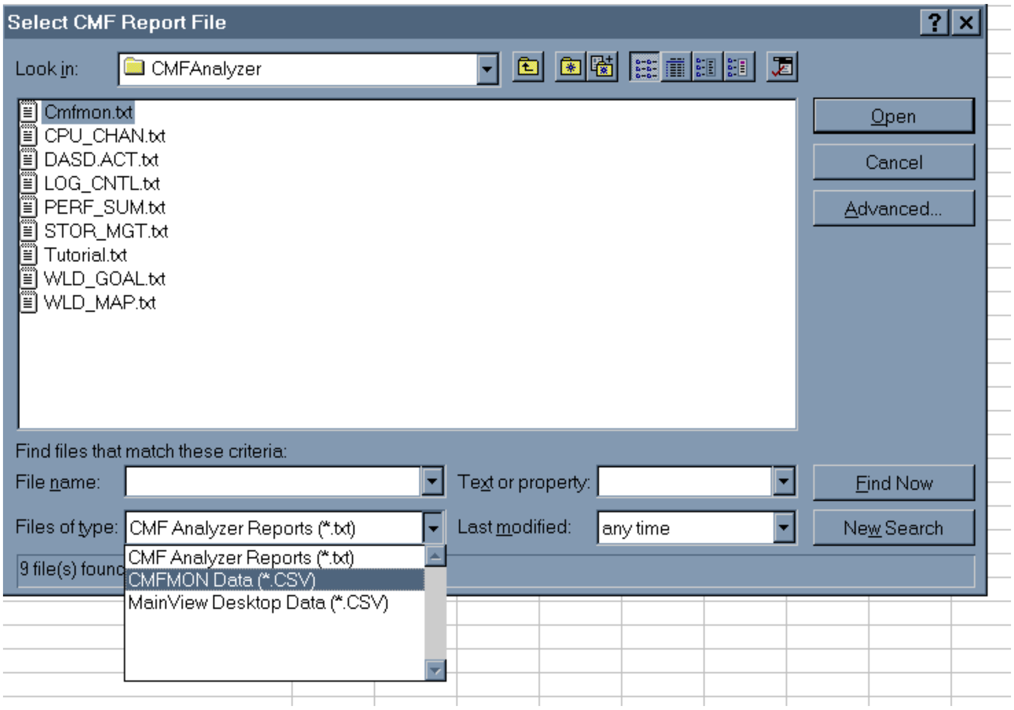

When you select a report, the Spreadsheet Converter

- creates a new workbook in which to store the results of the conversion
- automatically converts the reports in the file you select

NOTE

You can bypass the instructions by selecting the Convert CMF Reports button on the BMC Software toolbar, shown below.

Figure 22 BMC Software toolbar

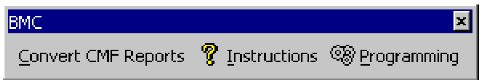

4 When the file contents have been converted, the Conversion Log is displayed, as shown in Figure 23 on page 111.

Figure 23 Conversion Log for converted reports

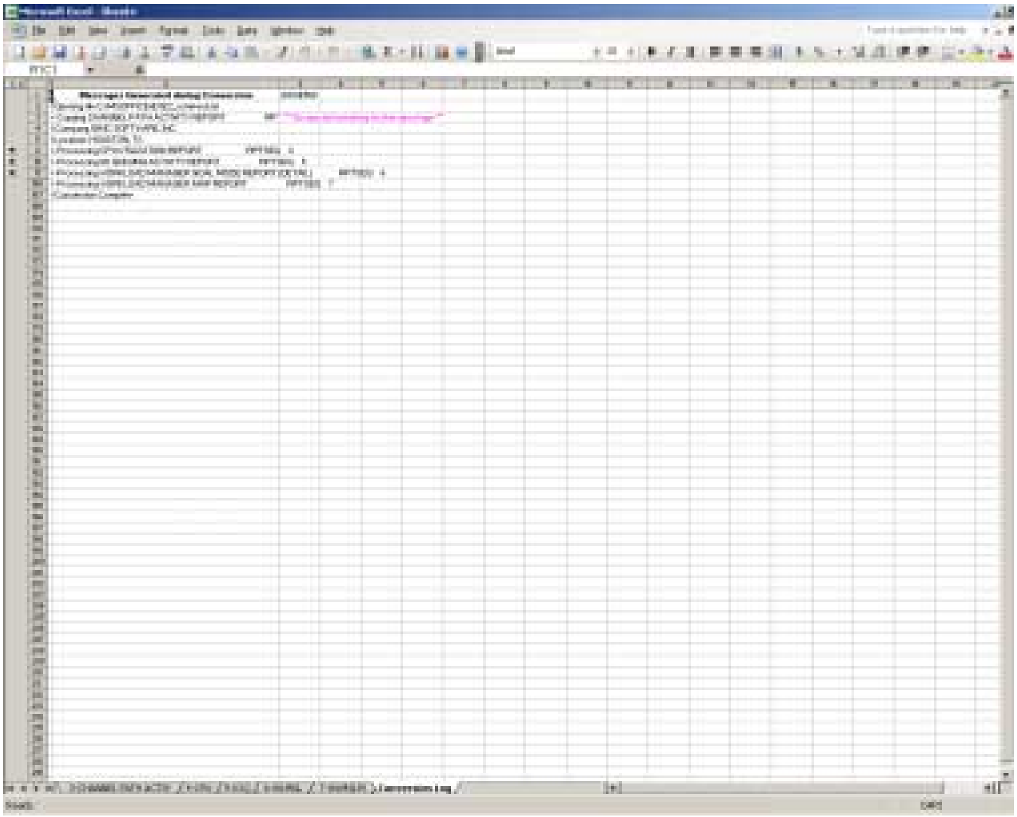

The Conversion Log is the last page in the workbook. The Conversion Log lists reports found in the input along with a description of how each one was processed. Any error messages produced during processing are also listed on this page. Additional information on any error messages you see can be displayed by selecting column 1 (hidden behind column 2).

5 Double-click on the mouse to completely display column 1. If any symbols were defined in a report, the names of the ranges defined can be found by expanding the outline (using the controls found in the left margin).

The Spreadsheet Converter can be invoked from Excel macros written in VBA, by coding a statement such as

```
Application.Run "cx98sscx!Main", "C:\Test\CMFRep.txt", "Test.xls"
```

where the two optional parameters are the file to be converted, and the name to be given to the resulting created workbook.

Spreadsheet converter output

When the Spreadsheet Converter is finished processing all of the reports in the input file, the results are stored in a set of pages in a new workbook. These reports can then be manipulated or exported to other programs, just like any other Excel spreadsheet.

For additional information about how to work with Excel spreadsheets, consult the Microsoft Excel documentation.

A brief tutorial

A file containing sample reports is included in the BBSAMP and UBBSAMP libraries. This tutorial uses this sample file to demonstrate how the Spreadsheet Converter changes your CMF Analyzer report data into spreadsheet format. The converted sample files also show how you can use Excel to manipulate the data to produce graphics or specialized reports.

1 Transfer the Spreadsheet Converter (CX98SSCX.XLA) to your PC.

   Using any file transfer method, be sure to download the converter program as binary. Refer to "Installing the Spreadsheet Converter on your PC" on page 106 for additional information.

2 Copy the sample file CX98REPG from the BBSAMP library to one of your data sets, or edit the copy in UBBSAMP.

3 Create a job card for CX98REPG, and modify the SYSUT2 DD statement to designate the data set that is to receive the output containing the CMF Analyzer reports.

4 Run the job.

   This task creates a file containing sample CMF Analyzer reports.

5 Download the output data set to your PC using any file transfer method.

   Be sure to transfer in ASCII mode and rename the file using the .txt extension. For this tutorial, you can name your file tutorial.txt.

6 Open the Spreadsheet Converter (CX98SSCX.XLA) in Excel by selecting **Add-ins** from the **Tools** menu.

7 Select **CMF Spreadsheet Converter** from the **Add-Ins Available** list box shown in Figure 24 on page 113.

Figure 24 Excel Add-Ins Available list box

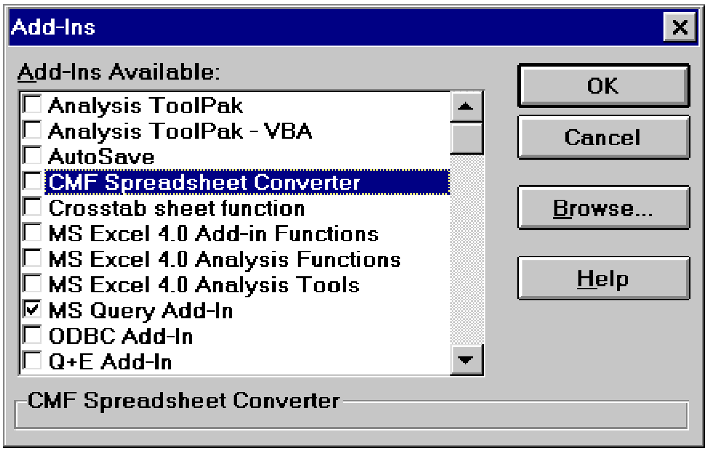

8 Run the Spreadsheet Converter by selecting **Convert CMF Reports**. A dialog box opens to display a list of your report files.

9 When you select your file (such as, tutorial.txt) from the list, it will automatically be converted into Excel spreadsheets using templates designed specifically for the CMF Analyzer reports.

10 When the conversion is complete, you should see the Conversion Log, with the names of the converted reports displayed on tabs at the bottom of the screen. The tabs can be scrolled using the arrows on the lower left of the screen.

11 Use your mouse to select the CMF Summary Report. The report displayed should look similar to Figure 25 on page 114. In this particular spreadsheet, a line graph was added to provide visual comparison between CPU, DASD, and TSO information.

Figure 25 Converted CMF Summary Report

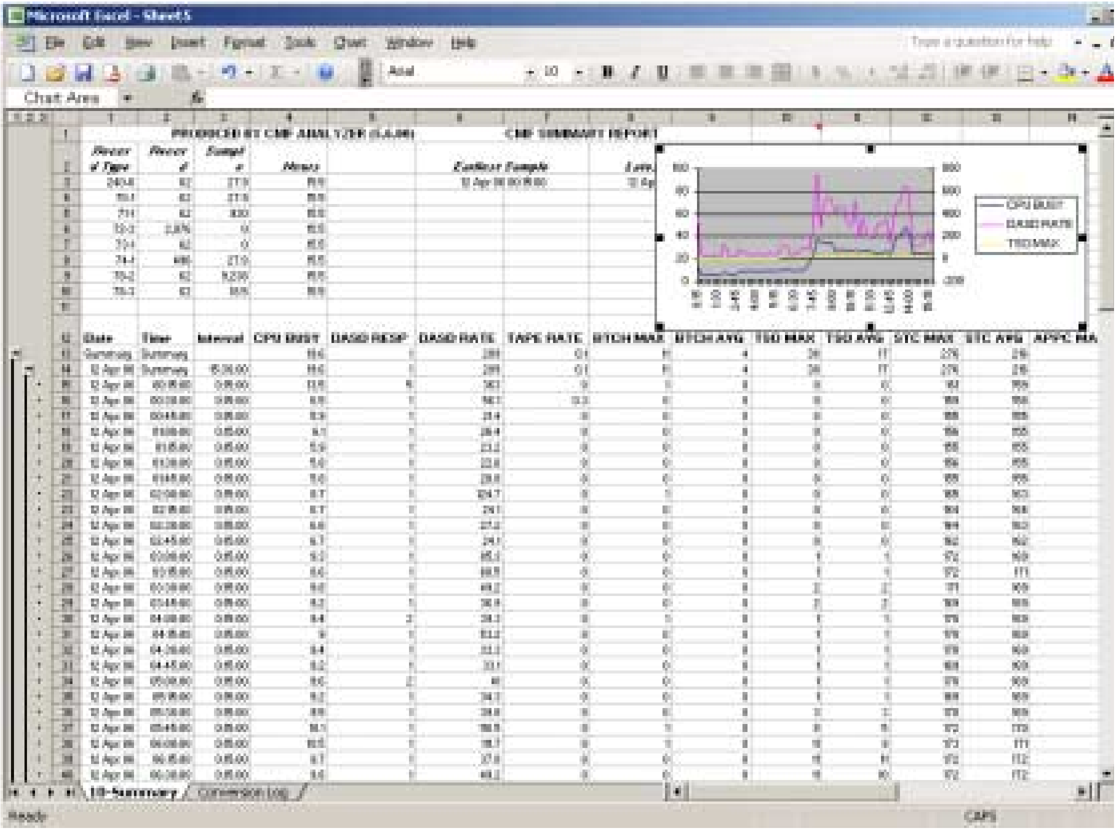

12 Select the tab for the CPU Utilization Report. The report displayed should look similar to Figure 26 on page 115. In this particular spreadsheet, a bar chart was generated to compare performance between systems.

Figure 26 CPU Utilization Report

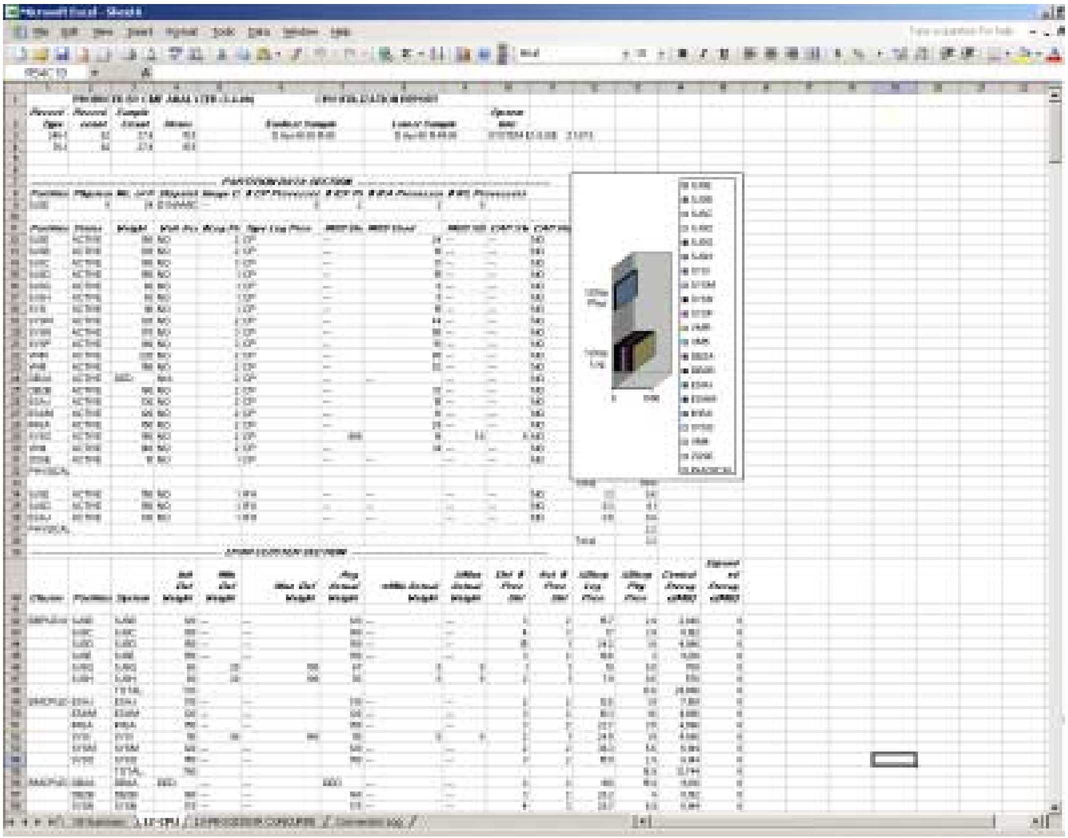

13 Select the tab for the Workload Manager Map Report. The report displayed should look similar to Figure 27 on page 116. It is difficult to display time-sensitive data as columns of figures. In this example, the polar plot graphic quickly shows the times when goals were not met.

Figure 27 Workload Manager Map Report

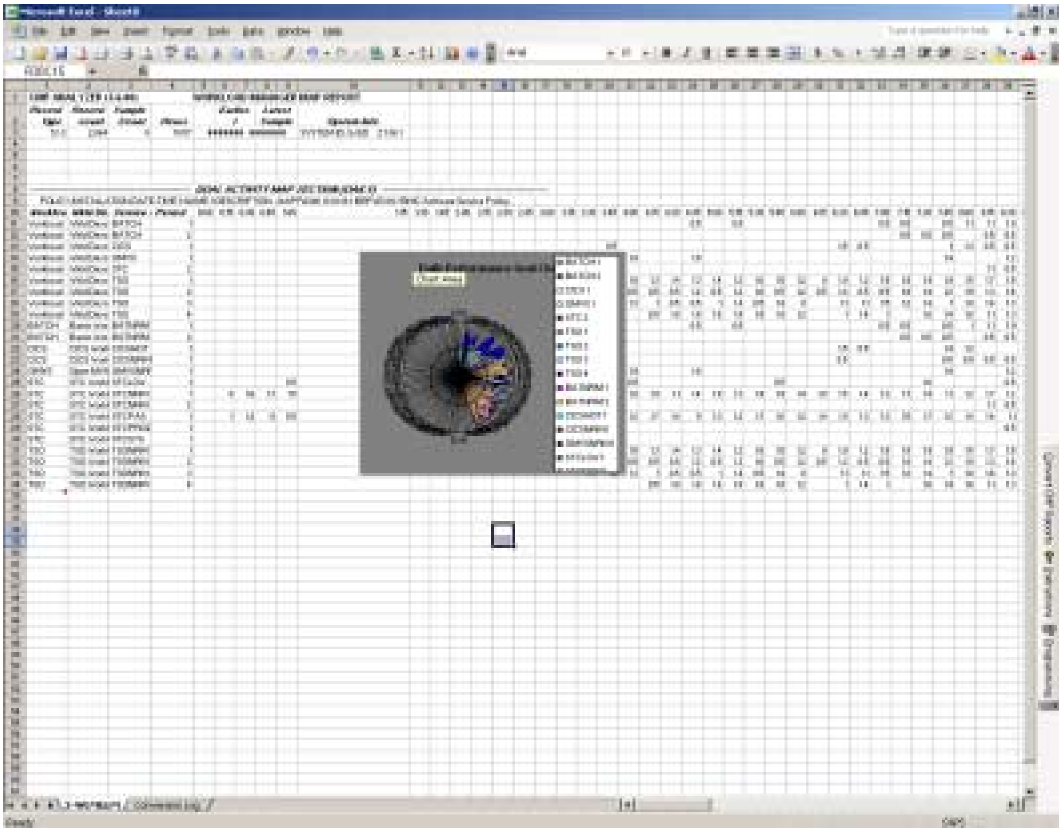

14 Select other reports to see how the data has been enhanced by using the Excel tools.

15 You can also use these sample reports to experiment with the Excel tools before converting your own reports.

Troubleshooting

Although the Spreadsheet Converter is very easy to use, this section describes the most common problems and how to fix them.

I cannot open the Spreadsheet Converter after I transfer it to my PC.

1 Make sure that the file CX98SSCX.XLA was transferred as a binary file.

2 Make sure that you are running Microsoft Excel 97 or later.

My CMF Analyzer or CMFMON reports do not convert.

1 Make sure that the SYSPRINT DD statement in your CMF Analyzer JCL points to the data set that is to contain the report

or

2 Make sure that the CMFPRINT DD statement in your CMF Analyzer JCL points to the data set that is to contain the report

or

3 Use the IBM Sysout Display and Search Facility (SDSF) to copy the data from the SYSOUT queue to a data set.

4 Make sure that you renamed your CMF Analyzer report file with the .txt extension.

5 Make sure that you saved your CMFMON report with comma separated values (CSV).

6 Make sure that you renamed your CMFMON report file with the .csv extension.

7 Make sure that you transferred the file to your PC in the ASCII format.

8 Display the transferred report on your PC to verify the following:

- There is carriage control in column 1.
- Data has been converted from EBCDIC to ASCII.
- Reports were produced by the CMF Analyzer.

9 Check column 1 of the Conversion Log to see if there are any exception messages. If there are, look in the Messages and Codes display for additional information.

I cannot format the reports in Excel the way I want to.

- Refer to the Microsoft Excel documentation, or contact Microsoft customer support.

Maintenance and support

Since the Spreadsheet Converter is distributed as a mainframe file, it will be maintained using standard SMP tools.

Maintenance

Updates to the Spreadsheet Converter will be distributed by way of the usual BMC Software Candidate and PUT mechanisms. Emergency fixes can be sent by standard BMC Software maintenance procedures. Whenever a Spreadsheet Converter PTF is sent, the HOLDDATA file will alert you to the availability of an updated version that needs to be downloaded again. Use the same procedures for downloading as described in “Installing the Spreadsheet Converter on your PC” on page 106.

Customer Support

In order to work on an incident for the Spreadsheet Converter, BMC Software Customer Support personnel need to have a copy of the reports that were being converted. It will also help if you send a copy of the resulting spreadsheet.

Since the report and spreadsheet files reside on your PC it should be possible to send them in either by e-mail, or by putting them on a diskette and mailing it in.

Part 2

CMF MONITOR reference

This part presents the following topics:

Chapter

6

Extractor control statements

This section describes the control statements used by the CMF MONITOR Extractor. The Extractor control statements and parameters are summarized in Table 16 and are cross-referenced to associated Analyzer control statements and report titles.

Following this table are separate sections that describe each Extractor control statement in detail; these sections are organized alphabetically.

Table 16 Extractor control statements (part 1 of 6)

Extractor command	Extractor function	Analyzer command and report title
ASMDATA See page 127 for more information.	samples auxiliary storage manager data produces SMF type 75 records; produces CMF type 240-02 and 240-09 user records	AUXSTOR (see page 332) Auxiliary Storage Report (see page 333) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491)
CACHE See page 129 for more information.	samples cache control unit data; produces SMF type 74-5 and 74-8 records (as of CMF MONITOR release 5.6, no longer produces 240-27 records)	CACHEACT (see page 216 for reports based on SMF type 74-5 records) Cache Subsystem Overview Report (see page 337) Cache Subsystem Activity Report (see page 340) when REPORT=SUBSYS is specified Cache Device Activity Report (see page 347) when REPORT=DEVICE is specified ESS (see page 251) ESS Statistics Report (see page 436)
CFDATA See page 132 for more information.	samples coupling facility activity; produces SMF type 74-4 records	CFACT (see page 219) Coupling Facility Activity Report (see page 378)

Table 16 Extractor control statements (part 2 of 6)

Extractor command	Extractor function	Analyzer command and report title
CHANNEL See page 134 for more information.	samples channel path activity data; produces SMF type 73 records	CMFSUM (see page 223) CMF Summary Report (see page 356) CHANNEL (see page 134) Channel Path Activity Report (see page 350) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491)
CPU See page 136 for more information.	samples CPU activity data produces SMF type 70-1 records; produces CMF type 240-01 user records	CMFSUM (see page 223) CMF Summary Report (see page 356) CPU (see page 230) CPU Utilization Report (see page 393) CPUCON (see page 232) Processor Concurrency Report (see page 500) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491) PROTKEY (see page 284) CPU Utilization by Protect Key Report (see page 411) PRSM (see page 285) Logical Partition Report (see page 476) SRM (see page 307) System Resources Manager Report (see page 515) TSOPERF TYPE=INT (see page 315) TSO Interval Summary Report (see page 534)
CRYPTO See page 140 for more information.	collects cryptographic hardware activity measurements; produces SMF type 70-2 records	CRYPTO (see page 233) Cryptographic Hardware Activity Report (see page 418)
CSMON See page 142 for more information.	activates the CMF MONITOR CSMS sampler, which gathers and formats data collected by COMMON STORAGE MONITOR, if active produces CMF type 240-29 user records	COMMSTOR (see page 227) Common Storage Usage Detail Report (see page 374) Common Storage Usage Summary Report (see page 377) CSMAPSAS Sample SAS report in the BBSAMP library (see Chapter 10, “Mapping CMF records created by CMF” for more information)

Table 16 Extractor control statements (part 3 of 6)

Extractor command	Extractor function	Analyzer command and report title
DEVICE See page 145 for more information.	samples whatever device classes are defined at the CLASS parameter produces SMF type 74-1 records; produces CMF type 240-05 user records	CMFSUM (see page 223) CMF Summary Report (see page 355) DASD (see page 237) Direct Access Report (see page 425) DASD,VOLSER (see page 237 and page 320) Direct Access Report Plot of Volume (see page 427) DEVACT (see page 243) Device Activity Report (see page 421) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491) SHARDEV (see page 297) Shared Device Activity Report (see page 505)
DISTIM See page 149 for more information.	monitors disabled time delay of CPU interrupts produces a CMF type 240-24 record	not applicable Disabled Delay Report (see page 429)
ENQUEUE See page 152 for more information.	collects enqueue contention data produces SMF type 77 records	ENQUEUE (see page 250) Enqueue Conflict Report (see page 434) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491)
EXTSUM See page 154 for more information.	automatically produces a spin-off report, called the Extractor Summary Report (see page 443), at each Extractor interval if SPINOFF=*class* is specified produces CMF type 240-06 and 240-07 user records, which are required for other reports	PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491)
FICONSW See page 158 for more information.	collects FICON Director configuration and activity data produces SMF type 74-7 records	FICONSW (see page 258) FICON Director Activity Report (see page 448)

Table 16 Extractor control statements (part 4 of 6)

Extractor command	Extractor function	Analyzer command and report title
HEADMOVE See page 159 for more information.	samples DASD head movement data produces CMF type 240-12, 240-13, and 240-14 user records	DASD (see page 237) Direct Access Report (see page 425) DASD,VOLSER (see page 237 and page 320) Direct Access Report Plot of Volume (see page 427)
HFS See page 165 for more information.	samples global, buffer, and file system statistics produces SMF type 74-6 records	HFS (see page 270) HFS Statistics Report (see page 454)
IOQ See page 167 for more information.	samples I/O queuing activity produces SMF type 78-3 records	CMFSUM (see page 223) CMF Summary Report (see page 355) IOQ (see page 272) I/O Queuing Activity Report (see page 464)
LINKMAP See page 169 for more information.	collects mapping data for link pack area produces CMF type 240-16 user records	LINKPACK (see page 273) Link Pack Area Report (see page 472)
OMVS See page 171 for more information.	collects OMVS kernel activity data produces SMF type 74-3 records	OMVS (see page 276) OMVS Kernel Activity Report (see page 488)
PAGING See page 173 for more information.	samples page data produces SMF type 71 records; produces CMF type 240-03 user records	CMFSUM (see page 223) CMF Summary Report (see page 355) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491) SRM (see page 307) System Resources Manager Report (see page 515) STORAGE (see page 308) Storage Management Report (see page 509)
PGDDLAY See page 175 for more information.	samples delay and storage utilization data for service classes produces SMF type 72-4 records	not applicable

Table 16 Extractor control statements (part 5 of 6)

Extractor command	Extractor function	Analyzer command and report title
REPORT See page 176 for more information.	required; defines recording mode and other Extractor operating characteristics produces CMF type 240-00 and 240-11 user records	not applicable LINKPACK (see page 273) Link Pack Area Report (see page 472)
TRACE See page 183 for more information.	invokes the trace facility produces CMF type 240-18 user records	TRACE (see page 312) Trace Report (see page 528)
TRACE76 See page 191 for more information.	samples selected fields produces SMF type 76 records	EXCEPTS TRACE=YES (see page 252) Exception Trace Detail Report (see page 442) GRAPH TYPE=TRACE (see page 259) Graphics Trace Detail Report (see page 451)
TSODATA See page 194 for more information.	monitors TSO command activity and gathers user statistics produces CMF type 240-20 and 240-21 user records	CMFSUM (see page 223) CMF Summary Report (see page 355) PERFSUM,PERFORM (see page 279 and page 277) Performance Summary Report (see page 491) TSOPERF TYPE=CMD (see page 315) TSO Command Summary Report (see page 532) TSOPERF TYPE=INT (see page 315) TSO Interval Summary Report (see page 534) TSOUSER (see page 317) TSO User Summary Report (see page 537), when USERS=YES is defined on the TSODATA Extractor control statement
USER See page 197 for more information.	provides an interface for user-written RMF exits User-written exits might or might not create their own record types.	not applicable
VSMDATA See page 200 for more information.	collects data on virtual storage, including SQA by subpool, CSA by subpool and key, and private area virtual storage by specific jobs produces SMF type 78-2 records	CMFSUM (see page 223) CMF Summary Report (see page 355) VIRTSTOR (see page 318) Virtual Storage Activity Report (see page 539)

Table 16 Extractor control statements (part 6 of 6)

Extractor command	Extractor function	Analyzer command and report title
WORKLOAD See page 203 for more information.	collects system workload running in service classes produces SMF type 72-3 records	CMFSUM (see page 223) CMF Summary Report (see page 355) PERFSUM (see page 279) Performance Summary Report (see page 491) WLMGL (see page 321) Workload Manager Goal Mode Report (see page 549)
XCFDATA See page 205 for more information.	samples Cross-System Coupling Facility (XCF) performance data produces SMF type 74-2 records	XCF (see page 323) Cross-System Coupling Facility Report (see page 413)

ASMDATA

```
ASMDATA
[SAMPLE={1000|nnnn}]
```

Overview

The ASMDATA control statement causes the CMF MONITOR Extractor to collect auxiliary storage management (ASM) data, including information about the I/O activity of page data sets and ASM data constants.

Monitoring modes

One ASMDATA statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The ASMDATA statement controls the ASMS sampler, which produces SMF type 75 records and CMF type 240-02 and 240-09 user records.

Analyzer statements and reports

AUXSTOR (see page 215)
Auxiliary Storage Report (see page 332)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

Other BMC Software products

The ASMDATA control statement must be defined in the CMF MONITOR Extractor JCL for MAINVIEW for z/OS. (See “MAINVIEW for z/OS” on page 35 for more information.)

Parameters

The parameter for the ASMDATA control statement is

SAMPLE= specifies the number of milliseconds between data gathering cycles for the ASM data being sampled

The default is 1000 or one millisecond. Acceptable values are 20 to 9999.

Example

```
ASMDATA  SAMPLE=2000
```

Instructs the CMF MONITOR Extractor to sample ASM data every 2000 milliseconds or once every two seconds.

CACHE

```
CACHE
[SUBSYS=(nnnn1,nnnn2,nnnn3,….nnnn16)]
[RECORDS={(}CACHE,ESSLINK,ESSRANK,ESS,ALL{)}
```

Overview

The CACHE control statement causes the CMF MONITOR Extractor to collect configuration and activity data for all or selected cache subsystems.

Monitoring modes

Only one CACHE statement can be defined for each monitoring mode, continuous (CPM) or intermittent (IPM). If more than one CACHE statement is specified, only the first statement will be accepted; all other statements will be discarded. Specifying multiple CACHE statements does not adversely impact initialization of the PAS.

Sampler and record types

The CACHE statement controls the Cache/ESS sampler, which produces SMF type 74-5 and 74-8 records.

As of version 5.6.00 of CMF MONITOR, SMF 240-27 records are no longer produced.

Analyzer statements and reports

CACHEACT (see page 216 for reports based on SMF type 74-5 records)
Cache Subsystem Overview Report (see page 337)
Cache Subsystem Activity Report (see page 340)
Cache Device Activity Report (see page 347)

ESS (see page 251 for reports based on SMF type 74-8 records)
ESS Statistics Report (see page 436)

NOTE

The ESS Statistics Report is created from SMF 74-8 records, which are created only if RECORDS=ESS, RECORDS=ESSLINK, RECORDS=ESSRANK, or RECORDS=ALL is specified on the CACHE control statement.

Other BMC Software products

The CACHE control statement must be defined in the Extractor JCL for MAINVIEW for z/OS. (See "MAINVIEW for z/OS" on page 35 for more information.)

Parameters

The parameters for this control statement are

SUBSYS=

specifies 1 to 16 cache subsystem IDs to be sampled

A cache subsystem ID is a four-digit hexadecimal string (for example, 03F5). When more than one subsystem ID is specified, they must be separated by commas and the entire string must be enclosed in parentheses.

When this parameter is omitted, CMF MONITOR detects and monitors all active cache subsystems automatically.

RECORDS=

causes the PAS to invoke a cache sampler; specify one of these values:

CACHE
creates only cache records (SMF 74-5); the default

Note: If you omit the RECORDS parameter from the CACHE control statement, the cache sampler default is to create *only* SMF 74-5 records, as if you had specified RECORDS=CACHE.

ESSLINK
creates SMF 74-8 records containing ESS Link Statistics

ESSRANK
creates SMF 74-8 records containing ESS Rank statistics and Extent Pool statistics

ESS
creates SMF 74-8 records containing ESS Link statistics, ESS Rank statistics, and Extent Pool statistics, as appropriate with your system configuration; the same as ESSLINK and ESSRANK combined

ALL
creates all possible records (that is, both SMF 74-5 and SMF 74-8 records), as appropriate with your system configuration

NOTE

If you specify parameters that are available in CMF MONITOR versions prior to 5.6, they are ignored and a warning message is issued.

Examples

```
CACHE    RECORDS=ALL
```

This example invokes the cache sampler and instructs it to collect both cache and ESS statistics (Link, Rank, and Extent Pool, as appropriate) from all detected cache and ESS systems.

```
CACHE    RECORDS=CACHE
```

This example invokes the cache sampler and instructs it to collect only cache statistics from all of the cache subsystems detected. No ESS statistics will be gathered.

```
CACHE
```

This example invokes the cache sampler with no RECORDS parameter, which will cause the cache sampler to assume the default RECORDS=CACHE parameter and therefore collect only cache statistics from all of the cache subsystems detected. No ESS statistics will be collected.

```
CACHE    SUBSYS=(CF00,0B00)
```

This example invokes the cache sampler and instructs it to sample just two cache subsystems—CF00 and 0B00. As no RECORDS parameter is specified, the cache sampler assumes the default RECORDS=CACHE.

```
CACHE    RECORDS=ESS
```

This example invokes the cache sampler and instructs it to collect only ESS statistics from all of the ESS subsystems detected; no cache statistics will be gathered.

```
CACHE    RECORDS=(CACHE,ESSLINK)
```

This example invokes the cache sampler and instructs it to collect cache statistics and ESS Link statistics; no ESS Rank statistics will be gathered.

CFDATA

```
CFDATA
[SAMPLE={10000|nnnnn}]
```

Overview

The CFDATA control statement causes the CMF MONITOR Extractor to collect data for all coupling facilities to which the system is connected, and write one SMF type 74-4 record for each coupling facility per interval.

To collect synchronous data from all systems to which the coupling facility is connected, BMC Software recommends that you use the SYNCH=SMF parameter of the CMF MONITOR Extractor REPORT control statement.

Monitoring modes

One CFDATA statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The CFDATA statement controls the CFTS sampler, which produces SMF type 74-4 records.

Analyzer statements and reports

CFACT (see page 219)
Coupling Facility Activity Report (see page 378)

Other BMC Software products

The CFDATA control statement is used by the CMF MONITOR Extractor only.

Parameters

The parameter for the CFDATA control statement is

SAMPLE= specifies the number of milliseconds between data-gathering cycles for coupling facility data

The default is 10000 milliseconds, or 10 seconds. Acceptable values are 5000 to 59000.

Examples

```
CFDATA
```

The CMF MONITOR Extractor collects coupling facility data once every ten seconds.

```
CFDATA SAMPLE=5000
```

The CMF MONITOR Extractor collects coupling facility data once every five seconds.

CHANNEL

```
CHANNEL
```

Overview

The CHANNEL control statement causes the CMF MONITOR Extractor to collect channel path busy counts. The system records channel path statistics in the channel path measurement table. CMF MONITOR inspects this table at the beginning and end of each recording interval, so the CHANNEL control statement does not need a default sample rate.

Monitoring modes

One CHANNEL statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The CHANNEL statement controls the CHNS sampler, which produces SMF type 73 records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 356)

CHANNEL (see page 134)
Channel Path Activity Report (see page 350)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

Other BMC Software products

The CHANNEL control statement is used by the CMF MONITOR Extractor only.

Parameters

There are no parameters for this control statement.

Example

```
CHANNEL
```

The CMF MONITOR Extractor collects channel path statistics from the channel path measurement table.

CPU

```
CPU
[SAMPLE={1000|nnnn}]
[,CAPMSG={YES|NO}]
[,CAPMSGRP={10|nn|NO}]
```

Overview

The CPU control statement causes the CMF MONITOR Extractor to collect CPU and CPU-dispatching data. This data includes information about CPU wait, busy, and idle status, online and offline times, queue depth, and processor concurrency.

Monitoring modes

One CPU statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The CPU statement controls the CPUS sampler, which produces SMF type 70-1 records and CMF type 240-01 user records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 355)

CPU (see page 230)
CPU Utilization Report (see page 393)

CPUCON (see page 232)
Processor Concurrency Report (see page 500)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

PROTKEY (see page 284)
CPU Utilization Report by Protect Key (see page 411)

PRSM (see page 285)
Logical Partition Report (see page 476)

SRM (see page 307)
System Resources Manager Report (see page 515)

TSOPERF TYPE=INT (see page 315)
TSO Interval Summary Report (see page 534)

Other BMC Software products

The CPU control statement must be defined in the CMF MONITOR Extractor JCL for MAINVIEW for z/OS. (See “MAINVIEW for z/OS” on page 35 for more information.)

Parameters

The parameters for the CPU control statement are

SAMPLE=	specifies the number of milliseconds between data-gathering cycles for CPU data The default is 1000 milliseconds or one second. Acceptable values are 20 to 9999.

CAPMSG=	specifies whether the MVS PAS issues the message notifying that the local logical partition is soft-capped by WLM
	
	If YES is specified, the MVS PAS issues the following message after detecting that the logical partition has been soft capped by WLM:
	
	`CMFCPU13 LPAR SOFT CAPPED BY WLM`
	
	When the logical partition is no longer capped, the following message is issued:
	
	`CMFCPU15 LPAR NO LONGER SOFT CAPPED BY WLM; CAPPED DURATION WAS hhh.mm.ss`

CAPMSG=
(*continued*)

Notes:

- If the weight of the local logical partition represents a CPU capacity below the defined capacity, WLM needs to turn capping on and off to keep the short-term CPU usage at the defined capacity. In this scenario, the previously listed CMFCPU13 and CMFCPU15 messages are issued repeatedly, normally a few minutes apart but could be as short as 10 seconds apart. When *hhh.mm.ss* is relatively short, the logical partition might be soft-capped again soon.

- The CPU sampler checks the capped status approximately every 10 seconds. When the message CMFCPU13 or CMFCPU15 is issued, the capped status has actually changed, on average, 5 seconds earlier.

- This parameter is applicable only for CPM mode and is ignored if specified for the IPM mode.

CAPMSGRP=

specifies whether the MVS PAS issues the message notifying that the local logical partition remains soft-capped by WLM

If a value *nn* from 1 to 60 is specified (the default is 10), the MVS PAS issues the following CMFCPU14 message every *nn* minutes:

```
CMFCPU14 LPAR SOFT CAPPED BY WLM SINCE hh:mm:ss
ddmmmyy FOR hhh.mm.ss
```

If NO is specified, the message CMFCPU14 is not issued.

Notes:

- If CAPMSG=NO is specified, either omit this parameter or specify CAPMSGRP=NO.

- This parameter is applicable only for CPM mode and is ignored if specified for the IPM mode.

Examples

```
CPU
```

The CMF MONITOR Extractor samples CPU data once per second, using the default SAMPLE=1000.

```
CPU  SAMPLE=1500
```

The CMF MONITOR Extractor samples CPU data once every 1 ½ seconds.

CRYPTO

```
CRYPTO
```

Overview

The CRYPTO control statement causes the CMF MONITOR Extractor to collect activity measurements of cryptographic hardware features, including the standard Cryptographic Coprocessor Facility (CCF), and the optional adjunct features: PCI Cryptographic Coprocessor (PCICC) and PCI Cryptographic Accelerator (PCICA).

Monitoring modes

One CRYPTO statement can be specified for each monitoring mode, CPM and IPM.

Sampler and record types

The CRYPTO statement controls the CRYS sampler, which produces SMF type 70-2 records.

Analyzer statements and reports

CRYPTO (see page 233)
Cryptographic Hardware Activity Report (see page 418)

Other BMC Software products

The CRYPTO control statement is used by the CMF MONITOR Extractor only.

Parameters

There are no parameters for this control statement.

Example

```
CRYPTO
```

The CMF MONITOR Extractor collects cryptographic hardware activity measurements.

CSMON

```
CSMON
```

Overview

The CSMON control statement causes the CMF MONITOR Extractor to gather and write data collected by COMMON STORAGE MONITOR (CSM). CSM data is gathered by CMF MONITOR Extractor only if CSM is active and the CSMON statement is defined.

CSM keeps its own data current, so data is collected once by the CMF MONITOR Extractor at the end of each defined interval; therefore, common storage obtained and subsequently released by a job within an interval is not reported. If CSM is not active, the CMF MONITOR Extractor bypasses processing until it becomes active.

Controlling COMMON STORAGE MONITOR

The CSMON control statement requires that the COMMON STORAGE MONITOR be active. The *CMF MONITOR Customization Guide* describes how to customize CSM as either a subsystem or a Started Task.

- If defined as a subsystem, CSM is started automatically at IPL.
- If defined as a Started Task, you can start CSM automatically after an IPL by adding this command to the SYS1.PARMLIB member COMMND*xx*:

  COM='START BB$CSMON[,SUB=MSTR]'

  where

BB$CSMON	is the name of the procedure containing the start command JCL (This default name is used by AutoCustomization when creating the procedure.)
SUB=MSTR	specifies that you want to start CSM before JES2 or JES3 so that you can track common storage requests from JES If SUB=MSTR is specified, the following must be true: ■ The JCL must not specify any JES data sets. ■ BB$CSMON must reside in SYS1.PROCLIB.

You can stop the CSM by executing the program BBXCSMON with the STOP option. A sample of the required JCL is shown here:

```
//STOPCSM EXEC PGM=BBXCSMON,PARM='STOP'
//STEPLIB DD  DISP=SHR,DSN=hilevel.BBLINK
<-- if not in linklist
```

Monitoring modes

A CSMON statement can be specified for CPM mode only; this statement is not valid for IPM mode.

Sampler and record types

The CSMON statement controls the CSMS sampler, which produces CMF type 240-29 user records.

Analyzer statements and reports

COMMSTOR (see page 227)
Common Storage Usage Summary Report (see page 377)

NOTE

The type 240-29 records can also be used to produce your own reports. See information about member CSMAPSAS, which is a sample SAS report in *hilevel*.BBSAMP, in Chapter 10, "Mapping CMF records created by CMF."

Other BMC Software products

The CSMON control statement is used by the CMF MONITOR Extractor only.

Parameters

There are no parameters for this control statement.

Example

```
CSMON
```

The CMF MONITOR Extractor activates the CSMS sampler, which gathers and formats COMMON STORAGE MONITOR data, if active.

DEVICE

```
DEVICE [ALL]
[,RANGE=(xxxx[:xxxx]1,...,xxxx[:xxxx]8)]
[,CLASS={DASD|TAPE|UNITR|CHRDR|GRAPH|COMM}]
[,EXCEPT=(xxxx[:xxxx]1,...,xxxx[:xxxx]8)]
[,OFFLINE={NO|YES}]
[,SAMPLE={1500|nnnn}]
```

Overview

The DEVICE control statement causes the CMF MONITOR Extractor to measure specified devices for busy and wait status, I/O activity, online and offline times, and volume activity.

NOTE

If you want SMF type 74-1 records that are compatible with those produced by RMF, you must specify the CLASS parameter.

Using DEVICE with HEADMOVE

If the HEADMOVE control statement (see “HEADMOVE” on page 159) is specified for only one mode (either CPM or IPM) and you also want to use the DEVICE control statement, both HEADMOVE and DEVICE must run in the same mode.

Monitoring modes

Any number of DEVICE control statements can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Different sample rates can be specified in each DEVICE statement; a high sampling rate can be specified for one device group, and a low rate can be specified for a less active group. However, a single device can be specified only once for each monitoring mode; subsequent duplicate specifications for a device are ignored.

Sampler and record types

The DEVICE statement controls the DEVS sampler, which produces SMF type 74-1 records and CMF type 240-05 user records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 355)

DASD (see page 237)
Direct Access Report (see page 425)

DASD,VOLSER (see page 237 and page 320)
Direct Access Report Plot of Volume (see page 427)

DEVACT (see page 243)
Device Activity Report (see page 421)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

SHARDEV (see page 297)
Shared Device Activity Report (see page 505)

Other BMC Software products

The DEVICE control statement must be defined in the CMF MONITOR Extractor JCL for MAINVIEW for z/OS. (See “MAINVIEW for z/OS” on page 35 for more information.)

Parameters

Parameters for the DEVICE control statement are as follows:

ALL= specifies that all online devices in the system are to be measured

RANGE= specifies a range of devices to be measured

The range can consist of a list of single device addresses, a list of device ranges (*xxxx*:*xxxx*), or a combination of both.

You can specify up to eight devices or device ranges. The first device specified for a device range must be a physically attached device.

CLASS= specifies a class of device to be sampled

Only one class can be specified per DEVICE statement. Only the EXCEPT, OFFLINE, and SAMPLE parameters are valid with CLASS. Parameters ALL and RANGE are ignored when CLASS is specified.

Note: This parameter must also be specified for CMF MONITOR Online and MAINVIEW for z/OS. If you want SMF type 74 records that are compatible with those produced by RMF, you must specify the CLASS parameter.

TAPE magnetic tape devices (UCBTBYT3 is X'80')

COMM communications equipment (UCBTBYT3 is X'40')

DASD direct access storage devices (UCBTBYT3 is X'20')

GRAPH graphics devices (UCBTBYT3 is X'10')

UNITR unit record devices (UCBTBYT3 is X'08')

CHRDR character reader devices (UCBTBYT3 is X'04')

Note: If ALL, RANGE, and CLASS are not specified, all online DASDs and tape devices are measured. If a device class other than TAPE or DASD is specified and RMF I/O monitoring is not active, the DEVICE sampler invokes start/stop I/O monitoring through BBXS.

EXCEPT= specifies a range of devices to be excluded from measurement activity

This exclusion applies only to the specification on this control statement and has no effect on prior or subsequent control statements. You can specify a maximum of eight devices or device ranges.

Note: The EXCEPT parameter of the HEADMOVE control statement (see “HEADMOVE” on page 159) has no impact on the data collected by the DEVICE sampler.

OFFLINE= specifies whether data will be collected for offline devices

The default is NO.

Activity for online devices is always collected. If OFFLINE=NO is specified, activity will be collected for devices that are varied online after the sampler has started.

SAMPLE= specifies the number of milliseconds between data gathering cycles for this group of devices

The default is 1500. Acceptable values are 20 to 9999.

Examples

```
DEVICE ALL,EXCEPT=(104,20F,313:41F,70A)
```

The CMF MONITOR Extractor samples all online devices (except device addresses 104, 20F, 313 through 41F, and 70A) at the default interval of every 1½ seconds.

```
DEVICE RANGE=(201,204:305),SAMPLE=1000
```

The CMF MONITOR Extractor samples device addresses 201 and 204 through 305 once every second.

```
DEVICE  SAMPLE=1000
```

All online DASD and tape devices are measured once per second.

```
DEVICE  SAMPLE=2000,CLASS=UNITR
```

The CMF MONITOR Extractor samples all unit record devices every two seconds.

```
DEVICE CLASS=DASD
```

The CMF MONITOR Extractor samples online DASDs only, at the default interval of every 1½ seconds.

```
DEVICE  CLASS=TAPE
```

The CMF MONITOR Extractor samples all magnetic tape devices, including all offline-to-online activity, at the default frequency of once every 1½ seconds.

DISTIM

```
DISTIM
[SPINOFF={A|class|NO}]
[,COPIES={1|nn}]
[,SAMPLE={157|nnnn}]
[,LODELAY={0|nnnn}]
[,HIDELAY={0|nnnn}]
[,SHIFT={0|nnnn}]
```

Overview

The DISTIM control statement allows you to monitor the disabled time delay of CPU interrupts. In each sampling interval, the disabled time sampler queues a TQE containing

TQEVAL	time that the interrupt should take place
TQEXIT	address of the exit

If an interrupted CPU runs disabled, the interrupt is stacked in the hardware until the running process enables itself. The exit performs a STCK to get the time that the routine actually did take control. The exit then calculates and records the difference as disabled time delay. If the delay falls inside specified limits, the ASID of the last address space in control is recorded as well.

Monitoring modes

One DISTIM statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The DISTIM statement controls the DITS sampler, which produces a CMF type 240-24 record.

Analyzer statements and reports

none

Other BMC Software products

The DISTIM control statement is used by the CMF MONITOR Extractor only.

Parameters

Parameters for the DISTIM control statement are as follows:

SPINOFF=	specifies the SYSOUT class to be dynamically allocated for printing the Disabled Delay Report Any valid SYSOUT class can be specified. The default is A. The SPINOFF=NO parameter eliminates the Disabled Delay Report.
COPIES=	specifies the number of copies of the Disabled Delay Report to be printed Up to 99 copies can be specified. The default is 1. This parameter is meaningful only if the SPINOFF parameter is defined with a valid SYSOUT class.
SAMPLE=	specifies the number of milliseconds between data gathering cycles; the default is 157 You should avoid specifying a sampling interval that is a divisor of an interval that is used by a different control statement. Because 1000, 1500, and 5000 are default sampling intervals for various control statements, you should avoid specifying a sampling interval that contains 0 as its final digit.
LODELAY=	specifies the lower limit of the range of delay time for which ASID counts are desired LODELAY must be below HIDELAY for the ASID count to be maintained. The default is 0.
HIDELAY=	specifies the upper limit of the range of delay times for which ASID counts are desired; the default is 0
SHIFT=	specifies the lower value of delay which is to appear on the graphical output It is the amount by which the graph is shifted to make the relevant portions fit on the page. The default is 0. The maximum is 380.

Example

```
DISTIM SPINOFF=R,COPIES=1,SAMPLE=293,LODELAY=30,HIDELAY=60
```

The CMF MONITOR Extractor produces one copy of a spin-off report in SYSOUT class R for each recording interval, sampling every 293 milliseconds. The report indicates which address spaces were last dispatched for any delay in the range of 30 to 60 microseconds.

ENQUEUE

```
ENQUEUE
[MAJOR=nnnnnnnn]
[,MINOR=nnnnnnnn]
```

Overview

The ENQUEUE control statement causes the CMF MONITOR Extractor to collect enqueue contention data.

Sampler and record types

The ENQUEUE statement controls the ENQS sampler, which produces SMF type 77 records.

Analyzer statements and reports

ENQUEUE (see page 250)
Enqueue Conflict Report (see page 434)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

Parameters

Parameters for the ENQUEUE control statement are as follows:

MAJOR=	specifies a major name to be monitored If this parameter is not specified, enqueue activity for all major and minor names is collected. The major name can be from 1 to 8 characters long.
MINOR=	specifies a minor name to be monitored If this parameter is not specified, enqueue activity for all minor names is collected. The minor name can be from 1 to 44 characters long. A MINOR name cannot be specified unless a MAJOR name is specified.

Examples

```
ENQUEUE
```

The CMF MONITOR Extractor collects all enqueue activity information.

EXTSUM

```
EXTSUM
[JOBCLASS=(JC=class1,JD=description1,...,JC=class36,JD=description36)]
[,SRVCLASS=(SC=serviceclass1,SCP=period1,...,SC=serviceclass64,
SCP=period64)]
[,JES={NO|YES}]
[,COPIES={1|n}]
[,SAMPLE={1000|nnnn}]
[,SPINOFF={A|class|NO}]
```

Overview

The EXTSUM control statement causes the CMF MONITOR Extractor to sample data on CPU and channel busy status, paging, I/O activity, TSO transaction rate, workload status, and service allocation percentages.

Monitoring modes

An EXTSUM statement can be specified for CPM mode only; this statement is not valid for IPM mode.

Sampler and record types

The EXTSUM statement controls the EXTS sampler, which produces CMF type 240-06 and 240-07 user records.

Analyzer statements and reports

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

Producing the Extractor Summary Report

The Extractor Summary Report is produced only if the SPINOFF parameter is defined with a valid SYSOUT class in the EXTSUM statement. The SPINOFF parameter causes the CMF MONITOR Extractor to dynamically allocate a JES spin-off data set at the end of each recording interval, without disrupting the continuous sampling functions of the CMF MONITOR Extractor. From this spin-off data, the CMF MONITOR Extractor automatically produces the Extractor Summary Report for every recording interval and sends it to the specified SYSOUT.

However, the other parameters and the WORKLOAD control statement also affect the Extractor Summary Report. If JES=NO is defined in the EXTSUM statement, the Job Class Activity section of the Extractor Summary Report is not produced.

Other BMC Software products

The EXTSUM control statement is used by the CMF MONITOR Extractor only.

Parameters

Parameters for the EXTSUM control statement are as follows:

JOBCLASS= specifies the job class queues to be monitored for activity and a description for each class to be reported in the Extractor Summary Report

You can specify a maximum of 36 classes and descriptions. Any class that is specified is reported separately on the Extractor Summary Report; all other classes are grouped together and described as OTHER. There are two values to define to the JOBCLASS parameter:

JC= name of the job class

JD= description of job class

For every JC= value specified, a job class description up to 15 characters long can be specified. The *description* specified as the JD value is displayed on the Extractor Summary Report.

SRVCLASS= specifies the service classes and periods to be monitored for the Extractor Summary Report and the Performance Summary Report

A maximum of 64 service classes can be specified. There are two values to define to the SRVCLASS parameter:

SC= name of the service class

You can specify the name of the service class, or you can use an asterisk (*) as a wildcard character to specify multiple service classes that use a pattern. For example, if you have two batch service classes (say, BATNRML and BATHOT), you can define SC=BAT* to specify both of those service classes.

Note: You can specify a maximum of 64 service classes. If use of the asterisk wildcard includes more than 64 service classes, only the first 64 service classes appear on the report.

SCP= number of the service class period

For every SC value specified, you can specify up to eight periods. If you do not specify SCP for a particular service class, all periods appear in the reports.

The SRVCLASS parameter can be used within the same control statement as the PERFORM parameter.

JES= specifies whether the job entry subsystem is to be sampled and controls whether the Job Class Activity section of the Extractor Summary Report is produced

The default is JES=NO.

For the Job Class Activity section of the Extractor Summary Report to be produced, this value must be defined as JES=YES, and the SPINOFF parameter must be defined with a valid SYSOUT class.

For JES2 sampling, code JES=YES.

For JES3 sampling, code JES=YES and make sure that the procedure for JES3INIT was assembled and linked during AutoCustomization or manual customization. If you have upgraded your JES3 release since you installed CMF MONITOR, you must reassemble the JES3INIT procedure to link to the new offsets. Use AutoCustomization or the instructions in the *hilevel*.BBILIB member @@YCH300 to reassemble the JES3INIT procedure.

COPIES= specifies the number of copies of the Extractor Summary Report to be printed

You can specify up to 99 copies. The default is 1. This parameter is meaningful only if the SPINOFF parameter is defined with a valid SYSOUT class.

SAMPLE= specifies the number of milliseconds between data gathering cycles

The default is 1000. The sample rate cannot be below 500 or over 9999.

SPINOFF= specifies the SYSOUT class to be dynamically allocated for printing the Extractor Summary Report

You can specify any valid SYSOUT class. The default is A. The SPINOFF=NO parameter eliminates all spin-off reports but allows Extractor summary records (CMF type 240-06 user records) to be written and later summarized by the Analyzer.

Examples

```
EXTSUM JOBCLASS=(JC=A,JD=PRODUCTION,JC=T,JD=TEST),
       SPINOFF=D,COPIES=3
```

The CMF MONITOR Extractor monitors and reports on job classes A and T. D is the SYSOUT class allocated to print the Extractor Summary Report. Three copies of the report are printed at Extractor recording intervals specified on the REPORT control statement.

Job class A has a description of PRODUCTION, and job class T has a description of TEST.

```
EXTSUM JOBCLASS=(JC=A,JD=PRODUCTION,JC=T,JD=TEST),
       SRVCLASS=(SC=*,SCP=1234)
```

The CMF MONITOR Extractor monitors and reports on job classes A and T and the first 4 periods on all service classes (up to 64 service classes).

FICONSW

```
FICONSW
```

Overview

The FICONSW control statement causes the CMF MONITOR Extractor to collect FICON Director configuration and activity data.

Sampler and record types

The FICONSW statement controls the FCSW sampler, which produces SMF type 74-7 records.

Analyzer statements and reports

FICONSW (see page 258)
FICON Director Activity Report (see page 448)

Parameters

There are no parameters for this control statement.

Examples

```
FICONSW
```

This example starts the FICON Director sampler, which collects FICON Director configuration and activity data.

HEADMOVE

```
HEADMOVE  [ALL]|
[,NUMBER=(c1,...,c15)]|
[,RANGE=(xxxx[:xxxx]1,...,xxxx[:xxxx]8)]
[,EXCEPT=(xxxx[:xxxx]1,...,xxxx[:xxxx]8)]
[,VTOC={NO|YES}]
[,DSNLIST=(dsn1,...,dsn101)]
[,DDNAME=ddn]
[,DSN='dsn']
[,ALTDSN='dsn']
[,SAMPLE={33|nn}]
[,OFFLINE={NO|YES}]
[,BUFSIZE={20|nnn}]
[,RESUME={0|n}]
[,SUSPEND={0|n}]
```

Overview

The HEADMOVE control statement causes the CMF MONITOR Extractor to measure specified DASDs for head movement activity.

HEADMOVE samples a maximum of 4,096 DASD devices. Parallel Access Volumes (PAVs) are not sampled.

Because its overhead can be greater than that of other samplers, BMC Software recommends that you run HEADMOVE only in IPM mode, with a sampling rate between 33 milliseconds and 25 milliseconds.

Using DEVICE with HEADMOVE

If the HEADMOVE control statement is specified for only one mode (either CPM or IPM), and you also want to use the DEVICE control statement, both HEADMOVE and DEVICE must run in the same mode. (See “DEVICE” on page 145 for more information.)

Monitoring modes

One HEADMOVE statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM); however, BMC Software recommends that this statement be run in IPM mode only.

Sampler and record types

The HEADMOVE statement controls the HMOV sampler, which produces CMF type 240-12, 240-13, and 240-14 user records.

Analyzer statements and reports

DASD (see page 237)
Direct Access Report (see page 425)

DASD,VOLSER (see page 237 and page 320)
Direct Access Report Plot of Volume (see page 427)

Other BMC Software products

The HEADMOVE control statement must be defined in the CMF MONITOR Extractor JCL for DSO. (See the *Data Set Optimizer User Guide and Reference* for more information.)

Parameters

Parameters for the HEADMOVE control statement are as follows:

ALL= specifies that all online DASDs in the system are to be measured

Note: If ALL is not specified, specify either NUMBER or RANGE, or both.

NUMBER= specifies a list of one or more numbers

Every online DASD having that number as its first digit is measured. You can specify up to 16 numbers. Each number that is specified must be a unique hexadecimal number in the range of 0 through F. If you specify more than one number, you must enclose the numbers in parentheses and separate them by commas.

RANGE= specifies a range of devices to be measured

The range can consist of a list of single device addresses, a list of device ranges (*xxxx*:*xxxx*), or a combination of both. You can specify a maximum of eight devices or device ranges.

Note: The NUMBER and RANGE parameters are ignored if ALL is specified.

EXCEPT= specifies a range of devices to be excluded from measurement activity

You can specify a maximum of eight devices or device ranges.

Note: The EXCEPT parameter of the DEVICE control statement (see “DEVICE” on page 145) has no impact on the data collected by the HEADMOVE sampler.

VTOC= specifies whether VTOC data is to be extracted from volumes being sampled for head movement

Data is collected at initialization. The HMOVRESCAN modify command can be used to collect and write VTOC records on request. (See “MODIFY command examples” on page 64 for more information.)

The default is NO. Specify YES only if the DSO option is used.

DSNLIST= specifies the names of up to 101 Extractor data sets where head movement data is to be written

The data set must be previously allocated and cataloged. This parameter can be used in place of DDNAME or DSN and ATLDSN. If you use the DSNLIST parameter, do not include the ddname CMFCDS*xx* or CMFIDS*xx* in the CMF MONITOR Extractor JCL.

You can specify DSNLIST=NULLFILE if you want the records to be discarded. This option might be useful if you are using IPM mode for running either CMF MONITOR Online or MAINVIEW for z/OS.

DDNAME= specifies the ddname of an Extractor data set (defined through JCL) to which the head movement data is written

This parameter is invalid if DSN or DSNLIST is used.

To support data set flipping, if DSO is writing to its own data set, specify ddnames CMFCDS*xx* for CPM and CMFIDS*xx* for IPM.

Note: If DDNAME, DSNLIST, or DSN is not specified, output goes to the file defined by the REPORT control statement.

DSN= specifies the data set name where head movement data is to be written

The data set must be previously allocated and cataloged. This parameter is invalid if DDNAME or DSNLIST is used. For DSN, do not include the ddname CMFCDS*xx* or CMFIDS*xx* in the CMF MONITOR Extractor JCL.

ALTDSN= specifies the data set name of a previously allocated, cataloged Extractor data set to be used for head movement data if the data set specified by DDNAME or DSN fills up

If this option is used, do not include the ddname CMFCDS*xx* or CMFIDS*xx* in the CMF MONITOR Extractor JCL. This parameter is invalid if the DSNLIST parameter is used.

SAMPLE= specifies the number of milliseconds between data gathering cycles for this group of devices

The default is 33. Acceptable values are 20 to 9999.

OFFLINE= specify OFFLINE=YES if entries should be reserved in the sample tables for selected devices that are marked offline

If these devices are then varied online, they are sampled. The default of NO eliminates unwanted dummy device entries.

BUFSIZE= specifies the maximum number of WKQEs (work queue elements) that the HEADMOVE sampler can acquire per sampling function

The default value is BUFSIZE=20. A WKQE is a unit of Extractor-managed CSA storage. There are approximately 3600 WKQEs available to the CMF MONITOR Extractor when CSA=512 is defined on the REPORT control statement.

Note: The BUFSIZE parameter should be defined if the SUSPEND and RESUME parameters are defined.

SUSPEND= defines a suspension threshold for CSA usage by the HEADMOVE sampler

The suspension threshold is the minimum number of WKQEs that must be available in CSA or HEADMOVE sampling suspends. The default is SUSPEND=0, which provides no suspension threshold. The suspension threshold is calculated as follows:

(value of SUSPEND=) * (value of BUFSIZE=) = (suspension threshold)

The SUSPEND, RESUME, and BUFSIZE parameters must always be defined together.

RESUME= defines a resumption threshold for CSA usage by the HEADMOVE sampler

The resumption threshold is the minimum number of WKQEs that must be available in CSA for HEADMOVE sampling to resume from a period of suspension. The default is RESUME=0, which provides no resumption threshold.

The RESUME value must always be greater than the SUSPEND value. The SUSPEND, RESUME, and BUFSIZE parameters must always be defined together.

The resumption threshold is calculated as follows:

(value of RESUME) * (value of BUFSIZE) = (resumption threshold)

Using the SUSPEND, RESUME, and BUFSIZE parameters

The following procedure should provide reasonable SUSPEND, RESUME, and BUFSIZE values for a system providing 512 K of CSA:

1. Specify CSA=512 on the REPORT control statement (see “REPORT” on page 176 for more information). The maximum CSA allowance provides the best environment.

2. Specify BUFSIZE=100 on the HEADMOVE control statement. If you already use a BUFSIZE value greater than 100, set it to 200.

3. Specify SUSPEND=4 on the HEADMOVE control statement.

4. Specify RESUME=8 on the HEADMOVE control statement.

If numerous suspension / resumption messages are logged while running the Extractor, you might need to adjust the SUSPEND and RESUME values. The following list provides some guidelines:

- Numerous suspension messages indicate that the SUSPEND value might be too large. Decrease the SUSPEND value by 1.
- Long suspension periods indicate that the RESUME value might be too large. Decrease the RESUME value by 1.
- Numerous short suspension periods indicate that the SUSPEND and RESUME values might be too close numerically. Increase the RESUME value by 1.
- If HEADMOVE terminates, the SUSPEND value is too small. Increase SUSPEND by 1 and reconsider the value of RESUME.

Example

```
HEADMOVE  ALL,EXCEPT=(142,145:147),OFFLINE=YES
```

The Extractor measures head movement activity for all online DASDs, except device addresses 142 and 145 through 147, reserving entries in the sample tables for devices marked offline. By default, the data is sampled every 33 milliseconds, and the output goes to the file defined in the REPORT control statement.

HFS

```
HFS
[FSNLIST=(fsname1,...,fsname101)
```

Overview

The HFS control statement causes the Extractor to collect global, buffer, and file system statistics for the hierarchical file system (HFS).

Monitoring modes

One HFS statement can be specified for each monitor mode, CPM and IPM.

Sampler and record types

The HFS statement controls the HFS sampler, which produces SMF type 74-6 records.

Analyzer statements and reports

HFS (see page 270)
HFS Statistics Report (see page 454)

Other BMC Software products

The HFS control statement is used by the CMF MONITOR Extractor only.

Parameters

The parameter for the HFS control statement is as follows:

FSNLIST= (*optional*) specifies the names of up to 101 hierarchical file systems

Each name is the fully qualified name of the MVS HFS data set (without quotation marks) that contains the file system.

This parameter is specified to produce the File System Statistics section of the HFS Statistics Report.

Example

```
HFS
```

The Extractor collects HFS global and buffer statistics.

```
HFS FSNLIST=SYSI.IBMOEM.SMS.OS270.GLDE.ETC
```

The Extractor collects HFS global and buffer statistics and the statistics for the file system SYSI.IBMOEM.SMS.OS270.GLDE.ETC.

IOQ

```
IOQ
[,IGNORE={NO|INACTIVE}]
[,CLASS=(class[,class]...)]
```

Overview

The IOQ control statement causes the Extractor to collect I/O configuration and activity data. Statistics are gathered for I/O processors (IOP) and Logical Control Units (LCU) and written to SMF type 78-3 records. These records are used by the Analyzer to produce the I/O Queuing Activity Report.

NOTE

When MVS runs under VM, do not specify the IOQ control statement. VM does not allow configuration data to be read, so the IOQ sampler does not function.

Monitoring modes

One IOQ statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The IOQ statement controls the IOQS sampler, which produces SMF type 78-3 records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 356)

IOQ (see page 272)
I/O Queuing Activity Report (see page 464)

Other BMC Software products

The IOQ control statement is used by the CMF MONITOR Extractor only.

Parameters

Parameters for the IOQ control statement are as follows:

IGNORE= specifies how LCUs that have no activity during an interval are to be treated

- **INACTIVE** specifies that no SMF 78-3 data be written for LCUs that had no activity for the reporting period
- **NO** specifies that SMF 78-3 data be written for all devices, whether or not they had any activity

  This value is the default.

CLASS= specifies the classes of devices to be sampled

The default is *all classes*.

- **TAPE** magnetic tape devices
- **COMM** communications equipment
- **DASD** direct access storage devices
- **GRAPH** graphics devices
- **UNITR** unit record devices
- **CHRDR** character reader devices

Examples

```
IOQ
```

This example invokes I/O queuing sampling for all device classes.

```
IOQ CLASS=DASD
```

This example invokes I/O queuing sampling for DASD devices only.

LINKMAP

```
LINKMAP
[INITRECS={FIRSTDS|EACHDS}]
```

Overview

The LINKMAP control statement causes the Extractor to produce mapping data on modules that are located in the pageable link pack area (PLPA).

Monitoring modes

One LINKMAP statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The LINKMAP statement controls the LPAM sampler, which produces CMF type 240-16 user records.

Analyzer statements and reports

LINKPACK (see page 273)
Link Pack Area Report (see page 472)

Other BMC Software products

The LINKMAP control statement is used by the CMF MONITOR Extractor only.

Parameters

The parameter for the LINKMAP control statement is as follows:

INITRECS= specifies when LINKMAP initial records are to be written

- If you specify INITRECS=FIRSTDS, LINKMAP records are written only when the Extractor begins recording. This selection is the default.
- If you specify INITRECS=EACHDS, LINKMAP records are written when the Extractor begins recording as well as whenever Extractor output flips to an alternate data set. In this case, the records have improved integrity, although slightly larger SMF files are produced.

Examples

```
LINKMAP
```

This example generates link pack mapping records only at Extractor initialization.

```
LINKMAP INITRECS=EACHDS
```

This example generates link pack mapping records at Extractor initialization and each time that recording flips to an alternate data set.

OMVS

```
OMVS
[SAMPLE={1000|nnnn}]
```

Overview

The OMVS control statement causes the Extractor to collect data on MVS OpenEdition kernel activity.

Monitoring modes

One OMVS statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The OMVS statement controls the OMVS sampler, which produces SMF type 74-3 records.

Analyzer statements and reports

OMVS (see page 276)
OMVS Kernel Activity Report (see page 488)

Other BMC Software products

The OMVS control statement is used by the CMF MONITOR Extractor only.

Restrictions

To run this control statement, the user ID under which the Extractor is running must be designated as a valid OMVS user ID. You can use the security system facilities for your system to define the user ID to OMVS. For more information about defining security for OMVS accounts, see the IBM manual *z/OS UNIX System Services Planning*.

Parameters

The parameter for the OMVS control statement is as follows:

SAMPLE= specifies the number of milliseconds between data gathering cycles for the paging data being gathered

The default value is 1000.

Example

```
OMVS  SAMPLE=5000
```

The Extractor samples OpenEdition MVS kernel activity data every five seconds.

PAGING

```
PAGING
[SAMPLE={5000|nnnn}]
```

Overview

The PAGING control statement causes the Extractor to collect paging and swapping data.

Monitoring modes

One PAGING statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The PAGING statement controls the PAGS sampler, which produces SMF type 71 records and CMF type 240-03 user records.

Analyzer statements and reports

AUXSTOR (see page 215)
Auxiliary Storage Report (see page 332)

CMFSUM (see page 223)
CMF Summary Report (see page 356)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

SRM (see page 307)
System Resources Manager Report (see page 515)

STORAGE (see page 308)
Storage Management Report (see page 509)

Other BMC Software products

The PAGING control statement must be defined in the Extractor JCL for MAINVIEW for z/OS. (See “MAINVIEW for z/OS” on page 35 for more information.)

Parameters

The parameter for the PAGING control statement is as follows:

SAMPLE=	specifies the number of milliseconds between data gathering cycles for the paging data being gathered The default and minimum value is 5000. Because this sampler causes high overhead, acceptable values are 5000 to 9999.

Example

```
PAGING  SAMPLE=5000
```

The Extractor samples paging data every five seconds.

PGDDLAY

```
PGDDLAY
```

Overview

The PGDDLAY sampler samples delay, storage utilization data, and all service class periods within each service class.

Monitoring modes

Only one PGDDLAY statement can be specified for a monitoring mode, either continuous (CPM) or intermittent (IPM), but not both.

Sampler and record types

The PGDDLAY statement controls the PGDS sampler, which produces SMF type 72-4 records for WLM.

Example

```
PGDDLAY
```

This example invokes the PGDS sampler, which samples delay and storage utilization data for service classes.

REPORT

```
REPORT {CPM|IPM}
[,INTERVAL={30|nn|HOUR|QTR|HALF|SMF}]
[,RUNTIME={1440|nnnn}]
[,SYNCH={00|nn|HOUR|QTR|HALF|SMF}]
[,CSA={80|nnn}]
[,SMFRECID={240|nnn}]
[,DSNLIST=(dsn1,...,dsn101)]
[,DSN='dsn']
[,ALTDSN='dsn']
[,SMF={NO|YES}]
[,DISP=NEW]
[,GBLS={YES|NO|1000...|9000}]
[, SRVCLASS={SYSSTC|WLM-RULE}]
```

Overview

This statement is required for all BMC Software products that use the CMF MONITOR Extractor. The REPORT control statement sets global parameters for an Extractor batch report job. It must be defined before any other Extractor control statement in the JCL. Only one REPORT statement can be defined in a single job.

REPORT statement controls include the following Extractor functions:

- monitoring mode
- runtime value
- recording interval
- location of where data is to be recorded

For example, Extractor output is written to SMF, if all of the following conditions exist:

- SMF=YES and SMFRECID parameters are specified.
- DSN, ALTDSN, and DSNLIST parameters are not specified.

Monitoring modes

One REPORT statement must be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The REPORT statement controls both the RECD and GBLS samplers, which produce CMF type 240-11 and 240-00 user records, respectively.

Analyzer statements and reports

none

NOTE

The GBLS=YES parameter must be defined on the REPORT control statement for

- LINKPACK (see page 273)
- Link Pack Area Report (see page 472)

Other BMC Software products

The REPORT control statement must be defined in the Extractor JCL, for BMC Software products that use the CMF MONITOR Extractor, which include

- DSO (see "DSO Analyzer" on page 31 for more information)
- MAINVIEW for z/OS (see "MAINVIEW for z/OS" on page 35 for more information)

Parameters

Parameters for the REPORT control statement are as follows:

CPM or IPM specifies whether the REPORT control statement and the succeeding control statements are for continuous (CPM) or intermittent (IPM) mode

This positional parameter must be specified.

INTERVAL=

specifies whether the REPORT control statement and the succeeding control statements are for continuous (CPM) or intermittent (IPM) mode

This positional parameter must be specified.

A number or a keyword parameter can be defined as the INTERVAL value. The default is 30. The maximum allowable numeric value is 60; the minimum allowable numeric value is 2.

Note: The INTERVAL parameter works in conjunction with the SYNCH parameter.

Keyword parameters are described in the following list:

HALF same as 30 minutes

HOUR same as 60 minutes

QTR same as 15 minutes

SMF same as the recording interval and synchronization value defined for SMF in the SMFPRM*xx* member of SYS1.PARMLIB

You can record data to CMF MONITOR data sets and still use INTERVAL=SMF.

Note: The INTERVAL parameter works in conjunction with the SYNCH parameter. If you define INTERVAL=SMF, do not define a value for the SYNCH parameter, unless you define SYNCH=SMF. It is not necessary to define both parameters, because defining one of these parameters with the SMF value causes the Extractor to use the same recording and synchronization values defined to SMF.

RUNTIME=

specifies the number of minutes that this monitoring mode is to remain active. This parameter is intended primarily for use with IPM mode. CPM mode should be allowed to run as long as possible so that long-term trending data can be collected.

The default is 1440 (untimed), which is also the maximum value. If the specified interval is greater than run time, the run time uses the default of 1440. The minimum value is 2.

Note: This time value is for elapsed time, not CPU time.

SYNCH= specifies the number of minutes into the hour to which the recording interval is synchronized

The recording interval is defined by the INTERVAL parameter.

The SYNCH parameter affects the INTERVAL parameter by determining the start time of recording, whereas the INTERVAL parameter defines the amount of time that passes before records are gathered by the Extractor.

The maximum SYNCH value is 60 minutes; the minimum value is 00, which is the default. Keyword parameters define the same time values as they do for INTERVAL.

Note: The SYNCH parameter works in conjunction with the INTERVAL parameter. If you define SYNCH=SMF, do not define a value for the INTERVAL parameter, unless you define INTERVAL=SMF. It is not necessary to define both parameters, because defining one of these parameters with the SMF value causes the Extractor to use the same recording and synchronization values defined to SMF.

CSA= specifies the amount of extended CSA required in K (1024 bytes)

The default is 80. The maximum value is 1024.

Note: A small amount of CSA is obtained below the line. The majority of CSA requested by the Extractor is satisfied from above the 16-megabyte line (ECSA).

SMFRECID= specifies an identification for the CMF MONITOR user subtype SMF records written by the Extractor

If you intend to write CMF MONITOR records to SMF, SMFRECID must be defined.

CMF MONITOR verifies that the number specified as the SMF record ID is a value between 128 and 255. The SMFRECID value must be unique from SMF record IDs used by other software on your system. The default is 240, if SMF=YES is not specified.

Note: If SMF=YES is specified, SMFRECID must also be specified.

DSNLIST=

specifies the names of up to 101 pre-allocated, cataloged data sets that contain the measurement data written by CPM or IPM mode

Data sets are dynamically allocated. If this option is used, do not include the ddnames CMFCPM*xx* or CMFIPM*xx* in the CMF MONITOR Extractor JCL, and do not define the SMF=YES parameter to the REPORT statement. You can use this parameter instead of the DSN and ALTDSN parameters.

You can specify DSNLIST=NULLFILE if you want the records to be discarded. This option can be useful if you are using IPM mode for running either CMF MONITOR Online or MAINVIEW for z/OS.

DSN=

specifies the name of a pre-allocated, cataloged data set that contains the measurement data written by CPM or IPM mode

The data set is dynamically allocated. If this option is used, do not include the ddnames CMFCPM*xx* or CMFIPM*xx* in the CMF MONITOR Extractor JCL, and do not define the SMF=YES parameter to the REPORT statement. Do not use this parameter if you use the DSNLIST parameter.

ALTDSN=

specifies the name of a pre-allocated, cataloged data set that is to be flipped to when the data set specified by DSN fills up with recorded data

Use this parameter only if you also use the DSN parameter. If ALTDSN is used, do not include the ddnames CMFCPM*xx* or CMFIPM*xx* in the CMF MONITOR Extractor JCL, and do not define the SMF=YES parameter to the REPORT statement.

SMF=

specifies whether data should be written to the SMF data sets

If SMF=YES is specified, SMFRECID must also be specified explicitly. NO is the default.

Notes:

- If SMF=YES is specified, do not define the DSNLIST, DSN, or ALTDSN parameters on the REPORT control statement, as they are mutually exclusive with SMF=YES.

- When SMF=YES is specified, the //CMFCPM1 DD, //CMFCPM2 DD, //CMFIPM1 DD, and //CMFIPM2 DD statements are ignored, if defined in the JCL.

- If data is to be recorded to SMF, the SYS and/or SUBSYS parameter of member SMFPRM*xx* in SYS1.PARMLIB must be specified.

DISP=NEW specifies at initialization that the primary CMF MONITOR Extractor output data set for the selected monitoring mode (CPM or IPM) is to be rewritten from the beginning

If you are recording to SMF, the DISP=NEW parameter is not required.

The DISP=NEW parameter applies only at CMF MONITOR Extractor initialization, when the primary output data set is first opened. When the primary data set becomes full and an alternate data set is switched to by the CMF MONITOR Extractor, the disposition automatically switches to MOD, and DISP=NEW has no effect.

Note: If you are using VSAM data sets and the first data set has been opened (by PERUSE, for example), the CMF MONITOR Extractor skips the first data set.

GBLS= gathers data needed for the link pack report

NO specifies that no global sampling takes place; CMF type 11 records are suppressed. However, the global sampler continues to perform some sampling functions under the DIE; these functions are used by other samplers.

The sample rate for GBLS can be controlled by specifying the number of milliseconds in multiples of 1000. The default is YES with a sampling rate of 1000.

SRVCLASS= specifies whether the MVS PAS is to run in the SYSSTC service class

If SYSSTC is specified, upon CPM mode initialization and at the end of every recording interval of CPM mode, CMF MONITOR assigns the MVS PAS to the SYSSTC service class if it does not run in this service class. SYSSTC is the default.

If WLM-RULE is specified, the MVS PAS runs in the service class determined by WLM classification rules. This service class should have an execution velocity goal of at least 80%.

Examples

```
REPORT IPM,INTERVAL=QTR,RUNTIME=30,SMFRECID=187,DSN='CMF.IPM.OUTPUT',
                  ALTDSN='CMF.IPM.ALT.OUTPUT'
```

The CMF MONITOR Extractor writes SMF type 70 and CMF type user records with an ID of 187 every quarter-hour over a 30-minute period. The CMF MONITOR Extractor is running in IPM mode. The data is written to data set CMF.IPM.OUTPUT, and if that data set is filled, an alternate data set, CMF.IPM.ALT.OUTPUT is allocated.

```
REPORT CPM,INTERVAL=30,SMFRECID=18
```

The CMF MONITOR Extractor writes SMF records and CMF type user records with an ID of 188 every 30 minutes (continuous mode). Defaults are used for CSA area (80K) and the runtime period of 1440 minutes, which is forever.

```
REPORT IPM,RUNTIME=120,SMFRECID=199,SMF=YES
```

The CMF MONITOR Extractor uses IPM mode to write SMF type 70 records and CMF type user records with an ID of 199. The IPM mode runs for 120 minutes. Data is written to the SMF data sets. Defaults are used for the sampling interval (30 minutes) and CSA area size (80 K).

TRACE

```
TRACE
[SAMPLE={60|nn}]
[,SRB={NO|YES}]
[,SYSEVENT=(n1,...,n6)]
[,ASID=asid]
[,ID=0|n]
[,NAME=jobname]
```

Overview

The TRACE control statement invokes the trace facility of the CMF MONITOR Extractor. Before using this facility, review the information explained in "SRB and SRM TRACE facilities" on page 186.

The trace facility uses two sampling methods:

SRB	schedules a user-written SRB routine at specified intervals A default SRB routine, CX10UMOD, is located in the CMF MONITOR BBSAMP data set.
SRM	allows you to trace specified SYSEVENTs This facility traces the job name for which the SYSEVENT was issued, along with parameter registers zero and one. For the TSEVENT (SYSEVENT 0), the name of the TSO command is traced rather than the parameter registers.

Monitoring modes

One TRACE statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The TRACE statement controls the TRCE sampler, which produces CMF type 240-18 user records.

Analyzer statements and reports

TRACE (see page 312)
Trace Report (see page 528)

Other BMC Software products

The TRACE control statement is used by the CMF MONITOR Extractor only.

Parameters

Parameters for the TRACE control statement are as follows:

SAMPLE= indicates the rate, in seconds, at which the user-written SRB routine receives control

Acceptable values are 1 to 999. The default is 60 seconds.

SRB= indicates whether the user-written trace routine is to be scheduled

The default is NO.

SYSEVENT= lists 1 to 6 optional SYSEVENTs that are to be traced

If nothing is specified, the SYSEVENT trace is suppressed. Valid SYSEVENT IDs are 0 through 79. A specification of SRB=NO and no SYSEVENTs results in an error.

Note: SYSEVENT data is traced for all address spaces, unless a NAME or ASID parameter is specified.

ID= specifies the trace ID to be assigned to the SYSEVENT trace records

The default is zero, which signifies that each SYSEVENT trace record is to be identified with the SYSEVENT ID. If ID is specified, all SYSEVENT trace records receive that ID.

Note: The ID parameter is ignored if no SYSEVENT parameter is defined.

NAME= specifies the name of the batch job, Started Task, or TSO logon ID for which SYSEVENT data is to be traced

The NAME parameter is invalid if ASID is defined.

Note: The NAME parameter is ignored if no SYSEVENT parameter is defined.

ASID= specifies the address space identifier for which SYSEVENT data is to be traced

ASID is invalid if NAME is used.

Note: The ASID parameter is ignored if no SYSEVENT parameter is defined.

Examples

```
TRACE   SAMPLE=1,SRB=YE
```

This example invokes the trace sampler; a user SRB is scheduled every second.

```
TRACE   SRB=NO,SYSEVENT=(0,8,9,15,20,21),NAME=TSOUSER
```

This example invokes the trace sampler with the SRM option. The following SYSEVENTs are examples of those traced for TSOUSER:

0	PPMODE
8	job select
9	job termination
15	swap-out complete
20	enqueue hold
21	enqueue release

Each trace entry has the ID of the SYSEVENT number.

```
TRACE   SAMPLE=10,SRB=YES,SYSEVENT=15,ID=77,ASID=10
```

This example invokes the trace sampler with both the SRB and SRM options. The user SRB is scheduled every 10 seconds. The swap-out complete system event (SYSEVENT 15) is traced for address space 10 and is stamped with an ID of 77.

SRB and SRM TRACE facilities

The CMF MONITOR trace facility uses service request block (SRB) and system resource manager (SRM) sampling methods.

SRB allows you to code trace routines for specialized system sampling

CMF MONITOR schedules global SRBs to perform many sampling functions. The trace facility permits you to interact with the SRB scheduling mechanism and introduce user-supplied trace routines.

At intervals specified by you, the SRB routine receives control and a data area of from 1 to 112 bytes is added as an entry in a trace record. Trace records vary in size up to a maximum of 4 K.

SRM allows selected SYSEVENTs to be traced

Using the SRM trace facility, you can trace selected SYSEVENTs as specified in the CMF MONITOR Extractor TRACE control statement (see “TRACE” on page 183 for more information). A trace for a SYSEVENT includes the name of the job for which the SYSEVENT was issued, along with the parameter registers, zero and one. TSEVENT (SYSEVENT zero) includes the TSO command name.

Implementing the user trace SRB

A sample user trace SRB routine is distributed with CMF MONITOR in two formats:

Load module
: contained in BBLINK member CX10UMOD

  The name of the trace SRB module must be CX10UMOD.

Assembly language source
: contained in BBSAMP member CX10UMOD

  The source for the SRB routine contains four macros that are used to code the user trace SRB. The SRB is composed of the following sections of code:

  ENTRY
  : contains standard linkage conventions that set up a save area and establish addressability to CMF MONITOR control blocks

    This action is performed by the CMFENTER MACRO.

  USER
  : is the section that you actually code

    After moving data to the trace data area, the CMFTRACE MACRO is issued to add that trace entry to the trace record.

  EXIT
  : performs standard-ending linkage conventions; provides an error routine that can terminate the trace, if necessary; and sets up a work area

    This action is performed by the CMFEXIT MACRO.

  DSECT
  : contains DSECTs for the necessary CMF MONITOR control blocks

    It is generated by the CMFDSECT MACRO.

NOTE

Before implementing a trace SRB sampler, BMC Software recommends that you assemble the sample SRB and study the source and comments. A typical CMF MONITOR trace SRB looks like the following:

CMFENTER
USER CODE
CMFTRACE
CMFEXIT
CMFDSECT

CMFTRACE MACRO

The CMFTRACE MACRO moves a data area into the trace buffer. This data area is then displayed as a trace entry in the CMF MONITOR trace record. The format of the CMFTRACE MACRO is

```
CMFTRACE  ID=99,LEN=U@LEN,ADDR=(R6),TYPE=SRB,MF=(E,(1))
```

ID=	identification of the trace entry Each trace entry has an ID that allows many trace entries to be created and selectively formatted. Any ID value from 0 to 99 is valid. The sample trace routine ID is 99. Because the SYSEVENT trace facility sets the default to ID for the trace entry of each SYSEVENT, BMC Software recommends that the SRB trace ID be restricted to a value between 73 and 99.
LEN=	length of the data area to be traced The maximum size of a data area is 112 bytes. The sample trace routine traces a 16-byte area.
ADDR=	address of the data area to be traced This address is set up with register 6 pointing to the 112-byte data area. Move the data to be traced to this area.
TYPE=	specifies the mode of this trace entry
MF=	specifies the address of a 16-byte remote parameter list used to maintain the trace SRB as re-entrant

User Trace SRB characteristics

The trace SRB routine is subject to the restrictions of all SRBs. SRBs have the following characteristics:

- All routines are entered in supervisor state, key zero, enabled, and unlocked.
- SVCs cannot be issued.
- The routines can take page faults.

The CMF MONITOR trace SRB must be re-entrant because it is possible for multiple CPUs belonging to the same multiprocessor to execute this routine at the same time. On entry to the SRB routine, register 13 points to a unique work area that contains a save area and a trace area. You should move all data to be traced to the trace area. The DSECT name, UWORK, maps this area; UWORK's address is in register 6.

Sample user trace SRB

The sample trace routine provided with CMF MONITOR traces the following fields:

- FLCEOPSW from the Prefixed Save Area (PSA)

  This field is the address from the interrupted external old PSW.

- ASCBJBNI / ASCBJBNS pointer to the current job name

To invoke the sample trace, specify **SRB=YES** in the CMF MONITOR Extractor TRACE control statement. The data generated is formatted by the Analyzer TRACE control statement.

User trace applications

The SRB routine is used as a pure trace function or as a sampler:

- A trace implies dumping a particular storage area repeatedly at specified intervals. This action is performed to key fields in system control blocks to track a special problem, such as a paging overload or excessive storage consumption.

- A sampler looks at various data areas and keeps an accumulator of the values, along with the number of observations. Average values can then be calculated.

To implement a sampler, follow this procedure:

1. Define a data area of up to 112 bytes.

2. Execute the required sampling function at every invocation of the SRB. Keep the sample values in the collection fields defined in the data area.

3. Issue a trace for this data area. At predetermined intervals, zero out the collection fields and start sampling again.

4. Write a program to analyze the data area written out. (CMF MONITOR Trace Report only formats the data.)

TRACE76

```
TRACE76
{FIELD=(name1,...,name120)}
[,SAMPLE={1000|nnnn}]
[,SAMPSET={60|nn}]
```

Overview

The TRACE76 control statement traces selected fields from MVS control blocks and summarizes the traced data in sets.

A set is a user-specified number of samples and contains maximum, minimum, and end snapshot values for the fields selected. The sum of values and the sum of squares are also included so that an average value and a standard deviation can be calculated.

Monitoring modes

One TRACE76 statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The TRACE76 statement controls the TRAS sampler, which produces SMF type 76 records.

Analyzer statements and reports

EXCEPTS (see page 252)
Exception Trace Detail Report (see page 442)

NOTE

The EXCEPTS Analyzer statement must have TRACE=YES defined.

GRAPH (see page 259)
Graphics Trace Detail Report (see page 451)

NOTE

The GRAPH Analyzer statement must have TYPE=TRACE defined.

Other BMC Software products

The TRACE76 control statement is used by the CMF MONITOR Extractor only.

Parameters

Parameters of the TRACE76 control statement are as follows:

FIELD=
name of the data field to be traced

This parameter is required. You can specify up to 120 fields. Valid data field names are described in "Traceable data fields" on page 630.

Note: Measures and trace values provided in "Traceable data fields" on page 630 contain variable information, such as device addresses, which reflects how the fields are reported in a graph. In defining measures, specify the full name of the measure—do not define the variable information.

SAMPLE=
rate, in milliseconds, between data gathering cycles

The default is 1000 (1 second).

SAMPSET=
number of samples per set

The default is 60.

SAMPLE and SAMPSET Usage Notes

These parameters directly impact the amount of ECSA that is required for the TRACE76 sampler. At initialization, the CMF MONITOR Extractor obtains a pool of ECSA in the amount specified in the CSA parameter of the REPORT control statement (see "REPORT" on page 176). Increasing the SAMPLE or SAMPSET values, or decreasing the INTERVAL value on the REPORT control statement, reduces the ECSA requirement.

Calculating CSA

Use this equation to calculate CSA required by the TRACE76 sampler:

((W * 20 * N) + (X * 32 * N) + J) * 40

N	(INTERVAL * 60000) / (SAMPLE *SAMPSET)
W	number of two-byte fields to trace
X	number of four-byte fields to trace
J	number of fields

This example assumes one two-byte and one four-byte field and uses the previously mentioned equation to calculate CSA:

```
N = (30 * 60000) / (1000 * 60) = 30

1 * 20 * 30        =    600
1 * 32 * 30        =    960
((600 + 960) + 2) * 40 =  62480
                       =  62K
```

Example

```
REPORT CPM,INTERVAL=30, CSA=50,SMFRECID=240,RUNTIME=60
TRACE76  SAMPLE=1000,SAMPSET=60,FIELD=(CCVUTILP,LSCTMTE)
```

The CMF MONITOR Extractor produces records at 30-minute intervals. The two-byte field **CCVUTILP** and the four-byte field **LSCTMTE** are being traced.

TSODATA

```
TSODATA
[SAMPLE={1000|nnnn}]
[,CMDS=(cl1,...,c32)]
[,LIMIT={32|nnn}]
[,USERS={YES|NO}]
```

Overview

The TSODATA control statement specifies the TSO commands to be monitored, and indicates whether TSO user data is collected.

TSO commands executed under ISPF do not issue SYSEVENT ZERO; therefore, these commands are not timed under ISPF. However, the EDIT and TEST subcommands under EDIT and TEST modes are counted and monitored. If FSE is being used and the subcommands are stacked, FSE replaced the last character of the subcommand with the character Z.

Monitoring modes

One TSODATA statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The TSODATA statement controls the TSOS sampler, which produces CMF type 240-20 and 240-21 user records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 355)

PERFSUM,PERFORM (see page 279 and page 277)
Performance Summary Report (see page 491)

TSOPERF (see page 315)
TSO Command Summary Report (see page 532) and TSO Interval Summary Report (see page 534)

TSOUSER (see page 317)
TSO User Summary Reports (see page 536), when TSOUSER USERS=YES is defined in the CMF MONITOR Extractor control statement set

Other BMC Software products

The TSODATA control statement is used by the CMF MONITOR Extractor only.

Parameters

Parameters for the TSODATA control statement are as follows:

SAMPLE= specifies the number of milliseconds between data gathering cycles for TSO data

The default is 1000. The sampling rate can be between 20 and 9999.

CMDS= identifies specific TSO commands (up to 32) that you want to monitor for response time information

If you do not specify CMD, the first *nnn* commands issued are monitored, where *nnn* is the value specified on the LIMIT parameter.

For standard TSO commands, TSO aliases (command short forms) are acceptable. Only one form of the command needs to be specified; for example, ALLOC also implies ALLOCATE. All commands are identified by their long names on the TSO Command Summary Report.

LIMIT= limits the total number of TSO commands (up to 251) that the CMF MONITOR Extractor can monitor

Any commands that are specified in the CMDS parameter are monitored along with any other commands that are detected, up to the LIMIT value. The default value for LIMIT is 32; the maximum value is 251. Specifying a value greater than 251 forces 251.

If LIMIT exceeds the number of commands that are specified in CMDS, the additional commands are the first ones issued.

If LIMIT is below the number of commands that are specified in CMDS, the limit is increased to the number of commands specified.

USERS= specifies whether statistics are to be collected for each TSO user at logoff

The default is YES. USERS=YES is required to produce the TSO User Summary report.

Examples

```
TSODATA
```

The first 32 TSO commands encountered are monitored at the default interval of once per second; according to default, data is gathered at each TSO user logoff.

```
TSODATA   SAMPLE=2000,CMDS=(TEST,STAT),LIMIT=50,USERS=NO
```

The TEST and STAT commands and up to 48 additional TSO commands are monitored every two seconds. No TSO user summary data is generated.

```
TSODATA   SAMPLE=500,LIMIT=25
```

Up to 251 TSO commands are monitored twice per second.

```
TSODATA   CMDS=EDIT,LIMIT=1,USERS=NO
```

The EDIT command is monitored at the default interval of once per second. No TSO user summary data is generated.

USER

```
USER
[SAMPLE={1000|nnnn}]
[,IEXIT=name]
[,DEXIT=name]
[,REXIT=name]
[,TEXIT=name]
```

Overview

The USER control statement provides the CMF MONITOR user the same user exit sampling capability as that provided by RMF. The USER control statement activates the CMF MONITOR user exits and specifies the sampling rate, which enables any user-written RMF exits to run under CMF MONITOR, unless they depend on RMF internal control blocks.

Default CMF MONITOR exit names and equivalent RMF exit names are listed in Table 17.

Table 17 Default CMF MONITOR exit names and equivalent RMF exit names

RMF exit name	CMF MONITOR exit name	Exit type
ERBMFIUC	CX10IXIT	initialization exit
ERBMFDUC	CX10DXIT	data gathering exit
ERBMFRUR	CX10RXIT	report writing exit
ERBMFTUR	CX10TXIT	termination exit

Monitoring modes

One USER statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The USER statement controls the USER sampler, which monitors user exit routines that can produce unique records types.

Analyzer statements and reports

none

Other BMC Software products

The USER control statement is used by the CMF MONITOR Extractor only.

How exits work with CMF MONITOR

CMF MONITOR supplies default exit routines that are equivalent to IBM utility module IEFBR14. These routines or user-supplied routines are loaded by the LEXICON module (CX10L020). Any combination of the exits can be used.

- If an exit is not present, CMF MONITOR returns control to the caller, passing back return code 0 to the calling program.
- IEXIT receives control from CSECT CX10USER in module CX10USER. REXIT and TEXIT receive control from CSECT USEREC in module CX10USER. The USER sampler, if loaded by IEXIT, receives control from CSECT USEDIE in module CX10USER.
- IEXIT, REXIT, and TEXIT are protected by an ESTAE. If an abend occurs in one of these routines, the appropriate CMF MONITOR error message identifies the abnormally ending routine as USER. DEXIT is protected by a functional recovery routine (FRR). If this routine ends abnormally, a message is issued for the USER sampler. Sampling for this function is terminated and a software error record is written to SYS1.LOGREC.

Parameters

Parameters for the USER control statement are as follows:

SAMPLE=	rate at which the sample routine (if loaded by CX10IXIT) receives control, in milliseconds; the default is 1000
IEXIT=	name of the initialization exit; the default is CX10IXIT
DEXIT=	name of the data gathering exit; the default is CX10DXIT
REXIT=	name of the report writing exit; the default is CX10RXIT
TEXIT=	name of the termination exit; the default is CX10TXIT

Example

```
USER   SAMPLE=2000,TEXIT=USERTXIT
```

This example invokes the USER sampling exits with control being given to the USER sampler every two seconds. All exits are loaded by using the default names except for the termination exit, which loads module name USERTXIT.

VSMDATA

```
VSMDATA
[SAMPLE={5000|nnnn}]
[JOBNAME=(job1,...,job32)]
```

Overview

The VSMDATA control statement causes the CMF MONITOR Extractor to collect information about the use of virtual storage, including the use of system queue area (SQA) by subpool, common storage area (CSA) by subpool and key, and private area virtual storage by specific jobs.

NOTE

Both CMF MONITOR VSMDATA sampler and RMF VSTOR sampler use the MVS VSMLIST service, which can have high CPU overhead.

Monitoring modes

One VSMDATA statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The VSMDATA statement controls the VSMS sampler, which produces SMF type 78-2 records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 355)

VIRTSTOR (see page 318)
Virtual Storage Activity Report (see page 539)

Other BMC Software products

The VSMDATA control statement is used by the CMF MONITOR Extractor only.

Usage notes for private area sampling

Private area sampling is most useful for long-running, nonswappable jobs. Sampling other jobs can be insufficient for the following reasons:

- Jobs are sampled only if they are active at the beginning of a recording interval. If a job starts after a recording interval has begun, it will not be sampled until the next recording interval.

  Similarly, if a job ends during a recording interval and another job of the same name begins, the data for the first job is flagged as *job terminated during interval* and sampling for the new job does not begin until the next recording interval. In addition, the data for the two like-named jobs is reported as though they were the same job.

- A job that begins and ends in the same recording interval is not sampled.

- If CMF MONITOR finds that a job is swapped out, it skips the current sample; thus, the sample count for a particular job might be below that of another job and below the overall sample count. If the job was swapped out for every sample during a recording interval, no data is recorded for that job.

Parameters

Parameters for the VSMDATA control statement are as follows:

SAMPLE=

specifies the number of milliseconds between data gathering cycles for the VSM data being gathered

The default is 5000, which is also the minimum value. The maximum sampling rate is 60,000.

Note: The virtual storage sampler is a high overhead sampler. For this reason, the lowest sampling rate allowed is 5000 milliseconds. If you specify a value below 5000, CMF MONITOR automatically changes it to 5000.

JOBNAME=

specifies the names of jobs for which private area storage usage data is to be collected

You can specify up to 32 job names. The sampling rate for jobs is the same as that for the common areas. If the JOBNAME parameter is not specified, no private area data is collected and the Private Area Storage Summary and Detail sections of the Virtual Storage Activity Report are unavailable.

Examples

```
VSMDATA  SAMPLE=5000
```

The CMF MONITOR Extractor samples virtual storage data once every five seconds (the minimum sampling rate). Private area data is not collected.

```
VSMDATA  SAMPLE=8000,JOBNAME=(RESOLVE,VAM2)
```

The CMF MONITOR Extractor samples virtual storage data for the common and private areas for jobs RESOLVE and VAM2 once every eight seconds.

WORKLOAD

```
WORKLOAD
```

Overview

The WORKLOAD control statement specifies that workload activity data is to be gathered.

Monitoring modes

One WORKLOAD statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The WORKLOAD statement controls the WLMS sampler, which produces SMF type 72-3 records.

Analyzer statements and reports

CMFSUM (see page 223)
CMF Summary Report (see page 355)

PERFSUM (see page 279)
Performance Summary Report (see page 491)

WLMGL (see page 321)
Workload Manager Goal Mode Reports (see page 549)

Other BMC Software products

The WORKLOAD control statement is used by the CMF MONITOR Extractor only.

Parameters

There are no parameters for this control statement.

Example

```
WORKLOAD
```

Workload activity data is to be collected.

XCFDATA

```
XCFDATA
[SAMPLE={1000|nnnn}]
```

Overview

The XCFDATA control statement causes the CMF MONITOR Extractor to collect performance data about the activity of the Cross-System Coupling Facility (XCF).

Monitoring modes

One XCFDATA statement can be specified for each monitoring mode, continuous (CPM) or intermittent (IPM).

Sampler and record types

The XCFDATA statement controls the XCFS sampler, which produces SMF type 74-2 records.

Analyzer statements and reports

XCF (see page 323)
Cross-System Coupling Facility Report (see page 413)

Other BMC Software products

The XCFDATA control statement is used by the CMF MONITOR Extractor.

Parameters

The parameter for the XCFDATA control statement is as follows:

SAMPLE=	specifies the sampling rate for collecting XCF data in milliseconds The default is 1000, which is also the minimum value.

Example

```
XCFDATA
```

Invokes collecting XCF measurement data once per second.

Chapter

7

Analyzer control statements

The CMF MONITOR Analyzer processes, analyzes, and formats data gathered by the CMF MONITOR Extractor or RMF into reports. The Analyzer uses general and report control statements to control its functions.

The Analyzer general control statements specify overall data selection and reporting parameters for CMF MONITOR reports. Report control statements cause specific reports to be produced. Report control statements use positional and keyword parameters to control the outcome of report data, just as Extractor control statements use parameters to control the input of report data.

The Analyzer also automatically produces some of its own reports that do not require control statements. These reports are Collection Phase Log reports, which describe various processing aspects of a particular batch report run.

General control statements

General control statements are valid only when placed at the beginning of the control statement set. These statements specify general data collection and reporting parameters for the Analyzer; the information from these control statements applies to all report requests.

General control statements can appear in any order within the general control statement group, except for positional general control statements. The DATETIME, CYCLE, and PERIOD control statements are positional, and must appear in that order.

NOTE

DATETIME is required if CYCLE and PERIOD are to be used.

A command error results if these commands are not entered in this sequence (see “CYCLE” on page 234, “DATETIME” on page 238, and “PERIOD” on page 281).

Table 18 summarizes all of the general control statements used by the Analyzer, the parameters for each statement, and function or report impact for each command.

Table 18 Analyzer general control statements (part 1 of 2)

Analyzer command	Parameter list	Function
CMFREC See page 221 for more information.	no parameters	writes only the CMF MONITOR records from an SMF input data set that were processed during a batch job to a data set that is defined in the Analyzer //CMXREC DD JCL statement
CYCLE See page 234 for more information.	{BIWEEKLY | CYCLE*nn* | CYCLE99 | DAILY | HOLIDAYS | MONTHLY | WEEKENDS | WEEKLY | WORKWKLY} [,(*sdate*:[*stime*],*edate*:[*etime*])1,... ,(*sdate*:[*stime*],*edate*:[*etime*])31]	specifies a subset of the DATETIME period that can be applied to any report by using the PERIOD statement
DATETIME See page 238 for more information.	(*sdate*:[*stime*],*edate*:[*etime*])	specifies an overall start and end date and time for the data that you want to use in your reports
DMSS See page 248 for more information.	{INIT | NOINIT} [,RESERVE={<u>400</u> | *nnn*}] [HIPERSP={<u>*</u> | nnn}]	specifies initialization of the DMSSMAIN work file, or the use and size of a hiperspace; also reserves storage for use by the DMSS component of the Analyzer
HEADERS See page 268 for more information.	[TITLE='title'] [,LOCATION='*location*'] [,LINES={<u>60</u> | *nn*}]	defines optional headline information and number of lines per report
PERIOD See page 281 for more information.	CYCLE=*cyclename* [,RPTS={<u>COMBINED</u> | SEPARATE}]	applies a date and time range defined by a CYCLE statement (CYCLE statements define a subset of the overall DATETIME value), to all or a single report control statement
RECTYPE See page 288 for more information.	[*xxx* | <u>240</u>] [,MONITOR={<u>CPM</u> | IPM | RMF}] [,STOPAFT={<u>EOF</u> | DATETIME}]	specifies characteristics of the Extractor data to be read; used if running RMF data or CMF MONITOR user type records with an SMF record ID other than 240

Table 18 Analyzer general control statements (part 2 of 2)

Analyzer command	Parameter list	Function
REPORTS See page 290 for more information.	REPORTS RPTGROUP=(*a*[,*b*[,*c*[,*d*]]]) [,DDGROUP=*e* [,SYSOUT=*class*]]	specifies how reports are to be grouped, whether separate DDs are to be used, and what default SYSOUT Class to use
SEVERITY See page 296 for more information.	nnn	specifies an error limit for processing
SHIFT See page 299 for more information.	[DINTV=(*starttime,intervals,duration*)] [(*startshift,endshift*)1 [,...,(*startshift,endshift*)96] | (080000,160000),(160000,240000), (000000,080000)] [,TYPE={GLOBAL | LOCAL}] [,DAYS={ALL | (MON | TUE | WED | THU | FRI | SAT | SUN) }] [,RPTS={SEPARATE | COMBINED | INTERVAL | DAILY}] [,UNIT=*unittype*]	specifies date and time range intervals by time of day and day of week, within the limitations set by the DATETIME, CYCLE, or PERIOD statements You can specify up to 96 intervals on the DINTV parameter, or a maximum of 96 startshift,endshift pairs is allowed.
SUBTITLE See page 309 for more information.	'text'	defines an additional headline
SYSPLEX See page 310 for more information.	[RPTS={COMBINED | SEPARATE}] [,SUBPLEX=(* | sysname*1*[,...,sysname*n*])}] [,TYPE={SYSID | SYSNAME}] [,INPUT={(dd*1*[,...,dd*n*])}]	provides a way to define the data set that is used in generating Analyzer reports

Report control statements

This section discusses the report control statements supported by the Analyzer. When a report control statement is used, the corresponding CMF MONITOR report is printed.

Table 19 summarizes the report control statements and parameters, and shows the corresponding report titles and associated Extractor control statements.

All statements follow the syntax rules in “Syntax conventions” on page 23.

Table 19 Analyzer report control statements (part 1 of 5)

Analyzer command	Report title	Extractor command
AUXSTOR See page 215 for more information.	Auxiliary Storage Report (see page 332)	ASMDATA (see page 127) PAGING (see page 173)
CACHEACT See page 216 for more information.	Cache Subsystem Overview Report (see page 337) with REPORT=OVERVIEW defined. Cache Subsystem Activity Report (see page 340) with REPORT=SUBSYS defined. Cache Device Activity Report (see page 347) with REPORT=OVERVIEW defined.	CACHE (see page 129)
CFACT See page 219 for more information.	Coupling Facility Activity Report (see page 378)	CFDATA (see page 132)
CHANNEL See page 220 for more information.	Channel Path Activity Report (see page 350)	CHANNEL (see page 134)
CMFSTAT See page 222 for more information.	CMF Record Statistics Report (see page 353)	none

Table 19 Analyzer report control statements (part 2 of 5)

Analyzer command	Report title	Extractor command
CMFSUM See page 223 for more information.	CMF Summary Report (see page 355)	CHANNEL (see page 134) CPU (see page 136) DEVICE (see page 145) IOQ (see page 167) PAGING (see page 173) TSODATA (see page 194) VSMDATA (see page 200) WORKLOAD (see page 203)
COMMSTOR See page 227 for more information.	Common Storage Usage Detail Report (see page 374) Common Storage Usage Summary Report (see page 377)	CSMON (see page 142)
CPU See page 230 for more information.	CPU Utilization Report (see page 393)	CPU (see page 136)
CPUCON See page 232 for more information.	Processor Concurrency Report (see page 500)	CPU (see page 136)
CRYPTO See page 233 for more information.	Cryptographic Hardware Activity Report (see page 418)	CRYPTO (see page 140)
DASD See page 237 for more information.	Direct Access Report (see page 425)	DEVICE (see page 145) HEADMOVE (see page 159)
DEVACT See page 243 for more information.	Device Activity Report (see page 421)	DEVICE (see page 145)

Table 19 Analyzer report control statements (part 3 of 5)

Analyzer command	Report title	Extractor command
DOMINO See page 246 for more information.	LOTUS DOMINO Server Summary Report (see Figure 75 on page 480) LOTUS DOMINO Server Detail Report (see Figure 76 on page 482) LOTUS DOMINO Database Activity Report (see page 485) LOTUS DOMINO User Activity Report (see page 487)	records created by LOTUS DOMINO server
ENQUEUE See page 250 for more information.	Enqueue Conflict Report (see page 434)	ENQUEUE (see page 152)
ESS See page 251 for more information.	ESS Statistics Report (see page 436)	CACHE (see page 129)
EXCEPTS See page 252 for more information.	Exception Trace Detail Report (see page 442) Exception Subreport (see page 440)	TRACE76 (see page 191) depends on measure or measures being reported
FICONSW See page 258 for more information.	FICON Director Activity Report (see page 448)	FICONSW (see page 158)
GRAPH See page 259 for more information.	Graphics Trace Detail Report (see page 451) Any graph type	TRACE76 (see page 191) depends on graph
HFS See page 270 for more information.	HFS Statistics Report (see page 453)	HFS (see page 165)
HTTP See page 271 for more information.	HTTP Server Summary Report (see page 459) HTTP Server Detail Report (see page 460)	records created by HTTP server
IOQ See page 272 for more information.	I/O Queuing Activity Report (see page 464)	IOQ (see page 167)

Table 19 Analyzer report control statements (part 4 of 5)

Analyzer command	Report title	Extractor command
LINKPACK See page 273 for more information.	Link Pack Area Report (see page 472)	LINKMAP (see page 169) Also, the Extractor REPORT control statement must have GBLS=YES defined (see page 176).
LPARCOMB See page 274 for more information.	CPU Utilization Report (see page 393)	none
OMVS See page 276 for more information.	OMVS Kernel Activity Report (see page 171)	none
PERFORM See page 277 for more information.	Performance Summary Report (see page 491)	none
PERFSUM See page 279 for more information.	Performance Summary Report (see page 491)	ASMDATA (see page 127) CHANNEL (see page 134) CPU (see page 136) DEVICE (see page 145) ENQUEUE (see page 152) EXTSUM (see page 154) PAGING (see page 173) TSODATA (see page 194) WORKLOAD (see page 203)
PROTKEY See page 284 for more information.	CPU Utilization by Protect Key Report (see page 411)	CPU (see page 136)
PRSM See page 246 for more information.	Logical Partition Report (see page 476)	CPU (see page 136)
SHARDEV See page 297 for more information.	Shared Device Activity Report (see page 505)	DEVICE (see page 145)

Table 19 Analyzer report control statements (part 5 of 5)

Analyzer command	Report title	Extractor command
SRM See page 307 for more information.	System Resources Manager Report (see page 515)	CPU (see page 136) PAGING (see page 173)
STORAGE See page 308 for more information.	Storage Management Report (see page 509)	PAGING (see page 173)
TRACE See page 312 for more information.	Trace Report (see page 528)	TRACE (see page 183)
TSOPERF See page 315 for more information.	TSO Command Summary Report (see page 532) TSO Interval Summary Report (see page 534)	CPU (see page 136) TSODATA (see page 194)
TSOUSER See page 317 for more information.	TSO User Summary Report (see page 536)	TSODATA with USERS=YES defined (see page 194)
VIRTSTOR See page 318 for more information.	Virtual Storage Activity Report (see page 539)	VSMDATA (see page 200)
VOLSER See page 320 for more information. Command requires DASD report control statement.	Direct Access Report Plot of Volume (see page 427)	none
WLMGL See page 321 for more information.	Workload Manager Goal Mode Reports (see page 549)	WORKLOAD (see page 203)
XCF See page 323 for more information.	Cross-System Coupling Facility Report (see page 413)	XCFDATA (see page 205)

AUXSTOR

```
AUXSTOR
```

Overview

The AUXSTOR control statement requests the Auxiliary Storage Report that is discussed in "Auxiliary Storage Report" on page 332.

The AUXSTOR statement uses data collected by the ASMDATA Extractor control statement (see "ASMDATA" on page 127) and the PAGING Extractor control statement (see "PAGING" on page 173).

Parameters

This statement has no parameters.

Example

```
AUXSTOR
```

The Auxiliary Storage Report is produced.

CACHEACT

```
CACHEACT
[REPORT={OVERVIEW|SUBSYS|DEVICE}]
[,MODEL={ALL|3990-03|3990-06|2105-20}]
[,ORDER={TOTAL|DASD|CACHE|ASYNC }]
[,SUBSYS=(nnnn,aaaa:bbbb,.....,dddd:eeee)16]
[,DEVICE=(nnnn,aaaa:bbbb,.....,dddd:eeee)16]
[,EXCLUDE=(nnnn,aaaa:bbbb,.....,dddd:eeee)16]
[,THRESHLD=nnnnnnnn]
```

Overview

The CACHEACT control statement causes the Analyzer to generate reports by using the SMF type 74-5 records that are generated by the CACHE Extractor control statement (see “CACHE” on page 129). The following reports can be produced:

Cache Subsystem Overview Report
This report contains a line of summary data for each cache subsystem, and a single page showing the DASD device activity. The last report page lists the most active DASD devices, by total and DASD I/O rates, and can have a variable number of lines, since filtering options and lack of device activity affect this section of the report. This report is the default.

Cache Subsystem Activity Report
This report contains detail data about every cache subsystem. For each subsystem, the report shows one page of detail information for the subsystem, plus one or more pages of summary data for every DASD device attached to that subsystem. These pages consist of one line per device, plus additional lines for RAID rank activity.

Cache Device Activity Report
This report contains detail data about every cache subsystem and detail data about every logical DASD device attached to a subsystem. For each cache subsystem, this report generates one page of detail subsystem information, and one or more pages of detail device information, one device per page.

NOTE

This report can be quite voluminous, so BMC Software suggests that you request it only for specific devices or ranges, or when troubleshooting cache subsystem problems.

Parameters

Parameters for this statement are as follows:

REPORT= specifies the type of report to generate, as follows:

- **OVERVIEW** produces the default Cache Subsystem Overview Report (see page 337)
- **SUBSYS** produces the Cache Subsystem Activity Report (see page 340)
- **DEVICE** produces the Cache Device Activity Report (see page 347)

MODEL= generates a report on particular cache subsystems by model ID:

- **ALL** reports all models; the default
- **3990-03** reports only model 3990-03 subsystems
- **3990-06** reports only model 3990-06 subsystems
- **2105-20** reports only model 2105-20 subsystems

ORDER= specifies the manner in which subsystems and devices in a report are ordered:

- **TOTAL** ordered by total I/O rates
- **DASD** ordered by DASD I/O rates
- **CACHE** ordered by Cache Hit rates
- **ASYNC** ordered by Async I/O rates

SUBSYS= generates a report on particular cache subsystems or ranges of subsystems by specifying a list of subsystem IDs or ranges of subsystem IDs

Any item in the list can be either a single subsystem ID (nnnn), or a range of subsystem IDs (aaaa:bbbb). They can appear in any combination and order, up to a maximum of 16. This parameter overrides any DEVICE or EXCLUDE filter, so devices that are specified by those keywords are ignored unless attached to the selected subsystem.

DEVICE= generates a report on particular cache DASD devices or ranges of DASD devices by specifying a list of device numbers or range of device numbers

Any item in the list can be either a device number (nnnn), or a range of device numbers (aaaa:bbbb). They can appear in any combination and order, up to a maximum of 16. This parameter is ignored when REPORT=OVERVIEW is specified or the default.

EXCLUDE= excludes particular cache DASD devices or ranges of DASD devices from all reports by specifying a list of device numbers or range of device numbers

Any item in the list can be either a device number (nnnn) or a range of device numbers (aaaa:bbbb). They can appear in any combination and order, up to a maximum of 16. If a device is specified in both the DEVICE= and EXCLUDE= lists, the former takes precedence. This parameter is ignored when REPORT=OVERVIEW is specified or the default.

THRESHLD= limits a report so that only those subsystems or devices meeting or exceeding the threshold value are included

The value is assumed to be the total I/O rate unless ORDER=DASD|CACHE|ASYNC is specified, in which case the threshold value is of the same type (DASD I/O rate, cache hit percentage, or Async I/O rate) as the order. Specify as an integer value, 0 to (2GB - 1). This parameter is ignored when REPORT=OVERVIEW is specified or the default.

Examples

```
CACHEACT
```

The Cache Subsystems Overview Report is produced.

```
CACHEACT    MODEL=2105-20,REPORT=SUBSYS
```

The Cache Subsystem Activity Report is produced for 2105 cache subsystems.

```
CACHEACT     REPORT=DEVICE,
             DEVICE=(AC00:AC1F,C108),ORDER=DASD
```

The Cache Device Activity Report is produced for devices AC00 to AC1F, and device C108, for a total of 33 devices. The report is ordered in descending DASD I/O rate sequence.

CFACT

```
CFACT
```

Overview

The CFACT control statement requests the Coupling Facility Activity Report. To collect synchronous data from all systems to which the coupling facility is connected, BMC Software recommends that you use the SYNCH=SMF parameter of the Extractor REPORT control statement.

NOTE

You must also specify the SYSPLEX control statement with RPTS=COMBINED to produce the Coupling Facility Activity Report.

Parameters

This statement has no parameters.

Example

```
SYSPLEX SUBPLEX=*,TYPE=SYSID,RPTS=COMBINED
CFACT
```

One Coupling Facility Activity Report is produced for each coupling facility.

CHANNEL

```
CHANNEL
[ZERO={YES|NO}]
```

Overview

The CHANNEL control statement requests the Channel Path Activity Report that is discussed in “Channel Path Activity Report” on page 350. If the Extractor data contains records from a system that has ESCON Multiple Image Facility (EMIF) support, the Channel Path Activity Report contains different fields that report on EMIF.

The CHANNEL statement uses data collected by the CHANNEL (see “CHANNEL” on page 134) Extractor control statements.

Parameters

The parameter for the CHANNEL control statement is as follows:

ZERO= controls whether channel paths with no activity during the reporting period are included in the report

The default is ZERO=YES, which includes all channel paths in the report, even those that were 0% busy.

When ZERO=NO is defined, reporting of channel paths with 0% busy is suppressed and does not appear on the report. All channel paths with busy activity are reported.

Examples

```
CHANNEL
```

The Channel Path Activity Report is produced and includes all channel paths, even those with no activity during the reporting period, because ZERO=YES is the default parameter.

```
CHANNEL   ZERO=NO
```

The Channel Path Activity Report is produced and includes only channel paths that had activity during the reporting period. Nonactive channel paths are not reported because of the ZERO=NO parameter.

CMFREC

```
CMFREC
```

Overview

The CMFREC control statement provides additional processing for CMF MONITOR SMF data. This control statement directs the Analyzer to identify the CMF MONITOR records saved to SMF, process them, and write only the CMF MONITOR records accepted for analysis to a data set defined by a //CMXREC DD statement. (See “Defining Analyzer JCL manually” on page 84.)

The CMFREC control statement is useful when processing a large amount of data. By removing extraneous data and saving only the processed CMF MONITOR records in a separate data set, subsequent runs against these records can be made without incurring additional overhead required to process those records not used by the Analyzer or excluded because of date, time, or system ID.

The CMFREC Analyzer statement is a general control statement and can process data from all Extractor control statements.

- The CMFREC control statement is ignored if no //CMXREC DD statement is defined in the Analyzer JCL.
- No report is generated by this control statement.

Parameters

This statement has no parameters.

Example

```
CMFREC
```

Only CMF MONITOR records are written to the data set that is defined on the //CMXREC DD statement.

CMFSTAT

```
CMFSTAT
```

Overview

The CMFSTAT control statement requests the CMF Records Statistics Report (see Figure 33 on page 353), which is a summary of date, time, and duration for all record types accepted for analysis.

This report includes information about all record types contained in the input data set.

Parameters

This statement has no parameters.

Example

```
CMFSTAT
```

The CMF Record Statistics Report is produced.

CMFSUM

```
CMFSUM
[INTERVAL={hh:mm:ss|Extractor interval|HOURLY|DAILY|
     WEEKLY|MONTHLY|QTRLY|SEMIANNL|FOREVER}]
[,MEASURE={ALL|RMF|(xxxxxxxx,...,xxxxxxxx)}]
```

Overview

The CMFSUM Analyzer control statement requests the CMF Summary Report that is discussed in Figure 34 on page 356.

The CMFSUM statement uses data collected by the following Extractor control statements:

- CHANNEL (see “CHANNEL” on page 134)
- CPU (see “CPU” on page 136)
- DEVICE (see “DEVICE” on page 145)
- IOQ (see “IOQ” on page 167)
- PAGING (see “PAGING” on page 173)
- TSODATA (see “TSODATA” on page 194)
- VSMDATA (see “VSMDATA” on page 200)

Parameters

Parameters for the CMFSUM control statement are as follows:

INTERVAL= specifies the time divisions on the graph

If no INTERVAL parameter is defined, the default used is the value defined to the INTERVAL parameter of the Extractor REPORT control statement (see “REPORT” on page 176). Do not define INTERVAL with a value greater than 24 hours; the interval will not be reported. The INTERVAL parameter can be specified with the following values:

hh:mm:ss

where

hh	Is the hour.
mm	Is the minutes.
ss	Is the seconds.

HOURLY Each line of the graph represents one hour.

INTERVAL= (*continued*)	**DAILY**	Each line of the graph represents one 24-hour period.
	WEEKLY	Each line of the graph represents one 168-hour period.
	MONTHLY	Each line of the graph represents one month.
	QTRLY	Each line of the graph represents a calendar year quarter.
	SEMIANNL	Each line of the graph represents a half year.
	FOREVER	Summarizes all input into one output line.

Note: If you define INTERVAL with the values WEEKLY, MONTHLY, QTRLY, SEMIANNL, or FOREVER, and you have MEASURE=INTERVAL defined, the CMF Summary Report does not report information in the INTERVAL *hh:mm:ss* field.

MEASURE= defines measures that are included in the report

The following values can be defined with this parameter:

ALL	Data on all measures are included in the report.
RMF	Data on the same measures that appear in the RMF Summary Report are included in the report.
xxxxxxxx	Is the name of a specific measure. Data on each defined measure is included in the report. You can define as many measures as you want. Measures must be separated by commas, and all of them enclosed within parentheses. Specific names for each measure that can be defined to the MEASURE parameter, as well as the corresponding report field name, are shown in Table 20 on page 225.

Note: If MEASURE=ALL is specified, or if more measures are specified than can be printed in one report, multiple reports are generated until all values have been reported upon.

Table 20 Valid measure names and corresponding report fields for MEASURE (part 1 of 2)

Measure name	Report field name
AFQUEUE	AVERAGE AVAILABLE FRAMES QUEUE
APPCAVG	APPC AVERAGE
APPCMAX	APPC MAXIMUM
AVGREADY	AVERAGE READY QUEUE
BATCHAVG	BATCH AVERAGE
BATCHMAX	BATCH MAXIMUM
CAPRATIO	AVERAGE CAPTURE RATIO
CHPUTIL	CHANNEL PATH UTILIZATION RATE
CHPBUSY	CHANNEL PATH BUSY
CPUBUSY	CPU BUSY
CPUBZMVS	MVS CPU BUSY
CPUSERV	CPU SERVICE RATE
CSALLOC	AVERAGE CSA ALLOCATED
DASDRATE	DASD RATE
DASDRESP	DASD RESPONSE TIME
DPAGING	DEMAND PAGING
ECSALLOC	AVERAGE ECSA ALLOCATED
EPGRATE	EXPANDED STORAGE PAGE RATE
ESFRAME	EXPANDED STORAGE FRAMES
ESQALLOC	AVERAGE ESQA ALLOCATED
EXCPRATE	AVERAGE EXCUSE RATE
FIXFRAME	AVERAGE FIXED FRAMES
HIGHUIC	HIGH UNREFERENCED INTERVAL COUNT
HSFRAME	AVERAGE HIPERSPACE FRAMES
INTERVAL	INTERVAL HH.MM.SS
IOSERV	I/O SERVICE RATE
LPARDISP	LPAR DISPATCH PERCENTAGE
MIGRAGE	MIGRATION AGE
MIGRATE	MIGRATION RATE
MSOSERV	MSO SERVICE RATE
OMVSAVG	OMVS AVERAGE
OMVSMAX	OMVS MAXIMUM
SPOOLUTL	SPOOL UTILIZATION
SQALLOC	AVERAGE SQA ALLOCATED
SRVRATE	SERVICE RATE
STCAVG	STC AVERAGE
STCMAX	STC MAXIMUM

Table 20 Valid measure names and corresponding report fields for MEASURE (part 2 of 2)

Measure name	Report field name
SWAPPAGS	SWAP PAGE RATE
SWAPRATE	SWAP RATE
TAPERATE	TAPE RATE
TRANRATE	TRANSACTION RATE
TSOAVG	TSO AVERAGE
TSOMAX	TSO MAXIMUM
TSOTRANS	TSO TRANSACTION RATE
VIOFRAME	AVERAGE VIO FRAMES
ZAAPBMVS	MVS ZAAP BUSY
ZAAPBUSY	ZAAP BUSY
ZIIPBMVS	MVS ZIIP BUSY
ZIIPBUSY	ZIIP BUSY

Examples

```
CMFSUM
```

A CMF Summary Report is produced with each row on the graph based on the period defined by the INTERVAL= valued defined to the Extractor REPORT control statement. (See “REPORT” on page 176 for more information.) The measures reported are identical to those reported in RMF Summary Report.

```
CMFSUM   INTERVAL=HOURLY,MEASURE=RMF
```

A CMF Summary Report is produced with each row on the graph representing a one-hour period, and the measures reported are the same as those in the RMF Summary Report.

```
CMFSUM   INTERVAL=00:30:00,MEASURE=(MIGRATE,MIGRAGE)
```

A CMF Summary Report is produced with each row on the graph representing a half-hour period, and the measures reported are the MIGRATE and MIGRAGE values.

```
CMFSUM   INTERVAL=DAILY
```

A CMF Summary Report is produced with each row on the graph representing a 24-hour period, and the measures reported are the same as those in the RMF Summary Report.

COMMSTOR

```
COMMSTOR
[REPORT={SUMMARY|DETAIL}]
[,SELECT=(name1,...,name32)]
[,SORT=([CSA<,AVG|MAX|MIN],|
    [CSA>,AVG|MAX|MIN],|
    [CSA,AVG|MAX|MIN],|
    [SQA<,AVG|MAX|MIN],|
    [SQA>,AVG|MAX|MIN],|
    [SQA,AVG|MAX|MIN])]
[,LIMIT=nnnn]
[TBLSIZE={20,000|nnnnnn}]
```

Overview

The COMMSTOR control statement requests either the Common Storage Usage Summary Report (see Figure 41 on page 377) or the Common Storage Usage Detail Report (see Figure 40 on page 375). If a summary report is requested, the average data values are printed. If a detail report is requested, the average, maximum, and minimum data values, and the date and time stamps of these data values are printed.

The COMMSTOR statement uses data that is collected by the CSMON Extractor control statement (see "CSMON" on page 142).

Parameters

Parameters for the COMMSTOR control statement are as follows:

REPORT= optionally selects either a summary report (default) that reports only the average CSA and SQA usage by job name, or a detail report that reports on the average, maximum, and minimum CSA and SQA usage by job name

SELECT= optionally selects from 1 to 32 specific job names or selection masks

Data for all other jobs is ignored. If SELECT is not specified, *all* job names are selected. The following wildcard characters can be used to define groups of job names:

+	any character
*****	any combination of characters

Note: Use the * wildcard character only at the end of a string. The + wildcard character can be used in any position within the string.

For example, the parameter SELECT=(+MF) would select only those job names that are three characters long and where the first character is any character, but the second and third characters are MF. On the other hand, the parameter SELECT=(+MF*) would select those job names that are from three to eight characters long and where the first character is any character, the second and third characters are MF, and the remaining characters are in any number and combination.

SORT= optionally selects the sort sequence for the report

The first entry defines what data area is to be the sort sequence.

If SORT is not specified, the sort is by job name.

CSA<	CSA below 16-MB line
CSA>	CSA above 16-MB line
CSA	all CSA
SQA<	SQA below 16-MB line
SQA>	SQA above 16-MB line
SQA	all SQA

The second SORT entry defines whether to use the average, maximum, or the minimum value as the sort sequence for the DETAIL report. These values define

MIN	minimum value; sort sequence ascending
MAX	maximum value; sort sequence descending
AVG	average value; sort sequence descending

LIMIT= specifies that only the first *nnnn* job name entries are displayed in the report

Job names are selected by alphabetical order.

TBLSIZE= specifies how much storage is allocated for an internal table used by the COMMSTOR report modules

The table is used to store job names and address space IDs (ASIDs) as ten-byte entries for each recording interval. The default size of the table is 20,000 bytes, which can store up to 2,000 job names and ASIDs.

Note: This parameter should be used only if error message CMF07385 is issued.

Examples

```
COMMSTOR  REPORT=DETAIL,SELECT=(ADMPRINT,ALT1),SORT=(CSA,AVG)
```

A detail Common Storage Usage Report is produced, containing CSA usage data for the jobs ADMPRINT and ALT1 only. The data is sorted by all CSA, regardless of whether it is above or below the 16M line, and the values are sorted by average in descending order.

```
COMMSTOR  REPORT=DETAIL,LIMIT=5
```

A detail Common Storage Usage Report is produced, containing CSA usage data for the first five jobs, based on job name by alphabetical order.

CPU

```
CPU
[MSUDIST=YES|NO]
[MSUDTAIL=4HOURMSU|INTVLMSU|INTVLWLM|TIME]
[MSUDTPCT=10|nnn]
[PRSMCHNG=MANYRPTS,ONERPT
```

Overview

The CPU control statement requests the CPU Utilization Report that is discussed in "CPU Utilization Report" on page 393.

To add combinations in the LPAR Combination Section or suppress this section altogether, see the LPARCOMB control statement that is discussed in "LPARCOMB" on page 274.

The CPU Analyzer statement uses data collected by the CPU Extractor control statement (see "CPU" on page 136).

Parameters

The parameters for the CPU control statement are as follows:

MSUDIST= specifies whether the Rolling 4-Hour MSU Usage Distribution section is produced; the default is YES

MSUDTAIL= specifies that the MSU Usage Detail section and the sorting order of the recording intervals are produced

- **4HOURMSU** The recording intervals are sorted in decreasing 4-hour MSU / hour (MSU / hour consumed in the four-hour period up to the end of the recording interval).
- **INTVLMSU** The recording intervals are sorted in decreasing interval MSU / hour (MSU / hour consumed during the recording interval).
- **INTVLWLM** The recording intervals are sorted in decreasing % capped by WLM during the recording interval.
- **TIME** The recording intervals are sorted in chronological order.

MSUDTPCT=
specifies that the top *nnn* % of recording intervals are formatted in the MSU Usage Detail section when 4HOURMSU, INTVLMSU, or INTVLWLM is specified for the MSUDTAIL parameter; *nnn* is in the range 1 to 100; the default is 10

MSUDTPCT is ignored when MSUDTAIL is not specified or when MSUDTAIL=TIME is specified.

PRSMCHNG=
specifies whether one or many reports are produced if changes are detected in PR/SM configuration

MANYRPTS
If changes are detected in PR/SM configuration, multiple reports are produced, one for each configuration. Examples of PR/SM configuration attributes include number of partitions, weight (share) of any partition, number of logical processors of any partition, and number of physical processors.

ONERPT
A single report is produced, ignoring all PR/SM configuration changes; the default.

Example

```
CPU MSUDTAIL=4HOURMSU
```

This example produces one CPU Utilization Report with MSU Usage Detail section containing the top 10% busiest of recording intervals sorted by four -hour MSU/hour consumed in the four-hour period up to the end of the recording interval.

```
CPU PRSMCHNG=MANYRPTS
```

This example produces multiple CPU Utilization Reports if changes are detected in PR/SM configuration.

```
CPU
```

This example produces one CPU Utilization Report, even if changes are detected in PR/SM configuration.

CPUCON

```
CPUCON
```

Overview

The CPUCON control statement requests the “Processor Concurrency Report” on page 500.

The CPUCON statement uses data collected by the CPU Extractor control statement (see “CPU” on page 136).

Parameters

This statement has no parameters.

Example

```
CPUCON
```

The Processor Concurrency Report is produced.

CRYPTO

```
CRYPTO
```

Overview

The CRYPTO control statement requests the “Cryptographic Hardware Activity Report” on page 418.

Parameters

This statement has no parameters.

Example

```
CRYPTO
```

The Cryptographic Hardware Activity Report is produced.

CYCLE

```
CYCLE
{BIWEEKLY|CYCLEnn|CYCLE99|DAILY|HOLIDAYS|
MONTHLY|WEEKENDS|WEEKLY|WORKWKLY}
[,(sdate:[stime],edate:[etime])1...,(sdate:[stime],edate:[etime])31]
```

Overview

The CYCLE control statement assigns a name to a unique date-time range. The period of the CYCLE range must be a subset of the range defined by the DATETIME control statement (see “DATETIME” on page 238 for more information). The CYCLE control statement requires that a DATETIME control statement also be defined.

A CYCLE range can be activated as a general default range to all report requests in a job or to individual report requests. Data is included in a report for any time range specified in a CYCLE statement, if the interval time for the data falls within that time range.

The CYCLE control statement only defines the name and time slice of a CYCLE range. A PERIOD statement (see “PERIOD” on page 281) is used to activate a CYCLE range. It is the placement of the PERIOD statement that determines whether one or all report requests are affected by a single CYCLE range.

When the CYCLE, DATETIME, and PERIOD control statements are used in the same statement set, they must appear in the following order, or else a command error results:

- DATETIME
- CYCLE
- PERIOD

Parameters

cycle name	is the first positional parameter in the CYCLE statement and must be specified The cycle name can be any of the values specified in Table 21 on page 235.

Table 21 Default values for the cycle name parameter

Cycle name	Default
BIWEEKLY	One report is generated for each two-week period defined within the cycle range.
CYCLE*nn*	One report is generated for each cycle range specified to the CYCLE*nn* value, where *nn* is a two-digit number between 01 and 98.
CYCLE99	One report is generated for the entire DATETIME period. The cycle range for the CYCLE99 value defaults to the DATETIME period.
DAILY	A separate cycle is generated for each day during the general DATETIME period, up to a maximum of 31 days.
HOLIDAYS	One report is generated for each cycle range specified to the HOLIDAYS value. If this information is desired, cycle ranges must be defined that specify the start and end dates and times covering the duration of any holiday. The value HOLIDAYS is printed on the report.
MONTHLY	A separate cycle is defined for each month during the general DATETIME period.
WEEKENDS	A separate cycle is defined for each weekend (Saturday and Sunday) during the general DATETIME period.
WEEKLY	A separate cycle is defined for each week (Monday through Sunday) during the general DATETIME period.
WORKWKLY	A separate cycle is defined for each work week (Monday through Friday) during the general DATETIME period.

NOTE

The BIWEEKLY, HOLIDAYS, and CYCLE*nn* values require that at least one cycle range be specified.

(*sdate*:[*stime*],*edate*:[*etime*])

The *s* and *e date* and *time* values define a specific cycle range to a *cycle name* value. Up to 31 cycle ranges can be defined to one CYCLE statement. The cycle range is defined by using the same format as that used in the DATETIME control statement (see “DATETIME” on page 238).

None of the ranges that are defined in a single CYCLE statement can overlap with one another. All ranges must be a subset of the range specified in the DATETIME control statement.

The cycle range is optional for all values except the BIWEEKLY, HOLIDAYS, and CYCLE*nn* values.

Examples

```
DATETIME    (03001:000000,03031:240000)
CYCLE       DAILY
PERIOD      CYCLE=DAILY
```

The reporting period is daily, within the DATETIME specified. Daily reports are generated for the 1st through the 31st days of 2003. Reports are not generated for any days past the 31st day in the range defined by the DATETIME control statement.

```
DATETIME     (02363:,03028:)
CYCLE        WEEKLY,(02364:,03006:),(03006:,03013:),
                    (03013:,03020:),(03020:,03027:)
PERIOD       CYCLE=WEEKLY
```

Data is collected weekly from day 364 of 2002 to day 6 of 2003; day 6 to day 13 of 2003; day 13 to day 20; and day 20 to day 27 (time default operates: 000000).

```
DATETIME    (03001:,03080:)
CYCLE       MONTHLY,(03001:,03032:)
PERIOD      CYCLE=MONTHLY
```

The collection period is the first month of 2003, using the time default of 000000.

NOTE

Collection Phase Log reports do not require cycle definition, but they *do* require DATETIME boundaries.

DASD

```
DASD
```

Overview

The DASD control statement requests the Direct Access Report that is discussed in “Direct Access Report” on page 425.

The DASD statement used in combination with the VOLSER control statement (see “VOLSER” on page 320 for more information), requests the Direct Access Plot of Volume Report that is discussed in “Direct Access Report Plot of Volume” on page 427. The VOLSER statement must immediately follow the DASD statement in the Analyzer JCL.

NOTE

Both reports require records written at CMF MONITOR Extractor initialization. To obtain the desired report, ensure that the input data includes records written when the extraction began. It is not necessary to include the period of these records in the DATETIME or SHIFT ranges.

The DASD statement uses data collected by the following Extractor control statements:

- DEVICE (see “DEVICE” on page 145)
- HEADMOVE (see “HEADMOVE” on page 159)

Parameters

This statement has no parameters.

Examples

```
DASD
```

This example produces the Direct Access Report.

```
DASD
VOLSER MVS001,SYSPK1,PAGE01
```

This example produces a Direct Access Report and a Direct Access Report Plot of Volume for DASD volumes MVS001, SYSPK1, and PAGE01.

DATETIME

```
DATETIME (sdate:[stime],edate:[etime])
```

Overview

The DATETIME control statement defines the overall date and time range for the data you want to use in your reports. A record is selected for a report only if the date and time of the start of its recording interval falls within the specified DATETIME range. To clarify, if you define DAYTIME as the following, any SMF record whose *start-of-interval* date and time falls between 8:00 A.M. and midnight of the specified days is used to produce records.

DATETIME (980322:080000,980333:240000)

Note that the data collection (or extraction) interval affects how reports are produced. For example, if REPORT CPM,SYNCH=59, INTERVAL=30 was specified when starting the CMF MONITOR Extractor, records will exist which have a *start-of-interval* time under a minute before 8:00 A.M. Thus, using the above DATETIME statement, these records will not be selected even though most of their data was gathered after 8:00 A.M.

A specific report time range can be smaller because there are other control statements, such as the CYCLE, SHIFT, and PERIOD statements that divide the overall DATETIME range into smaller ranges of time. When applied to a report control statement, the report is modified to include only data that falls between the defined period.

When using this statement, both a begin date-time and an end date-time must be specified.

Parameters

Parameters for the DATETIME control statement are as follows:

sdate

specifies the start date for the report data range

This parameter is required for the DATETIME statement. The colon (:) must be coded as the last character of the sdate parameter. This parameter is expressed in one of the following three formats:

Julian — *yydd*: or *yyddd*:, where

yy — The year can be represented as two digits. Two-digit years represent years in the range 1950-2049. Example: January 1, 2006 can be represented as 06001.

yyyy — The year can be represented as four digits. Example: January 1, 2006 can be represented as 2006001.

ddd — The number of days into the year can be represented as three digits and can be a value from 1 to 365 (366 during leap years).

Gregorian — *ddmmmyy*: or *ddmmmyyyy*: where

dd — represents the date of the month; must be a value between 1 and 31

(You are not required to insert a leading zero for values lower than 10.)

mmm — represents the month as three alphabetic characters

The acceptable values are as follows:

JAN	MAY	SEP
FEB	JUN	OCT
MAR	JUL	NOV
APR	AUG	DEC

yy — represents the year as two digits

Two-digit years represent years in the range 1950-2049. Example: January 1, 2006 can be represented as 01JAN06.

yyyy represents the year as four digits

Example: January 1, 2006 can be represented as 01JAN2006.

Relative *{-*nnn*}: where

* indicates today; if defined by itself without an -*nnn* value, the date specified is today

When an asterisk is used, the DATETIME card reads the current date from your system time.

[-*nnn*] (*optional*) indicates the number of days prior to today's date where the report data should start

For example, *-7 means the DATETIME range is today minus 7 days, or the date exactly one week ago (7 days) through today.

stime specifies the start time for the DATETIME range start date

The start time is optionally defined based on a 24-hour clock. This parameter is expressed as *hhmmss*{.th} where

hh number of hours; acceptable values are 00 to 24

mm number of minutes; acceptable values are 00 to 59

ss number of seconds; acceptable values are 00 to 59

.th thousandths of seconds; acceptable values are 00 to 99

This value is optionally defined. A period (.) must separate this value from the ss value, if defined.

Note: If the time is not defined, the default value midnight is used.

edate specifies the end date for the DATETIME range

This parameter is required and must be expressed in the same format as the sdate parameter.

etime specifies the end time for the DATETIME range end date

The end time is optionally defined. The same format used to define the optional stime parameter is used to define this parameter.

Syntax rules for the DATETIME statement

Syntax rules for the DATETIME statement are as follows:

- A colon must appear at the end of every date, regardless of what format is used.
- A comma must separate the start and end dates.
- The relative format is useful for running reports on a consistent basis because the dates defining the DATETIME range need not change each time the reports run.

  For example, if a site requires weekly reports to be produced each Monday, *-7 as the start range date and * as the end range date could be defined once; otherwise, the Julian day or Gregorian day or month values must be changed each week to obtain the current data.
- If the CYCLE or PERIOD statement is used, a DATETIME control statement must be used (see Table 21 on page 235 and "PERIOD" on page 281 for additional information). The order of these statements must be

  1. DATETIME
  2. CYCLE
  3. PERIOD

NOTE

If these commands are entered out of this sequence, a command error results.

Examples

```
DATETIME  (03294:104000,03294:110000)
```

This example specifies a record whose recording interval occurred on October 20th, 2003, any time from 10:40 A.M. until 11:00 A.M.

```
DATETIME  (03294:080000,03300:080000
```

This example specifies the date-time range from 0800 hours of day 294 (October 20th) of year 2003 to 0800 hours of day 300 (October 26th), 2003.

DEVACT

```
DEVACT
[TYPE={ALL|DASD|TAPE|COMM|GRAPH|UNITR|CHRDR}]
[,STORGRP={*|(aa,bb,dd:ff,ccc,....,gg:hh)16}]
[,RANGE=(aa:bb,cc,....,dd:ee)16]
[,ORDER={LCU|DEVICE|STORGRP}]
```

Overview

The DEVACT control statement requests the Device Activity Report (see Figure 55 on page 422). You can define a specific class of device for this report, or the Device Activity Report is generated for each class of device encountered in the input data.

The DEVACT statement uses data collected by the DEVICE Extractor control statement (see “DEVICE” on page 145).

Parameters

The parameters for the DEVACT control statement are as follows:

TYPE= specifies the class of device for which the Device Activity report is to be produced

The default is ALL, unless the STORGRP parameter is present, in which case the default is DASD.

STORGRP= specifies up to 16 SMS storage group names, or ranges of names, to be reported

Storage group names must be one to eight alphanumeric characters. Ranges of names are specified as two names separated by a colon (:). The first name in a range must precede the second in the standard collating sequence.

Single names and multiple ranges of names can be specified, separated by commas.

STORGRP=* can be used to indicate that ALL storage groups are to be reported.

RANGE= specifies up to 16 devices or device ranges to be reported

Device addresses must be entered as either three- or four-digit hexadecimal values. Device ranges are specified as two addresses separated by a colon (:).

The first address in a range must be less than the second address. Single devices and multiple ranges can be specified, separated by commas.

ORDER= specifies the order in which devices are reported

If LCU is specified, devices are reported by Logical Control Unit number and device number. If DEVICE is specified, the order is strictly by device number. If STORGRP is specified, the order is storage group name and device number.

The default is LCU, unless the STORGRP parameter is present, in which case the default is STORGRP.

Usage notes

The TYPE, STORGRP, and RANGE parameters are filters; that is, they are used to select the devices which are to be reported. The filtering is done from the least restrictive option (TYPE=) to the most restrictive (RANGE=). For example, if TYPE=TAPE,RANGE=(100:300) is specified, only tape devices whose numbers are between 100 and 300 are reported.

However, when STORGRP= is specified, any TYPE parameter is ignored and TYPE=DASD is forced.

When the STORGRP parameter is used, only those devices pertaining to the named storage groups are included in the report. If no storage group names are found in the input data, the following message is issued:

```
CMF07830: NO DATA AVAILABLE FOR THIS REPORT
```

When the RANGE parameter is used, only those devices specified are included in the report. If none of the specified device addresses are found in the input data, message CMF07830 is issued.

When both STORGRP= and RANGE= are specified, a device must meet both criteria in order to be reported. If none do, message CMF07830 is issued.

Example

```
DEVACT TYPE=DASD
```

This example produces the Device Activity Report for all DASD devices.

```
DEVACT  STORGRP=*
```

This example produces the Device Activity Report for all DASD devices that are in SMS storage groups.

```
DEVACT  STORGRP=(Y:Z,BMC)
```

This example produces the Device Activity Report for all DASD devices belonging to storage groups whose names are within the range Y to Z, and for the storage group named BMC. This report is ordered by storage group names.

```
DEVACT  STORGRP=(ADAMANT:ZZTOP),ORDER=DEVICE
```

This example produces the Device Activity Report for all DASD devices belonging to storage groups whose names are within the range ADAMANT to ZZTOP. This report is ordered by device number, not by storage group names.

DOMINO

```
DOMINO
[TYPE={SERVER|USER|DATABASE|ALL}]
```

Overview

The DOMINO control statement requests the LOTUS DOMINO activity reports. These reports provide information about the activities of Lotus Domino servers, users, and databases in order to analyze problem servers and to view performance data.

SMF type 108 records are used to generate the LOTUS DOMINO server report. These records are created by LOTUS DOMINO servers. CMF MONITOR does not create these records.

To produce a valid report, the SMF type 108 records must remain in the order in which they were written (SYSID, DATE, and TIME). The following sort control statement will sort the records in their original order:

```
SORT FIELDS=(15,4,CH,A,11,4,BI,A,7,4,BI,A),EQUALS
```

Parameters

The parameters for the DOMINO control statement are as follows:

TYPE= specifies the type of report desired, as one of the following:

SERVER	generates the server activity report (see “LOTUS DOMINO Server Report” on page 480)
DATABASE	generates the database activity report (see “LOTUS DOMINO Database Activity Report” on page 485)
USER	generates the user activity report (see “Figure 78LOTUS DOMINO User Activity Report” on page 487)
ALL	generates all of these reports

Example

```
DOMINO
```

This example produces the LOTUS DOMINO Server Activity report.

```
DOMINO  TYPE=USER
```

This example produces the LOTUS DOMINO User Activity report.

DMSS

```
DMSS  {INIT|NOINIT}
[,RESERVE=nnn]
[,HIPERSP={*|nnn}]
```

Overview

The DMSS control statement is optional. It can be specified to

- force the use of a hiperspace
- force initialization of the DMSSMAIN data set
- manually set the amount of storage used by the DMSS component of the Analyzer

BMC Software recommends that you omit this statement if you have also omitted the DMSSMAIN DD statement.

Parameters

Parameters for the DMSS control statement are as follows:

INIT specifies that the DMSSMAIN work file should be initialized

The Analyzer automatically determines whether a given DMSSMAIN data set was already initialized; it initializes a data set only if it was not previously initialized.

INIT should be specified if, for example, the data set was damaged by an Analyzer run that ended abnormally. This parameter is ignored when hiperspace is used.

NOINIT can be specified if the data set is old and was previously initialized by CMF MONITOR

Using this parameter saves I/O activity. NOINIT is the default.

RESERVE specifies the amount of storage (in K) to be reserved within the address space for the Analyzer

(See "Setting values of region, DMSS reserve, and CTRLSIZE" on page 91.) If this parameter is omitted, the Analyzer computes a default value based on the amount of available storage.

HIPERSP= specifies that a hiperspace is to be used, even when a DMSSMAIN data set is present, and 4 K hiperspace blocks are to be allocated as the initial size of the hiperspace

The number of blocks must be a numeric value between 10 and 520000.

If an asterisk (*) is specified instead of a numeric value, it indicates that the Analyzer default value is to be used. This default is currently 200 blocks.

Note that if the Analyzer runs out of hiperspace blocks, it attempts to extend the hiperspace by the number of blocks specified or defaulted to by this parameter.

Example

```
DMSS INIT,RESERVE=400
```

This example invokes the initialization of DMSSMAIN, reserving 400 K storage for the Analyzer.

ENQUEUE

```
ENQUEUE
[THRESHLD={0|nnn}]
[,TYPE={SUMMARY|DETAIL|ALL}]
```

Overview

The ENQUEUE control statement requests the Enqueue Conflict Report shown in Figure 60 on page 434.

The ENQUEUE Analyzer statement uses data collected by the ENQUEUE Extractor control statement (see “ENQUEUE” on page 152).

Parameters

THRESHLD= specifies a conflict percentage that causes the major name-minor name pair to appear on the report

The THRESHLD value has two decimal places implied. For example, if THRESHLD=250 is specified, the conflict percentage would be 2.5%. The percentage represents that portion of the total conflict for which any resource was enqueued with conflict. The default is zero.

TYPE= specifies the type of enqueue report to be produced

- **SUMMARY** lists the major and minor names that were in contention without listing their requestors
- **DETAIL** lists all requestors of a major and minor name
- **ALL** produces a summary listing followed by a detailed listing; the default

Example

```
ENQUEUE THRESHLD=150,TYPE=SUMMARY
```

This example produces an Enqueue Conflict report that shows major and minor name enqueue conflicts with a contention factor over 1.5%.

ESS

```
ESS
```

Overview

The ESS control statement requests the ESS Statistics Report shown in Figure 61 on page 437.

The ESS Analyzer statement uses data that is collected by the CACHE Extractor control statement (see "CACHE" on page 129).

Parameters

This statement has no parameters.

Example

```
ESS
```

This example produces the ESS Statistics Report.

EXCEPTS

```
EXCEPTS
[INTERVAL=hh:mm:ss]
[,MIN=0|n]
[,MAX=0|n]
[,ASSOC=(measname1,measname2,measname3)]
[,CPU={0|1|2|3|4|5|6|7|8|9|A|B|C|D|E|F|10|11|...|1F|ALL}]
[,MEASURE=measname]
[,TRACE={YES|NO}]
[,TRCETYPE=(T1=trace-type,...,T4=trace-type)]
```

Overview

The EXCEPTS control statement requests the Exception Subreport that is discussed in "Exception Subreport" on page 440 and the Exception Trace Detail Report shown in Figure 63 on page 442. The EXCEPTS statement uses data collected by the TRACE76 Extractor control statement (see "TRACE76" on page 191).

Parameters

Parameters for the EXCEPTS control statement are as follows:

INTERVAL= specifies the time divisions on the report

The default is the size of the recording interval. Specifying an interval that is less than the recording interval has the same effect as using the INTERVAL default. If TRACE=YES is specified, the default interval is the trace set size.

The EXCEPTS interval is used to gather data in multiples of the Extractor record interval. Data cannot be prorated over time. A data collection key is generated for each Extractor record encountered that meets date-time qualification requirements. The equation for this record key is

Extractor record start time / EXCEPTS interval

If the start of the Extractor record cycle and the size of the EXCEPTS interval parameter are known, you can determine which EXCEPTS record (graph line) contains which Extractor record group. The use of the EXCEPTS interval is illustrated in examples 1 and 2, following.

MIN= specifies the lower limit that can trigger the exception

It must be a one- to eight-digit number greater than or equal to zero. The default is zero, and there is no implied precision.

MAX= specifies the upper limit that can trigger the exception

It can be a one- to eight-digit number. The default is zero, and there is no implied precision.

ASSOC= specifies up to three additional measures (regardless of type) to be printed if the MEASURE value is less than the MIN value or greater than the MAX value

In the EXCEPTS statement, three or fewer ASSOC measures can be requested together, independent of type.

CPU= specifies the CPU ID (as a hexadecimal value) for which CPU-related measures are to be reported

Acceptable IDs are 0 through 1F, and ALL. ALL is the default.

MEASURE= specifies the measure that can be monitored for each exceptional condition

Only one measure can be monitored for each EXCEPTS command. Valid measures are the same as those used in the GRAPH statement; they are listed in Appendix C, "Measure and trace values." If you want to define more than one measure, use the ASSOC parameter.

TRACE= produces the Exception Trace Detail report, shown in Figure 63 on page 442

This report is for TRACE76 measures only. Any other graphics measures included on the measure list cause an error. One TRACE76 measure name can be on the MEASURE parameter, and three associative TRACE76 measure names can be included on the ASSOC parameter. TRACE76 measures are shown in Appendix C, "Measure and trace values."

Note: Look at member CMFANLTR in the UBBPARM data set for examples of using the TRACE, MEASURE, and TRCETYPE parameters. TRACE uses the ASSOC, MEASURE, and TRCETYPE parameters to format the trace report.

TRCETYPE= specifies a *trace-type* value to be associated with the entries on the MEASURE and ASSOC parameters

There are four values for the TRCETYPE parameter:

T1 specifies the *trace-type* value for the measure name on the MEASURE parameter

T2 specifies the *trace-type* value for the first measure name on the ASSOC parameter

T3 specifies the *trace-type* value for the second measure name on the ASSOC parameter

T4 specifies the *trace-type* value for the third measure name on the ASSOC parameter

The *trace-type* values can have any of the following values:

MIN minimum value captured during the specified time interval calculation: SMF76MIN

MAX maximum value captured during the specified time interval calculation: SMF76MAX

AVG average value captured during the specified time interval calculation: SMF76AVG / SMF76SSS; the default

END last value captured during the specified time interval calculation: SMF76ENV

STD standard deviation value captured during the specified time interval calculation:

(SMF76SSS * ∑(SMF76STD) -
(∑ (SMF76AVE) **2)) **.5 / SMF76SSS.

DIF difference between the starting end value and the ending end value calculation: SMF76ENV(last) - SMF76ENV(first)

Note: The *trace-types* collected must be selected in the RMF data collector. The CMF MONITOR Extractor always collects all *trace-types*. If a *trace-type* is selected for data that was not collected by RMF, it is marked N/A in the report output and a warning message is produced.

Examples

Determine which EXCEPTS interval contains Extractor record data for CPU busy, as of 10:15 A.M., on March 28, 2003, if

- Extractor start time is 00:15 A.M. on March 28, 2003
- Extractor REPORT control statement interval parameter value is QTR
- EXCEPTS parameter interval value is 02:00:00

From this information you know the following information:

- Data collection begins at 00:15:00 (from Extractor start time).

- Twelve EXCEPTS intervals cover the 24-hour period encompassing the data of March 28, 2003. (This value is calculated by dividing 24 hours by the EXCEPTS interval value of 2; the result is 12.)

- A maximum of eight Extractor intervals can be summed into each EXCEPTS data collection interval. (This value is calculated by dividing the EXCEPTS collection interval of 2 by the Extractor QTR hour (1/4) hour interval; the result is 8.)

Table 22 shows the EXCEPTS intervals and EXCEPTS interval time ranges for this situation. EXCEPTS interval 06 contains the Extractor record data for CPU busy at 10:15 A.M.

Table 22 EXCEPTS interval time range for twelve intervals

Interval number	EXCEPTS interval time range	Description
01	00:00:00 - 02:00:00	
02	02:00:00 - 04:00:00	
03	04:00:00 - 06:00:00	
04	06:00:00 - 08:00:00	
05	08:00:00 - 10:00:00	
06	10:00:00 - 12:00:00	
07	12:00:00 - 14:00:00	10:15:00 EXTR RECORD
08	14:00:00 - 16:00:00	
09	16:00:00 - 18:00:00	
10	18:00:00 - 20:00:00	
11	20:00:00 - 22:00:00	
12	22:00:00 - 24:00:00	

When using the SHIFT control statement with an EXCEPTS control statement that contains the interval parameter, use this equation to generate the record key:

```
Record Key = ((Extractor record start date)
             - (January 1, 2003)
             + (Extractor record start time))
             / EXCEPTS interval
```

Example 1

```
SHIFT (010000,090000),(090000,170000)
EXCEPTS MEASURE=CPU,INTERVAL=02:00:00
```

Determine which EXCEPTS intervals have a CPU busy less than 8% or greater than 20% if the following scenario is true:

- Extractor start time is 00:15 on March 28, 2003.
- Extractor REPORT control statement INTERVAL parameter value is QTR.
- EXCEPTS INTERVAL value is 02:00:00 (2 hours).
- SHIFT request time ranges are 010000 to 090000 (1:00 A.M. to 9:00 A.M.) and 090000 to 170000.

From this information, you know that the following statements are true:

- Data collection begins at 01:00:00 (from the SHIFT control statement request).
- There are four EXCEPTS intervals within each SHIFT time range definition defined on the SHIFT control statement. (This number is calculated by dividing the 8 hour time range by the EXCEPTS interval value of 2; the result is 4.)
- A maximum of eight Extractor intervals can be summed into each EXCEPTS data collection interval. (This number is calculated by dividing the EXCEPTS collection interval of 2 by the Extractor QTR (1/4) hour interval; the result is 8.)

Table 23 shows the EXCEPTS intervals and EXCEPTS interval time range for this situation. The EXCEPTS interval 05 contains the Extractor record data for CPU busy at 10:15 A.M.

Table 23 EXCEPTs interval time range for eight intervals

Interval number	EXCEPTS interval time range	Description
01	01:00:00 - 03:00:00	
02	03:00:00 - 05:00:00	
03	05:00:00 - 07:00:00	
04	07:00:00 - 09:00:00	
05	09:00:00 - 11:00:00	
06	11:00:00 - 13:00:00	
07	13:00:00 - 15:00:00	10:15:00 EXTR RECORD
08	15:00:00 - 17:00:00	

Example 2

```
EXCEPTS INTERVAL=00:30:00,MIN=8,MAX=20,
        MEASURE=PAGESEC,ASSOC=(PAGEINS,PAGEOUTS)
```

In this example, the INTERVAL parameter value causes the time division on the Exception Report generated to be 30 minutes; minimum / maximum trigger range is less than 8 or greater than 20; the average paging rate per second is measured, as well as the number of pages brought into and out of central storage.

Example 3

```
EXCEPTS  MEASURE=RCEAFC,
         TRACE=YES,
         ASSOC=(RCEPOOL,
                RCECOMPI,
                RCECOMPO),
         TRCETYPE=(T1=AVG,T2=DIF,T3=MAX,T4=MIN)
```

where

T1=AVG	specifies the trace value for exception-measure
T2=DIF	specifies the trace value for associate-measure-1
T3=MAX	specifies the trace value for associate-measure-2
T4=MIN	specifies the trace value for associate-measure-3

NOTE

VSDB and VSDA labels used in the calculations in Table 96 on page 503 are DSECT fields in VSDATA, which define areas within major SMF78RCD triplet areas associated with CSA and SQA data.

FICONSW

```
FICONSW
```

Overview

The FICONSW control statement produces the FICON Director Activity Report, shown in Figure 65 on page 449, which provides configuration and activity information about all FICON Directors (switches) to which the system was connected.

The FICONSW Analyzer statement uses data collected by the FICONSW Extractor control statement (see “FICONSW” on page 158).

Parameters

This statement has no parameters.

Example

```
FICONSW
```

This example produces the FICON Director Activity Report.

GRAPH

```
GRAPH
[TYPE={PLOT|KIVIAT|PIE|TAB|DISTRIB|
   PROFILE|TRACE}]
[,TRCETYPE=(T1=trace-type,...,T16=trace-type)]
[,INTERVAL=hh:mm:ss]
[,PLOTFILL={YES|NO}]
[,CIRCLE={YES|NO}]
[,AXES={YES|NO}]
[,LPI={6|n}]
[,CPI={10|n}]
[,LINES={60|n}]
[,FILLCHAR='*'

[,MEASURE=(measname1,...,measname16)]
[,LIMIT=(n1,...,n16)]
[,CPU={0|1|2|3|4|5|6|7|8|9|A|B|C|D|E|F|10|11|...|1F|ALL}]
[,NORMALIZ={YES|NO}]
```

Overview

The GRAPH control statement requests a Kiviat, pie, plot, tabular, distribution, or profile, and the Graphics Trace Detail Report that is discussed in “Graphics Trace Detail Report” on page 451. The valid parameters used for each graph type specified by the TYPE parameter are listed in Table 24 on page 260 and explained following the table.

The GRAPH statement uses data collected by the TRACE76 Extractor control statement (see “TRACE76” on page 191). The values must be specified in the MEASURE parameter of the GRAPH statement for reporting.

Table 24 Valid parameters for graph types

Parameter	Plot	KIVAT	PIE	TAB	DISTRIB	PROFILE	TRACE
INTERVAL	X	X	X	X	X	X	X
PLOTFILL	X	X				X	
CIRCLE		X	X				
AXES		X					
LPI		X	X				
CPI		X	X				
LINES		X	X				
FILLCHAR	X	X	X		X	X	
MEASURE	X	X	X	X	X	X	X
LIMIT	X	X			X	X	X
CPU	X	X	X	X	X	X	
IPS							
NORMALIZE					X		

Parameters

Parameters for the GRAPH control statement are as follows:

TYPE= specifies the type of graph to be produced

The default for TYPE is PLOT.

PLOT produces a horizontal bar graph (see Interval Bar Graph that is discussed in “Interval Bar Graph” on page 468)

KIVIAT produces a polygon within a circle (see “Kiviat Graph” on page 470)

Note: Special JCL is needed to produce this type of graph. (See “Producing graphics reports on a JES2 system” on page 91.)

PIE produces a segment graph within a circle (see “Pie Chart” on page 498)

Note: Special JCL is needed to produce this type of graph. (See “Producing graphics reports on a JES2 system” on page 91.)

TAB produces a tabular listing of up to 12 measures by interval (see “Tabular Subreport” on page 527)

TYPE= (*continued*)	**DISTRIB**	produces a distribution for each measure requested (see "Distribution Graph" on page 431)
	PROFILE	produces a horizontal bar graph based on time of day only (see "Profile Bar Graph" on page 502)
	TRACE	produces the Graphics Trace Detail Report This report is for TRACE76 measures only. Any other graphics measures included on the measure list will cause an error. Twelve TRACE76 measure names can be entered on the measure list. TRACE uses the MEASURE and TRCETYPE parameters to format the trace report. If a trace value is not specified for a measure name, the Graphics Trace Detail Report is produced, displaying a column for each of the six trace-type values of TRCETYPE. A separate graph is generated for each measure as requested. View member CMFANLTR in the UBBPARM data set for examples of using the TRACE, MEASURE, and TRCETYPE parameters. **Note**: For TYPE=TRACE, the only valid parameters are INTERVAL, MEASURE, and LIMIT.

TRCETYPE=

specifies a *trace-type* value to be associated with the entries on the MEASURE parameter

The trace-type values and the measure names are positional.

T*n*	specifies the trace-type value for the *n*th measure name

Up to 12 trace-type values and measure names can be entered. The default value is AVG. The trace-type values can be any of the following values:

MIN	minimum value captured during record interval calculation: smallest SMF76MIN value
MAX	maximum value captured during record interval calculation: greatest SMF76MAX value
AVG	average value captured during record interval calculation: SMF76AVG / SMF76SSS
END	last value captured during record interval calculation: SMF76ENV

TRCETYPE= (*continued*)	**STD**	standard deviation value captured during record interval calculation: (SMF76SSS * ∑ (SMF76STD) - (∑ (SMF76AVE) **2)) **.5 / SMF76SSS
	DIF	difference between the starting end value and the ending end value calculation: SMF76ENV(last) - SMF76ENV(first)

TRACE76 measures can be included on the measure list with other graphics measures for other graphics displays (PLOT, PROFILE, KIVIAT, TAB, and so on.). The default value, AVG, is assigned to any measure name that does not have a TRCETYPE value associated with it for any graph type except TRACE.

INTERVAL= specifies the time divisions on the graph

INTERVAL= can be specified with the following values:

hh:mm:ss

where

hh	Is the hour.
mm	Is the minutes.
ss	Is the seconds.
HOURLY	Each line of the graph represents one hour.
DAILY	Each line of the graph represents one 24-hour period.
WEEKLY	Each line of the graph represents one 168-hour period.
MONTHLY	Each line of the graph represents one month.
QTRLY	Each line of the graph represents a calendar year quarter.
SEMIANNL	Each line of the graph represents a half year.

Note: If INTERVAL is specified when TYPE=KIVIAT or PIE, one graph is produced for each occurrence of the interval.

The defaults are

TRACE	size of the trace set
PLOT, TAB, PROFILE	size of the Extractor record interval
KIVIAT, DISTRIBUTION, PIE	entire Extractor measurement interval

INTERVAL= (*continued*)

Specifying an interval that is less than the Extractor recording interval has the same effect as specifying the Extractor recording interval.

The GRAPH interval is used to gather data in multiples of the Extractor record interval. Data cannot be prorated over time. A data collection key is generated for each Extractor record encountered that meets date-time qualification requirements. The equation for this record key is

Record key = Extractor record start time / GRAPH interval

If the start of the Extractor record cycle and the size of the GRAPH interval parameter are known, you can determine which GRAPH record contains which Extractor record group.

PLOTFILL=

specifies the format of the plots in which two or more measures are shown

When you specify YES, a continuous string of characters leads from the plotted value on one line to the plotted value on the next line, making it easier to see the changes in values. When you specify NO, the point plotting format is produced. Whenever two or more characters occupy the same space, the character O is produced. The default is YES.

CIRCLE=

specifies whether a circle of asterisks should be printed to mark the outer boundary of the Kiviat graph or Pie chart

The default is YES.

AXES=

specifies whether periods should be printed to mark the location of the axis for each measure in the Kiviat graph

The default is YES.

CPI=

specifies the height and width of print characters

Set CPI to characters per inch. The default is CPI=10. If CPI is set incorrectly, the resulting Kiviat graph or Pie chart is elliptical.

LPI=

specifies the height and width of print characters

Set LPI to lines per inch. The default is LPI=6. If LPI is set incorrectly, the resulting Kiviat graph or Pie chart is elliptical.

LINES=

specifies the number of lines per page to be used for Kiviat or pie graphs

You cannot specify a value greater than 60. Kiviat and pie graphs are always printed one to a page, but unusual page sizes can be accommodated with this parameter. The default is 60. All other graph types print without page breaks until the graph is complete.

FILLCHAR= for Kiviat graphs, specifies the character to use to fill the interior of the polygon; the default is *

For plot graphs, specify one FILLCHAR for each measure requested; the default is *+=:-<.#12345678.

MEASURE= specifies the measures to be graphed

See Appendix C, “Measure and trace values,” for acceptable values. You can specify up to 16 measures for the Kiviat, Plot, Profile, and Distribution graphs and Pie chart. You can specify up to 12 measures for the Tabular Subreport (see “Tabular Subreport” on page 527 for more information) or Trace graph; if you specify more, only the first 12 are recognized.

LIMIT= specifies a limit value to be associated with each measure

There is a one-to-one correspondence between the MEASURE and LIMIT parameters for the Kiviat and distribution graphs. For the plot and profile graphs, the highest limit specified is used.

To establish useful graphs of measures that have a range of values exceeding 100, run the tab graph first, to determine a proper set of limit specifications. Two decimal places are implied for each LIMIT= value. The default for each value is 10000 (100.00). A value less than 1000 (10.00) causes results to be unpredictable.

The limit value should be high enough to accommodate the measure being graphed. For example, if the PAGEIO measure is being graphed and it is estimated that the value is in the thousands, specify a limit value of 1000000 (10,000.00). Erroneous graphs can result if the limit value is too small.

CPU= specifies the CPU ID (as a hexadecimal value) for which CPU data is to be displayed

Acceptable IDs are 0 through 1F, and ALL. ALL is the default.

NORMALIZ= specifies if the distribution graph is to be normalized

If normalized, the X-axis starts at the minimum value instead of zero (does not apply to percent), and the Y-axis ends at 50% (if no distribution exceeds 50%). The default is YES.

Examples

Determine which GRAPH interval contains Extractor record data for CPU busy, as of 10:15 A.M., on March 28, 2003:

- Extractor start time is 00:15 A.M. on March 28, 2003.
- INTERVAL value on the Extractor REPORT control statement is QTR.
- GRAPH INTERVAL value is 02:00:00.

From this information, you know that the following statements are true:

- Data collection begins at 00:15:00 (from Extractor start time).
- There are 12 GRAPH intervals that cover the 24-hour period encompassing the data of March 28, 2003. (This figure is calculated by dividing 24 hours by the GRAPH interval value of 2; the result is 12.)
- A maximum of eight Extractor intervals can be summed into each GRAPH data collection interval. (This figure is calculated by dividing the GRAPH collection interval of 2 by the Extractor QTR (1/4) hour interval; the result is 8.)

Table 25 shows the GRAPH intervals and interval time range for this situation. Interval Number 06 contains the Extractor record data for CPU busy at 10:15 A.M.

Table 25 GRAPH interval time range for twelve intervals

Interval number	GRAPH interval time range	Description
01	00:00:00 - 02:00:00	
02	02:00:00 - 04:00:00	
03	04:00:00 - 06:00:00	
04	06:00:00 - 08:00:00	
05	08:00:00 - 10:00:00	
06	10:00:00 - 12:00:00	10:15:00 EXTR RECORD
07	12:00:00 - 14:00:00	
08	14:00:00 - 16:00:00	
09	16:00:00 - 18:00:00	
10	18:00:00 - 20:00:00	
11	20:00:00 - 22:00:00	
12	22:00:00 - 24:00:00	

When using the SHIFT control statement with a GRAPH control statement containing the interval parameter, use the following equation to generate the record key:

```
Record Key = ((Extractor record start date)
             - (January 1, 2003)
             + (Extractor record start time))
             / GRAPH interval
```

For example

```
SHIFT (010000,090000),(090000,170000)
GRAPH MEASURE=CPU,INTERVAL=02:00:00
```

Determine which GRAPH interval contains Extractor record data for CPU busy, as of 10:15 A.M., on March 28, 2003, if the following statements are true:

- Extractor start time is 00:15 on March 28, 2003.
- INTERVAL value on the Extractor REPORT control statement is QTR.
- GRAPH parameter interval value is 02:00:00.
- SHIFT request time ranges are 010000 to 090000 and 090000 to 170000.

From this information you know that these statements are true:

- Data collection begins at 01:00:00 (from the SHIFT control statement request).
- There are four GRAPH intervals within each SHIFT time range definition defined on the SHIFT control statement. (This figure is calculated by dividing the 8 hour time range by the GRAPH interval value of 2; the result is 4.)
- A maximum of eight Extractor intervals can be summed into each GRAPH data collection interval. (This figure is calculated by dividing the GRAPH collection interval of 2 by the Extractor QTR hour (1/4) hour interval; the result is 8.)

Table 26 shows the GRAPH intervals and GRAPH interval time range for this situation. The GRAPH interval 05 contains the Extractor record data for CPU busy, at 10:15 A.M.

Table 26 GRAPH interval time range for eight intervals

Interval number	GRAPH interval time range	Description
01	01:00:00 - 03:00:00 03:00:00 - 05:00:00	
02	05:00:00 - 07:00:00 07:00:00 - 09:00:00	
03	09:00:00 - 11:00:00 11:00:00 - 13:00:00	
04	13:00:00 - 15:00:00 15:00:00 - 17:00:00	
05		
06		10:15:00 EXTR RECORD
07		
08		

```
GRAPH TYPE=PLOT,INTERVAL=01:00:00,MEASURE=(SUP,SYST)
```

This example instructs the Analyzer to produce a horizontal bar graph, using time divisions of one hour, displaying the percentage of CPU busy time during which the system was in supervisor and system test state.

```
GRAPH TYPE=KIVIAT,CPU=ALL,MEASURE=(DEV-13F,CHN-1,ERR-13F),
      LIMIT=(8000,,8000)
```

This example generates a Kiviat graph of all CPUs, displaying busy times for device 13F and channel 1, and the percentage of time device address 13F is in error recovery; measure limit value is 8000 for the first and third measures.

```
GRAPH TYPE=PIE,INTERVAL=04:00:00,MEASURE=(CPK-02,CPK-00,
      CPU,SUP,SYST,PPB)
```

This example produces a pie chart, using time divisions of 4 hours, to display CPU busy time for problem programs designated by protect key 02 and 00, total CPU capacity used, and percentage of time CPU was in supervisor, system task, and problem program states.

HEADERS

```
HEADERS
[TITLE='title']
[,LOCATION='location']
[,LINES={60|nn}]
```

Overview

The HEADERS control statement defines optional headline information for the CMF MONITOR reports. In addition, it can be used to adjust the maximum number of lines per report page.

Parameters

Parameters for the HEADERS control statement are as follows:

TITLE=	specifies a quoted-character string, which is centered on the second line of each CMF MONITOR report Up to 52 characters (excluding outside quotation marks) can be used. The enclosing quotation marks can be omitted if the string contains only alphabetic or numeric characters (no spaces). The default value is blanks.
LOCATION=	specifies a quoted-character string, which is centered on the third line of the CMF MONITOR reports Up to 40 characters (excluding outside quotation marks) can be used. The enclosing quotation marks can be omitted if the string contains only alphabetic or numeric characters (no spaces). The default value is blanks.
LINES=	specifies the maximum number of lines per page on the CMF MONITOR reports The maximum value is 99. The default value is 60. **Note**: Graphic reports ignore the Lines parameter. See “GRAPH” on page 259 for information about how graphic reports are printed.

Example

```
HEADERS TITLE='XYZ COMPANY, INC.',LOCATION='HOMETOWN, USA',LINES=55
```

CMF MONITOR reports are headed by XYZ COMPANY, INC. centered on the second line and HOMETOWN, USA centered on the third line; 55 lines are allotted for each page of this report.

HFS

```
HFS
```

Overview

The HFS control statement requests the HFS Statistics Report, which provides information about caching by HFS buffer pools and mounted file systems.

The HFS statement uses data collected by the HFS Extractor control statement.

NOTE

This report is not meaningful when data from multiple systems is combined into one report. Therefore, the control statement must not follow a SYSPLEX control statement specifying multiple systems and RPTS=COMBINED.

Parameters

This statement has no parameters.

Example

```
HFS
```

The HFS Statistics Report is produced.

HTTP

```
HTTP
```

Overview

The HTTP control statement requests the HTTP server report (see “HTTP Server Report” on page 458). This report provides information about the activities of HTTP servers. The information may be used to analyze problem servers and to view performance data.

SMF type 103 records are used to generate the HTTP server report. CMF MONITOR does not gather these records. HTTP servers create these records.

To produce a valid report, the SMF type 103 records must remain in the order in which they were written (SYSID, DATE, and TIME). The following sort control statement will sort the records in their original order:

```
SORT FIELDS=(15,4,CH,A,11,4,BI,A,7,4,BI,A),EQUALS
```

NOTE

This report differs from all other Analyzer reports due to the nature of the SMF type 103 records. A number of fields are not displayed in the report header, including the system name (only system ID is shown), the operating system version and release information, the duration of the report (both ACTUAL and RQSTD).

The following general control statements are not supported for this report and if present, are not applied to this report:

CYCLE, PERIOD, SHIFT, SYSPLEX

Parameters

This statement has no parameters.

Example

```
HTTP
```

The HTTP Statistics Report is produced.

IOQ

```
IOQ
[,IGNORE={NO|INACTIVE}]
```

Overview

The IOQ control statement requests the I/O Queuing Activity Report that is discussed in "I/O Queuing Activity Report" on page 464. Optionally, Logical Control Units (LCUs) that had no activity can be ignored; this setting might yield a more concise report.

Parameters

Parameters for the IOQ control statement are as follows:

IGNORE= specifies how LCUs that have no activity during the reporting period are to be treated

- **INACTIVE** specifies that the report not show LCUs that had no activity for the reporting period
- **NO** specifies that the report contain all LCUs, whether or not they had any activity

  This option is the default.

Example

```
IOQ    IGNORE=INACTIVE
```

The I/O Queuing Activity report is produced only for those LCUs that had some activity for the reporting period.

LINKPACK

```
LINKPACK
[THRESHLD={0|nnnnn}]
```

Overview

The LINKPACK control statement requests the Link Pack Area Report shown in Figure 73 on page 474.

NOTE

This report requires records that are written at CMF MONITOR Extractor initialization. To obtain the desired report, make certain that the input data includes the records that were written when the extraction began. It is not necessary to include the period of these records in the DATETIME or SHIFT ranges.

The LINKPACK statement uses data collected by the LINKMAP Extractor control statement (see “LINKMAP” on page 169) and the REPORT control statement must have GBLS=YES defined (see “REPORT” on page 176).

Parameters

The parameter for the LINKPACK control statement is as follows:

THRESHLD= specifies a one- to five-digit value used to determine an activity threshold value for any one link pack module

If the number of samples determines that a specific module had activity which exceeded the threshold value, the module is marked on the Link Pack Area Report with an action flag of ****. All active link pack modules are reported. The default is zero.

Example

```
LINKPACK    THRESHLD=6000
```

This example produces Link Pack Area report, with a threshold set to flag modules with over 6000 busy samples.

LPARCOMB

```
LPARCOMB
{[ALL=NO]
[, DESC='description',LPARS=( name1,...,name16)]}
```

Overview

The LPARCOMB control statement specifies a combination of LPARs to be included in the LPAR Combination Section of the “CPU Utilization Report” on page 393. One or more LPARCOMB control statements can be specified and they all must immediately follow a CPU control statement.

Parameters

The parameters for the LPARCOMB control statement are as follows:

ALL=	displays by default the combination with the description ALL LPARS (consisting of all LPARs) To suppress this combination, specify ALL=NO.
DESC=	specifies the description of the combination of LPARs This parameter can contain up to 32 characters. If it contains spaces, enclose the parameter in quotation marks; otherwise, quotation marks are not required. When this parameter is specified, the LPARS parameter is required.
LPARS=	specifies the names of up to 16 LPARs in the combination When this parameter is specified, the DESC parameter is required. If the reserved LPAR name PHYSICAL is specified, it is ignored.

Example

```
CPU
```

This example produces a CPU Utilization Report with an LPAR Combination Section containing only the default 'ALL LPARS' combination.

```
CPU
LPARCOMB  DESC='DB2 systems',LPARS=(DB2A,DB2B)
```

This example produces a CPU Utilization Report with an LPAR Combination Section containing the default 'ALL LPARS' combination consisting of all LPARs and the 'DB2 systems' combination consisting of LPARs DB2A and DB2B.

```
CPU
LPARCOMB  DESC='DB2 systems',LPARS=(DB2A,DB2B)
LPARCOMB ALL=NO
```

This example produces a CPU Utilization Report with an LPAR Combination Section containing only the 'DB2 systems' combination. The default 'ALL LPARS' combination is suppressed.

```
CPU
LPARCOMB ALL=NO
```

This example produces a CPU Utilization Report without the LPAR Combination Section.

OMVS

```
OMVS
```

Overview

The OMVS control statement requests the OMVS Kernel Activity Report that is discussed on page 488.

The OMVS Analyzer statement uses data collected by the OMVS Extractor control statement (see “OMVS” on page 171).

Parameters

This statement has no parameters.

Example

```
OMVS
```

The OMVS Kernel Activity Report is produced.

PERFORM

```
PERFORM
[,SCTYPE={CTL|RPT|BOTH}]
```

Overview

The PERFORM control statement is used with the PERFSUM control statement to modify report contents. PERFORM can also be used with the WLMGL control statement. This statement does not produce a report if used by itself.

The PERFORM control statement allows you to include or exclude service classes in the Workload Manager Goal Mode Report.

If used, the PERFORM statement must immediately follow the PERFSUM or WLMGL statement that it is modifying. Only one PERFORM statement can be defined after each PERFSUM or WLMGL statement; if multiple PERFORM statements are defined after a PERFSUM or WLMGL statement, only the last PERFORM statement affects the report.

NOTE

If PERFORM is used in conjunction with WLMGL TYPE=SUMMARY, the PERFORM card is ignored. (See the summary description of “Workload Manager Goal Mode Report” on page 549 for more information).

Parameters

Parameter meanings are as follows:

SCTYPE=

requests that service classes or report classes or both be included in the report

Acceptable values are as follows:

CTL — requests that only service classes be included in the report

This value is the default when defined with the PERFSUM statement.

RPT — requests that only report classes be included in the report.

SCTYPE= (*continued*)	**BOTH**	requests that both service and report classes be included in the report This value is the default when defined with the WLGML statement.

Example

```
PERFORM
```

When this example (without the SCTYPE parameter) follows a PERFSUM statement, only service classes are included in the PERFSUM report. When this example follows a WLGML statement, both service and report classes are included in the WLGML report.

```
PERFORM    SCTYPE=CTL
```

This example requests that only service classes be included in the PERFSUM or WLGML report.

```
PERFORM    SCTYPE=RPT
```

This example requests that only report classes be included in the PERFSUM or WLGML report.

```
PERFORM    SCTYPE=BOTH
```

This example requests that both service and report classes be included in the PERFSUM or WLGML report.

PERFSUM

```
PERFSUM
[BOTNECKS=([ENQUEUE={10|nn}][,CHANNEL={50|nn}]
    [,OVERUTIL={95|nn}][,UNDERUTL={20|nn}]
    [,PAGING={50|nn}][,TSO={15|nn}]
    [,SUPERVSR={50|nn}][,AUXSTORE={90|nn})]
```

Overview

The PERFSUM control statement requests the Performance Summary Report that is discussed in "Performance Summary Report" on page 491.

The PERFSUM statement uses data collected by the following Extractor control statements:

- ASMDATA (see "ASMDATA" on page 127)
- CPU (see "CPU" on page 136)
- CHANNEL (see "CHANNEL" on page 134)
- DEVICE (see "DEVICE" on page 145)
- ENQUEUE (see "ENQUEUE" on page 152)
- EXTSUM (see "EXTSUM" on page 154)
- PAGING (see "PAGING" on page 173)
- TSODATA (see "TSODATA" on page 194)
- WORKLOAD (see "WORKLOAD" on page 203)

Parameters

The value that constitutes a bottleneck can be overridden with the following parameters:

BOTNECKS= specifies the major parameter under which all of the following minor parameters fall

This parameter must be coded. Parentheses must enclose the minor parameters if more than one is coded; for example:

BOTNECKS=(CHANNEL=60,TSO=2)

ENQUEUE= specifies the percentage of total time that an excessive ENQUEUE conflict existed; the default is 10%

CHANNEL= specifies the percentage of overused channel path busy time; the default is 50%

OVERUTIL= specifies the percentage of overused CPU busy time; the default is 95%

UNDERUTL= specifies the percentage of underused CPU busy time; the default is 20%

PAGING= specifies a value for average pages per second that would be considered too high; the default is 50 pages per second

TSO= specifies an average TSO response time that would be considered too high; the default is 15 seconds

SUPERVSR= specifies a percentage of CPU busy time spent excessively in supervisor state; the default is 50%

AUXSTORE= specifies the percentage of overused auxiliary storage slots; the default is 90%

NOTE

- If the PERFORM control statement is not specified after the PERFSUM control statement, the Analyzer looks for PERFORM data written at Extractor initialization.
- If the PERFORM parameter was specified on the Extractor EXTSUM control statement (see “EXTSUM” on page 154 for more information), the data is only available to the Analyzer if the records, written at Extractor initialization, are included in the input data. It is not necessary to include the period of these records in the DATETIME or SHIFT ranges.

Example

```
PERFSUM   BOTNECKS=(OVERUTIL=98,CHANNEL=35,PAGING=20,TSO=6)
```

This example produces a Performance Summary Report designating bottleneck tolerance thresholds: the percentage of time the CPU is over 98% busy; the percentage of channel busy time over 35%; average pages per second over 20; and TSO response time slower than 6 seconds.

PERIOD

```
PERIOD CYCLE=cyclename
[,RPTS={COMBINED|SEPARATE}]
```

Overview

The PERIOD control statement can define a variable reporting period either to all reports or just one report specified in a single batch job. It also allows you to define how the data within the variable reporting period is presented, either as a single combined report or as a series of separate reports.

To use the PERIOD statement, a CYCLE statement must be defined first. A CYCLE statement (see "CYCLE" on page 234) is used to define the date-time range and assign a unique cycle name identifier to the variable reporting period. The CYCLE range can be equal to or a subset of the DATETIME range (see "DATETIME" on page 238 for more information). Then the PERIOD statement is used to apply the CYCLE range to a single or all reports in a control statement set.

The PERIOD statement can appear in the general control statement set, in the report control statement set, or in both sets. Within the general control statement set, PERIOD specifies a reporting period cycle name for all reports. Within the report control statement set, PERIOD specifies a reporting period cycle name for an individual report. Any time a PERIOD statement is defined directly to a specific report, it overrides any PERIOD statement defined in the general control statement set, but only for the one report.

When using the PERIOD statement, both the DATETIME and CYCLE control statements must be defined. The order of these statements must be

- DATETIME
- CYCLE
- PERIOD

If you do not enter the commands in this sequence, a command error results.

Parameters

Parameters for the PERIOD control statement are as follows:

CYCLE=
specifies the unique cycle name identifier defined in a preceding CYCLE control statement

This parameter is required and must be a recognized CMF MONITOR cycle name identifier from the following list:

DAILY
WEEKLY
WORKWKLY
WEEKENDS
BIWEEKLY
MONTHLY
HOLIDAYS
CYCLE*nn*, where ***nn*** is a two-digit number from 01 to 99

RPTS=
allows you to request individual reports or one report for the entire cycle; the default is COMBINED

COMBINED
one report is produced by combining all data gathered within the specified CYCLE range

SEPARATE
multiple reports are produced by separating all data gathered within the specified CYCLE range

Depending upon how your CYCLE range is defined, it can specify multiple subsets of a reporting period. When a CYCLE range consists of multiple subsets and RPTS=SEPARATE is specified, multiple reports are produced, each containing data from a separate subset of the reporting period.

For example, if the DATETIME range defines 1 ½ months' worth of data (and there is input data available for this period), a CYCLE range defined as DAILY produces one report for each whole or fractional day that falls within the DATETIME range.

Examples

```
PERIOD  CYCLE=WEEKLY,RPTS=SEPARATE
```

When defined in the general control statement set, this example produces separate reports for each weekly period.

```
IOQ
PERIOD  CYCLE=DAILY
```

When defined directly after a report control statement in the report control statement set, this example produces an I/O Queuing Activity report for each day’s data. This PERIOD statement has no affect on other report control statements defined in the report control statement set.

PROTKEY

```
PROTKEY   [CPU={0|1|2|3|4|5|6|7|8|9|A|B|C|D|E|F|10|11|...|1F|ALL}]
```

Overview

The PROTKEY control statement requests the CPU Utilization by Protect Key Report that is discussed in "CPU Utilization by Protect Key Report" on page 411.

The PROTKEY statement uses data collected by the CPU Extractor control statement (see "CPU" on page 136).

Parameters

The parameter for the PROTKEY control statement is as follows:

CPU= specifies the CPU ID (as a hexadecimal value) for which protect key information is to be produced

Acceptable IDs are 0 through 1F, or ALL. The default is ALL.

If more than one CPU is specified, Protect Key reports appear in the order specified in this parameter.

Example

```
PROTKEY   CPU=1
```

This example produces a CPU Utilization Report by Protect Key for CPU 1 only.

PRSM

```
PRSM
[MARKER=>]
[,DETAIL=(NONE,FULL,PHYSPROC,WAITCOMP,
    DISPINTV,STATUS,WEIGHT,LOGPROC)]
```

Overview

The PRSM Analyzer control statement requests the Logical Partition Report that is discussed in “Logical Partition Report” on page 476.

The Logical Partition Report shows the activity of logical partition (LPAR) management time incurred on a Processor Resource / Systems Manager (PR / SM).

The PRSM statement uses data collected by the CPU Extractor control statement (see “CPU” on page 136).

Parameters

Parameters for the PRSM Analyzer control statement are as follows:

MARKER=

placed to the right of fields to indicate where changes occurred to field values during the display interval

The default marker is > (greater-than sign); however, you can specify any other character as the marker.

DETAIL=

specifies the circumstances under which CMF MONITOR generates a new row of data for the partition or a new iteration of the report

If you select the default NONE, it must be the only value specified. All of the other values are field names and you may specify them in any combination.

If a change during the interval triggers more than one row of data for a partition, a subtotal row is generated to summarize the measurements.

Note: The characters in parentheses are abbreviations for the operands. They can be specified instead of the full operand name.

DETAIL= (*continued*)	**NONE (N)**	no partition splitting; the default
	FULL (F)	if you request FULL detail, generates a new report or a new row (depending on the change that occurs) when a change is detected in any one of the six fields described below
	PHYSPROC (P)	partition split on physical processor If you request DETAIL by physical processor, the report is started again on a new page each time there is a change in the number of physical processors during the display interval.
	WAITCOMP (WA)	partition split on wait completion If you request DETAIL by the wait completion flag, a new row is generated for that partition for each value of the wait completion.
	DISPINTV (D)	partition split on dispatch interval If you request DETAIL by dispatch interval, the report is started again on a new page each time there is a change during the display interval.
	STATUS (S)	partition split on status flag If you request DETAIL by the status flag, a new row is generated for that partition for each different status flag type encountered during the display interval.
	WEIGHT (WE)	partition split on weighting factor If you request DETAIL by the weighting factor, a new row is generated for that partition for each unique weighting factor encountered during the display interval.
	LOGPROC (L)	**p**artition split on logical processors If you request DETAIL by the number of logical processors, a new row is generated for that partition for each logical processor detected during the display interval.

Examples

These examples show what reports are generated given the parameters you specify in the control statement.

```
PRSM
```

This example produces the Logical Partition Report. No separation is performed for changed values.

```
PRSM  MARKER=*,DETAIL=(P,D)
```

This example produces a report that generates a new report for each combination of changes (identified with an *) in either the physical processor or the dispatch interval.

RECTYPE

```
RECTYPE  [xxx|240]
[,MONITOR={CPM|IPM|RMF}]
[,STOPAFT={EOF|DATETIME}]
```

Overview

The RECTYPE control statement specifies the characteristics of the Extractor data set to be read.

Parameters

Parameters for the RECTYPE control statement are as follows:

xxx specifies a number from 128 through 255

This number should be the same as the SMFRECID parameter value specified for the Extractor on the REPORT control statement. The default is 240.

MONITOR= specifies the monitoring mode used to produce data for the Analyzer run.

Only data from the specified monitoring mode is used. (Since CPM and IPM can run concurrently in the Extractor, statistics might overlap and produce erroneous results if the data were analyzed together.) The default is CPM.

CPM indicates that the Analyzer will process data that was produced by the Extractor Continuous Performance Monitoring mode

IPM indicates that the Analyzer will process data that was produced by the Extractor Intermittent Performance Monitoring mode

RMF indicates that the Analyzer will process data that was produced by the RMF monitor only

When you use this option, some reports are incomplete because RMF does not provide the additional data that CMF MONITOR produces in its user records.

STOPAFT= specifies when the Analyzer stops reading input data

The possible values are as follows:

- **EOF** the Analyzer stops reading input at the end of the file; the default
- **DATETIME** the Analyzer stops reading input at the end of the time range specified in the DATETIME statement

  If the input is out of time sequence, this parameter (STOPAFT=DATETIME) causes a portion of the data in the DATETIME range to be ignored.

  A DATETIME control statement must be defined to use STOPAFT=DATETIME parameter; otherwise, STOPAFT=EOF is assumed.

Example

```
RECTYPE    128,MONITOR=IPM,STOPAFT=DATETIME
```

This example specifies that data to be processed by the Analyzer is CMF MONITOR data extracted in IPM mode and CMF MONITOR user records with an SMF ID of 128. The Analyzer will stop reading records for the reports at the end of the DATETIME range.

REPORTS

```
REPORTS RPTGROUP=(a[,b[,c[,d]]])
[,DDGROUP=e [,SYSOUT=class]]
```

Overview

The REPORTS control statement specifies how reports are to be grouped, whether separate DDs are to be used, and what SYSOUT Class to use. For example, all reports for a system can be grouped together, or all reports requested by a report control statement (such as Performance Summary) can be grouped together. Reports can also be sent to separate DD statements for ease of distribution.

There can be only one REPORTS statement, and it must appear in the general control statement set.

Parameters

Parameters of the REPORTS control statement are as follows:

RPTGROUP= specify up to four levels of grouping, in any order, without repetitions

Possible selections are as follows:

SYSTEM	produces reports grouped by System The ID used for this determination are specified on the TYPE parameter of the SYSPLEX control card.
RPTSTMT	produces reports grouped by Report control statement
CYCLE	produces reports grouped by Cycle
SHIFT	produces reports grouped by Shift

DDGROUP= specifies the level at which reports are written to separate DD statements

You must specify the DDGROUP operand as one of the RPTGROUP operands.

When a report is produced that is different at the DDGROUP level from the previous report, a new DD is created for the new report.

If DDGROUP is not specified, reports are written to SYSPRINT.

SYSOUT= specifies the SYSOUT Class to which reports are written if DDGROUP is specified

The default class is the same as MSGCLASS.

Examples

The following examples illustrate the use of the REPORTS control statement:

```
REPORTS RPTGROUP=(SYSTEM,RPTNAME,CYCLE,SHIFT)
```

This example shows the default. It groups reports by System, and then REPORT control statement, and then Cycle, and finally Shift. All reports are written to the SYSPRINT DD.

```
SYSPLEX TYPE=SYSNAME,RPTS=SEPARATE
REPORTS RPTGROUP=(SYSTEM),DDGROUP=SYSTEM
```

This example groups reports by SYSTEM and generates DDNames based on SYSNAME.

```
SYSPLEX RPTS=COMBINED
REPORTS RPTGROUP=(SYSTEM),DDGROUP=SYSTEM
```

This example groups reports together (because of RPTS=COMBINED) and generates a DDName of CMFR0001.

```
REPORTS RPTGROUP=(RPTSTMT),DDGROUP=RPTSTMT,SYSOUT=F
```

This example groups reports by REPORT control statement and writes each report to a separate DD assigned to SYSOUT Class F.

```
REPORTS RPTGROUP=(SYSTEM,SHIFT),DDGROUP=SHIFT
```

This example groups reports by System and then Shift, and writes each System / Shift combination to a separate DD with names of CMFR*nnnnn* assigned to the same SYSOUT class as MSGCLASS.

For the following examples, assume that there are two systems: SYSA and SYSB.

Example 1

```
SYSPLEX SUBPLEX=*,RPTS=SEPARATE
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(SYSTEM,SHIFT,RPTSTMT)
PERFSUM
TSOUSER
```

The results of this example are as follows:

DDName	System	Shift	Report
SYSPRINT	SYSA	00:00-08:00	PERFSUM
SYSPRINT	SYSA	00:00-08:00	TSOUSER
SYSPRINT	SYSA	08:00-16:00	PERFSUM
SYSPRINT	SYSA	08:00-16:00	TSOUSER
SYSPRINT	SYSB	00:00-08:00	PERFSUM
SYSPRINT	SYSB	00:00-08:00	TSOUSER
SYSPRINT	SYSB	08:00-16:00	PERFSUM
SYSPRINT	SYSB	08:00-16:00	TSOUSER

Example 2

```
SYSPLEX SUBPLEX=*,RPTS=SEPARATE
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(SYSTEM,SHIFT,RPTSTMT)
DDGROUP=SYSTEM
PERFSUM
TSOUSER
```

The results of this example are as follows:

DDName	System	Shift	Report
SYSA	SYSA	00:00-08:00	PERFSUM
SYSA	SYSA	00:00-08:00	TSOUSER
SYSA	SYSA	08:00-16:00	PERFSUM
SYSA	SYSA	08:00-16:00	TSOUSER
SYSB	SYSB	00:00-08:00	PERFSUM
SYSB	SYSB	00:00-08:00	TSOUSER
SYSB	SYSB	08:00-16:00	PERFSUM
SYSB	SYSB	08:00-16:00	TSOUSER

Example 3

```
SYSPLEX SUBPLEX=*,RPTS=SEPARATE
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(SYSTEM,SHIFT,RPTSTMT),
DDGROUP=SHIFT
PERFSUM
TSOUSER
```

The results of this example are as follows:

DDName	System	Shift	Report
CMFR0001	SYSA	00:00-08:00	PERFSUM
CMFR0001	SYSA	00:00-08:00	TSOUSER
CMFR0002	SYSA	08:00-16:00	PERFSUM
CMFR0002	SYSA	08:00-16:00	TSOUSER
CMFR0003	SYSB	00:00-08:00	PERFSUM
CMFR0003	SYSB	00:00-08:00	TSOUSER
CMFR0004	SYSB	08:00-16:00	PERFSUM
CMFR0004	SYSB	08:00-16:00	TSOUSER

Example 4

```
SYSPLEX SUBPLEX=*,RPTS=SEPARATE
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(RPTSTMT,SYSTEM,SHIFT),
DDGROUP=RPTSTMT
PERFSUM
TSOUSER
```

The results of this example are as follows:

DDName	System	Shift	Report
PERFSUM	SYSA	00:00-08:00	PERFSUM
PERFSUM	SYSA	08:00-16:00	PERFSUM
PERFSUM	SYSB	00:00-08:00	PERFSUM
PERFSUM	SYSB	08:00-16:00	PERFSUM
TSOUSER	SYSA	00:00-08:00	TSOUSER
TSOUSER	SYSA	08:00-16:00	TSOUSER
TSOUSER	SYSB	00:00-08:00	TSOUSER
TSOUSER	SYSB	08:00-16:00	TSOUSER

Example 5

```
SYSPLEX SUBPLEX=*,RPTS=COMBINED
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(SYSTEM,SHIFT,RPTSTMT),
DDGROUP=SYSTEM
PERFSUM
TSOUSER
CFACT
```

The results of this example are as follows:

DDName	System	Shift	Report
CMFR0001	**ALL**	00:00-08:00	PERFSUM
CMFR0001	**ALL**	00:00-08:00	TSOUSER
CMFR0001	**ALL**	00:00-08:00	CFACT
CMFR0001	**ALL**	08:00-16:00	PERFSUM
CMFR0001	**ALL**	08:00-16:00	TSOUSER
CMFR0001	**ALL**	08:00-16:00	CFACT

Example 6

```
SYSPLEX SUBPLEX=*,RPTS=COMBINED
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(SHIFT,RPTSTMT),
DDGROUP=SHIFT
PERFSUM
TSOUSER
CFACT
```

The results of this example are as follows:

DDName	System	Shift	Report
CMFR0001	**ALL**	00:00-08:00	PERFSUM
CMFR0001	**ALL**	00:00-08:00	TSOUSER
CMFR0001	**ALL**	00:00-08:00	CFACT
CMFR0002	**ALL**	08:00-16:00	PERFSUM
CMFR0002	**ALL**	08:00-16:00	TSOUSER
CMFR0002	**ALL**	08:00-16:00	CFACT

Example 7

```
SYSPLEX SUBPLEX=*,RPTS=COMBINED
SHIFT DINTV=(000000,2,080000),RPTS=SEPARATE
REPORTS RPTGROUP=(RPTSTMT,SYSTEM,SHIFT),
DDGROUP=RPTSTMT
PERFSUM
TSOUSER
CFACT
```

The results of this example are as follows:

DDName	System	Shift	Report
PERFSUM	**ALL**	00:00-08:00	PERFSUM
PERFSUM	**ALL**	08:00-16:00	PERFSUM
TSOUSER	**ALL**	00:00-08:00	TSOUSER
TSOUSER	**ALL**	08:00-16:00	TSOUSER
CFACT	**ALL**	00:00-08:00	CFACT
CFACT	**ALL**	08:00-16:00	CFACT

SEVERITY

```
SEVERITY   nnn
```

Overview

The SEVERITY control statement specifies the maximum error severity that can be encountered for processing to continue. If the severity of an error exceeds the value defined in the SEVERITY control statement, the Analyzer terminates processing. If the error severity is equal to or less than the value defined in this statement, processing continues.

NOTE

If the SEVERITY statement is defined with an error value greater than 4, results can be unpredictable.

If the SEVERITY statement is omitted, the default of severity 0 (zero) is used, and processing terminates if any errors are encountered.

You can see the severity codes that are associated with each command error by typing **MSG** on the **COMMAND** line.

Parameters

The parameter for the SEVERITY control statement is as follows:

nnn	specifies a number from zero to 255 as the maximum error severity value; the default is zero

Example

```
SEVERITY   4
```

Processing continues if a control statement error of 4 or less is encountered.

SHARDEV

```
SHARDEV
[TYPE=(BOTH|DASD|TAPE)]
[,RANGE=(deva,devb,...devn)]
[,RANGE=(deva:devn)]
[,OFFLINE=(NO|YES)]
[REPORT=DETAIL|SUMMARY]
```

Overview

The SHARDEV control statement requests the Shared Device Activity Report.

The Shared Device Activity Report requires data from multiple MVS systems. To collect synchronous data from all systems sharing devices, BMC Software recommends using the SYNCH=SMF parameter of the Extractor REPORT control statement.

The SHARDEV statement uses data collected by the DEVICE Extractor control statement.

NOTE

- To specify which systems on which devices are shared, you must include the SYSPLEX control statement with RPTS=COMBINED. Issuing the SHARDEV control statement without the controlling SYSPLEX statement results in an error message.

- To produce interval type reports, you must specify each interval to be reported by using the *startshift*, *endshift* time pairs on the SHIFT control statement. You cannot specify RPTS=INTERVAL on the SHIFT statement.

- To match the logical device on various systems to a physical device, the Analyzer attempts to match devices using the 28-byte device ID, which is returned from self-describing devices. If the device ID contains only zeros, DASD devices are matched using the volume serial number and tape devices are matched using the device address. If a device cannot be matched by using any of these criteria, it does not appear in the report.

- BMC Software recommends that you use the DATETIME statement to generate a more meaningful report. Without this statement, all records are accepted from each system, regardless of their date/time, which could result in a report that contains many hours or days of data from one system and only a few minutes of data from another system.

Parameters

Parameters for the SHARDEV control statement are as follows:

TYPE=	specifies whether DASD, tape, or both DASD and tape devices should be included in the report; the default is BOTH
RANGE=	specifies which devices are reported You can specify either the device address or a range of device addresses on this parameter. The report contains a detail line for each system that shared the device during the reporting interval. Up to 16 devices or ranges of devices can be specified. If the RANGE parameter is not specified, all shared devices are reported.
OFFLINE=	specifies whether offline devices should be included in the report; the default is OFFLINE=YES
REPORT=	specifies the type of report If SUMMARY is specified, one line of information per device appears in the report; it summarizes activity from all of the systems. If DETAIL is specified, additional lines per device appear in the report; each line displays activity from an individual system. The default is DETAIL.

Examples

```
SYSPLEX SUBPLEX=*, RPTS=COMBINED
SHARDEV
```

This example produces a Shared Device Activity Report that lists all DASD and tape devices shared between all systems.

```
SYSPLEX SUBPLEX=*, RPTS=COMBINED
SHARDEV RANGE=(2140,2145)
```

This example produces a Shared Device Activity report for the two devices at addresses 2140 and 2145.

SHIFT

```
SHIFT
[DINTV=(starttime,intervals,duration)]
[(startshift,endshift1)1[,...,(startshift,endshift)96]|
 (080000,160000),(160000,240000),(000000,080000)]
[,TYPE={GLOBAL|LOCAL}]
[,DAYS={ALL|(MON|TUE|WED|THU|FRI|SAT|SUN)}]
[,RPTS={SEPARATE|COMBINED|INTERVAL|DAILY}]
[,UNIT=unittype]
```

Overview

The SHIFT control statement specifies data selection intervals by time of day and day of the week. For example, a single invocation of this command can produce CMF MONITOR reports for system activity during prime, swing, and graveyard shifts.

The SHIFT statement can appear in the general control statement set, the report control statement set, or both sets. In the general control statements, SHIFT defines the collection start and end times used to produce all the reports requested in the report control statements. A SHIFT statement for an individual report request overrides a global SHIFT statement for that one report request only.

The SHIFT statement can be used with the DATETIME, CYCLE, and PERIOD statements to produce monthly reports and basic history database trends.

Parameters

Parameters for the SHIFT control statement are as follows:

DINTV= allows you to generate up to 96 *startshift*, *endshift* pairs by specifying a start time, number of intervals, and a duration for data that is to be included in a report

This parameter specifies that reports be created from data obtained during multiple, evenly-spaced, sequential intervals. These three fields are all required, and commas must separate the three values.

starttime	contains the time you want the first interval to begin; this field must be specified in time format *hhmmss*
intervals	specifies the number of time pairs to be generated; this field must be specified as a number from 1 to 96
duration	specifies how long the interval lasts; this field must be specified in the format *hhmmss*

Notes:

- Use of the DINTV parameter or manually coded *startshift, endshift* pairs increases the storage requirements for the Analyzer. (See "Setting values of region, DMSS reserve, and CTRLSIZE" on page 91 for an explanation of Analyzer storage requirements.)
- While multiple *startshift, endshift* pairs may be used to specify any combination of report collection intervals, using the DINTV parameter eliminates the need to specify multiple *startshift, endshift* pairs when data from evenly spaced sequential intervals is required.
- Specifying both the DINTV parameter and one or more *startshift, endshift* pairs results in an error message.

(startshift,endshift) specifies a time-pair range that is equivalent to the duration of a shift

These variables are expressed in the time format *hhmmss*. Up to 96 different shift ranges can be specified in one SHIFT statement; the ranges must be separated by commas and enclosed within parentheses. Data is included in a report if its recording time is equal to or greater than the *startshift* value and less than the *endshift* value.

CMF MONITOR does not check to ensure that the time pairs specified correspond to default or actual shift periods. When time ranges are specified, they should not overlap. If they do and if RPTS=SEPARATE is specified, some data is used in more than one report shift.

If time-pair ranges are not specified, the following three default time pairs are used:

(080000,160000)—8:00 A.M. to 4:00 P.M., corresponding to shift 1

(160000,240000)—4:00 P.M. to 12:00 A.M. (midnight), corresponding to shift 2

(000000,080000)—2:00 A.M. (midnight) to 8:00 A.M., corresponding to shift 3

Note: The default for RPTS=DAILY is no SHIFT ranges. Time-pair ranges are mutually exclusive with RPTS=INTERVAL.

TYPE=	defines how the SHIFT statement is used	
	GLOBAL	valid in either the general or the report statement set This value defines the shift characteristics for all report requests that follow it, except single report requests associated with LOCAL SHIFT statements. If multiple GLOBAL SHIFT statements are used, each succeeding GLOBAL statement establishes new shift characteristics for the following report requests. GLOBAL is the default.
	LOCAL	valid only in the report statement set The LOCAL SHIFT statement defines the shift characteristics for the report request that immediately precedes it. A LOCAL SHIFT statement overrides any GLOBAL SHIFT statement, but only for that one report.
DAYS=	limits the scope of the input data by specifying from 1 to 7 days of the week The days can be listed individually or collectively by using the parameter values shown in Table 27. The default is ALL, except when RPTS=INTERVAL and RPTS=DAILY, because these parameters are mutually exclusive with DAYS=ALL.	

Table 27 Parameter values for DAYS

Parameter value	Abbreviations	Synonym (if applicable)
MONDAY	MON or MO	—
TUESDAY	TUE or TU	—
WEDNESDAY	WED or WE	—
THURSDAY	THU or TH	—
FRIDAY	FRI or FR	—
SATURDAY	SAT or SA	—
SUNDAY	SUN or SU	—
WEEKLY	none	ALL or (SU,MO,TU, WE,TH,FR,SA)
WORKWEEK	none	(MO,TU,WE,TH,FR)
WEEKENDS	none	(SA,SU)
ALL	none	WEEKLY or (SU,MO,TU,WE,TH, FR,SA)

RPTS= requests one report for the entire SHIFT definition set or individual reports for each time pair used to define the shift

SEPARATE One report is produced for each time pair of the corresponding SHIFT statement; SEPARATE is the default.

COMBINED One combined report is produced for all data gathered within the time ranges specified in the corresponding SHIFT statement.

INTERVAL A complete set of reports is produced for each recording interval. Reasonable limitations should be exercised with this parameter, because it might produce a large volume of output, which uses more memory and requires a larger DMSSMAIN data set or greater use of hiperspace.

Warning: All other SHIFT statements are ignored if this parameter is used anywhere in the control statement set.

When using the INTERVAL value, only one Control Card Log is printed to the //CMFLOG DD data set, and no Collection Phase Log reports are printed. Data from multiple images must not be merged, and the records must be sorted by date and time.

DAILY A complete set of reports is produced for each day where there is data in the input stream defined by the //EXTDATA DD statement.

Warning: All other SHIFT statements are ignored if this parameter is used anywhere in the control statement set.

When using the DAILY value, only one Control Card Log is printed to the //CMFLOG DD data set, and no Collection Phase Log reports are printed. Data from multiple images must not be merged, and the records must be sorted by date and time.

Note: This parameter does not function when RPTS=COMBINED is specified on the SYSPLEX control statement.

UNIT= defines a valid unit type for a temporary data set that is allocated when the RPTS=INTERVAL or RPTS=DAILY parameter is specified

For these RPTS= values, a temporary staging data set can be required and, if required, is dynamically allocated. It is possible that the default unit for your system does not allow the dynamic allocation of this data set. The UNIT value allows you to specify a valid unit type for this data set; for example, UNIT=SYSDA.

If your system still does not allow the dynamic allocation of the temporary staging data set after specifying the UNIT parameter, define //CMFSTAGE DD and //CMFSTAGO DD statements to your Analyzer JCL (see Table 12 on page 86, which lists JCL control statements for the CMF MONITOR Analyzer).

Examples

These examples illustrate the use of the SHIFT control statement in the general control statement set.

```
SHIFT
```

All defaults are in effect. This statement is equivalent to specifying:

```
SHIFT (080000,160000),(160000,240000),(000000,080000),
      TYPE=GLOBAL,DAYS=ALL,RPTS=SEPARATE
```

Regardless of the number of days encountered in the data set, three reports are generated for each report request, based on the default time pairs. The reports reflect an average profile of each shift time-pair range.

```
SHIFT  DAYS=MO
```

All defaults are in effect except DAYS. All Extractor data in the data set with a date of Monday is included for all report requests; all other days are excluded. The reports reflect a profile of an average Monday by each of the default shift time-pair ranges. If there is more than one Monday in the input data, a separate report is produced for each shift of each Monday.

```
SHIFT  DAYS=(MO,WE,FR),RPTS=COMBINED
```

All defaults are in effect except DAYS and RPTS. Data for Mondays, Wednesdays, and Fridays is collected into one report for each default shift time-pair range. Then all data collected under the default shift time-pair ranges is reported as a combined data group for all report requests. The result is one report for each report request containing data for Mondays, Wednesdays, and Fridays.

```
SHIFT  (100000,120000),(140000,160000)
```

All defaults are in effect except the shift time-pair ranges. Two reports for each report request are created. This feature can be used to break a shift into smaller reporting periods. Up to 96 ranges can be specified per SHIFT command, either as LOCAL shifts or as multiple GLOBAL shifts.

```
SHIFT  DINTV=(080000,32,001500)
```

All defaults are in effect except the shift time-pair ranges. For each report request, 32 reports are created. The first report contains data collected during a 15-minute interval beginning at 8:00 A.M. (from 8:00 to 8:15), the second report contains data collected from 8:15 A.M. to 8:30 A.M., and so on.

The following examples illustrate the use of the SHIFT control statement in the report control statement set.

```
SHIFT  (080000,120000),(130000,170000),RPTS=COMBINED
CPU
CHANNEL
DEVACT
SHIFT (000000,060000),(160000,220000),
       TYPE=GLOBAL,DAYS=WORKWEEK
PERFSUM
SRM
DASD
```

The second SHIFT statement that is displayed in this example is designated as TYPE=GLOBAL, redefining the shift characteristics for the report requests that follow.

The CPU, CHANNEL, and DEVACT report requests that precede the SHIFT TYPE=GLOBAL statement conform to the shift characteristics set by the first SHIFT statement in the control statement set.

```
IOQ
SHIFT TYPE=LOCAL
```

All defaults are in effect except TYPE. Because LOCAL is specified, this SHIFT statement applies only to the I/O Queuing Activity report request. Three I/O Queuing Activity reports are generated, one for each of the default time-pair ranges.

```
WLMGL
SHIFT TYPE=LOCAL,DAYS=WEEKENDS
```

All defaults are in effect except TYPE and DAYS. Because LOCAL is specified, this SHIFT statement applies only to the WLMGL report request. All workload summary-type data with a date of Saturday and Sunday is included in the Workload Summary report; all data from other days is excluded. Three reports are produced, one for each default time-pair range. Each report reflects an average profile for that shift for an average Saturday and Sunday.

The SHIFT statement can be used with the DATETIME, CYCLE, and PERIOD statements to produce monthly performance evaluation reports; for example

```
DATETIME (1MAR03:000000,31MAR03:235959)
CYCLE    DAILY
PERIOD   CYCLE=DAILY,
         RPTS=SEPARATE
SHIFT    (080000,163000)
DMSS     INIT,
         RESERVE=400
RECTYPE  240,
         MONITOR=CPM
CPU
LCU
AUXSTOR
PERFSUM
.
.
.
```

This request produces reports for each prime shift (8:00 A.M. to 4:30 P.M.) for each day in March 2003. Thirty-one sets of reports are produced.

The SHIFT statement can be used with DATETIME, CYCLE, and PERIOD to produce historical trend data for performance evaluation reporting; for example

```
DATETIME (1JAN03:000000,31MAR03:235959)
CYCLE    WEEKLY
PERIOD   CYCLE=WEEKLY,
         RPTS=SEPARATE
SHIFT    (080000,163000),
         DAYS=(MO,WE,FR)
DMSS     INIT,
         RESERVE=400
RECTYPE  240,
         MONITOR=CPM
CPU
TSOPERF
SRM
PERFSUM
WLMGL
.
.
.
```

This request produces 14 sets of reports, for the first 13 weeks of 2003 and one for the last 6 days of March. Each week comprises data collected from Monday, Wednesday, and Friday between 0800 and 1630.

Average profiling of a day, week, and weekend can be accomplished by using Tabular Summary reports. The PROFILE parameter on the GRAPH command statement enables graphic measures to be represented in a profile manner similar to the Tabular reports.

SRM

```
SRM
```

Overview

The SRM control statement requests the System Resources Manager Report that is discussed in “System Resources Manager Report” on page 515.

The SRM statement uses data collected by the following Extractor control statements:

- CPU (see “CPU” on page 136)
- PAGING (see “PAGING” on page 173)

Parameters

This statement has no parameters.

Example

```
SRM
```

This example produces the System Resource Manager Reports.

STORAGE

```
STORAGE
```

Overview

The STORAGE control statement requests the Storage Management Report that is discussed in "Storage Management Report" on page 509.

The STORAGE statement uses data collected by the PAGING Extractor control statement (see "PAGING" on page 173).

Parameters

This statement has no parameters.

Example

```
STORAGE
```

This example produces the Storage Management Report.

SUBTITLE

```
SUBTITLE  'text'
```

Overview

The SUBTITLE control statement specifies a text string to be printed on the fourth line of the requested report. This text string can also appear under the SUBTITLE column in the Report Table of Contents. The text string is used as a unique identifier when several similar reports are requested. The SUBTITLE control statement must follow the report control statement for the report to be subtitled.

Parameters

There is one parameter for this control statement:

'*text*'	The '*text*' is a quoted character string up to 52 characters long, excluding outside delimiting apostrophes. If the string contains no blank spaces and consists only of alphabetic and numeric characters, the enclosing apostrophes can be omitted.

Examples

```
IOQ
SUBTITLE   'SUMMARY REPORT'
```

One I/O Queuing Activity Report is produced for the entire reporting period, and the subtitle SUMMARY REPORT is printed on the fourth line of the report header.

```
IOQ
SUBTITLE  'DAILY REPORT'
SHIFT      RPTS=DAILY,TYPE=GLOBAL
```

Multiple I/O Queuing Activity Reports are produced, one for each whole or partial day within the reporting period, and the subtitle DAILY REPORT is printed on the fourth line of the report header.

SYSPLEX

```
SYSPLEX
[RPTS={COMBINED|SEPARATE}]
[,SUBPLEX={(*|sysname1[,...,sysnamen])}]
[,TYPE={SYSID|SYSNAME}]
[,INPUT={(dd1[,...,ddn])}]
```

Overview

With SYSPLEX, you select the data that is used in generating Analyzer reports. Data selection is based on SYSNAME or SYSID, and on DDNAME.

Parameters

Parameters for the SYSPLEX control statement are as follows:

RPTS=	specifies whether the data collected is combined into one report per group of SYSNAMEs or is divided into separate reports for each SYSNAME; the default is SEPARATE **Warning**: When using RPTS=COMBINED, BMC Software recommends that data from each MVS image covers identical date and time ranges. Averages are calculated based on the duration time of the data rather than the number of systems.
SUBPLEX=	specifies a list of SYSNAMEs that are used in selecting records for reporting You can specify up to 128 SYSNAMEs. The default is SUBPLEX=*, which specifies that records from all SYSNAMEs or SYSIDs (depending on what you specify in the TYPE parameter) found in the input data are accepted.
TYPE=	specifies whether the names used in the SUBPLEX parameter are SYSNAMEs or SYSIDs; the default is SYSNAME
INPUT=	specifies a list of DDNAMEs that are read You can specify up to 128 DDNAMEs, which are read in the order in which they are specified. The default is the same data sets specified by //EXTDATA in the Analyzer JCL (see “Defining Analyzer JCL manually” on page 84 for more information). **Note**: If you want to use data from the XDS data server for your reports, this parameter must be omitted.

Examples

```
SYSPLEX
```

This example specifies that separate reports are produced for each SYSNAME, reading every data set specified by //EXTDATA in the Analyzer JCL.

```
SYSPLEX RPTS=COMBINED
```

This example specifies that a single report is produced, combining all SYSNAMEs and reading every data set specified by //EXTDATA in the Analyzer JCL.

TRACE

```
TRACE
[ID={ALL|(0,1,2,...,99)}]
[,MODE={SRB|SRM}]
[,LIMIT=0|nnnn]
[,START=(yyddd:hhmmss)]
[,STOP=(yyddd:hhmmss)]
```

Overview

The TRACE control statement produces the Trace Report shown in Figure 90 on page 530. This statement causes the data in the trace records that are generated by the Extractor to be formatted. The trace formatter differs in architecture from other Analyzer reporters in that the trace formatter does not have a data collection phase. The trace records are formatted and written out as they are read. For this reason, a separate DCB is used to produce the Trace Report. The ddname for this DCB is allocated dynamically by the Analyzer, unless the optional DD statement, //CMXTRACE, is used to direct the output from the Trace Report.

The TITLE and LOCATION parameters from the HEADERS general control statement are not reproduced on the Trace Report. If a SUBTITLE statement is defined, it is displayed on the Trace Report.

NOTE

Records must be sorted by date and time when using the TRACE statement.

The TRACE statement uses data collected by the TRACE Extractor control statement (see “TRACE” on page 183).

Parameters

Parameters for the TRACE control statement are as follows:

ID= specifies IDs for which the trace data is to be formatted

Valid ID values are 0 through 99 and ALL; the default is ALL. If ALL is defined, all encountered trace IDs are formatted.

If you specify a value for this parameter, the corresponding values must be specified in the TRACE Extractor control statement.

MODE= specifies one of two types of trace entries to be formatted: SRB or SRM

If you do not specify a mode, all modes are formatted.

LIMIT= limits the number of entries to be processed

Valid values are any value from 0 through 9999; the default is 0. If 0 is defined, all entries encountered are formatted.

START= specifies the starting date and time for selecting trace records

If START is not specified, selection begins with the first trace record encountered. The starting date and time is defined in the same manner as the start date-time value for the DATETIME control statement (see “DATETIME” on page 238 for more information).

STOP= specifies the ending date and time for selecting trace records

If STOP is not specified, selection terminates with the last trace record encountered. The ending date and time is defined in the same manner as the end date-time value for the DATETIME control statement (see “DATETIME” on page 238 for more information).

Note: The DATETIME control statement affects the TRACE statement, but TRACE ignores all CYCLE, PERIOD, and SHIFT control statements. Data in the Trace Report is limited by the ranges defined in the DATETIME control statement, and in the START and STOP parameters of the TRACE control statement.

Examples

```
TRACE
```

This example produces the Trace report with all defaults in effect. These defaults are all modes and all IDs for the Analyzer date-time range, with no limit on the number of entries to be formatted.

```
TRACE  ID=(0,8,9),MODE=SRM,LIMIT=5000,START=(02145:113000),
       STOP=02145:143000
```

This example produces the Trace Report for only SRM-type entries with IDs of 0, 8, and 9. Only the first 5000 entries are formatted. Record selection begins at 1130 hours on May 25, 2003, and ends at 1430 hours on May 25, 2003.

```
TRACE  ID=ALL,MODE=SRB
```

This example produces the Trace Report for only SRB-type entries. All IDs are formatted.

```
TRACE  LIMIT=100,MODE=SRM,ID=0
```

This example produces the Trace Report for only SRM-type entries. Only ID 0 is selected and only the first 100 entries are formatted.

TSOPERF

```
TSOPERF
[TYPE={COMMAND|CMD|INTERVAL|INT|ALL}]
[,PLOT={RESPONSE|RSP|TSO}]
[,LIMIT={32|nn}]
[,SCALE={999|nnn}]
[,DARK={NO|YES}]
```

Overview

The TSOPERF control statement requests two reports:

- TSO Command Summary Report—discussed in “TSO Command Summary Report” on page 532
- TSO Interval Summary Report—discussed in “TSO Interval Summary Report” on page 534

The data for both of these reports is gathered by using the Extractor TSODATA control statement (see “TSODATA” on page 194). The TSODATA statement contains a LIMIT parameter that affects data extraction, which can also affect the information available for this report.

The TSOPERF statement uses data collected by the following Extractor control statements:

- CPU (see “CPU” on page 136)
- TSODATA (see “TSODATA” on page 194)

Parameters

Parameters for the TSOPERF control statement are as follows:

TYPE= specifies which of the two report formats is to be produced; the default is ALL

COMMAND	produces the Command Summary report; it can be abbreviated CMD
INTERVAL	produces the Interval Summary report; it can be abbreviated INT
ALL	produces both the Command and the Interval Summary reports

PLOT= Identifies which of two variables reported in the Interval Summary report is to be plotted

RESPONSE plots the average response time for the interval; it can be abbreviated RSP; the default

TSO plots the average CPU busy percent

LIMIT= limits the number of TSO commands that are reported

If 999 is specified, all TSO commands are reported. If a value less than 999 is specified, it defines the maximum number of TSO commands that can be reported. TSO commands are selected for the report by order of usage until the LIMIT value is reached. If TYPE=INT is specified, LIMIT is ignored. The default is 32.

SCALE= specifies the scaling value to be used on the response plot for the Command Summary report

The range of values is from 1 to 999. The default is 999: the maximum response time is used for the scale. If TYPE=INT is specified, SCALE is ignored.

DARK= specifies whether the average bar line on the plots will overstrike

Overstriking the print line highlights the average line. The default is NO.

Examples

```
TSOPERF
```

This example produces both the TSO Command Summary and the TSO Interval Summary Reports, showing the average response times of the 32 most-used commands. For the TSO Command Summary Report, the default scaling value is used to plot the maximum response times.

```
TSOPERF  TYPE=CMD,LIMIT=25,SCALE=10
```

This example produces a TSO Command Summary Report showing the average response times of the 25 most-used commands, using a scaling value of 10.

TSOUSER

```
TSOUSER
[LIMIT={32|nn}]
```

Overview

The TSOUSER control statement requests the TSO User Summary Report shown in Figure 93 on page 537. The report data includes only those active TSO users that were sampled. The most-active users are the ones with the greatest total service units.

The TSOUSER statement uses data collected by the TSODATA Extractor control statement, when USERS=YES is defined (see “TSODATA” on page 194).

Parameters

The parameter for the TSOUSER control statement is as follows:

LIMIT= limits the number of TSO users that are reported

If 999 is specified, all TSO users are reported. However, if a number less than 999 is specified, it defines the maximum number of TSO users that can be reported. TSO users are selected for the report by order of activity until LIMIT value is reached. The default is 32.

Examples

```
TSOUSER
```

This example produces a TSO User Summary Report that shows the activity data of the 32 most-active users who were sampled.

```
TSOUSER  LIMIT=999
```

This example produces a TSO User Summary Report that shows TSO activity for all users sampled.

```
TSOUSER  LIMIT=5
```

This example produces a report that shows activity for only the five most-active users who were sampled.

VIRTSTOR

```
VIRTSTOR
[MAP={YES|NO}]
[,DETAIL={YES|NO}]
[,JOBNAME=(job1,...,job25)]
```

Overview

The VIRTSTOR control statement requests the Virtual Storage Activity Report that is discussed in “Virtual Storage Activity Report” on page 539.

The VIRTSTOR statement uses data collected by the VSMDATA Extractor control statement (see “VSMDATA” on page 200).

Parameters

Parameters for the VIRTSTOR control statement are as follows:

MAP= specifies whether to generate the Virtual Storage Map; the default is YES

DETAIL= specifies whether to produce summary or detail reports

The detail report for common area data includes a breakdown of SQA by subpool and CSA by subpool and storage key. The detail report for private area data includes a breakdown of the private area data by subpool. The default is NO; only summary reports are produced.

JOBNAME= specifies the job name for which reports are to be produced

The default is to produce a report for every job for which data is encountered. Specifying job names in this field limits reporting to the specified job. To report on all jobs (the default), omit the JOBNAME parameter or leave it blank.

Note: If JOBNAME=ALL is specified, the Analyzer tries to print a report for a job named ALL and will suppress reporting for any other jobs.

Private area data is available for this report only for the job names specified on the VSMDATA sampler. If no job names are specified for the Extractor control statement, no private area data is available for this report. (See “VSMDATA” on page 200 for more information.)

Examples

```
VIRTSTOR
```

This example produces the Virtual Storage Map, Common Storage Summary Reports, and summary reports for all jobs sampled.

```
VIRTSTOR MAP=NO,DETAIL=YES
```

This example produces the Common Storage Summary and Detail Reports, and summary and detail reports for all jobs sampled.

```
VIRTSTOR DETAIL=YES,JOBNAME=(CICSPROD)
```

This example produces the Virtual Storage Map, Common Storage Summary and Detail Reports, and the Private Area Summary and Detail Reports for job CICSPROD.

```
VIRTSTOR JOBNAME=(########)
```

This example produces the Virtual Storage Map and Common Storage Summary Report. Because ######## is not a valid job name, private area reports are not produced.

NOTE

Any invalid or nonexistent job name can be used to suppress private area reports.

VOLSER

```
VOLSER  volser1,...,volser255
```

Overview

The VOLSER control statement can be used with the DASD control statement to request the Direct Access Report Plot of Volume that is discussed in "Direct Access Report Plot of Volume" on page 427. The VOLSER statement must immediately follow the DASD statement.

Parameters

The VOLSER control statement has one parameter:

volser*n*	You can specify up to 255 six-character volume serial numbers. If a VOLSER control statement is used without a volume serial list (the default), all volumes monitored during the measurement interval are listed.

NOTE

This plot requires records written at CMF MONITOR Extractor initialization. To obtain the desired report, make certain that the input data includes the records written when the extraction began. It is not necessary to include the reporting period of these records in the DATETIME or SHIFT ranges.

Example

```
DASD
VOLSER MVS001,SYSPK1,PAGE01
```

This example produces a volume activity plot of the Direct Access Report showing data for the volumes MVS001, SYSPK1, and PAGE01.

WLMGL

```
WLMGL
[UNUSED=(SUPPRESS|PRINT)]
[,TYPE=(DETAIL|ACTIVITY|DELAY|ALL)]
[,PERIOD=(DAY|WEEK|MONTH)]
```

Overview

The WLMGL control statement requests the Workload Manager Report that is discussed in "Workload Manager Goal Mode Report" on page 549. The PERFORM statement can be defined with the WLMGL control statement to include or exclude service classes. See "PERFORM" on page 277 for more information. To collect synchronous data from all systems in the sysplex, BMC Software recommends that you use the SYNCH=SMF parameter of the Extractor REPORT control statement.

Parameters

Parameters for the WLMGL control statement are as follows:

UNUSED= specifies whether reports are printed for service classes that report no work

- **SUPPRESS** prevents the generation of reports for service classes that report no work; the default
- **PRINT** requests that reports for all service classes be generated

TYPE= specifies the type or types of reports to be generated; allowable values are as follows:

- **DETAIL** generates the Detail section
- **ACTIVITY** generates the Goal Activity Map section
- **DELAY** generates the Goal Mode Delay Map section
- **ALL** generates all three reports

Note: If you do not specify the TYPE parameter, only the Detail section of the Goal Mode Report is produced.

PERIOD	specifies the granularity of the ACTIVITY and DELAY MAPS	
	DAY	requests that the map represent processing for each 15-minute period of a day If the data from more than one day is processed, the map shows the combined activity for all days. This value is the default.
	WEEK	requests that the map represent processing for each two-hour period of a week
	MONTH	requests that the map represent processing for each eight-hour period of a month

Examples

```
WLMGL
```

This example produces the Goal Activity Map section of the Workload Manager Goal Mode Activity Report.

```
WLMGL TYPE=ALL,PERIOD=WEEK
```

This example produces all sections of the Workload Manager Goal Mode Activity Report. The Goal Activity Map section and the Goal Mode Delay Map section are produced with values averaged over each two-hour period of the week.

```
WLMGL TYPE=DELAY,PERIOD=MONTH
```

This example produces the Goal Mode Delay Map section of the Workload Manager Goal Mode Activity Report. This map is produced with values averaged over each eight-hour period of the month.

XCF

```
XCF
[TYPE={SUMMARY|DETAIL|BOTH}]
[,ORDER=(SYSTEM, GROUP, MEMBER)]
```

Overview

The XCF Analyzer control statement requests the Cross-System Coupling Facility (XCF) Report that is discussed in “Cross-System Coupling Facility Report” on page 413. The XCF report shows the activity of the Cross-System Coupling Facility.

The information can be summarized for the system and presented individually by member.

The report is divided into three sections:

- **System Summary**

  This section contains information about XCF activity for each pairing of source and target systems, including activity within each system.

- **Path Utilization**

  This section contains information about XCF activity for each path between or within systems.

- **Detail Report**

  This section contains information about XCF activity for each member. You can sort the information in this section by using the ORDER parameter.

NOTE

Information that is not available for an inbound path is marked with <INBOUND> in that column for all inbound paths.

The XCF statement uses data collected by the XCFDATA Extractor control statement (see “XCFDATA” on page 205).

Parameters

Parameters for the XCF Analyzer control statement are as follows:

TYPE= specifies the type of reports to be generated as one of the following values:

- **SUMMARY** specifies a report of the system usage of XCF and buffer utilization
- **DETAIL** specifies a report of XCF usage by member
- **BOTH** specifies reports of the system usage of XCF, buffer utilization, and XCF usage by member; the default

ORDER= specifies the ordering of the information displayed in the DETAIL report

This operand has no effect on the SUMMARY reports. The DETAIL report contains a summary line that is calculated based on the first parameter of ORDER. For example, if you specify ORDER=GROUP,SYSTEM,MEMBER, the values on the summary line are calculated based on groups.

If all three values are specified, the order of display for the DETAIL report follows the sequence in which the values are defined. If one or two values are specified, the order of display for the DETAIL report starts with the value or values specified, and follows with the remaining values in the default order.

Examples

These examples show what reports are generated, given the parameters that are specified in the control statement:

```
XCF
```

This example produces SUMMARY and DETAIL reports. The DETAIL report is ordered SYSTEM, GROUP, and MEMBER, with summary lines for SYSTEM.

```
XCF  TYPE=DETAIL,ORDER=(GROUP,SYSTEM)
```

This example produces the DETAIL report. The DETAIL report is ordered GROUP, SYSTEM, and MEMBER, with summary lines for GROUP.

```
XCF  ORDER=MEMBER
```

This example produces SUMMARY and DETAIL reports. The DETAIL report is ordered MEMBER, SYSTEM, and GROUP, with summary lines for MEMBER.

Chapter

8

Analyzer reports

Whenever CMF MONITOR produces reports, it can also produce other preliminary reporting information. This information is optional and can be affected by Analyzer JCL. (See “Preliminary reporting information” on page 330 for more information.)

CMF MONITOR reports are listed alphabetically in this chapter. Samples of each CMF MONITOR report are illustrated, and field descriptions are explained. Information about the fields in the heading section of each report and how report data is formatted within report fields is found in Chapter 4, “Producing and using Analyzer reports.”

Table 28 provides an alphabetical listing of the CMF MONITOR reports. The table shows the required control statements for each report or log, and specifies the page number where each report is discussed.

Table 28 CMF MONITOR reports and control statements (part 1 of 3)

Report or log title	Page	Analyzer statement	Extractor statement
Auxiliary Storage Report	page 332	AUXSTOR	ASMDATA PAGING
Cache Subsystems Overview Report	page 337	CACHEACT	CACHE
Cache Subsystem Activity Report	page 340	CACHEACT REPORT=SUBSYS	CACHE
Cache Device Activity Report	page 347	CACHEACT REPORT=DEVICE	CACHE
Channel Path Activity Report	page 350	CHANNEL	CHANNEL
CMF Record Statistics Report	page 353	CMFSTAT	none required
CMF Summary Report	page 355	CMFSUM	CHANNEL, CPU, DEVICE, IOQ, PAGING, TSODATA, VSMDATA, WORKLOAD
Collection Phase Log	page 361	none required	none required
Common Storage Usage Detail Report	page 374	COMMSTOR REPORT=DETAIL	CSMON

Table 28 CMF MONITOR reports and control statements (part 2 of 3)

Report or log title	Page	Analyzer statement	Extractor statement
Common Storage Usage Summary Report	page 377	COMMSTOR	CSMON
Coupling Facility Activity Report	page 378	CFACT	CFDATA
CPU Utilization Report	page 393	CPU	CPU
CPU Utilization by Protect Key Report	page 411	PROTKEY	CPU
Cross-System Coupling Facility Report	page 413	XCF	XCFDATA
Cryptographic Hardware Activity Report	page 418	CRYPTO	CRYPTO
Device Activity Report	page 421	DEVACT	DEVICE
Direct Access Report	page 425	DASD	DEVICE, HEADMOVE
Direct Access Report Plot of Volume	page 427	DASD, VOLSER	DEVICE, HEADMOVE
Disabled Delay Report	page 429	none required	DISTIM
Distribution Graph	page 431	GRAPH TYPE=DISTRIB	depends on data to be graphed
Enqueue Conflict Report	page 434	ENQUEUE	ENQUEUE
ESS Statistics Report	page 436	ESS	CACHE
Exception Subreport	page 440	EXCEPTS	depends on data requested on report
Exception Trace Detail Report	page 442	EXCEPTS TRACE=YES	TRACE76
Extractor Summary Report	page 443	none required	EXTSUM
FICON Director Activity Report	page 449	FICONSW	FICONSW
Graphics Trace Detail Report	page 451	GRAPH TYPE=TRACE	TRACE76
HFS Statistics Report	page 453	HFS	HFS
HTTP Server Summary Report	page 459	HTTP	none required
HTTP Server Detail Report	page 460	HTTP	none required
I/O Queuing Activity Report	page 464	IOQ	IOQ, DEVICE
Interval Bar Graph	page 468	GRAPH TYPE=PLOT	depends on data to be graphed
Kiviat Graph	page 470	GRAPH TYPE=KIVIAT	depends on data to be graphed
Link Pack Area Report	page 472	LINKPACK	LINKMAP and GBLS=YES on REPORT statement
Logical Partition Report	page 476	PRSM	CPU
LOTUS DOMINO Server Report	page 480	DOMINO	none required
LOTUS DOMINO Database Activity Report	page 485	DOMINO	none required
OMVS Kernel Activity Report	page 488	OMVS	OMVS

Table 28 CMF MONITOR reports and control statements (part 3 of 3)

Report or log title	Page	Analyzer statement	Extractor statement
Performance Summary Report	page 491	PERFSUM, PERFORM	ASMDATA, CHANNEL, CPU, DEVICE, ENQUEUE, EXTSUM, IOQ, PAGING, TSODATA, WORKLOAD
Pie Chart	page 498	GRAPH TYPE=PIE	depends on data to be graphed
Processor Concurrency Report	page 500	CPUCON	CPU
Profile Bar Graph	page 502	GRAPH TYPE=PROFILE	depends on data to be graphed
Report Table of Contents	page 504	//RPTCONTS DD statement	none required
Shared Device Activity Report	page 505	SHARDEV	DEVICE
Storage Management Report	page 509	STORAGE	PAGING
System Resources Manager Report	page 515	SRM	CPU, PAGING
Tabular Subreport	page 527	GRAPH TYPE=TAB	depends on data to be graphed
Trace Report	page 528	TRACE	TRACE
TSO Command Summary Report	page 532	TSOPERF TYPE=CMD	TSODATA, CPU
TSO Interval Summary Report	page 534	TSOPERF TYPE=INT	TSODATA,CPU
TSO User Summary Report	page 536	TSOUSER	TSODATA
Virtual Storage Activity Report	page 539	VIRTSTOR	VSMDATA
Workload Manager Goal Mode Report	page 549	WLMGL	WORKLOAD

Preliminary reporting information

Whenever CMF MONITOR reports are requested, preliminary reporting information can be generated. This information includes log reports that are automatically generated about the system, the records, Analyzer processing, Extractor operation, and a table of contents.

Spinoff reports and preliminary reporting information that is generated include

- a table of contents listing the requested reports

NOTE

The Report Table of Contents is generated only if the //RPTCONTS DD JCL statement is specified in the Analyzer JCL. (See Chapter 4, "Producing and using Analyzer reports" for more information.)

- a copy of the Analyzer control statements used to request the reports, including any syntax error messages that are generated during control statement scanning (see "Control Card Log" on page 362 for more information)
- a collection of log reports that show extraction characteristics, installation performance specifications (IPS), SRM constants, and input record type counts for the collection phase (see "About extractor log reports" on page 331 for more information)
- a report that summarizes Extractor interval activity (see "Extractor Summary Report" on page 443 for more information)

NOTE

The Analyzer output for all spinoff reports can be directed to another data set by defining the //CMFLOG DD JCL control statement. (See Chapter 4, "Producing and using Analyzer reports" for more information.)

- to store the Analyzer report in a new data set, a data set allocated with
  - RECFM=FBM
  - LRECL=133
  - DSORG=PS

About extractor log reports

Extractor log reports can be directed to a specific location by using the Analyzer //CMFLOG DD statement.

There are five log reports about Extractor operation during data collection, the report records, SRM constants, IPS, and other processing data. Reports produced are as follows:

- "Data Distribution and DATETIME Chart" on page 371
- "Extraction Characteristics Report" on page 363
- "RMF/CMF Input Record Type Counts Report" on page 369
- "System Resources Manager Constants Report" on page 366

NOTE

If either RPTS=INTERVAL or RPTS=DAILY is specified on a SHIFT control statement, the above five reports are not produced.

If multiple systems are being reported or if PARM=NLOG is used in the //CMFRPTS EXEC statement of the Analyzer JCL, the Extraction Characteristics and System Resources Manager Constants reports are not produced.

Auxiliary Storage Report

The Auxiliary Storage Report provides information about the I/O activity of page data sets, and data constants for the Auxiliary Storage Manager.

The Auxiliary Storage Report is produced by using the AUXSTOR Analyzer control statement (see “AUXSTOR” on page 215). The data for this report is obtained by using the ASMDATA (see “ASMDATA” on page 127) and PAGING (see “PAGING” on page 173) Extractor control statements.

The report is divided into these sections:

- **Cross Reference**

  This section contains general information about the page data sets, and the devices on which they reside.

- **Page Data Set Data**

  This section lists more detailed information about the page data sets, such as start I/O rates and slots in use.

- **Page Data Set Slot Count**

  This section summarizes local page data set slot usage.

NOTE

The PAGING Extractor control statement must have been defined to the Extractor control statement set when data was collected for the reporting period; otherwise, this report section is not produced. The PAGING statement causes the Extractor to collect SMF type 71 records, which are required for this section of the report.

Each of the report sections is discussed on the following pages. A sample of the Auxiliary Storage Report is shown in Figure 28 on page 333.

Figure 28 Auxiliary Storage Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      AUXILIARY STORAGE REPORT                                  RPTSEQ 3 PAGE      6
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                                     REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 10.15.00  10 JUN YY 13.45.00                  HOUSTON, TX.                                        SYSTEM ID: SJSD     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-2/14/6,300/3.5  240-9/14/6,300/3.5  75-1/196/6,300/3.5  71-1/14/210/3.5
-------------------------------------------------------  CROSS REFERENCE SECTION  ---------------------------------------------------
                 INDEX         ADDR      DEVICE TYPE       VOLUME   VIO  ASM SERVICE  DATA SET NAME
                                                                          BURST PAGES
                   1           85B3         33903          SADPG1                 30  PAGE.SJSD.COMMON
                   2           4A1F         33903          SADPG4   YES           30  PAGE.SJSD.LOCALA
                   3           4A1F         33903          SADPG4   YES           30  PAGE.SJSD.LOCALB
                   4           4A1F         33903          SADPG4   YES           30  PAGE.SJSD.LOCALC
                   5           85B3         33903          SADPG1   YES           30  PAGE.SJSD.LOCAL1
                   6           85B3         33903          SADPG1   YES           30  PAGE.SJSD.LOCAL2
                   7           85B3         33903          SADPG1   YES           30  PAGE.SJSD.LOCAL3
                   8           A205         33903          SADPG2   YES           30  PAGE.SJSD.LOCAL4
                   9           A205         33903          SADPG2   YES           30  PAGE.SJSD.LOCAL5
                  10           A205         33903          SADPG2   YES           30  PAGE.SJSD.LOCAL6
                  11           4A10         33903          SADPG3   YES           30  PAGE.SJSD.LOCAL7
                  12           4A10         33903          SADPG3   YES           30  PAGE.SJSD.LOCAL8
                  13           4A10         33903          SADPG3   YES           30  PAGE.SJSD.LOCAL9
                  14           85B3         33903          SADPG1                 30  PAGE.SJSD.PLPA
-------------------------------------------------------  PAGE DATA SET DATA SECTION  ------------------------------------------------
                                                                                 DATASET  PERCENT     PERCENT      TOTAL   TOTAL
                        AVG I/O  AVG  XFER PAGES PER BURST    PAGES    TOTAL SSCH  SSCH  DATA SET SLOTS-IN-USE     SLOTS   BAD
      INDEX   CLASS    TIME MSEC TIME MSEC    AVG     MAX  TRANSFERRED   ISSUED   PER SEC  IN USE    AVG    MAX    ALLOC   SLOTS
        1     COMMON      ---       0.00      1.0       1          4          4       0.0     0.0     0.1    0.1    53999       0
        2     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
        3     LOCAL       0.68      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
        4     LOCAL       0.67      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
        5     LOCAL       0.68      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   134999       0
        6     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   134999       0
        7     LOCAL       0.74      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0    53999       0
        8     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
        9     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
       10     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
       11     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
       12     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
       13     LOCAL       1.36      0.00      0.0       0          0          0       0.0     0.0     0.0    0.0   179999       0
       14     PLPA        ---       5.38      1.0       2        372        368       0.0     0.0    88.2   88.2    17999       0
              TOTAL                                            376        372                                   2015986       0
-----------------------------------------------  PAGE DATA SET SLOT COUNT SECTION  ---------------------------------------------
                                 AVAILABLE      %        VIO      %      NON-VIO      %       BAD        TOTAL
                     MINIMUM -   1,943,988   100.0        0     0.0          0     0.0         0    1,943,988
                     MAXIMUM -   1,943,988   100.0        0     0.0          0     0.0         0    1,943,988
                     AVERAGE -   1,943,988   100.0        0     0.0          0     0.0         0    1,943,988
```

Cross Reference Section field descriptions

Each field in the Cross Reference Section of the Auxiliary Storage Report is described in Table 29.

Table 29 Field descriptions for the Cross Reference Section (part 1 of 2)

Field	Description
INDEX	internally generated index number for referencing data sets in this report
ADDR	address of the device on which the data set resides
DEVICE TYPE	type of the device on which the data set resides
VOLUME	serial number of the volume on which the data set resides
VIO	word YES is displayed if the local page data set accepts VIO pages

Table 29 Field descriptions for the Cross Reference Section (part 2 of 2)

Field	Description
ASM SERVICE BURST PAGES	number of pages per I/O burst
DATA SET NAME	name of the data set in question

Page Data Set Data Section field descriptions

Each field in the Page Data Set Data Section of the Auxiliary Storage Report is described in Table 30.

Table 30 Field descriptions for the Page Data Set Data Section (part 1 of 2)

Field	Description
INDEX	index number corresponding to the data sets in the Cross Reference section of this report The index is followed by an asterisk (*) if the data set was brought online during the sampling session.
CLASS	page space type, which can be one of the following values: **PLPA** pageable link pack area pages **COMMON** common system area pages **DUPLEX** duplicate PLPA pages; used only if duplexing is active **LOCAL** private area (address space) pages **SWAP** address space pages that were swapped out **NOGOOD** warning that the data set is unusable for paging or swapping by the ASM; the page data set must be reformatted **Note**: If marked with an asterisk (*), the volume for the page data set came online during the Extractor session.
AVG I/O TIME MSEC	average time in milliseconds to perform an I/O operation transferring a single page This field is applied only to LOCAL page data sets.
AVG XFER TIME MSEC	average time in milliseconds to transfer a page This value differs from AVG I/O TIME, in that it is the average of all of the paging I/O rather than the average of samples that transferred single pages. A value greater than AVG I/O TIME can indicate a high level of contention for access to the data set.
PAGES PER BURST – AVG	average number of pages that were transferred to the page data set in one I/O burst
PAGES PER BURST – MAX	maximum number of pages that were transferred to the page data set in one I/O burst
PAGES TRANSFERRED	number of pages that were transferred between storage and page data set

Table 30 Field descriptions for the Page Data Set Data Section (part 2 of 2)

Field	Description
TOTAL SSCH ISSUED	total number of SSCH requests by the ASM for the indicated data set
DATASET SSCH PER SEC	rate at which SSCH requests were issued to the data set
PERCENT DATA SET IN USE	percentage of time that the data set was observed to be in use by the ASM (an IORB for that data set was busy)
PERCENT SLOTS-IN-USE – AVG	average percentage of total slots allocated that were used for this data set
PERCENT SLOTS-IN-USE – MAX	maximum percentage of total slots allocated that were used for this data set
TOTAL SLOTS ALLOC	total number of slots that were allocated for this data set
TOTAL BAD SLOTS	total number of slots that encountered permanent I/O errors on this data set

Page Data Set Slot Count Section field descriptions

Table 31 describes each field in the Page Data Set Slot Count Section of the Auxiliary Storage Report. Minimum, maximum, and average values are shown for each field in this section of the report. Each row is labeled with its appropriate value description.

Table 31 Field descriptions for the Page Data Set Slot Count Section

Field	Description
AVAILABLE %	minimum, maximum, and average number of page data set slots that are available for use (which do not contain any data pages) This value is also expressed as a percentage of the total number of local page data set slots.
VIO %	minimum, maximum, and average number of page data set slots that contain pages for VIO data sets This value is also expressed as a percentage of the total number of local page data set slots.
NON-VIO %	minimum, maximum, and average number of page data set slots that contain pages for virtual storage associated with an address space This value is also expressed as a percentage of the total number of local page data set slots.
BAD	minimum, maximum, and average number of page data set slots that are unusable The Auxiliary Storage Manager (ASM) encountered a permanent I/O error accessing it.
TOTAL	minimum, maximum, and average number of all local page data set slots

Cache reports

The CACHEACT Analyzer control statement (see “CACHEACT” on page 216) is used to gather data for the following reports:

Report	Page	Description
Cache Subsystems Overview Report	page 337	shows an overview of caching activity at the subsystem level This report gives you a general idea of how your cache subsystems are performing, and provides information about where to look for specific performance problems. It consists of a summary line of information for each cache subsystem and for the most active devices (activity by total I/O rate and by DASD I/O rate).
Cache Subsystem Activity Report	page 340	shows detail information about every cache subsystem and an overview of activity on each subsystem’s attached DASD devices For each subsystem, the report contains one page of detailed subsystem information, plus one or more pages of device information consisting of a summary line of information for each device. Additionally, if applicable, there is detail information about RAID rank activity in the subsystem.
Cache Device Activity Report	page 347	contains detailed information about every cache subsystem and every logical DASD device attached to a subsystem For each cache subsystem, this report has one page of detailed subsystem information, and one or more pages of detailed device information, one device per page. The fields in this report are similar to those in the Cache Subsystem Activity report, except that this report shows the data for one particular device rather than for a subsystem. **Note**: This report can be quite voluminous, so BMC Software suggests that you request it only for specific devices or ranges, or when troubleshooting cache subsystem problems.
Ordering and filtering: The reports in this table can be ordered and filtered in different ways (see “CACHEACT” on page 216), so your report might not look exactly like the sample reports shown on the next few pages.		

Cache Subsystems Overview Report

This report is produced by the CACHEACT Analyzer control statement when the REPORT= parameter is either omitted or specified as REPORT=OVERVIEW. Use this report to see an overview of each subsystem's activity or to find out if a particular subsystem is either under- or overutilized.

By default, subsystems are ordered by subsystem ID (SSID on the report), but the order can be changed by the ORDER= parameter. The report can be filtered by cache subsystem model type or subsystem ID.

An example of the Cache Subsystems Overview Report is shown in Figure 29.

Figure 29 Cache Subsystems Overview Report (part 1 of 2)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    CACHE SUBSYSTEMS OVERVIEW REPORT                     RPTSEQ     3 PAGE      5
BMC SOFTWARE, INC.                                          BMC SOFTWARE, INC.                            REPORT DATE: DD MMM YY 16.44
REQD 10 JUN YY 09.50.00  10 JUN YY 10.20.00                    HOUSTON, TX.                              SYSTEM ID: SJSE     Z v.rr.n
ACTL 10 JUN YY 09.50.00  10 JUN YY 10.20.00                                                               REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-5/24/0/0.5

              CTRL        STORAGE(MB)  --I/O RATE-  --CACHE HIT RATE-  -----DASD I/O RATE-----------  ASYNC  -----% HITS------    %
SSID    TYPE  UNIT STATUS CACHE   NVS  CACHE   OFF   READ   DFW   CFW  STAGE DFWBP   ICL   BYP OTHER   RATE  TOTAL  READ WRITE  READ
----  ------- ---- ------ ----- -----  ----- -----  ----- ----- -----  ----- ----- ----- ----- -----  -----  ----- ----- ----- -----
0027  3990-03 5590 ACTIVE 1,024     8    0.4     0    0.4     0     0      0     0     0     0     0      0   99.9  99.9     0 100.0
0140  3990-06 8350 ACTIVE 28.1K   192  640.4     0  564.1  61.1     0    9.3     0   0.2   5.7     0   18.2   97.6  98.4  99.9  90.4
0144  3990-06 8733 ACTIVE 28.1K   192  192.5     0   97.9  94.4     0    0.2     0     0     0     0      0   99.9  99.8 100.0  51.0
0240  3990-06 8573 ACTIVE 28.1K   192  340.0     0  274.7  45.0     0    7.1     0     0  13.5     0   14.4   94.0  97.6  99.3  86.2
A300  2105-20 A333 ACTIVE 4,096   192    8.1     0    2.9   5.1     0      0     0     0     0     0    1.4  100.0  99.9 100.0  36.4
AC03  2105-20 C103 ACTIVE 4,096   192   25.8     0   25.6   0.1     0      0     0     0     0     0    0.1  100.0 100.0 100.0  99.4


PRODUCED BY CMF ANALYZER (v.r.mm)                    CACHE SUBSYSTEMS OVERVIEW REPORT                     RPTSEQ     3 PAGE      6
BMC SOFTWARE, INC.                                          BMC SOFTWARE, INC.                            REPORT DATE: DD MMM YY 16.44
REQD 10 JUN YY 09.50.00  10 JUN YY 10.20.00                    HOUSTON, TX.                               SYSTEM ID: SJSE     Z v.rr.n
ACTL 10 JUN YY 09.50.00  10 JUN YY 10.20.00                                                               REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-5/24/0/0.5

----------------------------------- MOST ACTIVE DEVICES -- IN DESCENDING ORDER BY TOTAL I/O RATE ----------------------------------

 DEV  VOLUME         ----STATUS---     %     I/O   -CACHE HIT RATE-  --------DASD I/O RATE--------  ASYNC  -------% HITS------     %
 NUM  SERIAL  SSID  CACHE  DASDFW    I/O    RATE   READ   DFW   CFW  STAGE DFWBP   ICL   BYP OTHER   RATE  TOTAL   READ  WRITE   READ
----  ------  ----  ------ ------  -----   -----  ----- ----- -----  ----- ----- ----- ----- -----  -----  -----  -----  -----  -----
8710  TSG337  0144  ACTIVE ACTIVE   87.7   168.9   84.5  84.4     0      0     0     0     0     0      0  100.0  100.0  100.0   50.0
8368  BAB319  0140  ACTIVE ACTIVE    6.9    44.4   43.8   0.1     0    0.5     0     0     0     0      0   98.9   98.9  100.0   99.7
83AE  TSG322  0140  ACTIVE ACTIVE    6.9    44.2   43.7   0.4     0      0     0     0     0     0    0.1   99.9   99.9  100.0   99.0
83AC  TSG320  0140  ACTIVE ACTIVE    4.2    27.0   26.6   0.1     0    0.2     0     0     0     0      0   99.3   99.3  100.0   99.5
83A2  TSG310  0140  ACTIVE ACTIVE    3.3    20.9   20.8     0     0    0.1     0     0     0     0      0   99.7   99.7  100.0  100.0
8525  SAZ12I  0240  ACTIVE ACTIVE    6.1    20.6   20.4     0     0    0.2     0     0     0     0      0   99.1   99.1      0  100.0
8712  TSG339  0144  ACTIVE ACTIVE   10.4    20.1   10.1  10.0     0      0     0     0     0     0      0  100.0  100.0  100.0   50.1
8526  SAZ12J  0240  ACTIVE ACTIVE    5.7    19.5   19.4     0     0      0     0     0     0     0      0   99.7   99.7      0  100.0
C11A  SHK026  AC03  ACTIVE ACTIVE   73.5    18.9   18.9     0     0      0     0     0     0     0      0  100.0  100.0      0  100.0
85E9  TSG324  0240  ACTIVE ACTIVE    5.0    17.0   14.3   2.7     0    0.1     0     0     0     0    0.1   99.6   99.6   99.8   84.2
8552  APD282  0240  ACTIVE ACTIVE    4.7    15.9    7.8   7.6     0      0     0     0   0.5     0    0.3   96.9   99.7   99.9   50.6
8330  TSG301  0140  ACTIVE ACTIVE    2.5    15.7   14.2   0.2     0    1.3     0     0     0     0    0.1   91.5   91.5   99.1   98.8
8360  BAB311  0140  ACTIVE ACTIVE    2.2    14.1   13.9   0.2     0    0.1     0     0     0     0      0   99.5   99.4  100.0   98.9
8311  SYM030  0140  ACTIVE ACTIVE    2.1    13.7   11.0   2.7     0      0     0     0     0     0    0.1   99.8   99.9   99.7   80.6
8367  BAB318  0140  ACTIVE ACTIVE    2.0    13.1   12.9   0.1     0    0.1     0     0     0     0    0.1   99.0   99.0   99.2   99.5
8306  SYSP14  0140  ACTIVE ACTIVE    1.9    12.4   11.9   0.5     0      0     0     0     0     0    0.1   99.9   99.9  100.0   96.1
8308  SYSP15  0140  ACTIVE ACTIVE    1.9    12.2   11.9   0.3     0      0     0     0     0     0    0.1   99.7   99.7   99.8   97.9
83A0  TSG308  0140  ACTIVE ACTIVE    1.9    12.1   11.5   0.6     0      0     0     0     0     0    0.2   99.9   99.9  100.0   95.3
```

Figure 29 Cache Subsystems Overview Report (part 2 of 2)

```
--------------------------------- MOST ACTIVE DEVICES -- IN DESCENDING ORDER BY DASD I/O RATE ----------------------------------

 DEV  VOLUME        ----STATUS---     %    I/O   -CACHE HIT RATE-   --------DASD I/O RATE--------  ASYNC  -------% HITS------     %
 NUM  SERIAL  SSID  CACHE  DASDFW    I/O  RATE   READ   DFW   CFW  STAGE DFWBP   ICL   BYP OTHER   RATE  TOTAL  READ WRITE   READ
----  ------  ----  ------ ------  -----  -----  ----- ----- -----  ----- ----- ----- ----- -----  -----  ----- ----- -----  -----
858C  SAEPG1  0240  ACTIVE ACTIVE    2.5    8.4    0.2   0.1     0      0     0     0   8.1     0    1.2    3.3  99.7 100.0   66.5
833B  SAEPG2  0140  ACTIVE ACTIVE    0.4    2.3    0.1     0     0      0     0     0   2.3     0    1.6    2.5  98.1     0  100.0
850F  SHRA06  0240  ACTIVE ACTIVE    0.6    2.1    0.3     0     0      0     0     0   1.8     0      0   16.4  99.3 100.0   96.2
8572  SANWK1  0240  ACTIVE ACTIVE    1.2    4.2    0.8   1.6     0      0     0     0   1.7     0    0.4   58.6  99.5  99.9   34.3
8330  TSG301  0140  ACTIVE ACTIVE    2.5   15.7   14.2   0.2     0    1.3     0     0     0     0    0.1   91.5  91.5  99.1   98.8
851B  BAB303  0240  ACTIVE ACTIVE    2.6    8.9    4.9   2.8     0    0.3     0     0   0.9     0    0.1   85.9  93.9 100.0   65.3
830B  SMFE36  0140  ACTIVE ACTIVE    0.4    2.6    0.1   1.6     0      0     0     0   0.9     0    0.9   63.1  97.8 100.0    3.1
8307  PAGF49  0140  ACTIVE ACTIVE    0.2    1.0    0.1     0     0      0     0     0   0.9     0    1.4    8.3 100.0 100.0   81.6
85C7  BAB367  0240  ACTIVE ACTIVE    1.3    4.5    3.6   0.2     0    0.7     0     0     0     0      0   84.4  83.8  99.4   96.0
85BE  SMS005  0240  ACTIVE ACTIVE    0.4    1.3    0.8     0     0    0.5     0     0     0     0      0   60.6  60.6 100.0  100.0
8552  APD282  0240  ACTIVE ACTIVE    4.7   15.9    7.8   7.6     0      0     0     0   0.5     0    0.3   96.9  99.7  99.9   50.6
8368  BAB319  0140  ACTIVE ACTIVE    6.9   44.4   43.8   0.1     0    0.5     0     0     0     0      0   98.9  98.9 100.0   99.7
85C8  BAB368  0240  ACTIVE ACTIVE    1.5    5.0    4.5   0.1     0    0.4     0     0     0     0      0   91.8  91.6  99.6   97.4
83AF  HSM300  0140  ACTIVE ACTIVE    1.6   10.0    9.3   0.4     0    0.4     0     0     0     0    0.3   96.0  95.9 100.0   96.3
83F3  FAT307  0140  ACTIVE ACTIVE    0.5    3.4    2.9   0.1     0    0.4     0     0     0     0      0   88.2  88.0 100.0   98.1
8369  BAB320  0140  ACTIVE ACTIVE    1.3    8.4    6.9   1.1     0    0.4     0     0     0     0    0.3   95.5  94.8 100.0   86.3
85C6  BAB366  0240  ACTIVE ACTIVE    1.4    4.8    4.4   0.1     0    0.4     0     0     0     0    0.1   92.6  92.5  99.2   98.6
856C  FAT312  0240  ACTIVE ACTIVE    1.5    5.0    4.5   0.1     0    0.3     0     0     0     0    0.1   93.2  93.1 100.0   97.4
```

Cache Subsystems Overview Report field descriptions

Table 32 describes each field that is displayed in the Cache Subsystems Overview Report.

Table 32 Field descriptions for the Cache Subsystems Overview Report (part 1 of 3)

Field	Description
Overview section	
SSID	Subsystem ID—a four-digit hexadecimal identifier that uniquely identifies the cache subsystem
TYPE	cache subsystem type
CTRL UNIT	physical control unit number of the caching subsystem; this number is equal to the lowest device number for the subsystem
STATUS	caching status of the entire subsystem
STORAGE (MB)	storage amounts, in megabytes:
CACHE	amount of storage configured for caching in the subsystem
NVS	amount of Non-Volatile Storage (NVS) storage in the subsystem
I/O RATE	I/O requests per second, by type:
CACHE	number of cacheable I/O requests to cache devices
OFF	number of I/O requests to non-cache devices
CACHE HIT RATE	I/O requests per second that completed without accessing DASD, by type:
READ	number of READ requests
DFW	number of DFW (DASD Fast Write) requests
CFW	number of CFW (Cache Fast Write) requests

Table 32 Field descriptions for the Cache Subsystems Overview Report (part 2 of 3)

Field	Description
DASD I/O RATE	I/O requests per second that accessed DASD, by type:
STAGE	number of normal or sequential requests
DFWBP	number of requests that caused DFW to be bypassed
ICL	number of ICL (Inhibit Cache Load) requests
BYP	number of requests that explicitly bypassed the cache
OTHER	number of requests that were either CFW BYPASS or DFW INHIBIT requests
ASYNC RATE	number of tracks asynchronously written to DASD in order to release space in cache or NVS storage A high rate indicates storage could be overcommitted.
% HITS	I/O requests serviced from the cache, as percentages of all I/O requests:
TOTAL	percentage of READ and WRITE requests serviced from the cache
READ	percentage of READ requests serviced from the cache
WRITE	percentage of WRITE requests serviced from the cache
READ	READ requests as a percentage of all I/O requests
MOST ACTIVE DEVICES sections IN DESCENDING ORDER BY TOTAL I/O RATE and BY DASD I/O RATE	
DEV NUM	four-digit device number that uniquely identifies the DASD device
VOLUME SERIAL	volume serial number
SSID	Subsystem ID—a four-digit hexadecimal identifier that uniquely identifies the cache subsystem
STATUS	cache status of the DASD device:
CACHE	caching status
DASDFW	DASD Fast Write (DFW) status
% I/O	I/O requests to this DASD volume, as a percentage of all I/O requests to the cache subsystem to which it belongs
I/O RATE	I/O requests per second to the volume
CACHE HIT RATE	I/O requests per second that completed without accessing DASD, by request type:
READ	number of READ requests
DFW	number of DFW (DASD Fast Write) requests
CFW	number of CFW (Cache Fast Write) requests
DASD I/O RATE	I/O requests per second that accessed DASD, by request type:
STAGE	normal or sequential requests
DFWBP	requests that caused DFW to be bypassed
ICL	ICL (Inhibit Cache Load) requests
BYP	requests that explicitly bypassed the cache
OTHER	requests that were either CFW BYPASS or DFW INHIBIT requests

Table 32 Field descriptions for the Cache Subsystems Overview Report (part 3 of 3)

Field	Description
ASYNC RATE	number of tracks asynchronously written to the DASD device in order to release space in the cache or NVS storage A high rate indicates that storage could be overcommitted.
% HITS	I/O requests serviced from the cache, as percentages of all I/O requests:
TOTAL	percentage of READ and WRITE requests serviced from the cache
READ	percentage of READ requests serviced from the cache
WRITE	percentage of WRITE requests serviced from the cache
READ	READ requests as a percentage of all I/O requests

Cache Subsystem Activity Report

This report is produced by the CACHEACT Analyzer control statement when REPORT=SUBSYS is specified.

By default, the report is ordered by subsystem ID, but the order can be changed by the ORDER= parameter. The report can be filtered by model, subsystem ID, and a threshold value. Specific subsystems can be selected or excluded.

An example of the Cache Subsystem Activity Report is shown in Figure 30 on page 341.

Figure 30 Cache Subsystem Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  CACHE SUBSYSTEM ACTIVITY REPORT                     RPTSEQ     4 PAGE     28
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                           REPORT DATE: DD MMM YY 16.44
REQD 10 JUN YY 09.50.00  10 JUN YY 10.20.00                  HOUSTON, TX.                              SYSTEM ID: SJSE     Z v.rr.n
ACTL 10 JUN YY 09.50.00  10 JUN YY 10.20.00                                                            REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-5/24/0/0.5

------------ SUBSYSTEM ID: A300   MODEL: 2105-20   CONTROL UNIT: A333   MFR: MCD   PLANT: 01   SERIAL: 00000131278E------------

   SUBSYSTEM STORAGE       NON-VOLATILE STORAGE       S U B S Y S T E M     S T A T U S         I/O TOTALS            HIT %
---------------------  ------------------------   ---------------------------------------   ---------------     ---------------
CONFIGURED      4.1GB  CONFIGURED          192MB   CACHING              - ACTIVE            CACHE     14,555     CACHE      100.0
AVAILABLE       3.4GB  PINNED                  0   NON-VOLATILE STORAGE - ACTIVE            OFFLINE        0     TOTAL      100.0
PINNED              0                              CACHE FAST WRITE     - ACTIVE            NONCACHE       0
OFFLINE             0                              IML DEVICE AVAILABLE - YES               TOTAL     14,555

CACHE I/O    -------------READ I/O REQUESTS------------     ---------------------WRITE I/O REQUESTS---------------------       %
REQUESTS       COUNT     RATE      HITS     RATE   HIT %       COUNT     RATE      FAST     RATE       HITS    RATE   HIT %     READ

NORMAL         4,851      2.7     4,851      2.7   100.0         223      0.1       223      0.1        223     0.1   100.0     95.6
SEQUENTIAL       442      0.2       437      0.2    98.9       9,039      5.0     9,039      5.0      9,039     5.0   100.0      4.7
CFW DATA           0        0         0        0       0           0        0         0        0          0       0       0        0

TOTAL          5,293      2.9     5,288      2.9    99.9       9,262      5.1     9,262      5.1      9,262     5.1   100.0     36.4

----------------------CACHE MISSES----------------------           ------------MISC------------          ------NON-CACHE I/O-----
REQUESTS      READ    RATE   WRITE    RATE  TRACKS    RATE                        COUNT   RATE                        COUNT   RATE

NORMAL           0       0       0       0       3       0           DFW BYPASS         0      0          ICL               0      0
SEQUENTIAL       5       0       0       0     463     0.3           CFW BYPASS         0      0          BYPASS            0      0
CFW DATA         0       0       0       0                           DFW INHIBIT        0      0
                                                                     ASYNC (TRKS)   2,480    1.4          TOTAL             0      0
TOTAL CACHE MISSES        5   RATE       0
TOTAL CACHE HITS     14,555   RATE    8.09                           --------DISK ACTIVITY-------         --HOST ADAPTER ACTIVITY--
                                                                              RESP   BYTES  BYTES                        BYTES  BYTES
---CKD STATISTICS---     ---RECORD CACHING---                                 TIME    /REQ   /SEC                         /REQ   /SEC
                                                                     READ   ------  ------  -----        READ        ------  -----
WRITE             0     READ MISSES       0                          WRITE  ------  ------  -----        WRITE       ------  -----
WRITE HITS        0     WRITE PROM     7,746


PRODUCED BY CMF ANALYZER (v.r.mm)                  CACHE SUBSYSTEM ACTIVITY REPORT                     RPTSEQ     4 PAGE     29
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                           REPORT DATE: DD MMM YY 16.44
REQD 10 JUN YY 09.50.00  10 JUN YY 10.20.00                  HOUSTON, TX.                              SYSTEM ID: SJSE     Z v.rr.n
ACTL 10 JUN YY 09.50.00  10 JUN YY 10.20.00                                                            REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-5/24/0/0.5

------------------------------------ SUBSYSTEM ID: A300   MODEL: 2105-20   CONTROL UNIT: A333 ------------------------------------

 DEV  VOLUME  RRID  ----STATUS---     %     I/O   -CACHE HIT RATE-  --------DASD I/O RATE--------  ASYNC  -------% HITS------     %
 NUM  SERIAL        CACHE  DASDFW    I/O   RATE   READ   DFW   CFW  STAGE DFWBP   ICL   BYP OTHER   RATE  TOTAL   READ  WRITE   READ
---- ------  ----  ------ ------  -----  -----  ----- ----- -----  ----- ----- ----- ----- -----  -----  -----  -----  -----  -----
A300  APD28G        ACTIVE ACTIVE   11.2    0.9      0   0.9     0      0     0     0     0     0    0.9  100.0  100.0  100.0    2.9
A301  SMFT01        ACTIVE ACTIVE    0.4      0      0     0     0      0     0     0     0     0      0  100.0  100.0      0  100.0
A302  SMFT02        ACTIVE ACTIVE    0.4      0      0     0     0      0     0     0     0     0      0  100.0  100.0      0  100.0
. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
A303  SMFT03  0100  ACTIVE ACTIVE   22.5    1.8      0   1.8     0      0     0     0     0     0    0.2  100.0  100.0  100.0    2.2
A313  SMFT28  0100  ACTIVE ACTIVE    0.4      0      0     0     0      0     0     0     0     0      0  100.0  100.0      0  100.0
A36F  --A36F  0101  ACTIVE ACTIVE    0.1      0      0     0     0      0     0     0     0     0      0  100.0  100.0      0  100.0

TOTALS FOR 112 DEVICES:           100.0    8.1    2.9   5.1     0      0     0     0     0     0    2.8  100.0   99.9  100.0   36.4

----------------------------------------------------- RAID RANK ACTIVITY -----------------------------------------------------

                          ---------READ REQUESTS---------   --------WRITE REQUESTS---------
  ID   TYPE      DA   HDD  COUNT  RATE AVG SZ BYTE/S RTIME   COUNT  RATE AVG SZ BYTE/S RTIME  ------ HIGHEST ACTIVITY VOLUMES--------
 ----  -------   --   ---  ----- ----- ------ ------ -----   ----- ----- ------ ------ -----
 0100  RAID-5    11     7     39     0  53.1K  1,151     4     139   0.1  76.3K  5,899    26   FTP109  SMFT03  APD28G  PAGC32  FTP103
 0101  RAID-5    11     7    101   0.1  38.3K  2,150    10       0     0      0      0     0   --A34E  FTP111  FTP112  --A338  --A339

 TOTAL               14    140   0.1  42.4K  3,301     8     139   0.1  76.3K  5,899    26
```

Cache Subsystem Activity Report field descriptions

Table 33 describes fields that are displayed in the Cache Subsystem Activity Report.

Table 33 Field descriptions for the Cache Subsystem Activity Report (part 1 of 6)

Field	Description	
	Subsystem Status section	
SUBSYSTEM ID	physical control unit number of the cache subsystem	
MODEL	hardware model ID	
CONTROL UNIT	physical control unit number of the caching subsystem; this number is equal to the lowest device number for the subsystem	
MFR	manufacturer code of the cache subsystem	
PLANT	plant code of the cache subsystem	
SERIAL	serial number of the cache subsystem	
SUBSYSTEM STORAGE	cache storage status, in megabytes:	
CONFIGURED	amount installed	
AVAILABLE	amount available	
PINNED	amount pinned in cache - DASD failure prevents data destaging	
OFFLINE	amount offline, possibly due to error	
NON-VOLATILE STORAGE	nonvolatile storage (NVS) status, in megabytes:	
CONFIGURED	amount installed	
PINNED	amount pinned in NVS - DASD failure prevents data destaging	
SUBSYSTEM STATUS	caching status of the entire subsystem:	
CACHING	ACTIVE	- caching online and usable
	PENDING ACTIVE	- cache being brought online
	INTERNAL ERROR	- internal error stopped caching
	HOST TERMINATION	- host deactivated caching
	CACHE SUSPENDED	- caching suspended
	PENDING OFFLINE	- cache deactivation is pending
	PENDING OFF FAIL	- cache deactivation failed
	IN MAINTENANCE	- cache disabled for maintenance
	?????	- unknown status
NON-VOLATILE STORAGE	ACTIVE	- NVS caching is online and usable
	INTERNAL ERROR	- internal error stopped NVS cache
	HOST TERMINATION	- host deactivated NVS caching
	DFW SUSPENDED	- DFW inhibited - battery failure
	PENDING OFFLINE	- NVS cache deactivation pending
	IN MAINTENANCE	- NVS disabled for maintenance
	?????	- unknown status
CACHE FAST WRITE	ACTIVE	- CFW is active
	DEACTIVATED	- CFW has been deactivated
IML DEVICE AVAILABLE	YES	- IML device is operational
	NO	- IML device is not operational

Table 33 Field descriptions for the Cache Subsystem Activity Report (part 2 of 6)

Field	Description
I/O TOTALS	number of I/O requests:
CACHE	cacheable I/O requests to cached devices
OFFLINE	I/O requests to non-cached devices
NONCACHE	sum of INHIBIT CACHE LOAD and CACHE BYPASS requests
TOTAL	sum of all I/O requests
HIT %	cache hits, as percentages of all I/O requests:
CACHE	percentage of cacheable I/O requests that were serviced from cache
TOTAL	percentage of all I/O requests that were serviced from cache
Subsystem Details section	
CACHE I/O REQUESTS	specified by the DEFINE EXTENT command; categorizes the way cache is managed, as one of three classes:
NORMAL	use the normal LRU algorithm to make cache space available
SEQUENTIAL	preload of tracks from DASD for sequential I/O requests
CFW DATA	handle WRITE and READ-AFTER-WRITE requests in cache
TOTAL	totals or averages for the three defined classes
READ I/O REQUESTS	I/O requests that had at least one SEARCH or READ command and no WRITE command; only requests to cached devices are counted:
COUNT	total number of SEARCH/READ requests
RATE	SEARCH/READ requests per second
HITS	SEARCH/READ requests handled without accessing DASD
RATE	SEARCH/READ requests handled without accessing DASD, per second
HIT %	SEARCH/READ cache hits, as a percentage of all such requests
WRITE I/O REQUESTS	I/O requests that had at least one WRITE command; only requests to cached devices are counted:
COUNT	total number of WRITE requests
RATE	WRITE requests per second
FAST	CACHE/DASD FAST WRITE (CFW and DFW) requests
RATE	CFW/DFW requests per second
HITS	CFW/DFW requests handled without accessing DASD
RATE	CFW/DFW requests processed without accessing DASD, per second
HIT %	CFW/DFW cache hits, as a percentage of all WRITE requests
% READ	READ requests as a percentage of all READ and WRITE requests

Table 33 Field descriptions for the Cache Subsystem Activity Report (part 3 of 6)

Field	Description
CACHE MISSES	occurs under the following situations: ■ Data to be read or written is not found in the cache. This problem causes the track or tracks containing the data to be read from DASD and placed in the cache. ■ A lack of NVS storage causes the data to be written to DASD immediately. ■ For some reason, DASD FAST WRITE (DFW) is inhibited.
READ	SEARCH / READ requests that accessed DASD
RATE	SEARCH / READ requests that accessed DASD, per second
WRITE	WRITE requests that accessed DASD
RATE	WRITE requests that accessed DASD, per second
TRACKS	number of tracks transferred from DASD to cache
RATE	number of tracks transferred from DASD to cache, per second
MISC	miscellaneous I / O requests
DFW BYPASS	number of DASD FAST WRITE (DFW) requests that might have been written to NVS, except that NVS was over-committed This situation forced the writes to be sent directly to DASD.
COUNT	total number of DFW BYPASS requests
RATE	DFW BYPASS requests per second
CFW BYPASS	number of I / O requests with the CACHE FAST WRITE (CFW) attribute that went directly to DASD because the cache was full No waiting for destaging is done for these requests.
COUNT	total number of CFW BYPASS requests
RATE	CFW BYPASS requests per second
DFW INHIBIT	if DASD FAST WRITE (DFW) is active, the number of WRITE requests that inhibited DFW If DFW is inactive, the number of WRITE requests that directly accessed the DASD is recorded in this field.
COUNT	total number of DFW INHIBIT requests
RATE	DFW INHIBIT requests per second
ASYNC (TRKS)	number of tracks asynchronously transferred from cache to DASD, in order to free space in the cache and the NVS For a duplex pair, this figure is the number of transfers from the primary cache to the secondary. A high number of ASYNC I / Os indicates an over-committed cache or NVS.
COUNT	number of tracks asynchronously transferred from cache to DASD
RATE	number of tracks asynchronously transferred, per second

Table 33 Field descriptions for the Cache Subsystem Activity Report (part 4 of 6)

Field	Description
NON-CACHE I/O	I/O requests that bypass the cache
ICL	inhibit cache load (ICL) Number of I/O requests that inhibited the loading of data into cache, even though the data was not found in the cache.
COUNT	total number of ICL requests
RATE	ICL requests per second
BYPASS	number of I/O requests that explicitly bypassed the cache, whether or not the data was in the cache:
COUNT	total number of BYPASS requests
RATE	BYPASS requests per second
TOTAL	total number of I/O requests that bypassed the cache:
COUNT	number of I/O requests that bypassed the cache
RATE	number of I/O requests that bypassed the cache, per second
TOTAL CACHE MISSES	count and rate of I/O requests that were not serviced from cache
TOTAL CACHE HITS	count and rate of I/O requests that were serviced from cache
DISK ACTIVITY	transfer activity from hard disk to cache and vice versa This information is available only for 2107 and 1750 subsystems; for all other subsystems, dashes will print in place of the missing data.
HOST ADAPTER ACTIVITY	I/O activity of normal, sequential, and CFW read and write requests This information is available only for 2107 and 1750 subsystems; for all other subsystems, dashes will print in place of the missing data.
RESP TIME	response time in milliseconds per read and write request
BYTES/REQ	average number of bytes transferred per read and write requests
BYTES/SEC	average number of bytes transferred per second per read and write requests
CKD STATISTICS	count-Key-Data (CKD) track format requests:
WRITE	number of CKD WRITE requests (also in total WRITE count)
WRITE HITS	CKD WRITE requests that were handled in cache
RECORD CACHING	dynamic record caching counts; normally, only whole tracks are cached:
READ MISSES	count of records to be read that were not in the cache
WRITE PROM	count of records to be written that needed to access DASD
Device Overview section	
DEV NUM	four-digit device number that uniquely identifies the DASD device
VOLUME SERIAL	volume serial number
XTNT POOL *or* **RRID**	ID of the extent pool to which the device belongs (for 2107s and 1750s only) RAID rank identifier (for systems other than 2107s and 1750s)

Table 33 Field descriptions for the Cache Subsystem Activity Report (part 5 of 6)

Field	Description
STATUS	cache status of the DASD device:
CACHE	caching status
DASDFW	DASD Fast Write (DFW) status
% I/O	I/O requests to this DASD volume, as a percentage of all I/O requests to the cache subsystem to which it belongs
I/O RATE	I/O requests per second to the volume
CACHE HIT RATE	I/O requests per second that completed without accessing DASD, by request type:
READ	number of READ requests
DFW	number of DFW (DASD Fast Write) requests
CFW	number of CFW (Cache Fast Write) requests
DASD I/O RATE	I/O requests per second that accessed DASD, by request type:
STAGE	normal or sequential requests
DFWBP	requests that caused DFW to be bypassed
ICL	ICL (Inhibit Cache Load) requests
BYP	requests that explicitly bypassed the cache
OTHER	requests that were either CFW BYPASS or DFW INHIBIT requests
ASYNC RATE	number of tracks asynchronously written to the DASD device in order to release space in the cache or NVS storage A high rate indicates that storage might be over-committed.
% HITS	I/O requests serviced from the cache, as percentages of all I/O requests
TOTAL	percentage of READ and WRITE requests that were serviced from the cache
READ	percentage of READ requests that were serviced from the cache
WRITE	percentage of WRITE requests that were serviced from the cache
READ	READ requests as a percentage of all I/O requests
RAID Rank Activity section (printed only for 2105 subsystems)	
ID	RAID rank identifier
TYPE	RAID rank type
DA	device adapter identifier
HDD	number of hard disk drives (HDD) in the RAID rank
READ REQUESTS	READ I/O requests:
COUNT	number of READ requests
RATE	number of READ requests, per second
AVG SZ	average number of bytes transferred per READ request
BYTE/S	average number of bytes read, per second
RTIME	average response time, in milliseconds, to an HDD READ I/O request

Table 33 Field descriptions for the Cache Subsystem Activity Report (part 6 of 6)

Field	Description
WRITE REQUESTS	WRITE I/O requests:
COUNT	number of WRITE requests
RATE	number of WRITE requests, per second
AVG SZ	average number of bytes that were transferred per WRITE request
BYTE/S	average number of bytes that were written, per second
RTIME	average response time, in milliseconds, to an HDD WRITE I/O request
HIGHEST ACTIVITY VOLUMES	top five volumes sorted in descending order by I/O activity

Cache Device Activity Report

This report is produced by the CACHEACT Analyzer control statement when REPORT=DEVICE is specified.

By default, the report is ordered by subsystem ID/device ID, but the order can be changed by the ORDER= parameter. The report can be filtered by model, subsystem ID, device ID, and a threshold value. Specific subsystems or devices can be selected or excluded.

An example of the Cache Device Activity Report is shown in Figure 31.

Figure 31 Cache Device Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    CACHE DEVICE ACTIVITY REPORT                        RPTSEQ     5 PAGE     35
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                             REPORT DATE: DD MMM YY 16.44
REQD 10 JUN YY 09.50.00  10 JUN YY 10.20.00                  HOUSTON, TX.                                SYSTEM ID: SJSE     Z v.rr.n
ACTL 10 JUN YY 09.50.00  10 JUN YY 10.20.00                                                              REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-5/24/0/0.5
------------- SUBSYSTEM ID: 0027   MODEL: 3990-03   CONTROL UNIT: 5590   MFR: MCD   PLANT: 01   SERIAL: 00000131278E--------------

   SUBSYSTEM STORAGE       NON-VOLATILE STORAGE        S U B S Y S T E M     S T A T U S         I/O TOTALS            HIT %
 ---------------------  ------------------------   ---------------------------------------    ---------------      ---------------
CONFIGURED         1GB  CONFIGURED           8.2MB  CACHING                     - ACTIVE          CACHE       782      CACHE       99.9
AVAILABLE          1GB  PINNED                   0  NON-VOLATILE STORAGE - ACTIVE          OFFLINE       0      TOTAL       99.9
PINNED               0                              CACHE FAST WRITE       - ACTIVE          NONCACHE      0
OFFLINE              0                              IML DEVICE AVAILABLE - YES             TOTAL       782

CACHE I/O    -------------READ I/O REQUESTS-------------     ---------------------WRITE I/O REQUESTS---------------------        %
REQUESTS       COUNT      RATE      HITS      RATE    HIT %      COUNT      RATE      FAST      RATE       HITS     RATE   HIT %      READ

NORMAL           782       0.4       781       0.4    99.9           0         0         0         0          0        0       0     100.0
SEQUENTIAL         0         0         0         0       0           0         0         0         0          0        0       0         0
CFW DATA           0         0         0         0       0           0         0         0         0          0        0       0         0

TOTAL            782       0.4       781       0.4    99.9           0         0         0         0          0        0       0     100.0

(continued on next page)
```

```
-----------------------CACHE MISSES----------------------            ------------MISC-----------          ------NON-CACHE I/O-----
REQUESTS      READ    RATE   WRITE    RATE  TRACKS    RATE                       COUNT   RATE                       COUNT   RATE

NORMAL           1       0       0       0       1       0            DFW BYPASS         0      0          ICL             0      0
SEQUENTIAL       0       0       0       0       0       0            CFW BYPASS         0      0          BYPASS          0      0
CFW DATA         0       0       0       0                            DFW INHIBIT        0      0
                                                                      ASYNC (TRKS)       0      0          TOTAL           0      0
TOTAL CACHE MISSES        1   RATE       0
TOTAL CACHE HITS     14,555   RATE    8.09                            --------DISK ACTIVITY-------        --HOST ADAPTER ACTIVITY--
                                                                               RESP   BYTES  BYTES                       BYTES   BYTES
---CKD STATISTICS---     ---RECORD CACHING---                                  TIME    /REQ   /SEC                        /REQ    /SEC
                                                                      READ   ------  ------  -----        READ      ------   -----
WRITE              0     READ MISSES        0                         WRITE  ------  ------  -----        WRITE     ------   -----
WRITE HITS         0     WRITE PROM         0

PRODUCED BY CMF ANALYZER (v.r.mm)                    CACHE DEVICE ACTIVITY REPORT                         RPTSEQ     5 PAGE    710
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                              REPORT DATE: DD MMM YY 16.44
REQD 10 JUN YY 09.50.00  10 JUN YY 10.20.00                  HOUSTON, TX.                                 SYSTEM ID: SJSE     Z v.rr.n
ACTL 10 JUN YY 09.50.00  10 JUN YY 10.20.00                                                               REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-5/24/0/0.5

-------------- VOLSER: APD28G    DEVICE NUMBER: A300     RAID RANK: 99    SUBSYSTEM ID: A300    MODEL: 2105-20 -------------------

      C A C H E   S T A T U S                         D U P L E X   S T A T U S                      I/O TOTALS          HIT %
-------------------------------------------     -------------------------------------------    ----------------   ---------------
CACHING         - ACTIVE                         DUPLEX PAIR      - NOT ESTABLISHED             CACHE      1,625   CACHE      100.0
DASD FAST WRITE - ACTIVE                         STATUS           - N/A                         OFFLINE        0   TOTAL      100.0
PINNED DATA     - NONE                           DUAL COPY VOLUME - N/A                         TOTAL      1,625

CACHE I/O    -------------READ I/O REQUESTS------------    ---------------------WRITE I/O REQUESTS---------------------       %
REQUESTS       COUNT     RATE      HITS     RATE   HIT %      COUNT     RATE      FAST     RATE      HITS    RATE   HIT %     READ

NORMAL            47        0        47        0   100.0          1        0         1        0         1       0   100.0     97.9
SEQUENTIAL         0        0         0        0       0      1,577      0.9     1,577      0.9     1,577     0.9   100.0        0
CFW DATA           0        0         0        0       0          0        0         0        0         0       0       0        0

TOTAL             47        0        47        0   100.0      1,578      0.9     1,578      0.9     1,578     0.9   100.0      2.9

-----------------------CACHE MISSES----------------------            ------------MISC-----------          ------NON-CACHE I/O-----
REQUESTS      READ    RATE   WRITE    RATE  TRACKS    RATE                       COUNT   RATE                       COUNT   RATE

NORMAL           0       0       0       0       2       0            DFW BYPASS         0      0          ICL             0      0
SEQUENTIAL       0       0       0       0       0       0            CFW BYPASS         0      0          BYPASS          0      0
CFW DATA         0       0       0       0                            DFW INHIBIT        0      0
                                                                      ASYNC (TRKS)   1,583    0.9          TOTAL           0      0
TOTAL CACHE MISSES        0   RATE       0
TOTAL CACHE HITS      1,625   RATE    0.90

---CKD STATISTICS---     ---RECORD CACHING---

WRITE              0     READ MISSES        0
```

Cache Device Activity Report field descriptions

The Cache Device Activity Report consists of the following sections and fields:

- The first page, for each subsystem, is the same as the first page of the Cache Subsystem Activity report. The fields in that report are described in Table 33 on page 342.
- The following pages, for each subsystem, contain device detail information, one page per device, for each device in the subsystem. This page consists of two sections: a status section and an activity section. The activity section uses the same format as the first page of the Cache Subsystem Activity Report, except that it is for activity that is specific to the device.

The fields in the status section of the Cache Device Activity Report are described in Table 34.

Table 34 Field descriptions for the Cache Device Activity Report

Field	Description
	CACHE STATUS
CACHING	**ACTIVE**—Caching is active. I/O requests to the device can be handled in cache. **DEACTIVATION PENDING**—Caching has been deactivated, but not all modified data has been transferred to DASD. **DEACTIVATED**—Caching has been deactivated for this device.
DASD FAST WRITE	**ACTIVE**—DASD FAST WRITE (DFW) I/O requests to the device can be processed. **DEACTIVATE PENDING**—DASD FAST WRITE has been stopped, but not all modified data has been transferred to DASD. **DEACTIVATED**—DASD FAST WRITE has been stopped for this device.
PINNED DATA	**NONE**—No data is pinned for this device. **EXISTS - DFW SUSPENDED**—Pinned data is present, and DASD FAST WRITE has been suspended. **EXISTS - DFW ACTIVE**—Pinned data is present, but DASD FAST WRITE has *not* been suspended.
	DUPLEX STATUS
DUPLEX PAIR	**NOT ESTABLISHED**—The device is not part of a duplex pair. **ACTIVE**—The device is active as the primary or secondary of a duplex pair. **PENDING**—The device is part of a duplex pair being established. **SUSPENDED**—The duplexing has been suspended due to an error in one or the other device of the duplex pair.
STATUS	**N/A**—The device is not part of a duplex pair. **PRIMARY**—The device is the primary device of a duplex pair. **SECONDARY**—The device is the secondary device of a duplex pair.
DUAL COPY VOLUME	**N/A**—The device is not part of a duplex pair. **nnnn**—Indicates the device address of the other device that is part of the duplex pair.
	I/O TOTALS
CACHE	This value is the number of cacheable I/O requests to the device.
OFFLINE	This value is the number of noncacheable I/O requests to the device.
TOTAL	This value is the number of I/O requests to the device.
	HIT %
CACHE	This value is the number of cache hits, as percentages of all cacheable I/O requests.
TOTAL	This value is the number of cache hits, as percentages of all I/O requests.

Channel Path Activity Report

The Channel Path Activity Report, in Figure 32 on page 350, shows channel path busy. It is produced by using the CHANNEL Analyzer control statement. The data is obtained by using CHANNEL Extractor control statements (see “CHANNEL” on page 134).

This report can have up to three major sections:

- **Channel Profile**

  This section shows specific information for each channel path.

- **Overview of DCM Channels**

  This section shows summary information for each channel type that is managed by DCM.

- **Hipersockets Channels**

  This section shows message rates and sizes for entire complex.

Figure 32 Channel Path Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  CHANNEL PATH ACTIVITY REPORT  (2064-00)            RPTSEQ    15 PAGE    105
BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                           SYSTEM ID: SJSE     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/28/25.1K/7  70-1/28/25.1K/7  73-1/28/0/7

----------------------------------------------------- OVERVIEW OF DCM CHANNELS -----------------------------------------------------
         CHANNEL                   UTILIZATION %     READ(MB/SEC)   WRITE(MB/SEC)
      GROUP  G    NO            PART  TOTAL    BUS      PART  TOTAL     PART  TOTAL
     CNC?M        2             0.01   0.02
CCSID=00----------------------------------------------------- CHANNEL PROFILE -----------------------------------------------------
 ID  CHANNEL G   SHR  HISTORY       UTILIZATION %     READ(MB/SEC)   WRITE(MB/SEC)            CHANNEL PATH
        TYPE                    PART  TOTAL    BUS      PART  TOTAL     PART  TOTAL             DESCRIPTION
 A6  CNC_S       YES            0.10   0.39                                             ESCON SWITCHED POINT TO POINT
 A7  CNC_S       YES            0.10   0.36                                             ESCON SWITCHED POINT TO POINT
 A8  CNC_S       YES            0.01   2.05                                             ESCON SWITCHED POINT TO POINT
 AA  CNC_S       YES            0.02   0.07                                             ESCON SWITCHED POINT TO POINT
 AB  CNC_S       YES            0.01   0.05                                             ESCON SWITCHED POINT TO POINT
 AE  CNC_S       YES            0.01   0.04                                             ESCON SWITCHED POINT TO POINT
 B0  CNC_S       YES            0.02   2.05                                             ESCON SWITCHED POINT TO POINT
 B2  CNC_S       YES            0.03   0.08                                             ESCON SWITCHED POINT TO POINT
 B8  CNC_S       YES            0.03   0.23                                             ESCON SWITCHED POINT TO POINT
 BD  CNC_S       YES            0.76   2.72                                             ESCON SWITCHED POINT TO POINT
 C6  CNC_P       YES            2.65   2.65                                             ESCON POINT TO POINT
 F8  OSD         YES            0.21   2.78   8.61      0.00   0.05     0.00   0.04     OSA DIRECT EXPRESS
 F9  OSD         YES            2.69   2.71   8.42      0.00   0.00     0.00   0.00     OSA DIRECT EXPRESS
 FD  CBP         YES            ----   0.00                                             INTEGRATED CLUSTER BUS PEER
 FE  CBP         YES            ----   0.00                                             INTEGRATED CLUSTER BUS PEER
-------------------------------------------------------- HIPERSOCKETS CHANNELS ---------------------------------------------------
 ID  CHANNEL G   SHR         --------------------- MESSAGES --------------------   ------ FAILED/SEC -------
        TYPE                 ---BYTES/SEC----  ----MSGS/SEC----  --AVG MSG SIZE--     SEND  ----RECEIVE-----
                                PART     TOTAL     PART     TOTAL     PART    TOTAL     PART      PART    TOTAL
 F7  IQD         YES           77.46     79.74        0         0      ---   143.4K        0         0        0
```

Channel Path Activity Report field descriptions

Table 35 describes each field in the Channel Path Activity Report.

Table 35 Field descriptions for the Channel Path Activity Report (part 1 of 2)

Field	Description
	Overview of DCM Channels Section
CHANNEL GROUP G NO	a summary line with the average values for all channels in this group; shown for each channel type that is managed by DCM G indicates the generation and is used to differentiate between channels of the same channel type, when one has significant differences from the other. The number of channels of the group is given in column NO.
	Channel Profile Section
CCSID	channel subsystem ID This field is available only when the processor has multiple channel subsystems.
ID	hexadecimal channel path ID
CHANNEL TYPE	type of channel A longer description of channel type is on the CHANNEL PATH DESCRIPTION column. If the last character is M, the channel is managed by Dynamic Channel Path Management (DCM) (for example, CNCSM is a managed channel of type CNC_S).
SHR	indication of whether the channel path is shared between logical partition This field contains YES if the channel path is shared; it is blank if the channel path is not shared.
STATUS	indication of whether the channel was offline or online at the end of the reporting interval This field is blank if online, it contains OFFLINE if offline.
HISTORY	indicates any changes to the status that occurred to the channel path The following list shows the values that can be displayed: **V** varied online or offline **M** modified **A** added **D** deleted
UTILIZATION % - PART	percent of time that a shared channel was busy for the home partition For FICON channels, this figure is a measure of how busy the channel processor was. For an unshared channel, the value is identical to that of the UTILIZATION % - TOTAL.

Table 35 Field descriptions for the Channel Path Activity Report (part 2 of 2)

Field	Description
UTILIZATION % - TOTAL	percent of time that a channel was busy for the whole complex For FICON channels, this figure is a measure of how busy the channel processor was. For an unshared channel, the value is identical to that of the UTILIZATION % - PART.
UTILIZATION % - BUS	utilization of the internal PCI bus of FICON channels This value is the percentage of the cycles per second used for I/O operations over the maximum internal bus cycles per second.
READ (MB/SEC) - PART	rate per second of megabytes read using the channel for the home partition **Note**: This column is applicable only to FICON channels.
READ (MB/SEC) - TOTAL	rate per second of megabytes read using the channel for the entire complex **Note**: This column is applicable only to FICON channels.
WRITE (MB/SEC) - PART	rate per second of megabytes written using the channel for the home partition **Note**: This column is applicable only to FICON channels.
WRITE (MB/SEC) - TOTAL	rate per second of megabytes written using the channel for the entire complex **Note**: This column is applicable only to FICON channels.
CHANNEL PATH DESCRIPTION	long description of channel type; also displayed by the system command D M=CHP
Hipersockets Channels Section	
MESSAGES **BYTES/SEC**	**PART**—rate of messages sent by programs for the partition **TOTAL**—rate of messages sent by programs for the entire complex
MESSAGES **MSGS/SEC**	**PART**—rate of messages sent by programs for the partition
MESSAGES **AVG MSG SIZE**	**PART**—average size of messages sent for the partition **TOTAL**—average size of messages sent for the entire complex
FAILED/SEC **SEND**	**PART**—rate of messages sent by the partition that failed, excluding the attempts failed due to unavailable buffers in the receiving partition
FAILED/SEC **RECEIVE**	**PART**—rate of messages sent by the partition that failed due to unavailable buffers in the receiving partition **TOTAL—**rate of messages sent by the entire complex that failed due to unavailable buffers in the receiving partition.

CMF Record Statistics Report

The CMF Record Statistics Report can be useful to verify the validity of the input data. The CMF Record Statistics Report is produced by using the CMFSTAT Analyzer control statement (see “CMFSTAT” on page 222). No Extractor control statement is required to collect data for this report.

The CMF Record Statistics Report records the following data for each record type accepted for processing:

- sum of the interval durations
- number of records
- sum of samples for all intervals
- earliest date and time interval
- latest date and time interval

An example of the CMF Record Statistics Report is shown in Figure 33.

Figure 33 CMF Record Statistics Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     CMF RECORD STATISTICS REPORT                          RPTSEQ    10 PAGE     68
BMC SOFTWARE, INC.                                             XYZ COMPANY                                  REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 08.17.19  10 JUN YY 16.00.00              WORLDWIDE HEADQUARTERS                             SYSTEM ID: SJSE     Z v.rr.n

    CMF/RMF/SMF          COUNT OF          TOTAL    TOTAL RECORD    START DATE  START TIME   END DATE    END TIME
     RECORD ID            RECORDS        SAMPLES      TIME VALUE    OF RECORD   OF RECORD    OF RECORD   OF RECORD
    MAJOR MINOR                                     DDD.HH.MM.SS    DD.MMM.YY   HH.MM.SS     DD.MMM.YY   HH.MM.SS
       70-01                   28         25,128        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       71-01                   28            420        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       72-03                1,176              0        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       73-01                   28              0        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       74-01                  112        100,380        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       74-02                   28            840        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       74-03                   28         25,128        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       74-04                   28          2,506        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       74-05                  308              0        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       74-06                   28              0        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       75-01                  112        100,512        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       76-01                   28         25,108        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       77-01                   28              0        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       78-02                   28          5,012        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
       78-03                   28          8,390        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-00                   28             28        07.00.04    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-01                   28         25,128        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-02                   28         25,128        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-03                   28            420        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-05                  112        100,380        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-06                   28         25,128        07.00.01    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-09                   28         25,128        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-11                  224         24,907        07.00.00    10 JUN YY   09.00.00     10 JUN YY   15.59.59
      240-12                    1              1        00.00.00    10 JUN YY   08.17.19     10 JUN YY   08.17.19
      240-13                  393         25,128        07.00.00    10 JUN YY   09.00.00     10 JUN YY   15.59.59
      240-16                    5              0        00.00.00    10 JUN YY   08.17.33     10 JUN YY   08.17.33
      240-19                  418          1,680        07.00.00    10 JUN YY   09.00.00     10 JUN YY   15.59.59
      240-20                   28         25,156        07.00.01    10 JUN YY   09.00.00     10 JUN YY   16.00.00
      240-21                   28              0        07.00.00    10 JUN YY   09.00.00     10 JUN YY   16.00.00
```

NOTE

The counts in this report represent what was selected by specifying the DATETIME range.

CMF Record Statistics Report field descriptions

Table 36 describes each field in the CMF Record Statistics Report.

Table 36 Field descriptions for the CMF Record Statistics Report

Field	Description
CMF/RMF RECORD ID MAJOR-MINOR	SMF record ID and subtype
COUNT OF RECORDS	count of record types encountered in the DATETIME range Multiple records for the same record type with the same date and time are counted as one record; for example, SMF type 72 workload records.
TOTAL SAMPLES	total number of samples that were encountered in the DATETIME range
TOTAL RECORD TIME VALUE	total elapsed measurement interval
START DATE OF RECORD	start date of the first record
START TIME OF RECORD	start time of the first record
END DATE OF RECORD	end date of the last record
END TIME OF RECORD	end time of the last record

CMF Summary Report

The CMF Summary Report shows an overview of system performance. The same information that is produced in the RMF Summary Report can appear in this report, as well as additional CMF MONITOR information that RMF does not report.

The CMF Summary Report is produced by using the CMFSUM Analyzer control statement (see "CMFSUM" on page 223). The data is obtained by using the following Extractor control statements (see Chapter 6, "Extractor control statements," for more information):

- CHANNEL
- CPU
- DEVICE
- IOQ
- PAGING
- TSODATA
- VSMDATA

An example of the CMF Summary Report is shown in Figure 34 on page 356.

Figure 34 CMF Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        CMF SUMMARY REPORT                                  RPTSEQ    11 PAGE     69
BMC SOFTWARE, INC.                                           XYZ COMPANY                                     REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00 10 JUN YY 16.00.00             WORLDWIDE HEADQUARTERS                               SYSTEM ID: SJSE     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-6/28/25.1K/7  240-20/28/25.2K/7  70-1/28/25.1K/7  71-1/28/420/7  72-3/1,176/0/7
73-1/28/0/7  74-1/112/25.1K/7  78-2/28/5,012/7  78-3/28/8,390/7
 DATE/     INTERVAL     CPU   DASD    DASD    TAPE  BTCH  BTCH   TSO   TSO   STC   STC  APPC  APPC  OMVS  OMVS  SWAP  DEMAND  SERVICE
  TIME     HH.MM.SS    BUSY   RESP    RATE    RATE   MAX   AVG   MAX   AVG   MAX   AVG   MAX   AVG   MAX   AVG  RATE  PAGING     RATE

10 JUN YY 07.00.00      6.2      9    72.3     2.1     5     0    31    25   125   117     0     0     1     1  0.01    1.67   290706

09.00.00  00.15.00      5.4      8    92.0    59.7     0     0    17    16   106   106     0     0     1     1  0.00    0.18   129951
09.15.00  00.15.00      4.7      8    52.0     0.0     0     0    20    19   107   106     0     0     1     1  0.00    0.03    95872
09.30.00  00.15.00      4.8      7    58.7     0.0     2     0    21    21   109   107     0     0     1     1  0.01    0.33   103837
09.45.00  00.15.00      4.6      9    51.2     0.0     0     0    19    19   106   104     0     0     1     1  0.00    0.04    99501
10.00.00  00.15.00      4.6     10    47.8     0.0     1     0    19    18   106   106     0     0     1     1  0.00    0.01    95111
10.15.00  00.15.00      5.5      7    72.5     0.0     0     0    24    21   113   109     0     0     1     1  0.01    0.07   113826
10.30.00  00.15.00      4.6     10    51.7     0.0     0     0    26    26   113   113     0     0     1     1  0.00    0.04    94459
10.45.00  00.15.00      6.2      9    89.3     0.0     1     0    29    28   120   116     0     0     1     1  0.01    0.29   139075
11.00.00  00.15.00      5.1     13    55.3     0.1     1     0    29    29   119   119     0     0     1     1  0.00    0.21   109092
11.15.00  00.15.00      4.7      7    57.1     0.0     0     0    30    29   121   119     0     0     1     1  0.00    0.12    94018
11.30.00  00.15.00      6.8      6    95.9     0.0     1     0    31    31   123   122     0     0     1     1  0.00    0.16   174335
11.45.00  00.15.00      5.7      9    49.7     0.0     0     0    30    30   123   123     0     0     1     1  0.00    0.02   124648
12.00.00  00.15.00      5.6      7    68.3     0.0     1     0    30    30   123   123     0     0     1     1  0.00    0.12   109128
12.15.00  00.15.00      5.7     10    57.5     0.0     1     0    30    30   123   122     0     0     1     1  0.00    0.31   114212
12.30.00  00.15.00      7.8     13    84.4     0.0     1     0    30    29   125   123     0     0     1     1  0.07   12.23   650506
12.45.00  00.15.00      6.6     16    58.8     0.0     0     0    29    29   121   120     0     0     1     1  0.00    8.68   356857
13.00.00  00.15.00      5.5     13    37.7     0.0     0     0    29    29   119   119     0     0     1     1  0.00    0.79    98832
13.15.00  00.15.00      7.0      7    56.5     0.0     0     0    30    29   119   119     0     0     1     1  0.00    1.35   638135
13.30.00  00.15.00      8.7      6   150.1     0.0     4     0    29    28   120   119     0     0     1     1  0.00    2.68   452641
13.45.00  00.15.00      8.6      9   135.1     0.0     1     0    29    28   123   121     0     0     1     1  0.05    9.82   725672
14.00.00  00.15.00      7.4      7    85.7     0.0     1     0    28    26   122   122     0     0     1     1  0.00    2.34   423713
14.15.00  00.15.00      6.9      6    90.8     0.3     1     0    26    25   122   122     0     0     1     1  0.00    0.90   542137
14.30.00  00.15.00      8.1     10    78.2     0.0     1     0    27    26   122   121     0     0     1     1  0.00    2.79   802765
14.45.00  00.15.00      7.9      7    72.1     0.0     1     0    26    25   120   120     0     0     1     1  0.00    0.30   777218
15.00.00  00.15.00      5.7     11    45.8     0.0     1     0    25    24   120   120     0     0     1     1  0.00    1.64   299051
15.15.00  00.15.00      8.4     14   114.9     0.0     5     2    23    22   121   118     0     0     1     1  0.00    0.59   286286
15.30.00  00.15.00      5.1     13    33.7     0.0     1     0    21    21   119   119     0     0     1     1  0.00    0.12   239000
15.45.00  00.15.00      5.6      8    83.0     0.0     2     0    22    21   121   120     0     0     1     1  0.00    0.45   249893
```

CMF Summary Report field descriptions

Table 37 on page 357 describes each field in the CMF Summary Report. (Abbreviated field names are shown in parentheses.)

Fields that appear in both CMF MONITOR and RMF are the default fields. Additional fields can be specified with the MEASURE parameter of the CMFSUM Analyzer control statement.

See “CMFSUM” on page 223 for more information about including these fields in the CMF Summary Report.

These fields appear only in CMF MONITOR reports; they do not appear in RMF reports.

Table 37 Field descriptions for the CMF Summary Report (part 1 of 5)

Field	Measure name	Description
DATE/TIME	default	day for the report data and start time of each interval being summarized Date is specified as *ddmmmyy* and time as *hh:mm:ss*.
INTERVAL HH.MM.SS	INTERVAL (*default*)	actual length of each interval included in the CMF Summary Report This field cannot report an interval length that spans more than 24 hours. No value is reported if MEASURE=INTERVAL is defined on the Analyzer CMFSUM report control statement, and INTERVAL= parameter specifies more than a 24 hour period, such as WEEKLY, MONTHLY, QTRLY, SEMIANNL, or FOREVER.
CHANNEL PATH BUSY (CP BUSY)	CHPBUSY	percent of time a channel was busy for the whole complex For FICON channels, this figure is a measure of how busy the channel processor was.
CPU BUSY	CPUBUSY (*default*)	**LPAR mode—**LPAR CPU busy percentage of standard CPs during the specified interval The formula is percentage of busy time / online time. **Basic mode or under VM—**average CPU busy percentage of standard CPs during the specified interval The formula is (interval - wait time) / interval. In this mode, this measure has the same meaning as CPUBZMVS.
MVS CPU BUSY	CPUBZMVS	MVS CPU busy percentage of the standard CPs during the specified interval The formula is as follows: **LPAR mode—**(online time - wait time) / online time **Basic mode or under VM—**(interval - wait time) / interval
ZAAP BUSY	ZAAPBUSY	LPAR CPU busy percentage of zAAPs (zSeries Application Assist Processors) during the specified interval The formula is percentage of busy time / online time.
MVS ZAAP BUSY	ZAAPBMVS	MVS CPU busy percentage of zAAPs (zSeries Application Assist Processors) during the specified interval The formula is as follows: **LPAR mode—**(online time - wait time) / online time **Under VM—**(interval - wait time) / interval

Table 37 Field descriptions for the CMF Summary Report (part 2 of 5)

Field	Measure name	Description
ZIIP BUSY	ZIIPBUSY	LPAR CPU busy percentage of zIIPs (zSeries Integrated Information Processors) during the specified interval The formula is percentage of busy time / online time.
MVS ZIIP BUSY	ZIIPBMVS	MVS CPU busy percentage of zIIPs (zSeries Integrated Information Processors) during the specified interval The formula is as follows: **LPAR mode—**(online time - wait time) / online time **Under VM—**(interval - wait time) / interval
LPAR DISPATCH PERCENTAGE (LPAR DISP)	LPARDISP	average dispatch percentage during the specified interval for all standard CPs for the logical partition on which records were extracted The formula is percentage of dispatch time / online time.
DASD RESPONSE TIME (DASD RESP)	DASDRESP (*default*)	average number of milliseconds required to complete an I/O request for all direct access storage devices (DASD) active during the specified interval; the range is from 0 to 99999
DASD RATE	DASDRATE (*default*)	rate of activity per second for all direct access storage devices (DASD) during a specified interval This measure is the accumulation of rate per second at which SSCH instructions were completed successfully for all DASDs. The range is from 0.0 to 999.9.
TAPE RATE	TAPERATE (*default*)	rate of activity per second for all magnetic tape devices during a specified interval This measure is an accumulation of rate per second at which SSCH instructions completed successfully for all magnetic tape devices. The range is from 0.0 to 999.9.
BATCH MAXIMIUM (BTCH MAX)	BATCHMAX (*default*)	maximum number of batch jobs active during a specified interval; the range is from 0 to 999
BATCH AVERAGE (BTCH AVG)	BATCHAVG (*default*)	average number of batch jobs active during a specified interval; the range is from 0 to 999
TSO MAXIMUM (TSO MAX)	TSOMAX (*default*)	maximum number of TSO sessions active during a specified interval; the range is from 0 to 999
TSO AVERAGE (TSO AVG)	TSOAVG (*default*)	average number of TSO sessions active during a specified interval; the range is from 0 to 999
STC MAXIMUM (STC MAX)	STCMAX (*default*)	maximum number of Started Tasks and mount tasks active during a specified interval; the range is from 0 to 999
STC AVERAGE (STC AVG)	STCAVG (*default*)	average number of Started Tasks and mount tasks active during a specified interval; the range is from 0 to 999
APPC MAXIMUM (APPC MAX)	APPCMAX (*default*)	maximum number of APPC address spaces active during a specified interval; the range is from 0 to 999

Table 37 Field descriptions for the CMF Summary Report (part 3 of 5)

Field	Measure name	Description
APPC AVERAGE (APPC AVG)	APPCAVG (*default*)	average number of APPC address spaces active during a specified interval; the range is from 0 to 999
SWAP RATE	SWAPRATE (*default*)	average rate per second of all swaps during a specified interval; the range is from 0.00 to 99.99
SWAP PAGE RATE (SWAP PAGES)	SWAPPAGS	rate per second of all swap-in and swap-out pages during a specified interval; the range is from 0.00 to 99.99
DEMAND PAGING	DPAGING (*default*)	rate per second of all demand paging requests during a specified interval The value reported is total number of non IO, non swap page-ins divided by the interval. The range is 0.00 to 99.99.
SERVICE RATE	SRVRATE (*default*)	rate per second at which I/O service, CPU service, and central storage service units were consumed during a specified interval; the range is 0 to 9999999
TRANSACTION RATE (TRANS RATE)	TRANRATE (*default*)	average number of transactions that ended during each second of a specified interval; the range is 0.000 to 99.999
OMVS AVERAGE (OMVS AVG)	OMVSAVG (*default*)	average number of OMVS address spaces active during a specified interval; the range is 0 to 999
OMVS MAXIMUM (OMVS MAX)	OMVSMAX (*default*)	maximum number of OMVS address spaces active during a specified interval; the range is 0 to 999
SPOOL UTILIZATION	SPOOLUTL	average JES spool space utilization expressed as a percentage Percentage is carried to one decimal place, in the format *nn.n*.
TSO TRANSACTION RATE (TSO TRNS)	TSOTRANS	average number of TSO transactions that ended during each second of a specified interval; the range is 0.00 to 99.99
CHANNEL PATH UTILIZATION RATE (CP RATE)	CHPUTIL	average rate per second at which I/Os used a channel path for all channel paths during a specified interval The value is calculated by dividing the path taken count by the interval; the range is 0.00 to 999.99.
AVERAGE CSA ALLOCATED (CSA AVG)	CSALLOC	average CSA allocated below 16 MB expressed as a percentage of the total CSA available below the 16-MB line Percentage is carried to one decimal place, in the format *nn.n*.
AVERAGE SQA ALLOCATED (SQA AVG)	SQALLOC	average SQA allocated below 16 MB expressed as a percentage of the total SQA available below the 16-MB line Percentage is carried to one decimal place, in the format *nn.n*.

Table 37 Field descriptions for the CMF Summary Report (part 4 of 5)

Field	Measure name	Description
AVERAGE ECSA ALLOCATED (ECSA AVG)	ECSALLOC	average CSA allocated above 16 MB expressed as a percentage of the total CSA available above the 1-MB line Percentage is carried to one decimal place, in the format *nn.n*.
AVERAGE ESQA ALLOCATED (ESQA AVG)	ESQALLOC	average SQA allocated above 16 MB expressed as a percentage of the total SQA available above the 16-MB line Percentage is carried to one decimal place, in the format *nn.n*.
AVERAGE READY QUEUE (READY AVG)	AVGREADY	average ready queue including jobs in and ready, out and ready, and logical out and ready during a specified interval; the range is 0 to 999
MIGRATION AGE (MIGR AGE)	MIGRAGE	average number of seconds an unreferenced page remained in expanded storage before being migrated to auxiliary storage during a specified interval; the range is 0.0 to 999999.9
MIGRATION RATE (MIGR RATE)	MIGRATE	average number of pages per second migrated from expanded storage to auxiliary storage duringa specified interval; the range is 0.00 to 999.99
HIGH UNREFERENCED INTERVAL COUNT (HIGH UIC)	HIGHUIC	High Unreferenced Interval Count (UIC) in seconds (time that a central storage frame has not been referenced) On z/OS systems 1.8 or later, the range is 1 to 65535; on z/OS systems 1.7 or earlier, the range is 1 to 2540.
AVERAGE EXCPs RATE (EXCP RATE)	EXCPRATE	average number of EXCPs per second during a specified interval; the range is 0 to 9999999
AVERAGE CAPTURE RATIO (CAP RATIO)	CAPRATIO	average capture ratio during a specified interval Capture ratio is an indication of the percentage of the CPU resources that are actually directed toward the processing of workloads. The ratio, expressed as a percentage, is calculated by determining the total amount of CPU time used to process workloads, and dividing it by the actual CPU time consumed. Percentage is carried to one decimal place, in the format 99.9.
AVERAGE FIXED FRAMES (FIX FRAME)	FIXFRAME	average number of fixed frames during a specified interval, expressed as a percentage of the average total number of frames Percentage is carried to one decimal place, in the format 99.9.
AVERAGE AVAILABLE FRAMES QUEUE (AVAIL FRAME)	AFQUEUE	average number of page frames available in central storage during a specified interval; the range is 0 to 999999

Table 37 Field descriptions for the CMF Summary Report (part 5 of 5)

Field	Measure name	Description
AVERAGE HIPERSPACE FRAMES (H.S. FRAME)	HSFRAME	average number of hiperspace page frames in expanded storage during a specified interval, expressed as a percentage of the total number of expanded storage frames online Percentage is carried to one decimal place, in the format 99.9.
AVERAGE VIO FRAMES (VIO FRAME)	VIOFRAME	average number of VIO page frames in expanded storage during a specified interval, expressed as a percentage of the total number of expanded storage frames online Percentage is carried to one decimal place, in the format 99.9.
EXPANDED STORAGE FRAMES (ESTOR FRAME)	ESFRAME	average number of available expanded storage frames during a specified interval, expressed as a percentage of the total number of expanded storage frames online Percentage is carried to one decimal place, in the format 99.9.
CPU SERVICE RATE (CPU SERVICE)	CPUSERV	rate per second at which TCB and SRB service units were consumed during specified interval; the range is 0 to 9999999
MSO SERVICE RATE (MSO SERVICE)	MSOSERV	rate per second at which central storage service units were consumed during a specified interval; the range is 0 to 9999999
I/O SERVICE RATE (I/O SERVICE)	IOSERV	rate per second at which I/O service units were consumed during a specified interval; the range is 0 to 9999999
EXPANDED STORAGE PAGE RATE (ESTOR PG RATE)	EPGRATE	rate per second at which pages moved from central storage to expanded storage during a specified interval; the range is 0 to 9999
SUMMARY	(*default*)	average value for all intervals for each field

Collection Phase Log

The Collection Phase Log immediately follows the Report Table of Contents, which is always the first page produced during an Analyzer job.

This log is divided into six parts:

- "Control Card Log"
- "Extraction Characteristics Report" on page 363
- "System Resources Manager Constants Report" on page 366
- "RMF/CMF Input Record Type Counts Report" on page 369
- "Data Distribution and DATETIME Chart" on page 371

Control Card Log

The Control Card Log lists all Analyzer control statements that are used for the report run. No Analyzer or Extractor control statements are required to produce the Control Card Log. It is produced automatically as part of the preliminary reporting information that CMF MONITOR supplies with each batch report job. (See “Preliminary reporting information” on page 330 for more information.)

An example of the Control Card Log is shown in Figure 35.

Figure 35 Control Card Log

```
 PRODUCED BY CMF ANALYZER (v.r.mm)                         CONTROL CARD LOG                                  RPTSEQ     1 PAGE      1
 BMC SOFTWARE, INC.                                                                                          REPORT DATE: DD MMM YY 13.40
   HEADERS  TITLE='XYZ COMPANY',LOCATION='WORLDWIDE HEADQUARTERS'
   REPORTS  RPTGROUP=RPTSTMT,DDGROUP=RPTSTMT
   AUXSTOR
   CACHE    TYPE=ALL
   CFACT
   CMFSTAT
   CMFSUM   - MEASURE=ALL
   COMMSTOR
   COMMSTOR REPORT=DETAIL
   CPU
   DASD
   VOLSER   FAT900,TSG900
   DEVACT
   ENQUEUE
 * EXCEPTS
   GRAPH    TYPE=TRACE,MEASURE=(CCVUTILP)
 * GRAPH    TYPE=TRACE,MEASURE=(ASMIORQC,ASMIORQR,CCVCPUCT,CCVUTILP,
 *          DMDTCMPL,DMDTRUA,MCVSTCRI,MCVMGAGE,RCEAFC,RCETOTFX,
 *          RCVPTR,RCVUICA,RMCATISC,SMCARCWT)
   IOQ
   LINKPACK
   OMVS
   PERFSUM
   PRSM
   SHARDEV
   SRM
   STORAGE
   TSOPERF
   TSOUSER
   VIRTSTOR
   WLMGL    TYPE=ALL
   XCF      TYPE=BOTH,ORDER=(GROUP,MEMBER,SYSTEM)
```

- Analyzer control statements and parameters are explained in Chapter 7, “Analyzer control statements.”
- Analyzer syntax error messages are explained in CMF MONITOR messages. (See “Related documentation” on page 19 for more information.)

Control Card Log field descriptions

Descriptions of the fields in the Control Card Log are listed in Table 38.

Table 38 Field descriptions for the Control Card Log

Field	Description
*** GENERAL CONTROL STATEMENTS**	Each general control statement used to specify global data collection and reporting parameters is listed in this area of the Control Card Log.
*** REPORT CONTROL STATEMENTS**	Each report control statement used to specify a desired report is listed in this area of the Control Card Log.
ERROR MESSAGES	Any errors encountered with the control card syntax are listed in this area of the Control Card Log.

Extraction Characteristics Report

The Extraction Characteristics Report is a Collection Phase Log report (see "Preliminary reporting information" on page 330 for more information) that lists the parameters used by the Extractor each time it is run. Since the Analyzer produces combined reports for multiple extractions, you must know the sample rates and types. Mixing Extractor data that contains unlike parameters can produce distorted results.

This report is produced only for extractions that occur within the overall DATETIME range. If no DATETIME is specified, the Analyzer produces one report for each extraction run. This report is not produced if RPTS=INTERVAL or RPTS=DAILY is specified on the SHIFT control statement.

No Analyzer or Extractor control statements are required to produce the Extraction Characteristics Report. It is produced automatically as part of the preliminary reporting information that CMF MONITOR supplies with each batch report job.

The Extractor characteristics that produced the data set read by the Analyzer are shown. One page is generated each time the Extractor is run.

An example of the Extraction Characteristics Report is shown in Figure 36.

Figure 36 Extraction Characteristics Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        COLLECTION PHASE LOG                            RPTSEQ     2 PAGE      3
BMC SOFTWARE, INC.                                                                                       REPORT DATE: DD MMM YY 13.40
                                                      EXTRACTION CHARACTERISTICS
                                         (IN EFFECT FROM 22 JUN YY 05.45.00 TO 22 JUN YY 10.15.00)
 EXTRACTOR VERSION NUMBER (v.r.mm)
 SU'S IN EFFECT =  0,1,4,5,6,7,8,9,10,11,13,15,16,17,21,24,25,27,29,30,32,33,37,39,42,43,47,
                   48,50,51,55,56,57,58,60,61,63,64,65,72,74,7
  SAMPLER NAME          SAMPLE RATE      SAMPLE COUNTS      SAMPLE COUNTS      SAMPLE COUNTS
                         IN MSECS          UNDER DIE          UNDER SRB          UNDER SRM
  ------------          -----------      -------------      -------------      -------------

      ASMS                 2,000             17,843             17,843                  0
      CA3H                 1,000             35,658                  0                  0
      CA5H                 1,000             35,658                  0                  0
      CA6H                 1,000                  0                  0                  0
      CFTS                     0              3,569                  0                  0
      CHNS                     0                 41                  0                  0
      CPUS                 2,000             17,843                  0                  0
      DEVS                 2,000             17,843                 40                  0
      DEVS                 2,000             17,843              8,955                  0
      EXTS                 2,000             17,843                  0              1,649
      GBLS                 1,000             35,343                  0                  0
      IOQZ                     0                597                556                  0
      PAGS                 6,000              5,952                597                  0
      RECD                   15M                 40                396                  0
      RUNT                 1,440M                 0                  0                  0
      VSMS                 6,000              5,952                 40                  0
      WLMS                 2,000             17,843                  0                  0
      XCFS                     0              1,191                  0                  0
      TOTALS                                231,059             28,427              1,649

EXTRACTOR CONTROL CARDS

                     REPORT  CPM,INTERVAL=15,SYNCH=00,CSA=512,SMFRECID=240,
                             RUNTIME=1440,SMF=NO 00140009
                     ASMDATA  SAMPLE=2000
                     CACHE
                     CFDATA
                     CHANNEL
                     CPU      SAMPLE=2000
                    *CRYPTO
                     DEVICE   SAMPLE=2000,CLASS=DASD
                     DEVICE   SAMPLE=2000,CLASS=TAPE,OFFLINE=YES
                     EXTSUM  SPINOFF=NO,JES=NO,SAMPLE=2000,
                             JOBCLASS=(JC=A,JD=CLASSA,
                                       JC=B,JD=CLASSB)
                     HEADMOVE ALL,SAMPLE=250,VTOC=YES
                     IOQ
                     LINKMAP
                     PAGING   SAMPLE=6000
                     PGDDLAY
                     TSODATA  LIMIT=50,USER=YES,SAMPLE=2000
                     VSMDATA  SAMPLE=6000
                     WORKLOAD
                     XCFDATA
```

Extraction Characteristics Report field descriptions

Descriptions of the fields in the Extraction Characteristics Report are listed in Table 39.

Table 39 Field descriptions for the Extraction Characteristics Report

Field	Description
EXTRACTOR VERSION NUMBER	version and release of the CMF MONITOR Extractor
SU'S IN EFFECT	selectable units on the system
SAMPLER NAME	names of samplers; valid names are shown in Table 40 on page 366
SAMPLE RATE IN MSECS	sample rate in milliseconds If the letter M follows sample rate, the value is in minutes.
SAMPLE COUNTS UNDER DIE	number of samples taken under DIE (Disabled Interrupt Exit)
SAMPLE COUNTS UNDER SRB	number of samples taken under a global SRB The sample count for the SRB is always equal to or less than the DIE sample count.
SAMPLE COUNTS UNDER SRM	number of samples taken running under SVC 95 (SYSEVENT)
EXTRACTOR CONTROL CARDS	extractor control cards used for this set of input data (See Chapter 6, "Extractor control statements," for more information.)

Sampler Names for the Extraction Characteristics Report

Sampler names that can be listed in this report are shown in Table 40.

Table 40 Sampler names for the Extraction Characteristics Report

Sampler	Description
ASMS	ASMDATA sampler
CA3H CA5H CA6H	CACHE sampler
CHNS	CHANNEL sampler
CPUS	CPU sampler
CSMS	CSMON sampler
DEVS	DEVICE sampler (more than one sampler can be present)
EQES	ENQUEUE sampler
EXTS	EXTSUM sampler
FCSW	FICON Director sampler
GBLS	REPORT sampler
HMOV	HEADMOVE SAMPLER
IOQS	IOQ sampler
LPAM	LINKMAP sampler
PAGS	PAGING sampler
PGDS	PGDDLAY sampler
RECD	REPORT sampler
TRAS	TRACE76 sampler
TRCE	TRACE sampler
TSOS	TSODATA sampler
USER	USER sampler
VSMS	VSMDATA sampler
WLMS	WORKLOAD sampler
XCFS	XCFDATA sampler

System Resources Manager Constants Report

The System Resources Manager Constants (SRM Constants) Report is a Collection Phase Log report (see “Preliminary reporting information” on page 330 for more information) that displays the control block-related symbolic name, the value, and the related function of each SRM constant.

No Analyzer or Extractor control statements are required to produce the System Resources Manager Constants Report. It is produced automatically as part of the preliminary reporting information that CMF MONITOR supplies with each batch report job. This report is not produced if RPTS=INTERVAL or RPTS=DAILY is specified on the SHIFT control statement.

By using the calculation methods of the System Resources Manager and the information provided in this report, installation performance standards can be tuned to improve performance.

An example of the SRM Constants Report is shown in Figure 37 on page 367.

Figure 37 SRM Constants Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        COLLECTION PHASE LOG                              RPTSEQ     2 PAGE      6
 BMC SOFTWARE, INC.                                                                                         REPORT DATE: DD MMM YY 13.40
                                                  SYSTEM RESOURCES MANAGER CONSTANTS
                                       (IN EFFECT FROM 22 JUN 06 05.30.00 TO 22 JUN 06 10.15.00)
 -------------------------------------------------------  CONSTANTS  ----------------------------------------------------------
 CONSTANT     VALUE           RELATED FUNCTION                       CONSTANT     VALUE           RELATED FUNCTION
 CCCAPMET       296           MIN USR EXEC INTVL FOR MTW COMP        CCCMNSIN         7           MIN INT COMPUTING SYS CPU UTIL
 CCCSIGUR        52           SIGNIF. MTW IN MS                      CCCAPDIV         5           RANGE OF SUBGRP WITHIN MTW GRP
 CCCAPGHI         0           APG HIGH VALUE                         CCCAPLDP         1           DP UNDISPATCHED APG USERS
 CCCAPGLO         0           APG LOW VALUE                          ICCTPILO        10           LOW THLD FOR PERCENT TPI
 ICCTPIHI        30           HIGH THLD FOR PERCENT TPI              MCCPLUS         10           AVL FRM Q DELTA FOR STEALING
 MCCASMT1        70           AUX STRGE SHORTAGE THRSHLD             MCCASMT2        85           CRITICAL ASM SHORTAGE THRSHLD
 MCCSTLCT         4           NO PAGES TO STEAL EACH ADRSPC          MCCSIPRT     10000           TIME BETWEEN PG-RATE CALC
 MCCSIWDL         3           TWSS DELTA % DECREASE                  MCCSIWDI         7           TWSS DELTA % INCREASE
 MCCSIETH         1           EXC TME THOLD PG RATE CALC             MCCAFCLO       400           AVAIL FRAME QUE LOW THOLD
 MCCAFCOK       600           AVAIL FRAME QUE OK THOLD               MCCFXTPR        80           % LOGICAL STORAGE THLD
 MCCFXEPR        92           % PHYSICAL STORAGE THLD                RCCUICTL         2           UIC LO THRSHLD
 RCCUICTH         4           UIC HI THRSHLD                         RCCCPUTL      2048           CPU LO THRSHLD
 RCCCPUTH      2048           CPU HI THRSHLD                         RCCPTRTL        50           PAGING RATE LO THLD
 RCCPTRTH        80           PAGING RATE HI THLD                    RCCFXTTH        72           % LOGICAL HIGH MPL THLD
 RCCFXTTL        66           % LOGICAL LOW MPL THLD                 RCCFXETH        88           % PHYSICAL HIGH MPL THLD
 RCCFXETL        82           % PHYSICAL LOW MPL THLD                LSCTUCTL         2           UIC LOW THRESHOLD
 LSCTUCTH        15           UIC HIGH THRESHOLD                     LSCTASTL      1000           ASM QUEUED REQ LOW THRESHOLD
 LSCTASTH      1000           ASM QUEUED REQ HIGH THRSHOLD           LSCTMTEL      2000           THINK TIME LOW THRESHOLD
 LSCTMTEH     60000           THINK TIME HIGH THRESHOLD              LSCTMTEI       500           THINK TIME INCREMENT
 LSCTMTED      1000           THINK TIME DECREMENT                   LSCTFTTL        58           LOGICAL LOW FIXED FRAME THLD
 LSCTFTTH        66           LOGICAL HIGH FIXED FRAME THLD          LSCTFETL        76           PHYSICAL LOW FIXED FRAME THLD
 LSCTFETH        82           PHYSICAL HIGH FIXED FRAME THLD         RMPTXCHT       256           EXCHG SWAP THRSHLD
 RMPTSAET      1000           SWAP EVALUATION THRSHLD                RMPTWMET         4           WLM EVALUATION THRESHOLD
 RMPTERV          9           ENQ RESIDENCY CPU SERVICE              RMPTTOM       3072           TOLERANCE FOR TIME PERIOD
 RMCTADJC       520           CPU RATE ADJUSTMENT                    RMCTTAPE                     TAPE SELECTION
```

SRM Constants Report Field Descriptions

Descriptions of the fields in the SRM Constants Report are listed in Table 41.

Table 41 Field descriptions for the SRM Constants Report

Field	Description
CONSTANT	symbolic name of the control block-related constant See Table 42 for more information.
VALUE	value of the constant
RELATED FUNCTION	function of the constant

Valid constants names for the SRM Constants Report

Each SRM constant value has a name. Constants that can appear in the SRM Constants Report are listed in Table 42.

Table 42 Names and functions of SRM constants values (part 1 of 2)

Constant	Function
CCCAPMET	minimum execution time before a new dispatching priority is computed for the user
CCCMNSIN	minimum interval for computing system CPU usage
CCCSIGUR	minimum mean-time-to-wait to be considered a heavy CPU user, in milliseconds
CCCAPDIV	range of a subgroup within the mean-time-to-wait group
CCCAPGHI	APG high value
CCCAPLDP	DP undispatched APG user
CCCAPGLO	APG low value
ICCTPILO	low threshold for the percentage of TPI
ICCTPIHI	high threshold for the percentage of TPI
MCCPLUS	available frame queue delta for stealing
MCCASMT1	shortage threshold for auxiliary storage
MCCASMT2	critical shortage threshold for auxiliary storage
MCCSTLCT	number of pages to steal from each address space
MCCSIPRT	time between page-rate calculation
MCCSIWDL	TWSS delta percent decrease
MCCSIWDI	TWSS delta percent increase
MCCSIETH	execution time threshold for page-rate calculation
MCCAFCLO	low threshold for available frame queue
MCCAFCOK	OK threshold for available frame queue
MCCFXTPR	percentage of logical storage threshold
MCCFXEPR	percentage of physical storage threshold
RCCUICTL	UIC low threshold
RCCUICTH	UIC high threshold
RCCCPUTL	low threshold for CPU usage
RCCCPUTH	high threshold for CPU usage
RCCPTRTL	low threshold for paging rate
RCCPTRTH	high threshold for paging rate
RCCFXTTH	Percentage of high logical MPL threshold
RCCFXTTL	percentage of low logical MPL threshold
RCCFXETH	percentage of high physical MPL threshold
RCCFXETL	percentage of low physical MPL threshold
LSCTUCTL	UIC low threshold value by which the system think time is decremented
LSCTUCTH	UIC high threshold value by which the system think time is decremented

Table 42 Names and functions of SRM constants values (part 2 of 2)

Constant	Function
LSCTASTL	low threshold for ASM queued requests
LSCTASTH	high threshold for ASM queued requests
LSCTMTEL	low think time threshold
LSCTMTEH	high think time threshold
LSCTMTEI	think time increment
LSCTMTED	think time decrement
LSCTFTTL	logical low fixed frame threshold
LSCTFTTH	logical high fixed frame threshold
LSCTFETL	physical low fixed frame threshold
LSCTFETH	physical high fixed frame threshold
RMPTXCHT	exchange swap threshold
RMPTSAET	swap evaluation threshold
RMPTWMET	WLM evaluation threshold
RMPTERV	enqueue residency for CPU service
RMPTTOM	minimum SRM invocation interval
RMCTADJC	CPU rate adjustment value
RMCTTAPE	tape selection

RMF/CMF Input Record Type Counts Report

The RMF/CMF Input Record Type Counts Report is a Collection Phase Log report (see “Preliminary reporting information” on page 330 for more information) that lists information about the input data set records.

No Analyzer or Extractor control statements are required to produce the RMF/CMF Input Record Type Counts Report. It is produced automatically as part of the preliminary reporting information that CMF MONITOR supplies with each batch report job. This report is not produced if RPTS=INTERVAL or RPTS=DAILY is specified on the SHIFT control statement.

A five-digit ID is used for records unique to CMF MONITOR. The first three digits of the ID match the user-specified SMFRECID value from the Extractor REPORT control statement; the last two digits are an internal ID. An example of this report is shown in Figure 38 on page 370.

Figure 38 RMF/CMF Input Record Type Counts Report

```
  PRODUCED BY CMF ANALYZER (v.r.mm)                        COLLECTION PHASE LOG                              RPTSEQ     2 PAGE      9
  BMC SOFTWARE, INC.                                                                                         REPORT DATE: DD MMM YY 13.40
  RMF/CMF INPUT RECORD TYPE COUNTS
     RECORD TYPE  COUNTS - RMF         CPM         IPM     RECORD DEFINITION
      SMF  70-01           0           28           0     CPU ACTIVITY
      SMF  71-01           0           28           0     PAGING ACTIVITY
      SMF  72-03           0        1,176           0     WORKLOAD ACTIVITY BY SERVICE CLASS
      SMF  73-01           0           28           0     CHANNEL PATH ACTIVITY
      SMF  74-01           0          112           0     DEVICE ACTIVITY
      SMF  74-02           0           28           0     XCF ACTIVITY
      SMF  74-04           0           28           0     COUPLING FACILITY ACTIVITY
      SMF  74-05           0          308           0     CACHE SUBSYSTEM DEVICE ACTIVITY
      SMF  75-01           0          112           0     PAGE DATA SET ACTIVITY
      SMF  78-02           0           28           0     VIRTUAL STORAGE DATA
      SMF  78-03           0           28           0     I/O QUEUEING DATA
      SMF 240-00           0           28           0     SRM CONSTANTS, IPS, EXTRACTOR CTL CARDS
      SMF 240-01           0           28           0     CPU DATA
      SMF 240-02           0           28           0     ASM DATA
      SMF 240-03           0           28           0     PAGING DATA
      SMF 240-05           0          112           0     DEVICE DATA
      SMF 240-06           0           28           0     EXTRACTOR SUMMARY
      SMF 240-09           0           28           0     ASM DATA
      SMF 240-50           0           28           0     OUTPUT WRITER STATISTICS DATA
      TOTAL RECORDS        0        5,099           0
  REPORTS WILL BE GENERATED USING ONLY CPM DATA
```

RMF/CMF Input Record Type Counts Report field descriptions

Descriptions of the fields in the RMF / CMF Input Record Type Counts Report are listed in Table 43.

Table 43 Field descriptions for the RMF/CMF Input Record Type Counts Report

Field	Description
RECORD TYPE	SMF record type
COUNTS - RMF	number of records in the input data set that is created by RMF
CPM	number of records in the input data set that is created by CPM mode
IPM	number of records in the input data set that is created by IPM mode
RECORD DEFINITION	type of data that the record contains
TOTAL RECORDS	total number of records of the type specified that are encountered in the input data set
REPORTS WILL BE GENERATED USING ONLY xxx DATA	monitoring mode that is reported on in this Analyzer run: CPM, IPM, or RMF

Data Distribution and DATETIME Chart

The Data Distribution and DATETIME Chart is a Collection Phase Log report (see “Preliminary reporting information” on page 330 for more information) that graphically reports the relationship between the date range specified in the Analyzer DATETIME statement and the distribution of data records. The date range in this chart is defined by the range specified in the DATETIME control statement, the data itself, or both.

No Analyzer or Extractor control statements are required to produce the Data Distribution and DATETIME Chart. It is produced automatically as part of the preliminary reporting information that CMF MONITOR supplies with each batch report job.

Data records excluded from reports (because they were extracted before or after the time of day constraints of the DATETIME range) still appear in the column for their day. No indication is given in this chart that these records are outside the DATETIME range. This chart reflects relationships based only on extraction and DATETIME dates, not on time.

An example of this report is shown in Figure 39 on page 372.

Figure 39 Data Distribution and DATETIME Chart

```
 PRODUCED BY CMF ANALYZER (v.r.mm)                        COLLECTION PHASE LOG                          RPTSEQ     2 PAGE     10
 BMC SOFTWARE, INC.                                                                                     REPORT DATE: DD MMM YY 13.40
                                                DATA DISTRIBUTION AND DATETIME CHART
          +         +
    2610  +         +    2610
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
    1914  +    C    +    1914
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
    1218  +    C    +    1218
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
     522  +    C    +     522
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
          |    C    |
        0 +----+----+
          |         |
 021JUN YY>+         +<23JUN YY
 LEGEND:  * DELIMITS THE DATETIME RANGE        R REPRESENTS RMF RECORDS         I REPRESENTS IPM RECORDS
          C REPRESENTS CPM RECORDS      EACH LINE REPRESENTS    58 RECORDS     EACH '+' COLUMN POSITION REPRESENTS   1 DAYS
```

Data Distribution and DATETIME Chart field descriptions

Descriptions of the fields in the Data Distribution and DATETIME Chart are listed in Table 44.

Table 44 Field descriptions for the Data Distribution and DATETIME Chart

Field	Description
*****	delimits the DATETIME range
C	represents CPM records, when present
R	represents RMF records, when present
I	represents IPM records, when present
+	plus sign (+) in each column position represents at least one day; the exact number of days for each + is shown in the legend information Each line represents the number of records, as indicated in the report legend.

The characters used in the chart are defined in the legend. The value of each line of data and the column positions vary depending on

- amount of data being reported
- DATETIME range
- date range of the data

The numbers along the vertical axis vary depending on the amount of data being reported; they are reference points for reading column totals.

Each + (plus sign) in the horizontal axis can represent a day or a group of days, as indicated in the legend. Dashes between plus signs are for spacing purposes only. The data records are aligned only over plus signs.

The columns of * (asterisks) at either side of the chart represent the start of the DATETIME period (left column) and the end of the DATETIME period (right column). A single column of asterisks indicates a DATETIME range that falls within one day. If no columns of asterisks appear, no DATETIME range was specified.

NOTE

The data symbols for records collected on the day of the start or end of the DATETIME range overlay the asterisks in that column.

Common Storage Usage Detail Report

The Common Storage Usage Detail Report provides a detailed view of the average, minimum, and maximum values of the common storage data collected in CMF-type 240-29 user records. This data is obtained by using the CSMON Extractor control statement (see "CSMON" on page 142).

This report is produced by using the COMMSTOR Analyzer control statement (see "COMMSTOR" on page 227) and specifying the REPORT=DETAIL parameter. There is also a Common Storage Usage Summary Report (see "Common Storage Usage Summary Report" on page 377), which is the default report that the COMMSTOR statement produces. The Common Storage Usage Summary Report provides a summary view of the common storage data collected in CMF type 240-29 user records.

The common storage records contain a very detailed level of data. A sample SAS routine is distributed in the BBSAMP data set member, CSMAPSAS, to use as a guide for writing customized in-house reports. The data in the CMF type 240-29 user record is documented in the BBSAMP data set member, CMFREC29. (See Chapter 10, "Mapping CMF records created by CMF," for more information.)

Figure 40 on page 375 is an example of the Common Storage Usage Detail Report.

Figure 40 Common Storage Usage Detail Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     COMMON STORAGE USAGE REPORT                      RPTSEQ    13 PAGE     79
 BMC SOFTWARE, INC.                                            XYZ COMPANY                              REPORT DATE: DD MMM YY 13.40
 ACTL 10 JUN YY 09.00.00 10 JUN YY 16.00.00             WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-29/1,342/0/7

 ---------------------------------------------- DETAIL REPORT SORTED BY JOBNAME          ------------------------------------------

        JOBNAME                          AVG STOR        MAX STOR   DATE      TIME          MIN STOR   DATE      TIME
        BCVKSR4     CSA BELOW 16M             136             136   1.215    9:15:00             136   1.215    9:15:00
                    CSA ABOVE 16M           5,440          10,264   1.215    9:15:00           2,224   1.215    9:44:59
                    CSA TOTAL               5,576          10,400   1.215    9:15:00           2,360   1.215    9:44:59
                    SQA BELOW 16M              38              96   1.215    9:15:00               0   1.215    9:44:59
                    SQA ABOVE 16M             112             280   1.215    9:15:00               0   1.215    9:44:59
                    SQA TOTAL                 150             376   1.215    9:15:00               0   1.215    9:44:59

        BCVSCAS     CSA BELOW 16M              69             944   1.215    9:15:00               0   1.215   10:30:00
                    CSA ABOVE 16M         196,068       2,907,856   1.215    9:15:00           4,096   1.215   10:30:00
                    CSA TOTAL             196,137       2,908,800   1.215    9:15:00           4,096   1.215   10:30:00
                    SQA BELOW 16M               0               0   1.215    9:15:00               0   1.215    9:15:00
                    SQA ABOVE 16M              22             640   1.215    9:15:00               0   1.215    9:15:00
                    SQA TOTAL                  22             640   1.215    9:15:00               0   1.215    9:15:00

        BCVSPASE    CSA BELOW 16M          16,239          54,392   1.215    9:15:00           1,680   1.215    9:15:00
                    CSA ABOVE 16M          72,621         526,896   1.215    9:15:00             256   1.215    9:15:00

                    CSA TOTAL              88,860         581,288   1.215    9:15:00           1,936   1.215    9:15:00
                    SQA BELOW 16M               0               0   1.215    9:15:00               0   1.215    9:15:00
                    SQA ABOVE 16M              93           2,672   1.215    9:29:59               0   1.215    9:15:00
                    SQA TOTAL                  93           2,672   1.215    9:29:59               0   1.215    9:15:00

        BKSPRMP4    CSA BELOW 16M               0               0   1.215   16:00:00               0   1.215   16:00:00
                    CSA ABOVE 16M             520             520   1.215   16:00:00             520   1.215   16:00:00
                    CSA TOTAL                 520             520   1.215   16:00:00             520   1.215   16:00:00
                    SQA BELOW 16M              64              64   1.215   16:00:00              64   1.215   16:00:00
                    SQA ABOVE 16M             256             256   1.215   16:00:00             256   1.215   16:00:00
                    SQA TOTAL                 320             320   1.215   16:00:00             320   1.215   16:00:00

        BKSPRMP5    CSA BELOW 16M               0               0   1.215   16:00:00               0   1.215   16:00:00
                    CSA ABOVE 16M             520             520   1.215   16:00:00             520   1.215   16:00:00
                    CSA TOTAL                 520             520   1.215   16:00:00             520   1.215   16:00:00
                    SQA BELOW 16M              64              64   1.215   16:00:00              64   1.215   16:00:00
                    SQA ABOVE 16M             256             256   1.215   16:00:00             256   1.215   16:00:00
                    SQA TOTAL                 320             320   1.215   16:00:00             320   1.215   16:00:00

        BMCPXK2     CSA BELOW 16M           1,013           4,232   1.215   14:45:00             136   1.215    9:15:00
                    CSA ABOVE 16M           3,827          10,264   1.215   14:45:00           2,072   1.215    9:15:00
                    CSA TOTAL               4,841          14,496   1.215   14:45:00           2,208   1.215    9:15:00
                    SQA BELOW 16M              85              96   1.215    9:15:00               0   1.215   15:29:59
                    SQA ABOVE 16M             293             408   1.215   15:29:59             280   1.215    9:15:00
                    SQA TOTAL                 379             408   1.215   15:29:59             376   1.215    9:15:00

        BMCPXK3     CSA BELOW 16M          22,654          28,280   1.215   10:30:00             136   1.215    9:15:00
                    CSA ABOVE 16M          88,740         119,648   1.215   10:30:00           2,112   1.215    9:15:00
                    CSA TOTAL             111,394         147,928   1.215   10:30:00           2,248   1.215    9:15:00
                    SQA BELOW 16M              85              96   1.215    9:15:00               0   1.215   15:29:59
                    SQA ABOVE 16M           1,053           1,464   1.215   10:30:00               0   1.215   15:29:59
                    SQA TOTAL               1,139           1,560   1.215   10:30:00               0   1.215   15:29:59
```

Common Storage Usage Detail Report column field descriptions

Table 45 describes each column field in the Common Storage Usage Detail Report.

Table 45 Column descriptions for the Common Storage Usage Detail Report

Field	Description
JOBNAME	name of address space that was in control when the common storage was obtained
AVG STOR	average amount of specified common storage that was held by the job name
MAX STOR	maximum amount of specified common storage that was held by job name in a single record interval
MAX STOR – DATE	date of record interval when maximum amount of specified common storage was held
MAX STOR – TIME	time of record interval when maximum amount of specified common storage was held
MIN STOR	minimum amount of specified common storage that was held by job name in a single record interval
MIN STOR – DATE	date of record interval when minimum amount of specified common storage was held
MIN STOR – TIME	time of record interval when minimum amount of specified common storage was held

Common Storage Usage Detail Report row field descriptions

Table 46 describes each row field in the Common Storage Usage Detail Report.

Table 46 Row descriptions for the Common Storage Usage Detail Report

Field	Description
CSA BELOW 16M	row containing the amount of CSA below the 16-MB line that the job name held
CSA ABOVE 16M	row containing the amount of CSA above the 16-MB line that the job name held
CSA TOTAL	row containing the amount of all CSA that the job name held
SQA BELOW 16M	row containing the amount of SQA below the 16-MB line that the job name held
SQA ABOVE 16M	row containing the amount of SQA above the 16-MB line that the job name held
SQA TOTAL	row containing the amount of all SQA that the job name held

Common Storage Usage Summary Report

The Common Storage Usage Summary Report provides a summary view of the average values of the common storage data collected in CMF type 240-29 user records. This data is obtained by using the CSMON Extractor control statement (see “CSMON” on page 142).

This report is produced by using the COMMSTOR Analyzer control statement (see “COMMSTOR” on page 227). There also is a Common Storage Usage Detail Report (see “Common Storage Usage Detail Report” on page 374), but the default COMMSTOR statement produces the Common Storage Usage Summary Report.

Figure 41 is an example of the Common Storage Usage Summary Report.

Figure 41 Common Storage Usage Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      COMMON STORAGE USAGE REPORT                               RPTSEQ    12 PAGE     73
BMC SOFTWARE, INC.                                             XYZ COMPANY                                       REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00              WORLDWIDE HEADQUARTERS                                  SYSTEM ID: SJSE     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-29/1,342/0/7

--------------------------------------------- SUMMARY REPORT SORTED BY JOBNAME            ----------------------------------------
```

JOBNAME	AVG CSA BELOW 16M	AVG CSA ABOVE 16M	AVERAGE CSA TOTAL	AVG SQA BELOW 16M	AVG SQA ABOVE 16M	AVERAGE SQA TOTAL
DC$RES	4,136	520	4,656	0	512	512
DC$RMM	8,864	285,584	294,448	2,312	10,355	12,667
DC$TCPIP	136	2,447,261	2,447,397	0	8,104	8,104
DC$ZSSI	288	0	288	0	0	0
DMRCASQA	56	180,512	180,568	0	0	0
DMRCA6KC	653	2,326,108	2,326,762	0	473	473
DMRCA6KP	56	180,296	180,352	0	0	0
DMRPASQA	25,176	61,336	86,512	0	0	0
DMRPA6KC	29,371	514,192	543,563	0	2,004	2,004
DMRPR1	0	13,653	13,653	0	0	0
DMRTMO1	0	4,096	4,096	0	0	0
DMRUPD3	0	4,096	4,096	0	0	0
DMRXCASE	864	3,070,672	3,071,536	0	640	640
DOMHHH1	9,288	953,626	962,914	4,480	2,457	6,937
DUMPSRV	0	4,664	4,664	64	89,457	89,521
GRS	0	0	0	48	12,631	12,679
I@61REPT	0	80	80	0	0	0
IEFSCHAS	0	0	0	0	16	16
IMSIR22	43,312	189,704	233,016	256	2,880	3,136
INIT	0	1,008	1,008	0	261	261
IOSAS	47	0	47	64	400	464
IRRDPTAB	0	59,976	59,976	0	0	0
IXGLOGR	0	32,608	32,608	64	2,352	2,416
I6AMP1	0	80	80	0	0	0
I6AM2PCT	608	69,632	70,240	1,496	8	1,504
I6AM32DL	0	0	0	304	0	304
JESXCF	0	329,560	329,560	64	8,904	8,968
JES2	28,240	539,224	567,464	64	1,616	1,680

Common Storage Usage Summary Report field descriptions

Table 47 describes each field in the Common Storage Usage Summary Report.

Table 47 Field descriptions for the Common Storage Summary Usage Report

Field	Description
JOBNAME	name of the address space in control when the common storage was obtained
AVG CSA BELOW 16M	average amount of CSA storage below the 16-MB line that is held by the specified job name
AVG CSA ABOVE 16M	average amount of CSA storage above the 16-MB line that is held by the specified job name
AVERAGE CSA TOTAL	average amount of all CSA storage that is held by the specified job name
AVG SQA BELOW 16M	average amount of SQA storage below the 16-MB line that is held by the specified job name
AVG SQA ABOVE 16M	average amount of SQA storage above the 16-MB line that is held by the specified job name
AVG SQA TOTAL	average amount of all SQA storage that is held by the specified job name

Coupling Facility Activity Report

The Coupling Facility Activity Report provides information about the following coupling facility components:

- activity against structures at both structure and system levels
- storage allocation and usage
- CPU utilization
- path and subchannel configuration

The Coupling Facility Activity Report is produced by using the CFACT Analyzer control statement (see “CFACT” on page 219). The data for this report is obtained by using the CFDATA Extractor control statement (see “CFDATA” on page 132). One report is produced for each coupling facility found in the input data.

This multipage report consists of four summary sections and three detail sections.

- summary sections

Processor	displays Processor utilization of the coupling facility
Storage	displays storage allocation and usage
Connectivity	displays the name of connected system members **Note**: This section allows you to verify that data from all systems that are connected to the coupling facility is accounted for and that the data from each connected system covers the same time window. If data from one or more connected systems is missing, or if the data from the connected systems does not cover the same time frame, then all coupling facility activity information is incomplete or skewed. If data from one connected system is missing, warning message CMF07440 is issued. If data from all systems is available but does not cover the same time frame, warning message CMF07441 is issued.
Structure	summarizes the attributes of each structure and its activities from all systems The structures are grouped by type in the following order: list, cache, and lock. Structures of the same type are ordered by name.

Figure 42 on page 380 shows an example of the four summary sections of the Coupling Facility Activity Report.

- detail sections

Subchannel Activity	displays the path and subchannel configuration of each connected system and its activity Figure 43 on page 386 shows an example of the Subchannel Activity section of the Coupling Facility Activity Report.
Structure Activity	displays request activity against each structure, viewed by system The structures are grouped by type in the following order: list, cache, and lock. Structures of the same type are ordered by name. See Figure 44 on page 388 for an example of the Structure Activity section of the Coupling Facility Activity Report.
CF to CF Activity	displays a summary of basic counts for duplexing related operations and information about the specific CF link types See Figure 45 on page 391 for an example of the CF to CF Activity section of the Coupling Facility Activity Report.

Beginning with Figure 43 on page 386, you can see examples of the detail sections of the Coupling Facility Activity Report.

Figure 42 Summary sections of the Coupling Facility Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                 COUPLING FACILITY ACTIVITY REPORT                     RPTSEQ     3 PAGE      6
BMC SOFTWARE, INC.                                         BMC ENGINEERING                              REPORT DATE: DD MMM YY 15.23
ACTL DD MMM YY 18.30.00  22 JAN 07 18.55.00                                                             SYSTEM ID: SJSE     Z 1.07.1


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-4/5/132/0.42
COUPLING FACILITY = CF16      ---------------------  PROCESSOR SUMMARY SECTION  ----------------------------------------------------
                                                                                 MODEL    VERSION    CFLEVEL    PLANT    SEQUENCE #
  # LOGICAL PROCESSORS    =     1                                               ------    -------    -------     --     ------------
  # PROCESSORS EFFECTIVE =    0.8                                               002094        S28         14      02    0000000422EA
  AVERAGE UTILIZATION    =   1.7%
  MAXIMUM UTILIZATION    =   1.8%   (PROCESSOR NUMBER = 00000000, FROM 22 JAN YY  18.35.00    TO 22 JAN YY  18.40.00)
---------------------------------------------------  STORAGE SUMMARY SECTION  -----------------------------------------------------
              --- SIZE ----  -% OF DUMP--                                       -- DEFINED --  -- ALLOCATED --
  STORAGE     AMOUNT  % OF   USED    MAX                         STORAGE         SIZE   % OF    SIZE   % OF
  ALLOCATION          CF             REQ                         ORGANIZATION           CF             DEFINED
  ----------  ------  -----  -----  -----                        ------------   ------  -----  ------  -------
  STRUCTURES    638M   32.2                                      CONTROL         1981M  100.0    642M     32.4
  DUMP            4M    0.2    0.0    6.3                        DATA           NOT DEFINED SEPARATELY FROM CONTROL
  AVAILABLE    1340M   67.6
  ----------  ------  -----
  TOTAL        1981M  100.0
-------------------------------------------------  CONNECTIVITY SUMMARY SECTION  ------------------------------------------------
   SYSTEM     DURATION        START TIME            END TIME
    NAME    DDD.HH.MM.SS
  --------  ------------  ------------------  ------------------
  SJSB      *** NO DATA ***
  SJSC      *** NO DATA ***
  SJSD      *** NO DATA ***
  SJSE          00.25.00  22 JAN YY  18.30.00  22 JAN YY  18.55.00
  SJSG      *** NO DATA ***
  SJSH      *** NO DATA ***
*WARNING* CMF07440 (IN SOR07440)  DATA FROM AT LEAST ONE SYSTEM OF COUPLING FACILITY IS MISSING

----------------------------------------------------  STRUCTURE SUMMARY SECTION  -------------------------------------------------
  TYPE        NAME       STATUS  --- SIZE --- -------------------- REQUESTS --------------------- LST/DIR DATA    LOCK
                                 AMOUNT % OF   ---- PER SEC -----  COUNT  % OF  -- SERVICE TIME -- ENTRIES ELEMENT ENTRIES  DIRECTORY
                                        CF                                 CF                     TOT/CUR TOT/CUR TOT/CUR  RECLAIMS
  ----- ---------------- ------  ------ ----- ------- ---------- ------- ----- ------- ---------- ------- ------- -------  -------
  LIST  APPC_STR1        ACTIVE      4M   0.2     0.0                  5   0.0   783.2 ++++++++       101    1507     N/A      N/A
                                                                                                        6     287     N/A
  LIST  DBGK_SCA         ACTIVE  16640K   0.8     0.0                  0   0.0     0.0              20728   41250     N/A      N/A
                                                                                                      210     859     N/A
  LIST  DFHXQLS_BCVCTS01 ACTIVE  20224K   1.0     0.0                  0   0.0     0.0              31033   28079     N/A      N/A
                                                                                                        0       0     N/A
  LIST  DSNDHH_SCA       ACTIVE      2M   0.1     1.7               2586   1.3   403.6 ++++          1664    3259     N/A      N/A
                                                                                                      161     340     N/A
  LIST  DSNDHO_SCA       ACTIVE      2M   0.1     1.7               2601   1.3   396.1 ++++          1664    3259     N/A      N/A
                                                                                                      160     339     N/A
  LIST  ISTMNPS          ACTIVE     10M   0.5     0.0                  0   0.0     0.0               7385   14717     N/A      N/A
                                                                                                        1       0     N/A
PRODUCED BY CMF ANALYZER (v.r.mm)                 COUPLING FACILITY ACTIVITY REPORT                     RPTSEQ    3 PAGE      7
BMC SOFTWARE, INC.                                         BMC ENGINEERING                              REPORT DATE: DD MMM YY 15.23
ACTL DD MMM YY 18.30.00 DD MMM YY  18.55.00                                                             SYSTEM ID: SJSE     Z 1.07.1


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-4/5/132/0.42
COUPLING FACILITY = CF16      ---------------------  STRUCTURE SUMMARY SECTION  ----------------------------------------------------
  TYPE        NAME       STATUS  --- SIZE --- -------------------- REQUESTS --------------------- LST/DIR DATA    LOCK
                                 AMOUNT % OF   ---- PER SEC -----  COUNT  % OF  -- SERVICE TIME -- ENTRIES ELEMENT ENTRIES  DIRECTORY
                                        CF                                 CF                     TOT/CUR TOT/CUR TOT/CUR  RECLAIMS
  ----- ---------------- ------  ------ ----- ------- ---------- ------- ----- ------- ---------- ------- ------- -------  --------
  LIST  IXCPATH1         ACTIVE     16M   0.8    66.0 ++++++++++   99068  51.1   626.5 ++++++        2866    2847     N/A      N/A
                                                                                                        2      24     N/A
  LIST  IXCPATH2         ACTIVE     16M   0.8     1.0               1528   0.8   632.5 ++++++        2866    2847     N/A      N/A
                                                                                                        1      16     N/A
  LIST  IXCPATH3         ACTIVE     16M   0.8    23.7 ++++         35527  18.3   366.9 ++++          2866    2847     N/A      N/A
                                                                                                        1      18     N/A
  LIST  IXCPATH4         ACTIVE     16M   0.8     2.6               3854   2.0   621.3 ++++++        2866    2847     N/A      N/A
                                                                                                        1      24     N/A
  LIST  LOG_DFHLOG_001   ACTIVE      4M   0.2     0.0                 57   0.0   993.7 ++++++++++    2018    9677     N/A      N/A
                                                                                                       94     506     N/A
  LIST  LOG_DFHSHUNT_001 ACTIVE      4M   0.2     0.0                 17   0.0   140.6 +              659   10882     N/A      N/A
                                                                                                       28     218     N/A
  LIST  RRS_DATA         ACTIVE      2M   0.1    11.7 ++           17595   9.1    94.1 +             2327    2449     N/A      N/A
                                                                                                      506     533     N/A
  LIST  RRS_LOGS         ACTIVE  25088K   1.2     0.0                 34   0.0   887.2 +++++++++    25932   25953     N/A      N/A
                                                                                                     7881    7892     N/A
  LIST  RRS_MAIN         ACTIVE      5M   0.3     0.3                517   0.3   661.8 ++++++        2000    2023     N/A      N/A
                                                                                                       15      64     N/A
  LIST  ISTGENERIC       ACTIVE     10M   0.5     9.2 +            13866   7.1   374.5 ++++         29284     585       4      N/A
                                                                                                        7       2       0
  LIST  MVIQA_EMHQ       ACTIVE   5376K   0.3     0.0                  0   0.0     0.0               1018    2006     256      N/A
                                                                                                        4       3       0
  LIST  MVIQA_MSGQ       ACTIVE  11008K   0.5     0.0                  0   0.0     0.0               4227    8395     256      N/A
                                                                                                     1117    1119       0
  LIST  MVI10A_EMHQ      ACTIVE      5M   0.3     0.0                  0   0.0     0.0               1018    2006     256      N/A
                                                                                                        6       5       0
  LIST  MVI10A_MSGQ      ACTIVE     10M   0.5     0.0                  0   0.0     0.0               4227    8395     256      N/A
                                                                                                       77     147       0
```

Field descriptions of the summary sections of the Coupling Facility Activity Report

Table 48 describes each field on the summary Sections of the Coupling Facility Activity Report.

Table 48 Field descriptions for the summary sections (part 1 of 5)

Field	Description
COUPLING FACILITY	coupling facility name **Note**: This field is displayed in the first section header on every page of this report.
# LOGICAL PROCESSORS	number of logical processors if the coupling facility is a partition in a PR/SM environment, or the number of physical processors if the coupling facility is standalone
# PROCESSORS EFFECTIVE	number of available logical processors, effectively, if the coupling facility is a partition in a PR/SM environment This value is a ratio of the total LPAR dispatch time (sum of CFCC execute time and wait time for all processors) over the interval duration. For a coupling facility that is ICMF LPAR, refer to "CPU Utilization Report" on page 393.
AVERAGE UTILIZATION	average CPU utilization of the coupling facility
MAXIMUM UTILIZATION, PROCESSOR NUMBER, FROM, TO	maximum utilization of any processor during all recording intervals PROCESSOR NUMBER identifies the processor, and FROM and TO identify the recording interval.
MODEL	coupling facility type
VERSION	coupling facility model
CFLEVEL	coupling facility microcode level
PLANT	two-character Manufacturer Plant Code; field will be blank for z/OS systems earlier than 1.8
SEQUENCE #	12-character Manufacturer Sequence Number; field will be blank for z/OS systems earlier than 1.8
STORAGE ALLOCATION	possible values for this field are as follows: **STRUCTURES** storage allocated and partitioned into structures The amount of storage allocated to structures reported here can exceed the TOTAL SIZE of all structures in the STRUCTURE SUMMARY SECTION because some amount of overhead is not attributable to any structure. **DUMP** storage reserved for dumping data **AVAILABLE** storage allocated to neither structures nor dump **TOTAL** total storage of the coupling facility

Table 48 Field descriptions for the summary sections (part 2 of 5)

Field	Description
SIZE – AMOUNT	storage in the coupling facility is allocated in blocks of 4-K bytes The unit is either kilobytes (K) or megabytes (M), if the size in K is a multiple of 1024 or greater than 99999.
SIZE – % OF CF	percentage of the total coupling facility storage allocated to structures, dump space, or available
% OF DUMP – USED	percentage of dump space in use, observed at end of interval
% OF DUMP – MAX REQ	percentage of the largest amount of dump space requested, compared with the dump space size This largest requested amount is a high water mark since coupling facility initialization or dump space reinitialization. If this value exceeds 100, at least 1 dump has been lost or truncated. The recommended size of dump space equals that of the largest structure or 5% of the coupling facility storage, whichever is larger.
STORAGE ORGANIZATION	structure and dump storage in the coupling facility occupies both CONTROL and DATA area
DEFINED – SIZE	for control area, the maximum amount it can occupy Data area, on the other hand, is not restricted to this size. When exhausted, data storage can be allocated from control area. The control area can be defined as equaling the total coupling facility storage. In other words, there is no restriction on the size of the control area; the entire coupling facility is shared between control and data areas. In this case, NOT DEFINED SEPARATELY FROM CONTROL is printed on the DATA line.
DEFINED – % OF CF	percentage of coupling facility storage allocated to control or data area
ALLOCATED - SIZE	amount of control or data area allocated to structures or dump The allocated size of data area can exceed the defined size (% OF DEFINED then exceeds 100), indicating that the data storage has been exhausted and data has been allocated from the control area.
ALLOCATED – % OF DEFINED	percentage of defined control or data area allocated to structures or dump
SYSTEM NAME	name of the system connected to the coupling facility If no records from a system are found in the input data, *** NO DATA *** is printed after the system name.
DURATION	total recording intervals of records produced on system
START TIME	start date and time of the first record
END TIME	end date and time of the last record

Table 48 Field descriptions for the summary sections (part 3 of 5)

Field	Description
TYPE	first user connecting to a structure specifies its type Possible values for this field are as follows: **LIST** allows users to share data organized as entries in a set of lists **CACHE** allows sharing of frequently referenced data **LOCK** allows serializing on shared resources, including list and cache structures If there is no activity against a structure from any connected system, its type is listed as UNK (for unknown).
NAME	name of the structure, up to 16 characters long A structure, and its attributes (maximum size, for example), is specified in the coupling facility resource management (CFRM) policy. The first subsystem or application that connects to the structure allocates it in the coupling facility. When there are multiple instances of the same structure, each instance is considered unique and listed separately in the Structure Summary section. In the Structure Activity section, the activity data of all instances of a structure is combined.
STATUS	status of the structure, observed at end of interval Possible values for this field are as follows: **ACTIVE** at the end of at least one interval, at least one system is connected to the structure; the ACTIVE values in the **Status** field can include either the PRIM or SEC qualifier **PRIM** rebuilt-old (primary) structure in a duplexing rebuild process; it is the first structure of the duplexed pair to be allocated **SEC** rebuilt-new (secondary) structure in a duplexing rebuild process; it is the second structure of the duplexed pair to be allocated **INACTV** structure is allocated in the coupling facility but is not connected to any system at the end of any interval **UNALLO** no system is connected to the structure at the end of interval and the structure is no longer allocated in the coupling facility; the existence of the structure is known because it was active earlier in the recording interval If the structure becomes active during the report interval, an asterisk is also printed.

Table 48 Field descriptions for the summary sections (part 4 of 5)

Field	Description
SIZE – AMOUNT	size of the structure, allocated in blocks of 4-K bytes The unit is in either kilobytes (K) or megabytes (M) if the size in K is a multiple of 1024 or greater than 99999. The structure size is either the value specified in the coupling facility resource management (CFRM) policy or the value specified at allocation time, whichever is greater.
SIZE – % OF CF	percentage of total coupling facility storage that is allocated to the structure
REQUESTS – PER SEC	rate per second of completed requests that are directed to the structure Graph is scaled such that the bar of the structure having the highest rate occupies the full width of the column.
REQUESTS – COUNT	number of completed requests that are directed to the structure
REQUESTS – % OF CF	percentage of completed requests against the structure, compared with the total number of requests against all structures
REQUESTS – SERVICE TIME	average number of microseconds that are required to complete a request (excluding queue time) Graph is scaled such that the bar of the structure having the longest service time occupies the full width of the column.
LST/DIR ENTRIES TOT/CUR	pair of numbers that provide information about list entries or cache directory entries For list or lock structures, the first number is the total number of list entries in the structure. For cache structures, the first number is the total number of directory entries in the structure. For list or lock structures, the second number is the average number of list entries in use at the end of the interval. For cache structures, the second number is the average number of directory entries in use at the end of the interval. If a field contains N / A, the data is not applicable to this structure type.
DATA ELEMENT – TOT/CUR	pair of numbers that provide information about list and cache data elements For list or lock structures, the first number is the total number of list elements in the structure. For cache structures, the first number is the total number of data elements in the structure. For list or lock structures, the second number is the average number of list elements in use at the end of the interval. For cache structures, the second number is the average number of data elements in use at the end of the interval. If a field contains N / A, the data is not applicable to this structure type.

Table 48 Field descriptions for the summary sections (part 5 of 5)

Field	Description
LOCK ENTRIES – TOT/CUR	pair of numbers that provide information about lock table entries for lock structures or serialized list structures The first number is the total number of lock table entries for the structure. The second number is the average non-zero lock table entry count at the end of the interval. If a field contains N / A, the data is not applicable to this structure type.
DIRECTORY RECLAIMS	number that provides the average number of directory reclaims for a cache structure Directory reclaim activity indicates a shortage of directory entries. A shortage of directory entries causes the coupling facility to *reclaim* already used directory entries that are associated with unchanged data. As a result, the copy of the data in the coupling facility is invalidated, and must be reread from DASD and registered to the coupling facility again. If a field contains N / A, the data is not applicable to this structure type.
SUMMARY	values on this line are either the total or the average of values of all instances of all structures If different structure instances exist at different ends of interval, totaling their sizes and percentages of coupling facility storage is meaningless. In this case, N / A is printed.

Figure 43 shows an example of the Subchannel Activity Section of the Coupling Facility Activity Report.

Figure 43 Subchannel Activity Section of the Coupling Facility Activity Report

```
 PRODUCED BY CMF ANALYZER (v.r.mm)                   COUPLING FACILITY ACTIVITY REPORT                   RPTSEQ     4 PAGE     96
 BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 16.58
 ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00              WORLDWIDE HEADQUARTERS                         SYSTEM ID: **ALL**  COMB-MVS


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-4/43/3,767/10.5
 COUPLING FACILITY = ICF91     ---------------------  SUBCHANNEL ACTIVITY SECTION  -------------------------------------------------
  SYSTEM  ---------------- REQUESTS -----------------  ------------- QUEUED REQUESTS --------------  REJECTED   -- CONFIG --
   NAME    TYPE    PER     COUNT  % OF  -SERVICE TIME--          #       %     -- DELAY TIME ---       REQUESTS   RESOURCE CNT
                   SEC             ALL     AVG   STD DEV         REQ     REQ   /REQ    S DEV   /ALL     RES  CNT
 -------- ------- ------- ------- ----- ------- -------  -------- ------- ----- ------- ------- -------  ---- ----  -------- ---

 SJSC     ALL       313.4 3948804  37.7   268.6   545.5                                                 PATH    0  PATH       3
           SYNC       2.4   30684   0.3    24.9    22.5  SYNC        3842  12.5      13      11       2  SCH  3842  SCH GEN    6
           ASYNC    186.4 2348065  22.4   271.8   548.3  ASYNC     939567  40.0    1102    1316     441            SCH USE    6
           CNVAS      0.0      63   0.0                  TOTAL     943409  39.7                                    SCH OPT    6
           UNSUC      0.0       0   0.0     0.0     0.0

 SJSE     ALL       250.5 6538033  62.3   531.4  2738.0                                                 PATH    0  PATH       2
           SYNC       7.3  189363   1.8    32.3    28.8  SYNC        6759   3.6      14      20       0  SCH  6759  SCH GEN    4
           ASYNC    144.7 3775349  36.0   556.4  2803.5  ASYNC    2766229  73.3    2055    7213    1506            SCH USE    4
           CNVAS      0.0      45   0.0                  TOTAL    2772988  69.9                                    SCH OPT    4
           UNSUC      0.0       0   0.0     0.0     0.0

 -------- ------- ------- ------- ----- ------- -------  -------- ------- ----- ------- ------- -------  ---- ----
 SUMMARY  ALL       542.0 1.0E+07 100.0   432.9  2193.9                                                 PATH    0
           SYNC      11.4  220047   2.1    31.2    28.1  SYNC       10601   4.8      14      17       1  SCH 10601
           ASYNC    316.5 6123414  58.4   447.3  2231.6  ASYNC    3705796  60.5    1813    6281    1097
           CNVAS      0.0     108   0.0                  TOTAL    3716397  58.6
           UNSUC      0.0       0   0.0     0.0     0.0
```

Field descriptions of the Subchannel Activity Section

Table 49 describes each field on Subchannel Activity section of the Coupling Facility Activity Report.

Table 49 Field descriptions for the Subchannel Activity Section (part 1 of 2)

Field	Description
SYSTEM NAME	name of the system connected to the coupling facility **Note**: If no records from a system are found in the input data, *** NO DATA *** is printed after the system name.
REQUESTS – TYPE	type of requests Possible values for this field are as follows: **SYNC** synchronous requests directed to structures **ASYNC** asynchronous requests directed to structures **CNVAS** requests directed to structures that originated as synchronous but converted to asynchronous A synchronous request is converted to asynchronous if the buffer size is greater than 4-K bytes or if multiple buffers are passed.

Table 49 Field descriptions for the Subchannel Activity Section (part 2 of 2)

Field	Description
REQUESTS – TYPE (*continued*)	**UNSUC** not completed due to hardware problems **ALL** total of all request types The ALL field value can exceed the sum of the values for SYNC, ASYNC, CNVS, and UNSUC because it also includes coupling-facility-managing requests and requests directed to structures that are no longer active at the end of the interval.
REQUESTS – PER SEC	rate per second of requests that originated from the system, viewed by request type **Note**: The rate and count of requests from all systems in the SUMMARY section exceeds the TOTAL rate and count in the Structure Summary section, because the former include unsuccessful requests and coupling-facility-managing requests.
REQUESTS – COUNT	number of requests that originated from a system
REQUESTS – % OF ALL	percentage of requests of a type that originated from a system, compared with the total number of requests from all systems
REQUESTS – SERVICE TIME - AVG, STD DEV	average number of microseconds that are required to complete a request (excluding queue time) and its standard deviation Service time of synchronous-converted-to-asynchronous (CNVAS) requests is included in that of asynchronous (ASYNC) requests. The service time of coupling-facility-managing (GLOBL) requests is not available.
QUEUED REQUESTS – # REQ	total number of requests to this coupling facility that were delayed in this interval
QUEUED REQUESTS – % REQ	percentage of requests to this coupling facility that were delayed in this interval
DELAY TIME / REQ	average delay time for each delayed request in microseconds
DELAY TIME S DEV	standard deviation of the average delay time for delayed requests
DELAY TIME / ALL	average delay time for all requests to the coupling facility in microseconds This number helps indicate how delays are affecting overall request activity to the coupling facility.
REJECTED REQUESTS – RES, CNT	number of requests that were rejected due to a temporary lack of resource: path (PATH) or subchannel (SCH)
CONFIG – RESOURCE, CNT	number of paths or coupling facility subchannels for each of the following resources: **PATH** number of paths connecting the system and the coupling facility **SCH GEN** number of GENed coupling facility subchannels **SCH USE** number of coupling facility subchannels that are currently in use **SCH OPT** number of coupling facility subchannels that this system can optimally use to satisfy its requests

Figure 44 shows an example of the Structure Activity section of the Coupling Facility Activity Report.

Figure 44 Structure Activity Section of the Coupling Facility Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                   COUPLING FACILITY ACTIVITY REPORT                     RPTSEQ     4 PAGE     98
BMC SOFTWARE, INC.                                              XYZ COMPANY                               REPORT DATE: DD MMM YY 16.58
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00              WORLDWIDE HEADQUARTERS                           SYSTEM ID: **ALL**  COMB-MVS


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-4/364/5,945/16.62
COUPLING FACILITY = ICF91 --------------------  STRUCTURE ACTIVITY SECTION  --------------------------------------------------
---------------------------------------------- STRUCTURE = DSNDB0S_SCA     --------------------------------------------------
 TYPE = LIST    SIZE =     9M
 SYSTEM  ---------------- REQUESTS -----------------  ------------- QUEUED REQUESTS --------------
   NAME    TYPE    PER     COUNT  % OF  -SERVICE TIME--  DELAY        #      %  -- DELAY TIME ---
                   SEC            ALL     AVG   STD DEV  REASON       REQ   REQ   /REQ    S DEV    /ALL
 -------- ------- ------- ------- ----- ------- -------  -------- ------- ----- ------- ------- -------
 SJSC     ALL        1.0   55100  33.9  1592.7  5419.3  NO SCH         2  0.0     215      59       0
           SYNC      0.0     310   0.2   890.7   792.3  PR WT          0  0.0       0       0       0
           ASYNC     1.0   54788  33.7  1596.7  5434.0  PR CMP       293  0.5    3225    2817      17
           CNVAS     0.0       2   0.0                  DUMP           0  0.0       0       0       0

 SJSE     ALL        1.0   52652  32.4  1928.3  3302.8  NO SCH         0  0.0       0       0       0
           SYNC      0.0     242   0.1  1158.9   747.0  PR WT          0  0.0       0       0       0
           ASYNC     1.0   52410  32.2  1931.8  3309.6  PR CMP        11  0.0   21653   11123       5
           CNVAS     0.0       0   0.0                  DUMP           0  0.0       0       0       0

 -------- ------- ------- ------- ----- ------- -------  -------- ------- ----- ------- ------- -------
 SUMMARY  ALL        3.0  162556 100.0  1734.3  4327.7  NO SCH         3  0.0    2851    3728       0
           SYNC      0.0     885   0.5  1016.9   803.0  PR WT          0  0.0       0       0       0
           ASYNC     3.0  161668  99.5  1738.2  4338.8  PR CMP       508  0.3    4538    1171      14
           CNVAS     0.0       3   0.0                  DUMP           0  0.0       0       0       0

--------------------------------------------------- STRUCTURE = DSNDB0X_SCA  ------------------------------------------------
 TYPE = LIST    SIZE =  9472K
 SYSTEM  ---------------- REQUESTS -----------------  ------------- QUEUED REQUESTS --------------
   NAME    TYPE    PER     COUNT  % OF  -SERVICE TIME--  DELAY        #      %  -- DELAY TIME ---
                   SEC            ALL     AVG   STD DEV  REASON       REQ   REQ   /REQ    S DEV    /ALL
 -------- ------- ------- ------- ----- ------- -------  -------- ------- ----- ------- ------- -------
 SJSC     ALL        1.0   53616  33.2  1600.6  5868.8  NO SCH         0  0.0       0       0       0
           SYNC      0.0     275   0.2   987.4   837.3  PR WT          0  0.0       0       0       0
           ASYNC     1.0   53341  33.0  1603.8  5883.5  PR CMP       279  0.5    4032    4947      21
           CNVAS     0.0       0   0.0                  DUMP           0  0.0       0       0       0

 SJSE     ALL        1.0   53888  33.4  1949.9  3858.1  NO SCH         0  0.0       0       0       0
           SYNC      0.0     249   0.2  1201.2   767.8  PR WT          0  0.0       0       0       0
```

Field descriptions of the Structure Activity Section

Table 50 describes each field on the Structure Activity section of the Coupling Facility Activity Report.

Table 50 Field descriptions for the Structure Activity Section (part 1 of 3)

Field	Description
STRUCTURE	name of the structure, up to 16 characters long **Note**: When there are multiple instances of the same structure, the activity data of all instances is combined.
TYPE	first user connecting to a structure specifies its type Possible values for this field are as follows: **LIST** allows users to share data organized as entries in a set of lists **CACHE** allows sharing of frequently referenced data **LOCK** allows serializing on shared resources, including list and cache structures If there is no activity against a structure from any connected system, its type is listed as UNK (for unknown).
SIZE	structure size in either K or M, if the size in K is a multiple of 1024 or greater than 99999
SYSTEM NAME	name of the system that is originating the requests against this structure
REQUESTS – TYPE	type of requests: synchronous (SYNC), asynchronous (ASYNC), originated as synchronous but converted to asynchronous (CNVAS) ALL is the total for all request types.
REQUESTS – PER SEC	rate per second of requests that originated from a system against this structure, viewed by request type
REQUESTS – COUNT	number of requests that originated from a system against this structure, viewed by request type
REQUESTS – % OF ALL	percentage of requests of a type that originated from a system, compared with the total number of requests from all systems against this structure
REQUESTS – SERVICE TIME – AVG, STD DEV	average number of microseconds that are required to complete a request (excluding queue time) and its standard deviation The service time of synchronous-converted-to-asynchronous (CNVAS) requests is included in that of asynchronous (ASYNC) requests.

Table 50 Field descriptions for the Structure Activity Section (part 2 of 3)

Field	Description
QUEUED REQUESTS – DELAY REASON	reason that a request gets queued Possible values for this field are as follows: **NO SCH** no coupling facility subchannel available **DUMP** structure being dumped **PR WT** duplexed request was holding a subchannel while waiting for a peer request to be started **PR CMP** one of the two duplexed operations is completed, but the completed subchannel remains unavailable for use until the peer operation is completed
QUEUED REQUESTS – # REQ	total number of requests to this structure that are delayed in this interval
QUEUED REQUESTS – % REQ	percentage of requests to this structure that are delayed in this interval
DELAY TIME / REQ	average delay time for each delayed request in microseconds
DELAY TIME S DEV	standard deviation of the average delay time for delayed requests
DELAY TIME / ALL	average delay time for all requests to the structure in microseconds This number indicates how delays are affecting overall request activity to the structure.
EXTERNAL CONTENTION – REQ TOTAL	total number of requests to the structure **Note**: This field is applicable only to serialized list or lock structures.
EXTERNAL CONTENTION – REQ DEFERRED	number of requests to the structure that were unable to be completed within the request issuers thread **Note**: This field is applicable only to serialized list or lock structures.
EXTERNAL CONTENTION – TRUE CONT	number of requests to the structure that were delayed due to lock contention This number is a subset of REQ DEFERRED. **Note**: This field is applicable only to lock structures.
EXTERNAL CONTENTION – FALSE CONT	number of requests to the structure that were delayed due to false lock contention False lock contention occurs when a user-supplied hash value is used to map a lock request to a lock entry and multiple unique lock requests hash to a single lock table entry. This number is a subset of -CONT, the total lock contention count. **Note**: This field is applicable only to lock structures. A high false lock contention count indicates either a poor hashing algorithm or an inadequate number of lock table entries.

Table 50 Field descriptions for the Structure Activity Section (part 3 of 3)

Field	Description
DATA ACCESS – READS	total number of read hits for a cache structure read request for the interval **Note**: This field is applicable only to cache structures.
DATA ACCESS – WRITES	total number of write hits for a cache structure write request for the interval **Note**: This field is applicable only to cache structures.
DATA ACCESS – CASTOUTS	total number of castouts for a cache structure for the interval Castouts occur when changed data is written to DASD. The castout count can indicate the volume of changed data being removed from a store-in cache structure, such as the DB2 global buffer pool structures. **Note**: This field is applicable only to cache structures.
DATA ACCESS – XI's	number of times that a data item residing in a local buffer pool was marked invalid by the coupling facility The XI count indicates the amount of data sharing among the users of the cache structure and the amount of write or update activity against the cache data. **Note**: This field is applicable only to cache structures.

Figure 45 shows an example of the CF to CF Activity section of the Coupling Facility Activity Report.

Figure 45 CF to CF Activity Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  COUPLING FACILITY ACTIVITY REPORT                     RPTSEQ     3 PAGE     51
Company X, Inc.                                             Company X, Inc.                              REPORT DATE: DD MMM YY 11.08
ACTL 23 NOV YY 23.44.00  24 NOV YY 14.59.00                  HOUSTON, TX.                                SYSTEM ID: **ALL**  Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-4/364/5,945/16.62
COUPLING FACILITY = ICF92     -------------------------  CF TO CF ACTIVITY  --------------------------------------------------------
 PEER    - CF LINKS -  ----------- REQUESTS ----------- -------------- QUEUED REQUESTS ---------------
    CF     TYPE    USE   PER     COUNT  -SERVICE TIME--             #      %  -- DELAY TIME -----
                         SEC              AVG   STD DEV             REQ   REQ    /REQ   S DEV    /ALL
 -------- -------  ---  ------- ------- ------- -------          ------- ----- ------- ------- -------
CF91    BLOCK      1     62.5 10,193K    31.5    60.5  SYNC          0     0       0       0       0
```

Table 51 describes each field on the CF to CF Activity section of the Coupling Facility Activity Report.

Table 51 Field descriptions for the CF to CF Activity Section

Field	Description
PEER CF	the Remote coupling facility name
CF LINKS	**TYPE**—type of CF to CF link **USE**—number of links used for coupling facility communications
REQUESTS - PER SEC	average number of requests per second This figure is the sum of the signal counters that have been sent from the subject CF to the remote CF. Signal counters used in this calculation are as follows: ■ ready-to-execute signal counter ■ ready-to-complete signal counter ■ halt-execution signal counter ■ request-for-suppression signal counter ■ request-for-suppression accepted signal counter
REQUESTS - COUNT	total number of requests in the interval
REQUESTS - SERVICE TIME	**AVG**—average service time in microseconds **STD DEV**—standard deviation to the average service time
QUEUED REQUESTS	queued requests are synchronous
QUEUED REQUESTS - # REQ	number of signals of all types that have experienced a delay in being sent from the subject CF to this remote CF
QUEUED REQUESTS - % REQ	percentage of requests that are delayed
QUEUED REQUESTS - DELAY TIME	**/REQ**—calculated by taking the average delay time in microseconds over all delayed requests **S DEV**—standard deviation of the average delay time **/ALL**—average delay time in microseconds for all requests—both delayed and not delayed

CPU Utilization Report

The CPU Utilization Report provides information about CPU busy times, CPU online and offline times, and CPU queue lengths.

The CPU Utilization Report is produced by using the CPU Analyzer control statement (see “CPU” on page 230). The data is obtained by using the CPU Extractor control statement (see “CPU” on page 136).

This report consists of these sections:

- **Summary**

  This section describes the activity of all CPUs in the system.

- **CPU**

  This section shows each CPU that was online during the measurement interval. It lists critical information by CPU and produces two graphs that visually represent the listed information. Each CPU ID and serial number are shown at the beginning of this section.

- **Partition Data**

  This section is produced if the Extractor records were gathered in a system running in a PR/SM environment. This section shows logical and physical processor utilization and configuration information by partition. (See “Partition Data Section for PR/SM environments” on page 399 for more information.)

- **LPAR Cluster**

  This section is produced if at least one LPAR cluster is defined on the physical processor. While repeating some of the information that is available in the Partition Data Section, this section provides additional information about partition attributes that are managed by Workload Manager and summarizes data at the LPAR cluster level.

- **LPAR Combination**

  This section displays combinations of LPARs and their highest rolling four-hour MSU per hour. This value helps to assess the license charge of Variable Workload License Charge (VWLC) products that are active either all the time or during periods of peak CPU utilization.

  To add combinations to the LPAR Combination Section or to suppress this section, see the LPARCOMB control statement discussed in “LPARCOMB” on page 274.

- **Rolling 4-Hour MSU Usage Distribution**

  This section helps to determine the defined capacity of an LPAR and evaluates the current defined capacity to help decide if capacity should be increased or decreased.

- **MSU Usage Detail**

  This section is produced if the parameter MSUDETAIL and MSUDIST=YES is specified. With Variable Workload License Charges (VWLC), the capacity for an LPAR is defined in terms of Millions of Service Units (MSUs). This section shows the distribution of four-hour MSU per hour usage.

  All CPU times reported in the CPU Utilization Report are based on measured values.

NOTE

Other CMF MONITOR reports that contain CPU times base their calculations on sampled values, so the CPU-related sample counts might be low in the CPU Utilization Report compared to the other reports. As the number of samples increases, however, the values in all the reports converge. (See Appendix A, “Statistical considerations,” for more information.)

Other reports that provide CPU times based on sampled values are as follows:

- “Channel Path Activity Report” on page 350
- “Extractor Summary Report” on page 443
- All graphics reports:
  - “Distribution Graph” on page 431
  - “Graphics Trace Detail Report” on page 451
  - “Interval Bar Graph” on page 468
  - “Kiviat Graph” on page 470
  - “Performance Summary Report” on page 491
  - “Pie Chart” on page 498
  - “Profile Bar Graph” on page 502
  - “Tabular Subreport” on page 527

Figure 46 shows an example of the Summary and CPU sections of the CPU Utilization Report.

Figure 46 CPU Utilization Report Summary and CPU sections

```
PRODUCED BY CMF ANALYZER (v.r.mm)                       CPU UTILIZATION REPORT                                 RPTSEQ     4 PAGE     12
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                                   REPORT DATE: DD MMM YY 11.09
REQD 27 JAN YY 10.30.00  27 JAN YY 10.59.59                  HOUSTON, TX.                                      SYSTEM ID: SJSD     Z v.rr.n
ACTL 27 JAN YY 10.30.00  27 JAN YY 10.59.59                                                                    REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/2/874/0.5  70-1/2/874/0.5
--------------------------------------------------------  SUMMARY SECTION  --------------------------------------------------------
                          CPU FAMILY: 2094  VERSION: 00  MODEL: 713
                          A CPU QUEUE EXISTED (EXECUTABLE ASIDS)(PCT OF EXT. TIME)                =   9.1
                                       (AVERAGE QUEUE DEPTH =  3.1, MAXIMUM QUEUE SIZE = 69)

                                                          CP                ZAAP                ZIIP
                          LPAR CPU UTILIZATION PERCENTAGE =   44.1               0.2                25.4
                          MVS  CPU UTILIZATION PERCENTAGE =   48.1               0.1                27.8

                                                       BATCH                TSO        STARTED TASK
                          CPU UTILIZATION DISTRIBUTION    =    0.3               2.1                97.6
--------------------------------------------  CPU 00 (TYPE:CP   SERIAL:0422EA) SECTION  --------------------------------------------
CPU UTILIZATION
  ONLINE              99.9
  BUSY PERCENT        44.1                                        CPU BUSY PERCENT                        CPU BUSY DISTRIBUTION
                     % OF TOTAL TIME      % OF CPU BUSY
  TCB MODE                 17.9                 40.7     .****************                         100 |
     SUPVR ENABLED     3.4               7.8             .+++                                            |
     SUPVR DISABLED   13.7              31.1             .++++++++++++                            P      |
     PROB PROG          .8               1.8             .                                        C      |
  LOCAL SRB MODE           26.1                 59.3     .************************                T      |
     ENABLED           7.1              16.1             .++++++                                   75    |
     DISABLED         19.0              43.2             .+++++++++++++++++                       O      |
  GLOBAL SRB MODE            .0                   .0     .                                        F      |
     ENABLED            .0                .0             .                                               |
     DISABLED           .0                .0             .                                        O      |
                                                         .                                        C  50  |
  TOTAL ENABLED TIME       11.3                 25.7     .**********                              C      |
  TOTAL DISABLED TIME      32.7                 74.3     .******************************          U      |
  TOTAL PROB PROG TIME       .8                  1.8     .                                        R      |
  TOTAL SUPVR TIME         43.3                 98.2     .***************************************  R      |
                                                         .                                        E  25  |    *    * * *
% USED FOR BATCH             .2                   .5     .                                        N      |    *    * * *
% USED FOR TSO              1.4                  3.3     .*                                       C      |    *    * * *
% USED FOR STARTED TASKS   42.4                 96.2     .**************************************  E      |    *    * * *
                                                         .                                               |    *    * * *
  AVERAGE DISABLE DELAY (SECONDS)                        -----------------------------------------       ---------------------
          DIE DELAY           =    .001239               0       2       4       6       8       1          0   1   3   5   7   9
          SIG DELAY           =    .000000                       0       0       0       0       0              8   6   4   3   1
          AVG FOR THIS CPU    =    .001239                                                       0
                                                                                                                BUSY PERCENT
  VECTOR FACILITY IS OFF                                                                                   EACH '*' REPRESENTS 5 %
```

Summary Section field descriptions

Table 52 describes each field in the Summary Section of the CPU Utilization Report.

Table 52 Field descriptions for the Summary Section of the CPU Utilization Report

Field	Description
CPU FAMILY VERSION MODEL	CPU family, version, and model
A CPU QUEUE EXISTED	percentage of time that address spaces were found dispatchable and waiting for CPU
AVERAGE QUEUE DEPTH	average number of dispatchable address spaces that were waiting for CPU
MAXIMUM QUEUE SIZE	maximum number of dispatchable address spaces that were waiting for CPU
LPAR CPU UTILIZATION PERCENTAGE	LPAR CPU busy percentage of standard CPs, zAAPs, and zIIPs The formula is the percentage of busy time / online time.
MVS CPU UTILIZATION PERCENTAGE	MVS CPU busy percentage of standard CPs, zAAPs, and zIIPs The formula is **LPAR mode:** (online time - wait time) / online time **Basic mode or under VM:** (interval - wait time) / interval
CPU UTILIZATION DISTRIBUTION	distribution of CPU busy time among batch, TSO, and Started Task address spaces, totaling 100%

CPU Section field descriptions

Two graphs are included in the CPU section of the CPU Utilization Report:

- The graph at the center of the CPU section shows the amount of CPU time spent in states, such as TCB, Local SRB, Global SRB, Supervisor, Problem program, Enabled, Disabled, and others.
- The graph at the left of the CPU section is a distribution graph showing the frequency of occurrence of each percentage busy. The horizontal axis represents the percent busy for this CPU. The vertical axis represents the frequency of that condition. This graph can be used to determine the stability of CPU usage during the sampling period.

Table 53 on page 397 describes each field in the CPU section of the CPU Utilization Report.

Table 53 Field descriptions for the CPU Section (part 1 of 2)

Field	Description
TYPE	type of processor: ■ CP for standard CPs ■ ZAAP for zAAPs ■ ZIIP for zIIPs
SERIAL	serial number of the CPU
ONLINE	percentage of time that this CPU was online; this value is always 100% in a single processor environment
BUSY PERCENT	percentage of time that this CPU was busy
% OF TOTAL TIME	portion of the total time that elapsed when the CPU was online and in the indicated state
% OF CPU BUSY	portion of the CPU busy time that was spent in the indicated state
TCB MODE	sum of the percentages of time that was spent in supervisor state enabled, disabled, or in problem program state while running under a TCB
SUPVR ENABLED	percentage of time that was spent in supervisor state enabled while running under a TCB
SUPVR DISABLED	percentage of time that was spent in supervisor state disabled while running under a TCB
PROB PROG	percentage of time that was spent in problem program state while running under a TCB
LOCAL SRB MODE	sum of the percentages of time that the system was enabled or disabled while running under a local SRB
LOCAL SRB MODE - ENABLED	percentage of time that the system was enabled while running under a local SRB
LOCAL SRB MODE - DISABLED	percentage of time that the system was disabled while running under a local SRB
GLOBAL SRB MODE	sum of the percentages of time that the system was enabled or disabled while running under a global SRB
GLOBAL SRB MODE - ENABLED	percentage of time that the system was enabled while running under a global SRB
GLOBAL SRB MODE - DISABLED	percentage of time that the system was disabled while running under a global SRB
TOTAL ENABLED TIME	total percentage of time that the system was enabled
TOTAL DISABLED TIME	total percentage of time that the system was disabled; all measured disabled time values are elevated in a PR/SM environment due to CPU switching
TOTAL PROB PROG TIME	total percentage of time that the system spent in a problem program state
TOTAL SUPVR TIME	total percentage of time that the system spent in a supervisor state
% USED FOR BATCH	percentage of time on this CPU that was spent on batch
% USED FOR TSO	percentage of time on this CPU that was spent on TSO
% USED FOR STARTED TASKS	percentage of time on this CPU that was spent on Started Tasks

Table 53 Field descriptions for the CPU Section (part 2 of 2)

Field	Description
DIE DELAY	average time that the Extractor was delayed in getting control of this CPU because another disabled function was running.
SIG DELAY	average time that the Extractor was delayed in getting control of another CPU in the system because a disabled function was running on that CPU This value is produced only in a multiple processor environment.
AVG FOR THIS CPU	average time that the Extractor was delayed in beginning and completing a sample while executing on this CPU, where Delay in Beginning—DIE delay time on this CPU Delay in Completing—SEGL delay time waiting to sample the status of another CPU; this value is displayed only in a multiple processor environment The following are examples of delays that can occur in the system: ■ The Extractor gains control under the DIE on CPU 0 after a delay of 300 microseconds. ■ CPU 1 is sampled by RISGNL. The sampling on CPU 1 is delayed by 100 microseconds. ■ The Extractor later gains control under the DIE on CPU 1 after a delay of 50 microseconds. ■ CPU 0 is sampled by RISGNL. The sampling on CPU 0 is delayed by 280 microseconds. **Note**: CMF MONITOR considers a CPU to be disabled in any of three cases: 1. A CPU is physically disabled whenever it masks off external and I/O interruptions. 2. A CPU is logically disabled whenever it holds a global spin lock. 3. A CPU is switched out of partition and is unavailable to service the interrupt.
VECTOR FACILITY IS	ON indicates Vector Facility is online; OFF indicates Vector Facility is offline
VECTOR FACILITY AFFINITY	amount of Vector Facility used by the CPU for the specified interval

Partition Data Section for PR/SM environments

The CPU Utilization Report provides a Partition Data Section if your records were generated in a PR/SM environment. It shows logical and physical processor utilization by partition. The partitions are grouped by the type of physical processors that they use: standard CP, ICF, zAAP (zSeries Application Assist Processor), zIIP (zSeries Integrated Information Processor), and so on. A partition having multiple processor types appears multiple times, once for each processor type that it has. The last partition reported for each processor type is always PHYSICAL. PHYSICAL is the PR/SM overhead that cannot be attributed to any individual partition.

NOTE

If a partition has a mixture of shared and dedicated processors, in addition to an existing line for each partition, two additional lines show data similar to the partition level line but separately for shared and dedicated processors.

An example of the Partition Data Section of the CPU Utilization Report for a PR/SM environment is shown in Figure 47 on page 400.

Figure 47 Partition Data Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        CPU UTILIZATION REPORT                              RPTSEQ      3 PAGE       5
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                                REPORT DATE: DD MMM YY  7.50
ACTL DD MMM YY 02.00.46  DD MMM YY 03.15.00                   HOUSTON, TX.                                   SYSTEM ID: SJSD      Z 1.06.1


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/5/4,447/1.24  70-1/5/4,447/1.24
------------------------------------------------------ PARTITION DATA SECTION -----------------------------------------------------
                        MVS PARTITION NAME   =  SJSD                    NO. OF PARTITIONS    =       24
                        IMAGE CAPACITY       =   122                    DISPATCH INTERVAL    =  DYNAMIC
                        PHYSICAL PROCESSORS  =    17
                                 CP          =    13
                                 ZAAP        =     1
                                 IFL         =     0
                                 ICF         =     2
                                 ZIIP        =     1

                           WAIT  -LOG PRCR-  ----- MSU ------  -CAPPING-   -% DISPATCHED-
PARTITION  STATUS  WEIGHT  COMP   CNT  TYPE   DEF  USED  %DEF  %WLM  HARD     LOG     PHY              CPU DISPATCH PERCENT
---------  ------  ------  ----  ----  ----  ----  ----  ----  ----  ----  ------  ------  -----------------------------------------
SJSD       ACTIVE     MIX  NO     1.0  CP     ---    15   ---   ---  NO      23.7     1.8  .
SJSB       ACTIVE     100  NO     2.0  CP     ---    10   ---   ---  NO       8.2     1.3  .
SJSC       ACTIVE     120  NO     2.0  CP     ---    15   ---   ---  NO      12.4     1.9  .
SJSG       ACTIVE      60  NO     1.0  CP     ---     5   ---   ---  NO       7.8     0.6  .
SJSH       ACTIVE      60  NO     1.0  CP     ---     5   ---   ---  NO       7.9     0.6  .
SYSI       ACTIVE     100  NO     1.0  CP     ---    33   ---   ---  NO      54.3     4.2  .PP
SYSM       ACTIVE     150  NO     2.0  CP     ---    18   ---   ---  NO      14.4     2.2  .P
SYSN       ACTIVE     170  NO     2.0  CP     ---    22   ---   ---  NO      17.9     2.8  .P
SYSP       ACTIVE     MIX  NO       2  CP     ---    69   ---   ---  NO      56.8     8.7  .PPP
VMR        ACTIVE     220  NO       2  CP     ---    49   ---   ---  NO      39.9     6.1  .PP
VM5        ACTIVE     200  NO       2  CP     ---    55   ---   ---  NO      45.0     6.9  .PP
DB2A       ACTIVE     DED  N/A      2  CP     ---   ---   ---   ---  NO     100.0    15.4  .PPPPPP
DB2B       ACTIVE     MIX  NO     2.0  CP     ---    66   ---   ---  NO      53.9     8.3  .PPP
ESAJ       ACTIVE      90  NO     2.0  CP     ---    19   ---   ---  NO      15.3     2.3  .P
ESAM       ACTIVE      90  NO     2.0  CP     ---     6   ---   ---  NO       4.5     0.7  .
IMSA       ACTIVE     100  NO     2.0  CP     ---    10   ---   ---  NO       7.8     1.2  .
SJSE       ACTIVE     MIX  NO     2.0  CP     ---    33   ---   ---  NO      26.9     4.1  .PP
SYSO       ACTIVE     150  NO     2.0  CP     999    59   5.9   0.0  NO      48.3     7.4  .PPP
VM4        ACTIVE     155  NO       2  CP     ---    26   ---   ---  NO      21.1     3.2  .P
ZOSE       ACTIVE       1  NO       1  CP     ---   ---   ---   ---  NO       0.0     0.0  .
PHYSICAL                                                                          8.3  .PPP
                                                                               ------  .
                                                                      TOTAL      88.0  .PPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPP
                                                                                       .
SJSD       ACTIVE     MIX  NO       1  ZAAP   ---   ---   ---   ---  NO       0.2     0.2  .
SYSM       ACTIVE     150  NO       1  ZAAP   ---   ---   ---   ---  NO       0.9     0.9  .
SYSP       ACTIVE     MIX  NO       1  ZAAP   ---   ---   ---   ---  NO       0.2     0.2  .
DB2B       ACTIVE     MIX  NO       1  ZAAP   ---   ---   ---   ---  NO       0.2     0.2  .
ESAJ       ACTIVE      90  NO       1  ZAAP   ---   ---   ---   ---  NO       2.1     2.1  .P
IMSA       ACTIVE     100  NO       1  ZAAP   ---   ---   ---   ---  NO       0.2     0.2  .
SJSE       ACTIVE     MIX  NO       1  ZAAP   ---   ---   ---   ---  NO       1.4     1.4  .
SYSO       ACTIVE     150  NO       1  ZAAP   ---   ---   ---   ---  NO       0.3     0.3  .
PHYSICAL                                                                          6.8  .PP
                                                                               ------  .
                                                                      TOTAL      12.3  .PPPPP
                                                                                       .
CF11       ACTIVE     500  NO       1  ICF    ---   ---   ---   ---  NO      12.5     6.2  .PP
CF12       ACTIVE     DED  N/A      1  ICF    ---   ---   ---   ---  NO     100.0    50.0  .PPPPPPPPPPPPPPPPPPPP
CF13       ACTIVE     100  NO       1  ICF    ---   ---   ---   ---  NO       0.3     0.1  .
CF16       ACTIVE     375  NO       1  ICF    ---   ---   ---   ---  NO      11.0     5.5  .PP
PHYSICAL                                                                         11.8  .PPPP
                                                                               ------  .
                                                                      TOTAL      73.6  .PPPPPPPPPPPPPPPPPPPPPPPPPPPPPP
                                                                                       .
SJSD       ACTIVE     MIX  NO       1  ZIIP   ---   ---   ---   ---  NO       0.2     0.2  .
SYSM       ACTIVE     150  NO       1  ZIIP   ---   ---   ---   ---  NO       0.1     0.1  .
SYSN       ACTIVE     170  NO       1  ZIIP   ---   ---   ---   ---  NO       0.2     0.2  .
DB2B       ACTIVE     MIX  NO       1  ZIIP   ---   ---   ---   ---  NO       1.0     1.0  .
ESAJ       ACTIVE      90  NO       1  ZIIP   ---   ---   ---   ---  NO       0.2     0.2  .
IMSA       ACTIVE     100  NO       1  ZIIP   ---   ---   ---   ---  NO       0.2     0.2  .
SJSE       ACTIVE     MIX  NO       1  ZIIP   ---   ---   ---   ---  NO       0.3     0.3  .
PHYSICAL                                                                          5.1  .PP
                                                                               ------  .
                                                                      TOTAL       7.3  .PPP
                                                                                       .
                                                                                       -----------------------------------------
                                                                                       0       2       4       6       8       1
                                                                                               0       0       0       0       0
```

Partition Data Section field descriptions

Table 54 describes each field in the Partition Data section of the CPU Utilization Report.

Table 54 Field descriptions for the Partition Data Section (part 1 of 2)

Field	Description
MVS PARTITION NAME	name of the home partition
NO OF PARTITIONS	number of PR/SM partitions that are defined for this processor complex
IMAGE CAPACITY	CPU capacity in millions of service units (MSU) per hour
DISPATCH INTERVAL	time slice interval in milliseconds The word DYNAMIC indicates a variable time slice interval.
PHYSICAL PROCESSORS	number of physical processors in this processor complex If processor types are identified by hardware (CP, ICF, IFL, ...), the next lines list the | number of physical processors of each type.
PARTITION	up to eight-character partition name **Note**: Partition PHYSICAL, if present, is actually not a logical partition; rather, it is a placeholder for PR/SM overhead that cannot be attributed to any individual partition.
STATUS	state of the partition, either ACTIVE or INACTIVE, at the end of the report interval If the partition was inactive at the end of the report interval but had been previously active, the rest of the line is blank except for the AVG % DISPATCHED columns.
WEIGHT	dispatcher weight that is assigned to a partition Possible values for this field are as follows: ***nnn*** — all logical processors not dedicated; they have same weight of *nnn* **MIX** — all logical processors not dedicated; they do not have same weights **DED** — all logical processors dedicated **DNE** — a mixture of dedicated and nondedicated processors; nondedicated processors have equal weights **DNN** — a mixture of dedicated and nondedicated processors; nondedicated processors do not have equal weights

Table 54 Field descriptions for the Partition Data Section (part 2 of 2)

Field	Description	
WAIT COMP	wait completion attribute of nondedicated logical processors that are assigned to a partition Possible values for this field are as follows:	
	YES	processor dispatched to a partition remains dispatched until the time slice expires
	NO	processor dispatched to a partition is returned to PR/SM and becomes available to other partitions as soon as the partition enters an enabled wait state
	MIX	partition has a mixture of nondedicated logical processors with wait completion attributes of YES and NO
	N/A	partition has only dedicated processors
LOG PRCR	**CNT**	if partition is managed by WLM, the average number of online processors; if partition is not managed, the number of logical processors defined for the partition
	TYPE	processor types displayed if identified by hardware **Note:** zAAPs are displayed as ZAAP (on z9 or later CPUs) or ICF (on pre-z9 CPUs). zIIPs are displayed as ZIIP.
MSU	**DEF**	defined capacity of the partition in MSU/hour Defined capacity can be specified for a partition running z/OS on a z900 machine. Three dashes (---) appear if the partition does not have defined capacity.
	USED	MSU per hour actually used
	%DEF	percentage of used MSU/hour over defined capacity Three dashes (---) appear if the partition does not have defined capacity.
These three columns are applicable only for standard CPs:		
CAPPING	**%WLM**	percentage of time that WLM capped the partition Three dashes (---) appear if the partition does not have defined capacity.
	DEF	YES indicates that hard cap was specified for the partition; NO indicates that the partition was not hard capped
% DISPATCHED	**LOG**	percentage of dispatch time over online time of all logical processors that are defined for the partition
	PHY	percentage of dispatch time over online time of all physical processors in the complex
CPU DISPATCH PERCENT	a graphic representation of the logical and physical processor dispatched percentage, where L shows the logical value and P shows the physical value	

LPAR Cluster Section

The CPU Utilization Report provides an LPAR Cluster Section if at least one LPAR cluster is defined on the central processor complex (CPC). An LPAR cluster is a set of logical partitions that are resident on the same CPC and in the same parallel sysplex. Workload Manager (WLM) manages processor access based on workload goals by dynamically adjusting the weights of the partitions within an LPAR cluster and optimizing the number of online processors for each partition.

The LPAR Cluster section contains CPU utilization and the partition attributes that are managed by Workload Manager (weight and number of online processors), and it summarizes data at the LPAR cluster level.

NOTE

The LPAR Cluster section contains only information about standard CPs because WLM manages only this processor type.

An example of the LPAR Cluster Section of the CPU Utilization Report is shown in Figure 48.

Figure 48 LPAR Cluster Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        CPU UTILIZATION REPORT                                 RPTSEQ    14 PAGE     99
BMC SOFTWARE, INC.                                             XYZ COMPANY                                      REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00              WORLDWIDE HEADQUARTERS                                 SYSTEM ID: SJSE     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  70-1/56/25.1K/7
-----------------------------------------------------  LPAR CLUSTER SECTION  -------------------------------------------------------
                               ----------- WEIGHT ------------  -------- PROCESSOR --------  - STORAGE (MB) --
                               -- DEFINED ---  --- ACTUAL ----  --- NUMBER ----  - % BUSY -  CENTRAL  EXPANDED
CLUSTER   PARTITION  SYSTEM    INIT  MIN  MAX  AVG  %MIN  %MAX  DEFINED  ACTUAL   LOG   PHY

BBPLEX01  SJSD       SJSD        30   10   40   30   0.0   0.0        5     4.0   3.6   1.0      768         0
          SJSE       SJSE        30   10   40   30   0.0   0.0        3     3.0   6.2   1.3      768         0
-----------------------------  ------------------------------  --------------------------  -----------------
                     TOTAL       60                                                     2.4    1,536         0
                                                                                                                                  0
```

LPAR Cluster Section field descriptions

Table 55 describes each field in the LPAR Cluster Section of the CPU Utilization Report.

Table 55 Field descriptions for the LPAR Cluster Section

Field	Description	
CLUSTER	name of the LPAR cluster	
PARTITION	name of the partition	
SYSTEM	MVS system name	
WEIGHT - DEFINED	INIT	initial weight that the partition has when the system is IPLed The TOTAL line shows the weight of the LPAR cluster.
	MIN, MAX	minimum and maximum weight that the partition can have Three dashes (---) appear if the partition is not WLM-managed.
WEIGHT - ACTUAL	AVG	average weight of the partition
	%MIN	percentage of time that the actual weight was in the minimum bandwidth of minimum weight to minimum weight plus 10% Three dashes (---) appear if the partition is not WLM-managed.
	%MAX	percentage of time that the actual weight was in the maximum bandwidth of maximum weight minus 10% to maximum weight Three dashes (---) appear if the partition is not WLM-managed.
PROCESSOR - NUMBER	DEFINED	maximum number of logical processors that can be online for the partition
	ACTUAL	average number of logical processors that were actually online for the partition
PROCESSOR - % BUSY	LOG	percentage of busy time over online time of logical processors
	PHY	percentage of busy time over online time of all physical processors The TOTAL line shows the CPU utilization of the LPAR cluster.
STORAGE (MB) - CENTRAL EXPANDED	central and expanded storage, in megabytes, assigned to the partition The TOTAL line shows the total storage for the LPAR cluster.	

LPAR Combination Section

This section displays the combination of LPARs that are specified on the LPARCOMB statement and their highest rolling four-hour MSU per hour. This value helps assess the license charge of those Variable Workload License Charge (VWLC) products that are active either all the time or during periods of peak CPU utilization. The combination with the description ALL LPARS (consisting of all LPARs), when present, is displayed first. Other combinations are displayed in the same order as the LPARCOMB statements that specify the combinations.

An example of the LPAR Combination Section of the CPU Utilization Report is shown in Figure 49. This sample section is produced by these control statements:

```
CPU

LPARCOMB DESC='DB2 SYSTEMS',LPARS=(DB2A,DB2B)

LPARCOMB DESC='SAN JOSE SYSTEMS',LPARS=(SJSD,SJSE)
```

Figure 49 LPAR Combination Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                       CPU UTILIZATION REPORT                          RPTSEQ     3 PAGE      8
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                            REPORT DATE: DD MMM YY 17.01
ACTL 10 JUN YY 00.25.26  10 JUN YY 17.00.00                  HOUSTON, TX.                               SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/67/29.8K/16.58  70-1/67/29.8K/16.58
----------------------------------------------------  LPAR COMBINATION SECTION  -----------------------------------------------------
                                  -HIGH 4-HOUR MSU/HOUR-  ------------------------ LPAR ------------------------  --EXCLUDED LPARS--
    COMBINATION DESCRIPTION        MSU   TIME    DATE     LPAR      MSU  LPAR      MSU  LPAR      MSU  LPAR      MSU    LPAR     REASON
--------------------------------  ----  -----  ---------  ------------  ------------  ------------  ------------  --------  --------
ALL LPARS                          440  17:00  12 DEC 02  SYSI      41  SYSN      20  ESAJ      38  VMR       74  CF04      ICF PROC
                                                          DB2A      43  ESAM      21  SJSE      61  SYSP      79  CF03      ICF PROC
                                                          IMSA      18  DB2B      34  SJSD      11

DB2 SYSTEMS                         77  17:00  12 DEC 02  DB2A      43  DB2B      34

SAN JOSE SYSTEMS                    72  17:00  12 DEC 02  SJSD      11  SJSE      61                              SJSZ      ABSENT
```

LPAR Combination Section field descriptions

Table 56 describes each field in the LPAR Combination Section of the CPU Utilization Report.

Table 56 Field descriptions for the LPAR Combination Section

Field	Description	
COMBINATION DESCRIPTION	description of the LPAR combination An asterisk (*) following this description indicates that the description is not unique among the listed LPAR combinations.	
HIGH 4-HOUR MSU/HOUR	MSU	highest four-hour MSU / hour consumed by the LPAR combination If the input records do not cover any four-hour periods or no LPAR member of the combination had been active for four contiguous hours, NO 4-HOUR PERIODS is displayed.
	TIME, DATE	time and date when the highest four-hour MSU / hour value occurred
LPARS	LPAR	name of the LPARs in the combination (including LPARs that were always inactive)
	MSU	four-hour MSU / hour of this LPAR at the time when the highest four-hour MSU / hour of the LPAR combination occurred If three dashes (---) appear, this LPAR had not been active for four hours.
EXCLUDED LPARS	LPARs that were excluded from the calculation of the highest four-hour MSU / hour	
LPAR	name of LPAR	
REASON	ABSENT	LPAR was not found in the input SMF type 70 records
	ICF PROC	LPAR used ICF processors and hence is not MVS partitions
	DED PROC	a remote LPAR had dedicated processors; there is not enough information in the SMF type 70 CPU record to calculate MSU usage
	WAIT YES	a remote LPAR had the attribute Wait Completion set to Yes; there is not enough information in the SMF type 70 CPU record to calculate MSU usage
	Note: If only one remote LPAR is excluded because of DED PROC or WAIT YES, generate this report on the SMF records produced on the excluded LPAR. The LPAR will then be the local LPAR, and MSU usage can then be calculated.	

Rolling 4-Hour MSU Usage Distribution Section

This section shows the distribution of four-hour MSU per hour usage for the busiest 10% of rolling four-hour periods.

If a partition does not have a defined capacity, this section helps you determine the partition value. For the listed MSU values, CMF MONITOR calculates the expected percent of time that soft capping is in effect for the partition if the defined capacity is set to one of these MSU values. To use this section, decide what percentage of time is acceptable for Workload Manager to cap the partition, and then look for that percentage on the PROJECTED % CAPPED BY WLM line. The MSUs line shows the MSU value to be used for the defined capacity. For example

```
MSUs                                                   33
                                                     -----
% OF 4-HOUR PERIODS USING MORE THAN THE MSU AMOUNT    0.4
PROJECTED % CAPPED BY WLM                             0.8
```

In this example, 0.4% of four-hour periods consumed more than 33 MSU per hour. If the defined capacity is set at 33 MSU per hour, it is projected that the partition is capped by Workload Manager 0.8% of the time. In other words, if you would like the partition to be capped 0.8% of the time, set the defined capacity of the partition at 33 MSU per hour.

If a partition already has defined capacity, you can adjust its value by using this section and the % Capped by WLM column (**% WLM**) in the Partition Data Section.

Calculation method for PROJECTED % CAPPED BY WLM

Every five minutes, Workload Manager calculates the MSU per hour usage averaged over the previous four-hour period. If the average exceeds the defined capacity, the partition is capped for the next five minutes; that is, it can only use as much CPU as set by defined capacity. At the end of this (next) five-minute period, Workload Manager calculates the four-hour average again. If it drops below defined capacity, the capping is removed; otherwise, the partition stays capped.

Similarly, CMF MONITOR calculates the four-hour MSU per hour average at the end of every recording interval *n*. If the average exceeds the tentative defined capacity, the partition is considered capped for the entire length of the next recording interval *n*+1. (If the partition used more than the defined capacity in this next interval, the excess usage is charged to the following interval *n*+2.) At the end of the recording interval *n*+1, CMF MONITOR calculates the four-hour average again. If it drops below tentative defined capacity, the partition is no longer considered capped.

An example of the Rolling 4-Hour MSU Usage Distribution Section of the CPU Utilization Report is shown in Figure 50.

Figure 50 Rolling 4-Hour MSU Usage Distribution Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                       CPU UTILIZATION REPORT                              RPTSEQ     3 PAGE      6
BMC SOFTWARE, INC.                                           XYZ COMPANY                                    REPORT DATE: DD MMM YY 14.46
ACTL 10 JUN YY 04.00.00  10 JUL YY 17.00.00             WORLDWIDE HEADQUARTERS                              SYSTEM ID: SJSB     02.08.00

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  70-1/2,645/475K/659.7
-----------------------------------------  ROLLING 4-HOUR MSU USAGE DISTRIBUTION SECTION  ------------------------------------------
CPU FAMILY: 9672  VERSION: A5  MODEL: Z57               CEC CAPACITY =    149
PARTITION = SJSC       NUMBER OF 4-HOUR PERIODS =   2,481
   CONSUMED MSU                                            45     44     43     42     41     40     39     38     37     36
                                                        -----  -----  -----  -----  -----  -----  -----  -----  -----  -----
   % OF 4-HOUR PERIODS USING MORE THAN THE CONSUMED MSU    0.2    0.2    0.4    0.6    0.8    0.8    1.1    2.1    3.1    3.3
   PROJECTED % CAPPED BY WLM                               0.2    0.2    0.4    0.5    0.7    0.8    1.1    2.2    3.1    3.7

   CONSUMED MSU                                            35     34     33     32     31     30     29     28     27     26
                                                        -----  -----  -----  -----  -----  -----  -----  -----  -----  -----
   % OF 4-HOUR PERIODS USING MORE THAN THE CONSUMED MSU    3.7    4.4    5.5    6.3    7.2    8.1    8.8    9.2    9.7   10.4
   PROJECTED % CAPPED BY WLM                               4.7    5.2    6.3    7.1    8.0    9.1   10.0   11.0   12.1   13.3

PARTITION = SJSD       NUMBER OF 4-HOUR PERIODS =   2,480
   CONSUMED MSU                                            19     18     17     16     15     14     13     12     11     10
                                                        -----  -----  -----  -----  -----  -----  -----  -----  -----  -----
   % OF 4-HOUR PERIODS USING MORE THAN THE CONSUMED MSU    0.4    0.6    1.0    1.1    1.3    1.5    1.7    2.3    3.7    5.5
   PROJECTED % CAPPED BY WLM                               0.3    0.6    0.9    1.1    1.3    1.5    1.7    2.4    3.8    6.8

   CONSUMED MSU                                             9      8
                                                        -----  -----
   % OF 4-HOUR PERIODS USING MORE THAN THE CONSUMED MSU    8.8   12.8
   PROJECTED % CAPPED BY WLM                              11.5   18.5

PARTITION = SJSG       NUMBER OF 4-HOUR PERIODS =   2,497
   CONSUMED MSU                                             5      4      3
                                                        -----  -----  -----
   % OF 4-HOUR PERIODS USING MORE THAN THE CONSUMED MSU    0.8    6.4   28.2
   PROJECTED % CAPPED BY WLM                               0.9    9.3   43.0
```

Rolling 4-Hour MSU Usage Distribution Section field descriptions

Table 57 describes each field in the Rolling 4-Hour MSU Usage Distribution Section of the CPU Utilization Report.

Table 57 Field descriptions for the Rolling 4-Hour MSU Usage Distribution Section

Field	Description
CPU FAMILY VERSION MODEL	CPU family, version, and model
CEC CAPACITY	capacity of all CPs in millions of service units (MSUs) per hour CPs are those available for MVS partitions.
PARTITION	partition name
NUMBER OF 4-HOUR PERIODS	number of four-hour periods
CONSUMED MSU	MSUs per hour consumed by rolling four-hour periods
% OF PERIODS USING MORE THAN THE CONSUMED MSU	percentage of rolling four-hour periods that are consuming more than the MSU amount **Note**: CMF MONITOR displays the MSU per hour values of the top 10% periods, not to exceed 20 values.
PROJECTED % CAPPED BY WLM	percentage of time projected to be capped by Workload Manager if the defined capacity equals the MSU amount **Note**: This field is shown only if the partition does not have defined capacity.

MSU Usage Detail Section

This optional section displays MSU usage at the recording interval level in order to identify peak CPU usage occurrences and patterns such as time of day, day of month, and so on. The intervals are ordered depending on the value of the MSUDTAIL parameter specified on the Analyzer CPU control statement:

- in decreasing four-hour MSU per hour (MSU per hour consumed in the four-hour period up to the end of the recording interval) if MSUDTAIL=4HOURMSU is specified
- in decreasing interval MSU per hour (MSU per hour consumed during the recording interval) if MSUDTAIL=INTVLMSU is specified
- in decreasing % capped by WLM during the recording interval if MSUDTAIL=INTVLWLM is specified
- in increasing time order if MSUDTAIL=TIME is specified

An example of the MSU Usage Detail Section of the CPU Utilization Report is shown in Figure 51.

Figure 51 MSU Usage Detail Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        CPU UTILIZATION REPORT                                RPTSEQ     3 PAGE      9
BMC SOFTWARE, INC.                                             XYZ COMPANY                                     REPORT DATE: DD MMM YY 14.08
ACTL 10 JUN YY 16.15.00  11 JUN YY 16.30.00              WORLDWIDE HEADQUARTERS                                SYSTEM ID: SJSE     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  70-1/192/17.2K/23.87
PARTITION: DB2A     ------------------------------------- MSU USAGE DETAIL ----------------------------------------------------------
 --- INTERVAL ---  INTVL   4-HR   %CAP     --- INTERVAL ---  INTVL   4-HR   %CAP     --- INTERVAL ---  INTVL   4-HR   %CAP
   DATE      TIME    MSU    MSU    WLM       DATE      TIME    MSU    MSU    WLM       DATE      TIME    MSU    MSU    WLM
 11 JUN 03  16:30   43.6   40.0    ---     11 JUN 03  16:15   43.9   38.1    ---     11 JUN 03  16:00   49.0   36.5    ---
 11 JUN 03  15:45   40.0   34.5    ---     11 JUN 03  15:30   35.3   32.7    ---     11 JUN 03  15:15   44.0   31.2    ---
 11 JUN 03  15:00   45.9   29.8    ---     11 JUN 03  14:45   43.6   27.6    ---     11 JUN 03  14:30   44.7   26.1    ---
 11 JUN 03  11:15   22.3   25.0    ---
```

MSU Usage Detail Section field descriptions

Table 58 describes each field in the MSU Usage Detail section of the CPU Utilization Report.

Table 58 Field descriptions for the MSU Usage Detail Section

Field	Description
PARTITION	partition name
INTERVAL - DATE, TIME	date and time at the end of the recording interval
INTVL MSU	MSU / hour consumed during the recording interval
4-HR MSU	MSU / hour consumed in the four-hour period up to the end of the recording interval
INTVL %WLM	percentage of time capped by WLM during the recording interval Three dashes (---) appear if the partition does not have defined capacity.

CPU Utilization by Protect Key Report

The CPU Utilization by Protect Key Report shows CPU status by storage protection key.

The CPU Utilization by Protect Key Report is produced by using the PROTKEY Analyzer control statement (see “PROTKEY” on page 284). The data is obtained by using the CPU Extractor control statement (see “CPU” on page 136).

There are three sections in the CPU Utilization by Protect Key Report:

- **Sub System**

  This section gives information about protect keys 0 through 7.

- **V = V**

  This section shows all jobs running under key 8 for V=V.

- **V = R or Special V = V**

  This section shows all protect keys that had either nonpageable jobs that were running under them or pageable jobs that were not in key 8. Keys 9 through F can appear in this section as well.

If dashes (---) appear for a given protect key under the CPU column, no activity information was encountered during sampling. If dashes appear for all protect keys under a given CPU, that CPU does not exist. For purposes of this report, you must specify the ID of the CPU to be examined.

Figure 52 on page 412 shows an example of the CPU Utilization by Protect Key Report.

Figure 52 CPU Utilization by Protect Key Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                CPU UTILIZATION BY PROTECT KEY REPORT                RPTSEQ     4 PAGE     25
 BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 14.08
 ACTL 10 JUN YY 16.15.00  11 JUN YY 16.30.00            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/96/17.2K/23.87
 ---------------------------------------------------   CPU  0: SUB SYSTEM SECTION  -------------------------------------------------
                                    KEY        SUB SYSTEM          BUSY        SUPRV        PROBL

                                     0         SUPERVISOR           5.5        100.0           .0
                                     1         SCHEDULER,JES         .0        100.0           .0
                                     3         RESERVED              .0           .0        100.0
                                     4         RESERVED              .5         98.8          1.2
                                     5         DATA MANAGEMENT       .0        100.0           .0
                                     6         TCAM/VTAM             .1        100.0           .0
                                     7         IMS/DB2               .2        100.0           .0
-------------------------------------------------------  CPU  0: V = V SECTION  ----------------------------------------------------
                                    KEY        SUB SYSTEM          BUSY        SUPRV        PROBL

                                     8         V = V                 .2         19.4         80.6
-*************************************************************************************************************************************
----------------------------------------------------   CPU  1: SUB SYSTEM SECTION  -------------------------------------------------
                                    KEY        SUB SYSTEM          BUSY        SUPRV        PROBL

                                     0         SUPERVISOR           5.4        100.0           .0
                                     1         SCHEDULER,JES         .0        100.0           .0
                                     4         RESERVED              .7        100.0           .0
                                     5         DATA MANAGEMENT       .0        100.0           .0
                                     6         TCAM/VTAM             .1        100.0           .0
                                     7         IMS/DB2               .2        100.0           .0
```

CPU Utilization by Protect Key Report field descriptions

NOTE

The same columns of data appear in each section of this report.

Table 59 describes each field in the CPU Utilization by Protect Key Report.

Table 59 Field descriptions for the CPU Utilization by Protect Key Report

Field	Description
KEY	protect key being measured; its value is 0 through F
SUB SYSTEM	name of the system function that usually resides in this protect key
BUSY	percentage of CPU time in this protect key
SUPRV	percentage of CPU busy time in this protect key in supervisor state
PROBL	percentage of CPU busy time in this protect key in a problem state

Cross-System Coupling Facility Report

The Cross-System Coupling Facility Report shows the activity of the Cross-System Coupling Facility (XCF). The information can be presented in a summarized format for the system and individually by the member.

The Cross-System Coupling Facility Report is produced by using the XCF Analyzer control statement (see “XCF” on page 323). The data for this report is obtained by using the XCFDATA Extractor control statement (see “XCFDATA” on page 205).

There are three sections to this report:

- System Summary
- Path Utilization Section
- Detail Report

Figure 53 shows an example of the Cross-System Coupling Facility Report.

Figure 53 Cross-System Coupling Facility Report (part 1 of 2)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                CROSS-SYSTEM COUPLING FACILITY REPORT                  RPTSEQ     3 PAGE      5
 BMC SOFTWARE, INC.                                             XYZ COMPANY                             REPORT DATE: DD MMM YY 17.41
 ACTL 10 JUN YY 09.00.00  10 JUN YY 12.30.00             WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSC     02.09.00

 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-2/14/420/3.5
-------------------------------------------------------- SYSTEM SUMMARY ---------------------------------------------------------

                                                          -- BUFFER EFFICIENCY PERCENTAGES --
                                                                                               MESSAGES
  FROM     TO        TRANSPORT     BUFFER    REQUESTS                 MSG GT   MSG LT   MSG =    MIGRATED TO     MESSAGES
 SYSTEM   SYSTEM       CLASS       LENGTH   SATISFIED     DEGRADED    BUFFER   BUFFER   BUFFER    ALT PATH       REJECTED

 SJSA     SJSC        <INBOUND>   <INBOUND>          0          0.0      0.0      0.0      0.0            0             0
 SJSB     SJSC        <INBOUND>   <INBOUND>          0          0.0      0.0      0.0      0.0            0             0
 SJSC     SJSA        DEFAULT           956     41,574         93.0     14.8      0.0     85.2            0             0
 SJSC     SJSB        DEFAULT           956     34,791         90.4     17.6      0.0     82.4            0             0
 SJSC     SJSC        DEFAULT           956     34,438        100.0      7.8      0.0     92.2            0             0
 SJSC     SJSD        DEFAULT           956     43,061         97.1     15.9      0.0     84.1            0             0
 SJSC     SJSE        DEFAULT           956     70,362         93.8     17.8      0.0     82.2            0             0
 SJSC     SJSG        DEFAULT           956     38,687         96.4     12.6      0.0     87.4            0             0
 SJSC     SJSH        DEFAULT           956     22,166         95.0     15.9      0.0     84.1            0             0
 SJSD     SJSC        <INBOUND>   <INBOUND>          0          0.0      0.0      0.0      0.0            0             0
 SJSE     SJSC        <INBOUND>   <INBOUND>          0          0.0      0.0      0.0      0.0            0             0
 SJSG     SJSC        <INBOUND>   <INBOUND>          0          0.0      0.0      0.0      0.0            0             0
 SJSH     SJSC        <INBOUND>   <INBOUND>          0          0.0      0.0      0.0      0.0            0             0

 ---------------------------------------------------- PATH UTILIZATION SECTION -----------------------------------------------------

                                                                      -- RETRY --   - PERFORMANCE PERCENTAGES -
  FROM     TO       TRANSPORT  FROM-TO DEVICE/  REQUESTS    AVERAGE
 SYSTEM   SYSTEM      CLASS    STRUCTURE        SATISFIED  QUEUE LEN  LIMIT COUNT   REFUSED  APPENDED IMMEDIATE   MEMBER STATUS

 SJSE     SJSA      DEFAULT    IXCPATH1            14,441       0.00     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSA      DEFAULT    IXCPATH2            14,798       0.00     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSA      DEFAULT    IXCPATH3            12,026       0.00     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSA      DEFAULT    IXCPATH4            14,474       0.00     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSB      DEFAULT    IXCPATH1            11,826       0.01     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSB      DEFAULT    IXCPATH2            14,428       0.01     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSB      DEFAULT    IXCPATH3            13,156       0.01     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSB      DEFAULT    IXCPATH4            16,173       0.01     10     0       0.0       0.1      99.9   ACTIVE
 SJSE     SJSC      DEFAULT    IXCPATH1            49,274       0.00     10     0       0.0       0.0     100.0   ACTIVE
 SJSE     SJSC      DEFAULT    IXCPATH2            41,618       0.00     10     0       0.0       0.0     100.0   ACTIVE

(continued on next page)
```

Figure 53 Cross-System Coupling Facility Report (part 2 of 2)

SJSE	SJSC	DEFAULT	IXCPATH3	29,388	0.00	10	0	0.0	0.1	99.9	ACTIVE
SJSE	SJSC	DEFAULT	IXCPATH4	62,073	0.00	10	0	0.0	0.1	99.9	ACTIVE
SJSE	SJSD	DEFAULT	IXCPATH1	15,368	0.00	10	0	0.0	0.0	100.0	ACTIVE
SJSE	SJSD	DEFAULT	IXCPATH2	10,030	0.00	10	0	0.0	0.0	100.0	ACTIVE
SJSE	SJSD	DEFAULT	IXCPATH3	18,364	0.00	10	0	0.0	0.0	100.0	ACTIVE
SJSE	SJSD	DEFAULT	IXCPATH4	12,882	0.00	10	0	0.0	0.1	99.9	ACTIVE
SJSE	SJSG	DEFAULT	IXCPATH1	12,702	0.01	10	0	0.0	0.1	99.9	ACTIVE
SJSE	SJSG	DEFAULT	IXCPATH2	19,809	0.01	10	0	0.0	0.0	100.0	ACTIVE
SJSE	SJSG	DEFAULT	IXCPATH3	19,089	0.01	10	0	0.0	0.1	99.9	ACTIVE
SJSE	SJSG	DEFAULT	IXCPATH4	17,295	0.01	10	0	0.0	0.2	99.8	ACTIVE
SJSE	SJSH	DEFAULT	IXCPATH1	10,613	0.02	10	0	0.0	0.0	100.0	ACTIVE
SJSE	SJSH	DEFAULT	IXCPATH2	11,668	0.02	10	0	0.0	0.1	99.9	ACTIVE
SJSE	SJSH	DEFAULT	IXCPATH3	4,002	0.02	10	0	0.0	0.0	100.0	ACTIVE
SJSE	SJSH	DEFAULT	IXCPATH4	7,624	0.03	10	0	0.0	0.0	100.0	ACTIVE
SJSG	SJSE	<INBOUND>	IXCPATH1	41,745	0.00	10	0	0.0	0.0	0.0	ACTIVE
SJSG	SJSE	<INBOUND>	IXCPATH2	49,324	0.00	10	0	0.0	0.0	0.0	ACTIVE

-- DETAIL REPORT --

			MESSAGES RECEIVED		MESSAGES SENT		
SYSTEM	GROUP	MEMBER	COUNT	PERCENT OF TOTAL	COUNT	PERCENT OF TOTAL	MEMBER STATUS
SJSB	COFVLFNO	SJSB	2,740	7.9	6,919	21.7	ACTIVE
SJSB	DC$CAS	SJSB	19,584	56.3	19,738	61.9	ACTIVE
SJSB	EDML	CCD00001	0	0.0	0	0.0	ACTIVE
SJSB	EDML	EDML	0	0.0	0	0.0	ACTIVE
SJSB	EDML	EDMS	0	0.0	0	0.0	ACTIVE
SJSB	EZBTCPCS	SJSBDC$TCPIP	0	0.0	0	0.0	ACTIVE
SJSB	IGWXSGIS	N200121522435437	0	0.0	0	0.0	ACTIVE
SJSB	IMFOTMA	IMS51T	0	0.0	9	0.0	ACTIVE
SJSB	IXCLO03B	M5153	25	0.1	12	0.0	ACTIVE
SJSB	IXCLO038	M508	0	0.0	0	0.0	ACTIVE
SJSB	IXCLO039	M5114	0	0.0	0	0.0	ACTIVE
SJSB	SYSATB02	M416	0	0.0	0	0.0	ACTIVE
SJSB	SYSENF	SJSB	10	0.0	0	0.0	ACTIVE
SJSB	SYSGRS	SJSB	663	1.9	663	2.1	ACTIVE
SJSB	SYSIGW00	IGWCLM01SJSB	10	0.0	10	0.0	ACTIVE
SJSB	SYSIGW01	IGWCLM01SJSB	0	0.0	0	0.0	ACTIVE
SJSB	SYSIGW02	SMLS1SJSB	238	0.7	238	0.7	ACTIVE
SJSB	SYSIGW03	IGWSHC00SJSB	0	0.0	0	0.0	ACTIVE
SJSB	SYSMCS	SJSB	7,558	21.7	1,764	5.5	ACTIVE
SJSB	SYSMCS2	SJSB	0	0.0	0	0.0	ACTIVE
SJSB	SYSWLM	SJSB	3,831	11.0	2,411	7.6	ACTIVE
SYSTEM TOTAL			34,788	6.8	31,881	6.2	

System Summary section field descriptions

Table 60 describes each field in the System Summary section of the Cross-System Coupling Facility Report.

Table 60 Field descriptions for the System Summary section

Field	Description
FROM SYSTEM	system from which the message originated (for internal messages this is the same as the TO SYSTEM)
TO SYSTEM	system to which the message is sent (for internal messages this is the same as the FROM SYSTEM)
TRANSPORT CLASS	transport class name as defined to XCF **Note**: Information not available for an inbound path is marked with <INBOUND>.
BUFFER LENGTH	size of the message buffer
REQUESTS SATISFIED	number of messages successfully sent where a TRANSPORT CLASS received a TO SYSTEM message that originated from its own FROM SYSTEM
BUFFER EFFICIENCY PERCENTAGES	information under this section head is related to buffer efficiency: **DEGRADED** — percentage of MSG GT BUFFER that suffered degradation due to overhead incurred by XCF in finding a suitably large buffer **MSG GT BUFFER** — percentage of the messages satisfied that required a larger buffer length than defined for the Transport Class **MSG LT BUFFER** — percentage of the messages satisfied that required less buffer length than defined for the Transport Class **MSG = BUFFER** — percentage of the messages satisfied that required the same buffer length as defined for the Transport Class
MESSAGES MIGRATED TO ALT PATH	number of messages migrated to an alternate path because the designated primary path was not available
MESSAGES REJECTED	number of messages rejected due to an XCF failure to obtain sufficient buffer space

Path Utilization Section field descriptions

Table 61 describes each field in the Path Utilization Section of the Cross-System Coupling Facility Report.

Table 61 Field descriptions for the Path Utilization Section

Field	Description
FROM SYSTEM	system from which the message originated (for internal messages this is the same as the TO SYSTEM)
TO SYSTEM	system to which the message is sent (for internal messages this is the same as the FROM SYSTEM)
TRANSPORT CLASS	transport class name as defined to XCF Information not available for an inbound path is marked with <INBOUND> in that column for all inbound paths.
FROM-TO DEVICE/ STRUCTURE	address of the CTC devices from which and to which messages are sent along this path, or the name of the coupling facility structure used to send messages
REQUESTS SATISFIED	number of messages successfully sent on this path
AVERAGE QUEUE LEN	average number of messages that had to wait to be sent along this path
RETRY	information under this section is related to the number of times that XCF had to retry to send messages: **LIMIT** maximum number of times XCF can retry to send a message **COUNT** number of times XCF had to retry to send messages
PERFORMANCE PERCENTAGES	information under this section is related to XCF message sending performance: **REFUSED** percentage of all message requests that were refused because the maximum message limit had been reached (inbound only) **APPENDED** percentage of all message requests that were appended to a message or messages already in the process of being sent over the path (outbound only) **IMMEDIATE** percentage of all message requests that were sent over a path that was suspended (not busy) at the inception of the path request (outbound only)
MEMBER STATUS	status of the member during the recording interval Members can be Active, Partially Activated, Re-activated, or De-activated.

Detail Report section field descriptions

Table 62 describes each field in the Detail Report section of the Cross-System Coupling Facility Report.

Table 62 Field descriptions for the Detail Report section

Field	Description
SYSTEM	system on which this path, group, and member is defined
GROUP	group to which the member is assigned
MEMBER	member name
MESSAGES RECEIVED	information under this section header is related to the number of XCF messages that were received by members: **COUNT** number of messages received for this member **PERCENT OF TOTAL** percentage of all messages received that were for this member
MESSAGES SENT	information under this section header is related to the number of XCF messages that were sent for members: **COUNT** number of messages sent for this member **PERCENT OF TOTAL** percentage of all messages sent that were for this member
MEMBER STATUS	status of the member during the recording interval Members can be Active, Partially Activated, Re-activated, or De-activated.

Cryptographic Hardware Activity Report

The Cryptographic Hardware Activity Report shows the activity of cryptographic hardware features, including the standard Cryptographic Coprocessor Facility (CCF), and the optional adjunct features: PCI Cryptographic Coprocessor (PCICC), PCI-X Cryptographic Coprocessor (PCIXCC), and PCI Cryptographic Accelerator (PCICA).

The Cryptographic Hardware Activity Report is produced by using the CRYPTO Analyzer control statement (see “CRYPTO” on page 233). The data for this report is obtained by using the CRYPTO Extractor control statement (see “CRYPTO” on page 140).

NOTE

The data in the ICSF SERVICES section of this report reflects the activity on the system where the SMF records were produced. The data in the rest of the sections reflects the activity from all systems in the CPC.

Figure 54 shows an example of the Cryptographic Hardware Activity Report.

Figure 54 Cryptographic Hardware Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)           CRYPTOGRAPHIC HARDWARE ACTIVITY REPORT                       RPTSEQ     3 PAGE      5
BMC SOFTWARE, INC.                                            XYZ COMPANY                                REPORT DATE: DD MMM YY 14.11
                                                        WORLDWIDE HEADQUARTERS                           SYSTEM ID: SYSE      Z v.rr.n
 ACTL 22 JAN YY 00.15.00  22 JAN YY 00.30.00

 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  70-2/4/0/0.25

-------------------------------------------------  PCI CRYPTOGRAPHIC ACCELERATOR ---------------------------------------------------
 TYPE = PCICA

          ------- TOTAL ---------  ----- 1024-BIT ME -----  ----- 2048-BIT ME -----  ----- 1024-BIT CRT ----  ----- 2048-BIT CRT ----
      ID   RATE  EXEC TIME  UTIL%   RATE  EXEC TIME  UTIL%   RATE  EXEC TIME  UTIL%   RATE  EXEC TIME  UTIL%   RATE  EXEC TIME  UTIL%
      --  -----  ---------  -----  -----  ---------  -----  -----  ---------  -----  -----  ---------  -----  -----  ---------  -----
       2  165.2        1.3   21.5  107.1        1.1   11.8    0.0        0.0    0.0   58.1        1.7    9.7    0.0        0.0    0.0
       3  269.9        1.8   48.6    0.0        0.0    0.0    0.0        0.0    0.0    0.0        0.0    0.0  269.9        1.8   48.6
          -----  ---------  -----  -----  ---------  -----  -----  ---------  -----  -----  ---------  -----  -----  ---------  -----
 SUMMARY  435.1        1.6   35.0  107.1        1.1    5.9    0.0        0.0    0.0   58.1        1.7    4.8  269.9        1.8   24.3

-------------------------------------------------  PCI CRYPTOGRAPHIC COPROCESSOR ---------------------------------------------------
 TYPE = PCIXCC

          -------- TOTAL --------  RSA KEY
      ID   RATE  EXEC TIME  UTIL%  GEN RATE
      --  -----  ---------  -----  --------

       0    0.1     3205.0   32.1       0.1
       1   83.0        1.1    8.8       0.0
          -----  ---------  -----      ----
 SUMMARY   83.1        1.5   20.4       0.1

  ------------------------------------------------------------ ICSF SERVICES -----------------------------------------------------------
 TYPE = PCIXCC

          TOTAL   DES ENCRYPTION   DES DECRYPTION   ------ MAC -----   ---- HASH ----   ------ PIN ------
                  SINGLE  TRIPLE   SINGLE  TRIPLE   GENERATE  VERIFY   SHA-1  SHA-256   TRANSLATE  VERIFY
          -----   ------  ------   ------  ------   --------  ------   -----  -------   ---------  ------
    RATE    2.3      0.0     1.0      0.3     1.0        0.0     0.0     0.0      0.0         0.0     0.0
    SIZE    ---      0.0   209.3    176.0   297.3        0.0     0.0     0.0      0.0         ---     ---
```

Cryptographic Hardware Activity Report field descriptions

Table 63 describes each field in the Cryptographic Hardware Activity Report.

Table 63 Field descriptions for the Cryptographic Hardware Activity Report (part 1 of 3)

Field	Description	
	PCI CRYPTOGRAPHIC ACCELERATOR **Note**: This section will appear only when PCI Cryptographic Accelerator hardware is installed.	
TYPE	type of cryptographic accelerator: ■ PCICA—PCI Cryptographic Accelerator ■ CEX2A—Cryptographic Express2 Accelerator	
ID	PCICA index	
TOTAL	**RATE**	rate of all RSA operations per second
	EXEC TIME	execution time, in milliseconds, of all RSA operations
	UTIL%	ratio of execution time of all RSA operations and elapsed time
1024-BIT ME	rate, execution time, and utilization for all RSA operations in 1024-bit ME format as defined	
2048-BIT ME	rate, execution time, and utilization for all RSA operations in 2048-bit ME format as defined	
1024-BIT CRT	rate, execution time, and utilization for all RSA operations in 1024-bit CRT format as defined	
2048-BIT CRT	rate, execution time, and utilization for all RSA operations in 2048-bit CRT format as defined	
SUMMARY	**RATE**	rate per second of all RSA operations for all processors
	EXEC TIME	execution time, in milliseconds, of all RSA operations for all processors
	UTIL%	ratio of execution time to elapsed time of all RSA operations for all processors
	PCI CRYPTOGRAPHIC COPROCESSOR **Note**: This section will appear only when PCI Cryptographic Coprocessor hardware is installed.	
TYPE	type of cryptographic coprocessor: ■ PCICC—PCI Cryptographic Coprocessor on pre-z990 processor ■ PCIXCC—PCIX Cryptographic Coprocessor introduced on the IBM z990 processor ■ CEX2C—Crypto Express2 Coprocessor, combining PCICA and PCIXCC functions into a single feature on z890 and z990 processors	
ID	PCICC index	

Table 63 Field descriptions for the Cryptographic Hardware Activity Report (part 2 of 3)

Field	Description	
TOTAL	**RATE**	rate of all operations per second
	EXEC TIME	average execution time, in milliseconds, of all operations
	UTIL%	ratio of execution time for all operations to elapsed time
RSA KEY GEN RATE	rate for RSA-key-generation operations per second	
SUMMARY	**RATE**	rate per second of all operations for all processors
	EXEC TIME	execution time, in milliseconds, of all operations for all processors
	UTIL%	ratio of execution time to elapsed time of all operations for all processors
	RSA KEY GEN RATE	rate per second of all RSA-key-generation operations for all processors
ICSF SERVICES		
TYPE	type of cryptographic hardware that is using ICSF services: ■ PCIXCC for z990 processors ■ PCICC for non-z990 processors	
TOTAL	**RATE**	rate of all the calls per second
DES ENCRYPTION	**SINGLE RATE**	single DES: number of calls to encipher per second
	SINGLE SIZE	single DES: average number of bytes of data enciphered
	TRIPLE RATE	triple DES: number of calls to encipher per second
	TRIPLE SIZE	triple DES: average number of bytes of data enciphered
DES DECRYPTION	**SINGLE RATE**	single DES: number of calls to decipher per second
	SINGLE SIZE	single DES: average number of bytes of data deciphered
	TRIPLE RATE	triple DES: number of calls to decipher per second
	TRIPLE SIZE	triple DES: average number of bytes of data deciphered
MAC	**GENERATE RATE**	rate of calls to MAC generate per second
	GENERATE SIZE	average number of bytes of data MAC generated
	VERIFY RATE	rate of calls to MAC verify per second
	VERIFY SIZE	average number of bytes of data MAC verified

Table 63 Field descriptions for the Cryptographic Hardware Activity Report (part 3 of 3)

Field	Description	
HASH SHA-1 **HASH SHA-256**	**RATE**	rate of calls to hash per second
	SIZE	average number of bytes of data hashed The measures represent two hashing algorithms: Secure Hash Algorithm-1 (SHA-1) and -256 (SHA-256). SHA-256 was introduced with the z9-109 processor and is supported on z/OS 1.6 or later.
PIN	**TRANSLATE RATE**	rate of calls to PIN translate per second
	VERIFY RATE	rate of calls to PIN verify per second

Device Activity Report

The Device Activity Report displays I/O activity by device number (see Figure 55). The report is formatted by device class. Device classes can be specifically requested or can default to whatever classes are found in the data. Within a device class, devices are sorted as specified by the ORDER parameter on the DEVACT control statement. When devices are sorted by LCU or storage group, there is a summary line for each LCU or storage group.

The Device Activity Report is produced by using the DEVACT Analyzer control statement (see “DEVACT” on page 243). The data for this report is obtained by using the DEVICE (see “DEVICE” on page 145) Extractor control statement.

NOTE

Data is displayed by device class, regardless of how the device data was requested on the Extractor DEVICE control statement.

A sample of this report is shown in Figure 55 on page 422.

Figure 55 Device Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        DEVICE ACTIVITY REPORT                         RPTSEQ     3 PAGE      6
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                           REPORT DATE: DD MMM YY 11.45
ACTL 30 JUN YY 00.33.22  30 JUN YY 11.45.00                   HOUSTON, TX.                              SYSTEM ID: SJSE     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-1/315/19.9K/11.19
                                              --------- AVERAGE TIME IN MSEC ----------  --------- % OF TIME ---------
                                                         INIT   DEV                                DEV                       AVG
DEV    DEVICE  VOLUME STORAGE  PAV    SSCH    TOTAL      COMND  BUSY                        DEV    IN   DEV  DEV  MOUNT   DSETS
NUM     TYPE   SERIAL  GROUP   MX    PER SEC   RESP  IOSQ  RESP DELAY  PEND  DISC  CONN    CONN   USE  RESV ALLOC  PEND   ALLOC
---- -------- ------ -------- ---  --------- ----- ----- ----- ----- ----- ----- -----   ----- ----- ----- ----- -----   ------
4100  33903    APD28G            3      0.007   0.6   0.0   0.1   0.0   0.2   0.0   0.4     0.0   0.0   0.0 100.0   0.0     0.0
4101  33903    DMGN20            3      0.007   0.5   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4102  33903    DMGS10 SGSVCDMP   3      0.015   4.8   0.0   0.1   0.0   0.2   0.0   4.6     0.0   0.0   0.0 100.0   0.0     0.0
4103  33903    PAGC32            3      0.007   0.5   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4104  33903    SAFPG1            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4105  33903    SADPG2            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4106  33903    SAZO4U            5      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4107  33903    SAZO4V            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4108  33903    SAZO4W            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4109  33903    --4109            2      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
410A  33903    --410A            2      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
410B  33903    --410B            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
410C  33903    PSTR01            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
410D  33903    PSTR02            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
410E  33903    PSTR03            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
410F  33903    PSTC01            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4110  33903    PSTS01            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4111  33903    PSTW01            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4112  33903    PSTW02            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4113  33903    SAEPG4            3      4.300   1.9   0.0   0.1   0.0   0.3   0.1   1.5     0.7   0.7   0.0 100.0   0.0     3.0
4114  33903    PSTU01            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4115  33903    SA39AU            9      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4116  33903    SA39AV            2      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4117  33903    SA39AW            2      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4118  33903    --4118            2      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
4119  33903    --4119            2      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
411A  33903    --411A            3      0.007   0.4   0.0   0.1   0.0   0.2   0.0   0.2     0.0   0.0   0.0 100.0   0.0     0.0
...
...
LCU 0045 ( 56 ACTIVE)                   4.705   1.8   0.0   0.1   0.0   0.3   0.1   1.4     0.0   0.0   0.0 100.0   0.0     3.0
```

Device Activity Report field descriptions

Table 64 describes each field in the Device Activity Report.

Table 64 Field descriptions for the Device Activity Report (part 1 of 3)

Field	Description
DEV NUM	four-digit hexadecimal number that identifies this device An asterisk (*) following this number indicates that the online/offline status of the device changed during the report interval.
DEVICE TYPE	type of the physical I/O device on which the volume is mounted
VOLUME SERIAL	volume serial number of the last volume that was mounted on this device; applicable only to DASDs and tape drives
STORAGE GROUP	name of the storage group that the DASD is assigned to if it is managed by SMS
PAV MX	number of exposures (base and aliases) of a Parallel Access Volume (PAV) at the end of the report duration; applicable only to DASDs An asterisk (*) following this number indicates that the number of exposures changed during the report duration.

Table 64 Field descriptions for the Device Activity Report (part 2 of 3)

Field	Description
SSCH PER SEC	rate per second at which SSCH instructions issued for this device were completed successfully OFFLINE is printed under this column if the device was offline during the entire report interval. If the online/offline status of a device changed during a recording interval, the device is considered offline and its data is skipped for that particular interval. **Note**: In the LCU Summary Line, this field contains the average SSCH per second for all devices that are online to the LCU.
TOTAL RESP	average number of milliseconds that are required by the device to complete an I/O request (including IOSQ, pending, connect, and disconnect time); commonly known as response time
IOSQ	average number of milliseconds of delay that an I/O request encountered because the device was busy performing an I/O from the local system
INIT COMND RESP	average number of milliseconds of delay that an I/O request encountered, beginning from when the first command of the channel program is sent to the device until the device indicates that it has accepted the command; this delay is part of pending time Data is available only on a z990 or later processor.
DEV BUSY DELAY	average number of milliseconds of delay that an I/O request encountered because the device was busy due to I/O from another system; this delay is part of pending time
PEND	average number of milliseconds that an I/O request must wait for hardware, such as an available channel path or control unit, as well as time between the SSCH pending at the channel and the device active on the subchannel This time also includes delays caused by another processor that is reserving this device in a shared DASD environment.
DISC	average number of milliseconds during which the device was processing an SSCH instruction but not transferring data
CONN	average number of milliseconds during which the device was processing an SSCH instruction and transferring data
DEV CONN	percent of time during the measurement interval that the device was connected to a channel path
DEV IN USE	percent of time during the measurement interval that the device was in use; this number includes device connect and disconnect time
DEV RESV	percent of time during the measurement interval that this device was reserved by the processor on which the CMF MONITOR Extractor was executing
DEV ALLOC	percent of time that the device was allocated to one or more data sets DASD devices that are marked as permanently resident in the UCB always show 100% allocation. Not applicable to tape devices.
MOUNT PEND	percent of time that the device had a mount pending request outstanding; this number is reported only for DASD devices

Table 64 Field descriptions for the Device Activity Report (part 3 of 3)

Field	Description
AVG DSETS ALLOC	average number of data sets allocated on this device; this field reports only on DASDs
NUM OF MOUNTS	number of mounts for the tape device for the reporting interval An asterisk (*) to the left of this value indicates that a mount pending condition existed at the start of the recording interval. An asterisk (*) to the right of the value indicates that a mount pending condition existed at the end of the recording interval. For the LCU, the field contains the total number of mounts for all devices on the LCU for the reporting interval.
AVG MOUNT TIME	average mount pending time for the tape device For the LCU, the field contains the average mount pending time for all devices on the LCU. The field is reported as *hh:mm:ss*; the maximum time reported in this field is 99:59:59.
TIME DEVICE ALLOC	total time that the tape device was allocated during the reporting interval For the LCU, the field contains the average allocation time for all active devices on the LCU during the reporting interval. The field is reported as *hh:mm:ss*; the maximum time reported in this field is 99:59:59.
LCU	total/summary line for all devices on this LCU
***n* ACTIVE**	number of devices on this LCU that had some I/O activity
NOT READY	percent of time that the device was not ready for use; not applicable to DASD or tape devices

Direct Access Report

The Direct Access Report summarizes activity on each volume. The amount of activity of each volume directly indicates the effect of each volume on system performance.

The Direct Access Report is produced by using the DASD Analyzer control statement (see “DASD” on page 237). The data for this report is obtained by using the DEVICE (see “DEVICE” on page 145) and HEADMOVE (see “HEADMOVE” on page 159) Extractor control statements.

Extractor overhead for this report is high because a high resolution sampling interval is required. However, information in the report can be used to reduce contention and improve the use of DASD resources.

A sample report is shown in Figure 56.

Figure 56 Direct Access Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        DIRECT ACCESS REPORT                              RPTSEQ    17 PAGE    113
BMC SOFTWARE, INC.                                            XYZ COMPANY                                  REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00             WORLDWIDE HEADQUARTERS                             SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-12/1/1/0  240-13/393/25.1K/7  74-1/112/25.1K/7
-------------------------------------------------------  ACTIVITY BY VOLUME  ------------------------------------------------------
 VOLUME    DEVICE    TOTAL HEAD    MOVEMENTS TO      HEAD MOVEMENT TIME         ERROR       MOUNTED      ALLOC      DEVICE
  NAME      TYPE     MOVEMENTS    ALTERNATE CYLS     TOTAL (SEC)     AVG (MSEC)  RECOVERY     PERCENT      PERCENT    NUMBERS
 BAB310    33903             4         0              0.112          28.00        0.00       100.00       100.00     835F
 BAB311    33903             2         0              0.020          10.00        0.00       100.00       100.00     8360
 BAB312    33903             0         0              0.000           0.00        0.00       100.00       100.00     8361
 BAB313    33903             4         0              0.072          18.00        0.00       100.00       100.00     8362
 BAB314    33903            23         0              0.285          12.39        0.00       100.00       100.00     8363
 BAB315    33903             4         0              0.032           8.00        0.00       100.00       100.00     8364
 BAB316    33903             7         0              0.057           8.14        0.00       100.00       100.00     8365
 BAB317    33903             9         0              0.166          18.44        0.00       100.00       100.00     8366
 BAB318    33903             0         0              0.000           0.00        0.00       100.00       100.00     8367
 BAB319    33903             0         0              0.000           0.00        0.00       100.00       100.00     8368
 BAB320    33903             2         0              0.056          28.00        0.00       100.00       100.00     8369
 BAB321    33903             9         0              0.078           8.67        0.00       100.00       100.00     836A
 BAB322    33903            19         0              0.286          15.05        0.00       100.00       100.00     836B
 BAB323    33903            71         0              0.518           7.30        0.02       100.00       100.00     836C
 BAB324    33903            49         0              0.697          14.22        0.00       100.00       100.00     836D
 BAB325    33903            29         0              0.510          17.59        0.00       100.00       100.00     836E
 BAB326    33903            66         0              0.766          11.61        0.00       100.00       100.00     836F
 BAB327    33903             1         0              0.024          24.00        0.00       100.00       100.00     8370
 BAB328    33903            33         0              0.689          20.88        0.01       100.00       100.00     8371
 BAB329    33903            10         0              0.098           9.80        0.00       100.00       100.00     8372
 BAB330    33903             0         0              0.000           0.00        0.00       100.00       100.00     8373
 BAB331    33903            59         0              0.313           5.31        0.24       100.00       100.00     8374
 BAB332    33903           249         0              2.694          10.82        0.00       100.00       100.00     8375
```

In this report, the time spent in error recovery includes both retries and permanent I/O errors. A bit in the IOSB is turned on whenever an error recovery procedure (ERP) gains control from IOS. This bit is sampled, and a 1 is added to a counter each time it appears.

ERPs issue messages IEA00A through IEA005I: `INTERVENTION REQUIRED, PATH INOPERATIVE,` and `PERMANENT I/O ERROR`. All ERPs write records to the SYS1.LOGREC data set.

An ERP is entered from a channel end appendage (CHE) or abnormal end appendage (ABE) if an exception condition is indicated in the IOSB and the device is accessible.

A CHE or ABE appendage is entered if the following conditions exist:

- The CSW status bits in the IOSB (IOSCSW) contain no status information other than a PCI, channel end, device end, attention, unit exception, or wrong length record indication.
- The CSW status bits contain a unit check, channel data check, channel control check, or interface control check.

NOTE

If error counts exist but no ERP messages are issued, one of three causes can be responsible: logical error recovery, ISAM overflow, or successful retries.

Direct Access Report field descriptions

Table 65 describes each field in the Direct Access Report.

Table 65 Field descriptions for the Direct Access Report

Field	Description
VOLUME NAME	name of the volume that was measured
DEVICE TYPE	device type on which the volume was mounted
TOTAL HEAD MOVEMENTS	total number of head movements observed
MOVEMENTS TO ALTERNATE CYLS	total number of head movements observed to alternate cylinders
HEAD MOVEMENT TIME	total number of seconds that the device spent in head movement time, and average number of milliseconds devoted to each head movement
ERROR RECOVERY	percentage of time that the device on which the volume was mounted spent in error recovery
MOUNTED PERCENT	percentage of time that the device was mounted
ALLOC PERCENT	percentage of time that at least one data set on the volume was allocated, or percentage of time that the volume was marked permanently resident
DEVICE NUMBERS	four-digit addresses of the devices on which the volume was mounted

Direct Access Report Plot of Volume

The Direct Access Report Plot of Volume graphically summarizes head movement activity for DASD volumes. A histogram is given for each volume. Each histogram shows where movement on a direct access volume is occurring. By knowing the region of the volume that is associated with a large amount of head movement, you can reorganize the data sets on the volume and decrease head movement time.

The Direct Access Report of Volume is produced by using the DASD (see “Direct Access Report” on page 425) and VOLSER (see “VOLSER” on page 320) Analyzer control statements. The data for this report is obtained by using the DEVICE (see “DEVICE” on page 145) and HEADMOVE (see “HEADMOVE” on page 159) Extractor control statements. A sample of this report is shown in Figure 57. Column headings are explained in the legend.

Figure 57 Direct Access Report Plot of Volume

```
PRODUCED BY CMF MONITOR (5.5.00)                   DIRECT ACCESS REPORT PLOT OF VOLUME                 RPTSEQ     3 PAGE     46
BMC SOFTWARE, INC.                                            XYZ COMPANY                              REPORT DATE: DD MMM YY 10.08
ACTL 10 JUN YY 13.49.13  10 JUN YY 10.21.33            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SYSB      Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74/57/289K/21.44   240-12/10/10/.00   240-13/8021/288K/21.70

--------------------------------------------------  VOLUME BAB003 ACTIVITY PATTERN  ------------------------------------------------

    TOTAL HEAD MOVEMENT TIME    (56.688 SEC)    AVERAGE HEAD MOVEMENT TIME  (12.860 MSEC)     PHYSICAL HEAD MOVEMENTS     (4408)
       39.6 +                                            *
            I                                            *
            I                                            *
            I                                            *
            I                                            *
            I                                            *
            I                                            *
            I                                            *
       29.7 +                                            *
            I                                            *
    S       I                                            *
    A       I                                            *
    M       I                                            *
    P       I                                            *
    L       I                                            *
    E       I                                            *
       19.8 +                                            *
    P       I                                            *
    E       I                                            *
    R       I                                            *
    C       I                                            *
    E       I                                            *
    N       I                                            *                           *
    T       I                                            *                           *
        9.9 +                                            *                           *
            I                                            *                          **
            I                                            *                          **
            I                                            *                          **    *
            I                                     * *    *                          **    *
            I           *                         * *    *                          **    *
            I           *                         * *    *   *                      **    *
            I       *  **        *            *   * *    *   *      *       *   *   **    *
             ---------+---------+---------+---------+---------+---------+---------+---------+--------- (MAX =   886)

                      1         2         3         4         5         6         7         8
                      0         0         0         0         0         0         0         0
                      0         0         0         0         0         0         0         0
                                                   ACTIVITY DISTRIBUTION BY CYLINDER CLASS
                                                         (CLASS WIDTH = 10 CYLINDERS)
```

The horizontal axis shows all possible cylinders on which head movement can occur. The vertical axis shows the percentage of time that head movement activity occurred. The upper limit of this axis is self-scaling; it represents the highest percentage of time that is observed for head movement activity on this particular volume.

Head movement activity is represented by columns of asterisks. The asterisks appear above cylinder classes that had head movement activity; one cylinder class is equal to 10 cylinders. Each vertical asterisk represents one percentage unit of head movement activity for that cylinder class.

Direct Access Report Plot of Volume field descriptions

Table 66 describes each field in the Direct Access Report Plot of Volume.

Table 66 Field descriptions for the Direct Access Report Plot of Volume

Field	Description
TOTAL HEAD MOVEMENT TIME	total number of seconds spent on head movements for this volume
AVERAGE HEAD MOVEMENT TIME	average number of milliseconds per head movement for this volume
PHYSICAL HEAD MOVEMENTS	total number head movements observed for this volume

Disabled Delay Report

The Disabled Delay Report is a two-page report that is produced by using the DISTIM Extractor control statement. No Analyzer control statement is required. This report is produced at the end of every interval, as specified in the REPORT Extractor control statement (see "REPORT" on page 176).

The first page of the report shows the values of the LODELAY and HIDELAY values used by the DISTIM control statement as well as a total count for the values below, within, and above these parameters. It also displays a graph of the frequency of occurrence for each value of disabled delay.

The horizontal axis of the graph shows the possible values (in microseconds) of the disabled delay. The vertical axis shows the number of times each particular value occurred. Each asterisk (*) represents one occurrence of a disabled delay of a particular value.

The shift parameter eliminates all values that fall below the value that you specify. If you specify SHIFT=20, for example, the values from 1 to 19 do not appear on the graph.

The second page of the report shows the number of times that each address space was the last one dispatched before a sample is taken, where the delay time falls in the range of LODELAY to HIDELAY.

A sample of the Disabled Delay Report is shown in Figure 58.

Figure 58 Disabled Delay Report

```
PRODUCED BY CMF MONITOR (5.5.00)                     DISABLED DELAY REPORT                         SYSTEM ID: SYSB      PAGE     1
BMC SOFTWARE, INC.                                        FOR  5 MINUTES                           INTERVAL: 10 JUN 03 17:50/17:55
---------------------------------------------------< CONSTANTS  AND COUNTERS >----------------------------------------------------
  LOW DELAY  HI DELAY   0 TO LO  LO TO HI   OVER HI          CPU BUSY %                                CT OVFL   DL OVFL
         70       140       471      1032       367                  34                                      0       131
|
|
|
+
|
|
|                                                                         *
|                                                                         *
|                                                                         *
|                                                                         *
|                                                                       * *
|                                                                       * *
|                                                                       * *
+                                                                       * *
|                                                                       * *
|                                                                       * *
|                                                                       * *
|                                                                       * *
|                                                                       * *
|                                                                       * *
|                                                                    ** *
|                                                               *    ** *   *
|                                                               **   ** * *   **
+                                                               **   ** * *   **
|                                                               **  *** ***   **
|                          * *                                  **  *** ***   **
|                          * *                                  **  *** ****  **
|                      *   * * *                               **** ********* **
|                      *   * * *                               **** ********* **
|                      * * * * *                               ************** **
|                     ** * * * *                               ************** **
|                     ** * *** *                               ***************** *
|                     **** *****                               ***************** *   *
+                     **********                             * ***************** *   *
|                     **********  *                          * ***************** *  **
|                    ***********  *                          *************************
|                    ***********  *                          *************************  *
|                  * ***********  *                          *************************  *
|                  *************  *                         **************************  *
|                  ************* **                        **************************** *
|                  ****************                       *******************************  *     *     *
|                  *****************                      ******************************** * *   *     *
|                 ******************                      ******************************** * *   **   ** *
+              *  *******************                     ******************************** ***   **   ** *     *
|              *  *******************  *                  ******************************** ***   **   ** *     * *
|              ** *******************  *                  ************************************* ****  ** **    * *
|             ***********************  **                ***********************************************  **    * * *     *
|             ***********************  ** **     *     *************************************************  **    * * *   * *  *
|             *********************** ******     * *   ******************************************************  *** *  ** *  *
|             ********************************   * *   ******************************************************  ****** ** ****   *
|             ***************************************  ******************************************************* ************** * **
|        *   **************************************************************************************************************** * **
| SHIFT     20  *******************************************************************************************************************
+---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+

PRODUCED BY CMF MONITOR (5.5.00)                    DISABLED DELAY REPORT                           STEM ID: SYSB      PAGE     2
BMC SOFTWARE, INC.                                       FOR  5 MINUTES                           INTERVAL: 10 JUN 03 17:50/17:55
---------------------------------------------------< LAST ASID IN LO TO HI RANGE >------------------------------------------------
      ASID   COUNT        ASID   COUNT        ASID   COUNT        ASID   COUNT        ASID   COUNT        ASID   COUNT
         1     255          10       6          14       2          18       4          23       1          24       1
        29       7          31       1          38      12          60       2          81       1          89       1
        92       1          99       4         111       1         116       2         151       2         172       4
       177       1         179       1         202       1         213       1         245       2         259       1
       284       1         288       2         362       4         365       1         367       1         509       3
```

Disabled Delay Report field descriptions

Table 67 describes each field in the Disabled Delay Report.

Table 67 Field descriptions for the Disabled Delay Report

Field	Description
LOW DELAY	LODELAY value specified on the control card
HI DELAY	HIDELAY value specified on the control card
0 to LO	number of samples that fell below the LODELAY value specified on the control card
LO TO HI	number of samples that fell into the range of LODELAY and HIDELAY values specified on the control card
OVER HI	number of samples that occurred above the HIDELAY value specified on the control card
CPU BUSY %	percentage of samples during which the CPU was running
CT OVFL	number of samples that went over 212, which is the maximum number of samples for a particular delay value
DL OVFL	number of delay samples that were over 511 microseconds
ASID	ID number of the address space
COUNT	number of times that a particular address space was the last one dispatched before a sample is taken, where the delay time is in the range of LODELAY to HIDELAY

Distribution Graph

The Distribution Graph shows how frequently selected conditions exist during the measurement interval. The GRAPH TYPE=DISTRIB Analyzer control statement produces this graph (see “GRAPH” on page 259). Up to 16 measures can be specified in each GRAPH statement; each measure specified generates a new graph. No simultaneous graphing is permitted in distribution graphics.

The Extractor control statements required for the Distribution Graph are dependent upon the information to be graphed. (See “Numeric list of record types” on page 47 for information about the record types generated by each Extractor control statement, and a description of the specific type of data collected by each Extractor control statement.)

In data processing, frequency of occurrence of conditions generally cluster around a certain set of values. For this reason, distribution can be normalized; this is the default. The horizontal axis shows the percentage of CPU busy time. The vertical axis shows the percentage of time that the CPU had a specific percentage of activity. A range of 0 to 50%, in 5% increments, is displayed if the greatest percentage of occurrence is below 50%.

The first fill character is used for all measures, no matter how many measures are requested.

A sample Distribution Graph for CPU usage is shown in Figure 59.

Figure 59 Distribution Graph

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      DISTRIBUTION GRAPH OF CPU                        RPTSEQ    27 PAGE     62
BMC SOFTWARE, INC.                                            XYZ COMPANY                               REPORT DATE: DD MMM YY 9.33
ACTL 10 JUN YY 06.45.01  10 JUN YY 07.45.00             WORLDWIDE HEADQUARTERS                          SYSTEM ID: SYSB    Z v.rr.n


                                                              P E R C E N T S   O F   A C T I V I T Y
                               0        10        20        30        40        50        60        70        80        90       100
                               +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
                             50.         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
 CPU ENVIRONMENT          P    .         .         .         .         .         .         .         .         .         .         .
       REQUESTED:  ALL       45.         .         .         .         .         .         .         .         .         .         .
                          E    .         .         .         .         .         .         .         .         .         .         .
 STANDARD                      .         .         .         .         .         .         .         .         .         .         .
 DEVIATION =     18.8     R    .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
   AVERAGE =       37     C  40.         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
    MEDIAN =       32     E    .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
      MODE =       25     N    .         .         .         .         .         .         .         .         .         .         .
                             35.         .         .         .         .         .         .         .         .         .         .
                          T    .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                          O  30.         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                          F    .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                               .         .         .         .         .         .         .         .         .         .         .
                             25.         .         .         .         .         .         .         .         .         .         .
                          O    .         .         .    *****.         .         .         .         .         .         .         .
                               .         .         .    *****.         .         .         .         .         .         .         .
                          C    .         .         .    *****.         .         .         .         .         .         .         .
                               .         .         .    *****.         .         .         .         .         .         .         .
                          C  20.         .         .    *****.         .         .         .         .         .         .         .
                               .         .         .    *****.         .         .         .         .         .         .         .
                          U    .         .         .    **********     .         .         .         .         .         .         .
                               .         .         .    **********     .         .         .         .         .         .         .
                          R    .         .         .    **********     .         .         .         .         .         .         .
                             15.         .         .    **********     .         .         .         .         .         .         .
                          R    .         .         .    **********     .         .         .         .         .         .         .
                               .         .         .    **********     .         .         .         .         .         .         .
                          E    .         *****     .    **********     .         .         .         .         .         .         .
                               .         *****     .    **********     .         .         .         .         .         .         .
                          N  10.         *****     .    **********     .         .         .         .         .         .         .
                               .         *****     .    **********     .         .         .         .         .         .         .
                          C    .         *****     .    **********     .         .         .         .         .         .         .
                               .         *****     .    **********     .         .         .         .         .         .         .
                          E    .         *************************     ********************.         **********.         .         .
                              5.         *************************     ********************.         **********.         .         .
                               .         *************************     ********************.         **********.         .         .
                               .         *************************     ********************.         **********.         .         .
                               .         *************************     ********************.         **********.         .         .
                               .         *************************     ********************.         **********.         .         .
                              0+---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
```

As shown, the CPU was never below 10% busy during the reporting period nor greater than 80% busy. A quarter of the time, the CPU was 25% to 30% busy.

A plateau in the graph represents an accumulated percentage of occurrence. The CPU ran at an accumulated busy rate of 40% to 60%, as shown in the example. The CPU ran at this rate 24% of the time, not 6% of the time.

For an explanation of mean, mode, and standard deviation, see Appendix A, "Statistical considerations."

Distribution Graph field descriptions

Table 68 describes each field in the Distribution Graph.

Table 68 Field descriptions for the Distribution Graph

Field	Description
CPU ENVIRONMENT REQUESTED:	CPU number or numbers selected in the GRAPH control statement
STANDARD DEVIATION =	standard deviation arrived at in calculating percentage of CPU busy time
AVERAGE	average value arrived at in calculating percentage of CPU busy time
MEDIAN =	median value arrived at in calculating percentage of CPU busy time
MODE =	indicates the recording mode under which the records were extracted **Note**: An asterisk (*) printed before MODE= indicates that more modes were identified than could be reported. Up to the first three modes are reported.
PERCENT OF OCCURRENCE	axis showing percentage of time that the measurement was at the indicated value
PERCENTS OF ACTIVITY	axis representing percentage of measured activity during the interval Each column represents a normalized value, showing a grouping of measurement values (see breakdown in the description of the report).

Enqueue Conflict Report

The Enqueue Conflict Report provides information about resources that build up queues greater than the threshold specified on the ENQUEUE Analyzer control statement (see “ENQUEUE” on page 250). Data is collected for this report by using the ENQUEUE Extractor control statement (see “ENQUEUE” on page 152). The information in this report can be helpful in locating bottlenecks.

The Enqueue Conflict Report can be requested in detailed or summary formats. A sample including both formats is shown in Figure 60. The summary format does not include the **OWNER OR REQUESTOR** field.

Figure 60 Enqueue Conflict Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      ENQUEUE CONFLICT REPORT                              RPTSEQ    19 PAGE    136
BMC SOFTWARE, INC.                                           XYZ COMPANY                                    REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                               SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  77-1/28/0/7
GRS MODE: STAR ------------------------------------ ENQUEUE CONFLICT SUMMARY REPORT ---------------------------------------------
                                                     (THRESHOLD PERCENT =    .00 )
  CONFLICT   DURATION    (MAJOR)    (MINOR)                                                                    AVG        MAX
  PERCENT    HH.MM.SS     QNAME      RNAME                                                      SCOPE        DEPTH      DEPTH

     .02     00.00.04    DSPURI01   AAO.IMS710.RECON1                                          SYSTEM        1.0          1
     .14     00.00.35    ECMCONFL   A..O.                                                      SYSTEM        1.0          1
                                    C07D3
                                    12560
     .13     00.00.33    IGDCDSXS   SYSI.IBMSMS.R150.COMMDS                                    SYSTEM        1.0          1
     .00     00.00.00    SYSDSN     BMCPXK.BBIJRNLL                                            SYSTEM        1.0          1
     .06     00.00.15    SYSIEA01   DMPDSENQ                                                   SYSTEM        1.0          1
     .08     00.00.20    SYSIEA01   SDDSQ                                                      SYSTEM        1.0          1
     .43     00.01.48    SYSIEA01   SDUMPENQ                                                   SYSTEM        3.1         4+
     .00     00.00.00    SYSIGGV2   CATALOG.ICF.MSTRA39A                                       SYSTEM        1.4          3
    1.50     00.06.17    SYSIGGV2   ICFUCAT.VSYSP14                                            SYSTEM        1.2         4+
     .09     00.00.23    SYSIGGV2   ICFUCAT.VSYSP15                                            SYSTEM        1.2          2
     .00     00.00.00    SYSIGGV2   SYS1.VOLCAT.VA                                             SYSTEM        1.0          1
     .00     00.00.00    SYSVSAM    AAO.IMS710.RECON1.DATAICFUCAT.VSYSP14...I                  SYSTEM        1.0          1
                                    CCD4CDEFFF4DCCDDF4CCECCCCECCE4EEEEDFF210C
                                    116B942710B953651B41319364313B528271496F9

GRS MODE: STAR ------------------------------------ ENQUEUE CONFLICT DETAIL REPORT ----------------------------------------------
                                                     (THRESHOLD PERCENT =    .00 )
  CONFLICT   DURATION    (MAJOR)    (MINOR)                             OWNER OR      SYSTEM                AVG     MAX
  PERCENT    HH.MM.SS     QNAME      RNAME                             REQUESTER      NAME        SCOPE    DEPTH   DEPTH

     .13     00.00.33    IGDCDSXS   SYSI.IBMSMS.R150.COMMDS             MIMGR   -O/E   SJSE       SYSTEM     1.0      1
                                                                        SMS     -W/E   SJSE       SYSTEM
     .06     00.00.15    SYSIEA01   DMPDSENQ                            DUMPSRV -O/E   SJSE       SYSTEM     1.0      1
                                                                        DUMPSRV -W/E   SJSE       SYSTEM
     .08     00.00.20    SYSIEA01   SDDSQ                               DUMPSRV -O/E   SJSE       SYSTEM     1.0      1
                                                                        DUMPSRV -W/E   SJSE       SYSTEM
     .42     00.01.46    SYSIEA01   SDUMPENQ                            DUMPSRV -O/E   SJSE       SYSTEM     3.2     4+

     .01     00.00.01    SYSZMCS    SYSMCS#CL1                          CONSOLE -O/E   SJSD       SYSTEMS    1.0      1
                                                                        AAOSN61 -W/E   SJSC       SYSTEMS
     .00     00.00.00    SYSZMCS    SYSMCS#CL1                          CONSOLE -O/E   SJSC       SYSTEMS    1.0      1
                                                                        AAOSN62 -W/E   SJSC       SYSTEMS
     .00     00.00.00    SYSZMCS    SYSMCS#CL1                          CONSOLE -O/E   SJSC       SYSTEMS    1.0      1
                                                                        AAOSSEK2-W/E   SJSC       SYSTEMS
     .01     00.00.01    SYSZMCS    SYSMCS#CL1                          CONSOLE -O/E   SJSD       SYSTEMS    1.0      1
                                                                        I5APAS  -W/E   SJSD       SYSTEMS
```

Enqueue Conflict Report field descriptions

Table 69 describes each field in the Enqueue Conflict Report.

Table 69 Field descriptions for the Enqueue Conflict Report (part 1 of 2)

Field	Description
GRS MODE	displayed at the end of the most recently recorded interval; possible values are STAR RING NONE
CONFLICT PERCENT	percentage of time that a conflict was observed for this enqueue request To calculate the conflict percent, use this formula: conflict time x 100 / measurement interval
DURATION	length of time that the conflict existed—reported in hours, minutes, seconds (*hh.mm.ss*) The duration of the conflicts reflects actual time. For example, if job B waits 10 minutes for a resource held exclusively by job A, and job C starts 5 minutes after job B but waits for the same resource held by job A, the duration reported would be 15 minutes (assuming that jobs B and C could share the resource).
(MAJOR) QNAME	major name used for the enqueue request; an * means that the resource was still held when the measurement interval ended This field works in combination with (MINOR) RNAME values.
(MINOR) RNAME	minor name used for the enqueue request Three messages can be produced by this report: ■ no contention occurred ■ no contention occurred above threshold ■ warning - CMF table full, data was lost during extraction The last warning message indicates that enqueue conflict events during the interval exceeded the CMF MONITOR Extractor table capacity. The data in the report is accurate for the QNAME-RNAME combinations printed, but other QNAME-RNAME combinations occurred that were not printed.
THRESHOLD PERCENT	value supplied in the ENQUEUE control statement, used to determine when the enqueue conflict occurred
OWNER OR REQUESTOR	job name or Started Task name of the owner or requestor of the resource During the contention event, CMF MONITOR selects the job or Started Task when the wait queue is the longest. Codes appear to the right of the owner or requestor name. Valid codes are O requestor owned the resource W requestor was waiting for the resource S requestor made a shared request E requestor made an exclusive request

Table 69 Field descriptions for the Enqueue Conflict Report (part 2 of 2)

Field	Description
SYSTEM NAME	name of the system on which job is running; if a name is not available, this field contains NONAME
SCOPE	scope of the enqueue request. There are two values for this field: SYSTEM indicates that a serially reusable resource is being shared among address spaces on this system SYSTEMS indicates that the resource is being shared among systems; enqueue requests with a scope of STEP are not reported
AVG DEPTH	average of the longest wait queue from each contention event
MAX DEPTH	maximum queue depth observed for this major- or minor-name request **Note**: The upper limit for values in this field is 4, because there are only four fields.

ESS Statistics Report

The ESS Statistics Report displays the adapters and their I/O activity for each Enterprise Storage Server (ESS) that is connected to the system. The statistics are kept inside the hardware and, as a result, include I/O activity from all connected systems.

The ESS Statistics Report is produced by using the ESS Analyzer control statement (see “ESS” on page 251). The data is obtained by using the CACHE Extractor control statement, specifying RECORDS=ESS or RECORDS=ALL (see “CACHE” on page 129).

A sample ESS Statistics Report is shown in Figure 61 on page 437.

Figure 61 ESS Statistics Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        ESS STATISTICS REPORT                                RPTSEQ     3 PAGE      6
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                                 REPORT DATE: DD MMM YY 11.49
ACTL 22 SEP YY 12.22.57  22 SEP YY 18.44.26                    HOUSTON, TX.                                   SYSTEM ID: S0W1     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-8/15/5/1.05
------------------------------------ ESS SERIAL NUMBER: 00000AMAA1   TYPE-MODEL: 002107-921 ----------------------------------------

------------------------------------------------ ESS LINK STATISTICS SECTION --------------------------------------------------

----- ADAPTER -----   LINK TYPE          BYTES      BYTES PER    OPERATIONS     RESP TIME       I/O
 ID   TYPE                              PER SEC     OPERATION      PER SEC     /OPERATION     INTENSITY

0030  FIBRE  2Gb      ECKD READ         632,963          854         741.1            0          12.8
                      ECKD WRITE        111,891        2,815          39.7            0           2.0
                                        -------      -------       -------      -------       -------
                      SUMMARY           744,854          954         780.9            0          14.7

0031  FIBRE  2Gb      ECKD READ           8,964        1,981           4.5            0           0.1
                      ECKD WRITE          2,596        1,579           1.6          0.1           0.1
                                        -------      -------       -------      -------       -------
                      SUMMARY            11,559        1,874           6.2            0           0.3

0100  FIBRE  2Gb      ECKD READ         620,988          843         736.8            0          12.8
                      ECKD WRITE        106,319        2,867          37.1            0           1.8
                                        -------      -------       -------      -------       -------
                      SUMMARY           727,307          940         773.9            0          14.6

0101  FIBRE  2Gb      ECKD READ         774,583        2,582         300.0            0           8.1
                      ECKD WRITE        227,623        1,680         135.5            0           4.7
                                        -------      -------       -------      -------       -------
                      SUMMARY            1,002K        2,301         435.5            0          12.8

0230  FIBRE  2Gb      ECKD READ         626,560          854         733.5            0          12.6
                      ECKD WRITE        117,151        2,904          40.3            0           1.9
                                        -------      -------       -------      -------       -------
                      SUMMARY           743,711          961         773.8            0          14.6

0300  FIBRE  2Gb      ECKD READ         790,745        2,526         313.0            0           8.2
                      ECKD WRITE        226,689        1,748         129.7            0           4.6
                                        -------      -------       -------      -------       -------
                      SUMMARY            1,017K        2,298         442.7            0          12.8


------------------------------------------------ ESS EXTENT POOL STATISTICS SECTION ------------------------------------------------

 --- EXTENT POOL ---    ------- REAL EXTENTS -------
  ID    ---- TYPE ---    CAPACITY    EXTENT     ALLOC
                         (GBYTES)     COUNT    EXTENTS

 0000  CKD    1Gb          1,560     1,771     1,771

 0001  CKD    1Gb          1,560     1,771     1,771


---------------------------------------------------- ESS RANK STATISTICS SECTION ----------------------------------------------------

                      RAID   ------ READ OPERATIONS -------     ------ WRITE OPERATIONS ------     --------- RANK ATTRIBUTES -------
--- EXTENT POOL ---   RANK      OPS   RTIME   BYTES   BYTES        OPS   RTIME   BYTES   BYTES     --ARRAY--  MIN    CAP    - RAID -
 ID    ---- TYPE ---   ID      /SEC     /OP    /SEC     /OP       /SEC     /OP    /SEC     /OP     NUM  WDTH  RPM  GBYTES     TYPE

0000  CKD     1Gb     0000   1.6722  0.0065  68,110  40,731     2.4952  0.0165  763.6K    306K      1     6  N/A    1,800  RAID 5

0001  CKD     1Gb     0001   0.0045  0.0216   207.7  46,261     0.1909  0.1569  11,906  62,363      1     6  N/A    1,800  RAID 5

0002  CKD     1Gb     0002   0.0005  0.0080       0       0     0.0747  0.0185   519.1   6,947      1     6  N/A    1,800  RAID 5

0003  CKD     1Gb     0003   0.0404  0.0120   2,111  52,258     0.0275  0.2380   242.3   8,822      1     6  N/A    1,800  RAID 5
```

ESS Statistics Report field descriptions

Table 70 describes each field in the ESS Statistics Report.

Table 70 Field descriptions for the ESS Statistics Report (part 1 of 2)

Field	Description
ESS SERIAL NUMBER	serial number of the primary control unit
TYPE-MODEL	type and model of the primary control unit
ESS Link Statistics Section	
	Note: z/OS does not return any statistics to CMF for ESCON adapters that are connected to Enterprise Storage Servers (ESS). As a result, no ESCON adapters can appear in this section.
ADAPTER	ID adapter ID TYPE adapter type and link speed; for example, FIBRE 2Gb **Note**: ??? is displayed if type is unknown.
LINK TYPE	type of I/O handled by the adapter: ECKD extended count key data (statistics not available for 2105 ESS model) PPRC Peer-to-Peer Remote Copy SCSI small computer system interface **Note**: CMF MONITOR cannot identify link type if there is no activity; in this case, NO DATA OR ALL DATA IS ZERO is displayed. The SUMMARY line displays the summary of activity in both directions (read and write or send and receive).
OPERATIONS PER SEC	average number of operations completed per second; the meaning of operation depends on the link type ECKD one count per command PPRC one count per command SCSI one count per read or write
RESP TIME/OPERATION	average response time per operation, in milliseconds
BYTES PER SEC	average number of bytes transferred per second
BYTES PER OPERATION	average number of bytes transferred per operation
I/O INTENSITY	measurement of how busy a link or adapter is I/O INTENSITY is the response time of all operations divided by elapsed time and measured in milliseconds/second. Because a FICON channel can carry concurrent operations, it is possible for this value to exceed 1000.
ESS Extent Pool Statistics Section	
EXTENT POOL	pool of allocation units for logical volumes ID extent pool identifier TYPE extent pool type, for example, FIBRE 1 Gb or CKD 1 Gb

Table 70 Field descriptions for the ESS Statistics Report (part 2 of 2)

Field	Description
CAPACITY (GBYTES)	capacity, in gigabytes, of physical storage for real extents in the extent pool; the available capacity for the operating system
EXTENT COUNT	number of extents defined in an extent pool; a discrete number of extents can be used to create volumes
ALLOC EXTENTS	number of extents allocated from the extent pool
	ESS Rank Statistics Section
EXTENT POOL	pool of allocation units for logical volumes ID — extent pool identifier TYPE — extent pool type, for example, FIBRE 1 Gb or CKD 1 Gb
RAID RANK ID	RAID rank identifiers in the extent pool **Note**: If RAID RANK ID=POOL, that line represents a summary of all rank values of the entire extent pool. If the EXTENT POOL has only one RAID rank, no summary (POOL) line is generated.
READ OPERATIONS	summary of read operations: OPS / SEC — number of read operations per second RTIME / OP — average response time of read operations in milliseconds BYTES / SEC — average number of bytes read per second BYTES / OP — average number of bytes per read operation
WRITE OPERATIONS	summary of write operations: OPS / SEC — number of write operations per second RTIME / OP — average response time of write operations in milliseconds BYTES / SEC — average number of bytes written per second BYTES / OP — average number of bytes per write operation
ARRAY	NUM — number of arrays on the rank WDTH — sum of disk drive modules (DDMs) of a rank, excluding spares of the rank For example, for a RAID-5 array with 6 data disks and 1 parity disk, ARRAY WDTH is 7; for a RAID-10 with 3 mirrored disks, ARRAY WDTH is 6.
MIN RPM	slowest drive of the rank in units of 1000 RPM (revolutions per minute)
RANK CAP	sum of bytes of a rank
RAID TYPE	RAID type found for the rank, for example, RAID-5 or RAID-10 **Note**: In the line displaying the average values for the entire extent pool (where RRID=POOL), “MIXED” is shown if different RAID types have been encountered for the individual ranks in the extent pool.

Exception Subreport

The Exception Subreport lists the exceptional conditions defined by the EXCEPTS Analyzer control statement (see “EXCEPTS” on page 252). The listing shows the exceptional condition being reported and the time that the condition was detected. Up to four fields can be measured for an exceptional condition.

The Extractor control statements that are required for the Exception Subreport are dependent upon the information to be reported. (See “Numeric list of record types” on page 47 for information about the record types generated by each Extractor control statement and a description of the specific type of data collected by each Extractor control statement.)

A sample of the Exception Subreport is shown in Figure 62.

Figure 62 Exception Subreport

```
PRODUCED BY CMF ANALYZER (v.r.mm)                         EXCEPTION SUBREPORT                                 RPTSEQ     5 PAGE      8
BMC SOFTWARE, INC.                                            XYZ COMPANY                                     REPORT DATE: DD MMM YY 16.53
ACTL 10 JUN YY 09.00.00  10 JUN YY 12.45.00             WORLDWIDE HEADQUARTERS                                SYSTEM ID: SJSC     02.09.00


    EXCEPTIONAL CONDITION WHEN PAGESEC   LT         8  OR GT        20

   CPU ENVIRONMENT REQUESTED:  ALL CPU'S
   DATE     TIME                                  VALUE
                            PAGESEC        PAGEINS        PAGEOUTS

 10JUN 03    09.00.00           0.21         381.00           0.00
             09.30.00           0.44         791.00           0.00
             10.00.00           0.20         358.00           0.00
             10.30.00           0.06         103.00           0.00
             11.00.00           0.06         100.00           0.00
             11.30.00           0.02          28.00           0.00
             12.00.00           0.06          87.00          25.00
```

Exception Subreport field descriptions

Table 71 describes each field in the Exception Subreport.

Table 71 Field descriptions for the Exception Subreport

Field	Description
EXCEPTIONAL CONDITION WHEN xxx	measure name from EXCEPTS control statement (xxx)
LT	minimum threshold value from EXCEPTS control statement or the default
GT	maximum threshold value from EXCEPTS control statement or the default
CPU ENVIRONMENT REQUESTED	CPU in a multiprocessor system for which CPU-related measures are to be reported; from the EXCEPTS control statement
DATE TIME	interval start date and time during which the exceptional condition occurred
VALUE	selected exception measures (from EXCEPTS control statement) are printed as column headings; values of the exception measures fill the columns Associated measures are printed when the exception measure is printed, if the ASSOC parameter of the EXCEPTS control statement is used.

Exception Trace Detail Report

The Exception Trace Detail Report is produced by using the Analyzer EXCEPTS control statement (see “EXCEPTS” on page 252). Data for this report is gathered by using the TRACE76 Extractor control statement (see “TRACE76” on page 191).

The report lists the exceptional conditions defined on the EXCEPTS control statement. The listing shows the exceptional condition being reported and the time that the condition was detected. Up to four fields can be measured during an exceptional condition.

One line of values is produced for each date and time interval that is specified on the EXCEPTS command. If the default time interval is used, one line is produced for each sample cycle that is encountered in the input data record. These lines of data are printed only if they fall outside the exception limits that are specified in the control statement. Large amounts of input data require large amounts of storage for DMSS. Storage is defined by the REGION parameter of the //CMFRPTS EXEC control statement (see “DMSS” on page 248 for more information).

If the TRCETYPE parameter is not used in parallel with the TRACE=YES parameter, the value AVG is substituted in the list for each trace measure that does not have a corresponding characteristic value.

A sample of the Exception Trace Detail Report is shown in Figure 63.

Figure 63 Exception Trace Detail Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    EXCEPTION TRACE DETAIL REPORT                    RPTSEQ     3 PAGE     14
BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 11.54
ACTL 10 JUN YY 11.25.46  10 JUN YY 11.45.01            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SYSB     Z v.rr.n


    EXCEPTIONAL CONDITION WHEN ASMIORQR  LT         0  OR GT        0
    CPU ENVIRONMENT REQUESTED:  CPU  0
   DATE     TIME                                  VALUE
                         ASMIORQR       RCEAFC         CCVUTILP       RCVPTR
                            AVG           AVG            AVG            AVG
 10JUN 03    11.29.46    1.44E+06        185.80         102.50          33.70
             11.30.00    1.45E+06        251.95         102.77          34.17
             11.31.00    1.45E+06        207.83          93.45          31.85
             11.32.00    1.46E+06        285.03          47.82          22.83
             11.33.00    1.47E+06        246.98          89.60          49.45
             11.34.00    1.48E+06        277.43          86.28          30.28
             11.35.00    1.48E+06        272.10          99.18          23.85
             11.36.00    1.48E+06        393.67          89.78          14.23
             11.37.00    1.48E+06        327.67          87.17           9.93
```

Exception Trace Detail Report field descriptions

Measure values for the Exception Trace Detail Report are described in Appendix C, “Measure and trace values.”

Extractor Summary Report

The Extractor Summary Report shows a synopsis of system activity during the Extractor interval. This report is produced dynamically by the Extractor according to a user-definable interval.

There is no Analyzer control statement for the Extractor Summary Report. To produce this report, the required Extractor control statement is EXTSUM with the SPINOFF=*a*, where *a* is a valid output class (see “EXTSUM” on page 154).

The following control statements provide data used by this report:

- CHANNEL (see page 134)
- CPU (see page 136)
- IOQ (see page 167)
- PAGING (see page 173)
- TSODATA (see page 194)

The Extractor Summary Report is divided into five sections:

- **CPU Utilization**

  This section lists the CPU busy percentages for each processor on the system.

  NOTE

  This information is accurate only if PR/SM is not active.

- **System Utilization**

  This section shows the five busiest channel paths, along with spool space usage, page-in rate, swap rate, I/O rate, and TSO transaction rate.

- **Job Class Activity**

  This section shows the average number of active jobs in a given job class, along with the average amount of CPU used by that job class and the number of jobs active in that class at the end of the interval. The total number of jobs and job steps completed during the interval also are displayed. The job activity is also summarized by batch (all job classes), TSO, and Started Tasks.

NOTE

The job class and job class descriptions are taken from the user specification on the EXTSUM statement. If JES=NO is specified, no job class activity section is produced and no spool space usage is calculated

- **Service Class Activity**

  This section summarizes, by service class and period during the Extractor interval, a variety of goal mode performance measures, including performance index and execution velocity. The service class and periods to be summarized are taken from the SRVCLASS parameter on the EXTSUM statement. If no SRVCLASS parameter is specified on the EXTSUM statement, this section is not produced. This section is produced only for systems running in goal mode.

An example of the Extractor Summary Report is shown in Figure 64.

Figure 64 Extractor Summary Report (part 1 of 2)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     EXTRACTOR SUMMARY REPORT                         SYSTEM ID: SJSE    PAGE     1
BMC SOFTWARE, INC.                                        FOR  5 MINUTES                               INTERVAL: 10 JUN 03 17:30/17:35
-------------------------------------------------------<  CPU UTILIZATION  >-----------------------------------------------------
            CPU ID        -0-    -1-    -2-
            % BUSY        7.7    7.4    6.9
------------------------------------------------------<  SYSTEM UTILIZATION  >---------------------------------------------------
     PERCENT UTILIZATION       CH PATH    AVG                                                      AVG      MIN      MAX
     FIRST  BUSIEST PATH          F6     18.63                   SPOOL SPACE USAGE    ( PERCENT )    51.5     51.5     51.5
     SECOND BUSIEST PATH          EC     17.79                   PAGE-IN RATE         (PAGES/SEC)     3.3      0.0     14.4
     THIRD  BUSIEST PATH          B4      1.06                   SWAP SEQUENCES       (SWAPS/MIN)     0.0      0.0     46.0
     FOURTH BUSIEST PATH          59      0.23                   I/O ACTIVITY RATE    (IO'S /SEC)     ---      ---      ---
     FIFTH  BUSIEST PATH          F8      0.08                   TSO TRANSACTION RATE (TRANS/MIN)     ---      ---      ---
-----------------------------------------------------<  JOB CLASS ACTIVITY  >----------------------------------------------------
                               JOB       JOB CLASS          ACTIVE  JOBS                 COMPLETED
                              CLASS      DESCRIPTION         AVG    END    %CPU            JOBS  STEPS
                              A          CLASSA              0.0      0     0.0               0      0
                              B          CLASSB              0.0      0     0.0               0      0
                              OTHER      ALL_OTHERS          5.1      5     6.2               3      3
                              BATCH      ALL_BATCH           5.1      5     6.2               3      3
                              TSO        ALL_TSO            14.0     14     0.0               0      0
                              STC        ALL_STC            24.9     29    93.8               0      0
                              TOTAL                         44.0     48   100.0               3      3

----------------------------------------------------<  SERVICE CLASS ACTIVITY  >-------------------------------------------------

                POLICY: BBPLEX01              ACTIVATED:  06/10/YY AT 00:50:20              INTERVAL:   5 MINUTES


 SERVICE CLASS: BATNRM   DESCRIPTION: Batch Normal Jobs                  WORKLOAD: BATCH      RESOURCE GROUP:           PERIODS: 2

PER  IMP   PERF    AVG NO     TRANSACTIONS           RESPONSE TIME     EXECUTION VELOCITY     CPU     EXE      TOTAL SERVICE UNITS
           INDX   ADR SPCS   PER/SEC / TOTAL        GOAL         AVERAGE   GOAL     ACTUAL    USING   DELAY     PER/SEC      SU/1000
 1    4    0.21        2.1      0.01       3                        8.01   10%      47.9%       57      62       3,159          947
 2    5    0.05        3.0                                                   1%      20.3%       16      63       4,260        1,278


 SERVICE CLASS: STCLOW   DESCRIPTION: Low Priority STC's                 WORKLOAD: STC        RESOURCE GROUP:           PERIODS: 1

PER  IMP   PERF    AVG NO     TRANSACTIONS           RESPONSE TIME     EXECUTION VELOCITY     CPU     EXE      TOTAL SERVICE UNITS
           INDX   ADR SPCS   PER/SEC / TOTAL        GOAL         AVERAGE   GOAL     ACTUAL    USING   DELAY     PER/SEC      SU/1000
 1    D                1.0                                                                       0       0           0            0

 (continued on next page)
```

Figure 64 Extractor Summary Report (part 2 of 2)

```
 SERVICE CLASS: STCNRM   DESCRIPTION: Normal STC's                       WORKLOAD: STC       RESOURCE GROUP:           PERIODS: 1

PER  IMP   PERF    AVG NO     TRANSACTIONS            RESPONSE TIME     EXECUTION VELOCITY     CPU     EXE       TOTAL SERVICE UNITS
           INDX   ADR SPCS   PER/SEC / TOTAL         GOAL         AVERAGE   GOAL     ACTUAL     USING   DELAY     PER/SEC      SU/1000
 1    4    1.16       39.0                                                   25%     21.6%      133     482      56,864       17,059


 SERVICE CLASS: STCPAS   DESCRIPTION: PAS STC's                          WORKLOAD: STC       RESOURCE GROUP: PASSTC    PERIODS: 1

PER  IMP   PERF    AVG NO     TRANSACTIONS            RESPONSE TIME     EXECUTION VELOCITY     CPU     EXE       TOTAL SERVICE UNITS
           INDX   ADR SPCS   PER/SEC / TOTAL         GOAL         AVERAGE   GOAL     ACTUAL     USING   DELAY     PER/SEC      SU/1000
 1    3    2.16        2.0                                                   60%     27.8%       20      52      26,066        7,819


 SERVICE CLASS: STCPROD  DESCRIPTION: Production STC's                   WORKLOAD: STC       RESOURCE GROUP:           PERIODS: 1

PER  IMP   PERF    AVG NO     TRANSACTIONS            RESPONSE TIME     EXECUTION VELOCITY     CPU     EXE       TOTAL SERVICE UNITS
           INDX   ADR SPCS   PER/SEC / TOTAL         GOAL         AVERAGE   GOAL     ACTUAL     USING   DELAY     PER/SEC      SU/1000
 1    3                1.0                                                   40%                  0       0           0            0


 SERVICE CLASS: TSONRM   DESCRIPTION: Normal TSO Users                   WORKLOAD: TSO       RESOURCE GROUP:           PERIODS: 4

PER  IMP   PERF    AVG NO     TRANSACTIONS            RESPONSE TIME     EXECUTION VELOCITY     CPU     EXE       TOTAL SERVICE UNITS
           INDX   ADR SPCS   PER/SEC / TOTAL         GOAL         AVERAGE   GOAL     ACTUAL     USING   DELAY     PER/SEC      SU/1000
 1    2    0.50       13.9      0.11       33      .50  90%           .03           25.0%        1       3         174           52
 2    3    0.50        0.0      0.02        5     2.00  80%           .07            0.0%        0       2           7            2
 3    4                0.0
 4    4                0.0
```

Extractor Summary Report field descriptions

Descriptions of the fields in the Extractor Summary Report are listed in Table 72. For more information, see Appendix B, “Workload measurement.”

Table 72 Field descriptions for the Extractor Summary Report (part 1 of 4)

Field	Description
CPU ID	CPU ID for each CPU assigned to the partition from which the records were extracted
% BUSY	percentage of CPU busy for each CPU
CH PATH	channel path number
AVG	average channel path utilization
SPOOL SPACE USAGE	average / minimum / maximum percent of JES spool space used
PAGE-IN RATE	average / minimum / maximum page-ins per second; the minimum / maximum values represent a period of one minute
SWAP SEQUENCES	average / minimum / maximum address space swap sequences per minute
I/O ACTIVITY RATE	average / minimum / maximum number of SSCH instructions per second; the minimum / maximum values represent a period of one minute **Note**: Data for this field is not available on 3090 or later processors. When data is not available, this field contains dashes (---).
TSO TRANSACTION RATE	average / minimum / maximum number of TSO transactions per minute
JOB CLASS	job class as specified on the EXTSUM statement

Table 72 Field descriptions for the Extractor Summary Report (part 2 of 4)

Field	Description
JOB CLASS DESCRIPTION	job class description as specified on the EXTSUM statement
ACTIVE JOBS AVG	average number of jobs active in a given class during the Extractor interval
ACTIVE JOBS END	number of jobs active in a given job class at the end of the Extractor interval
ACTIVE JOBS % CPU	average percentage of CPU busy time that was used by a given job class during the Extractor interval
COMPLETED JOBS	total number of jobs that were completed for a given job class during the Extractor interval
COMPLETED STEPS	total number of job steps that were completed for a given job class during the Extractor interval
SUMMARY PERIODS	performance periods as specified on the EXTSUM PP parameter
WORKLOAD DESCRIPTION	description as specified on the EXTSUM PD parameter
RESPONSE TIME	average response time for all transactions that ended in the specified period or periods
TRANSACTIONS PER/SEC	average number of transactions per second during the Extractor interval for the specified period or periods
TRANSACTIONS TOTAL	total number of transactions during the Extractor interval for the specified period or periods; M indicates units of 1000
EXECUTION VELOCITY	rate at which transactions are executing
CPU SERVICE PER/SEC	average number of CPU service units per second that were used during the Extractor interval for the specified period or periods
CPU SERVICE SU/1000	total number of CPU service units, in thousands, that were used during the Extractor interval for the specified period or periods
POLICY	named set of performance goals that the workload manager uses as a guideline to match resources to work
ACTIVATED	date and time that the service policy became active
INTERVAL	workload interval This value is usually the same as the Extractor recording interval, but it might be shorter if the policy changed in the middle of an Extractor recording interval.
SERVICE CLASS	group of work with the same performance goals, resource requirements, or business importance Service class names are taken from the workload policy. For a service class to show up in this report, it must ■ be listed in the service policy ■ be listed in the SRVCLASS parameter of the EXTSUM Extractor control statement ■ have at least one active address space during the interval.

Table 72 Field descriptions for the Extractor Summary Report (part 3 of 4)

Field	Description
DESCRIPTION	description given for the service class in the active policy, which is specified when the service class is defined; this field is blank if the service class description in the policy is blank
WORKLOAD	name given to the workload to which this service class belongs A workload is a group of work that is tracked, managed, and reported as a unit. It is usually a grouping of similar service classes.
RESOURCE GROUP	resource group defined for the service policy The resource group determines the amount of processing capacity across one or more MVS images assigned for this service class. If no resource group is defined, this field is blank.
PERIODS	number of periods defined for this service class A service class can have up to eight periods defined for it. This field shows the actual number of periods that the service class contains, even if you did not specify all of those periods on the SRVCLASS parameter of the EXTSUM Extractor control statement.
PER	service class period being reported on for the row of information that follows it If an individual period has no active address spaces for the interval, it contains a zero, which is its average number of address spaces. If none of the periods for a service class has an active address space, the service class is not listed. For more information about service class periods, see Appendix C, “Measure and trace values.”
IMP	importance level assigned to this goal
PERF INDX	performance index for this service class; the performance index is a relative calculation to determine how well your service classes are meeting their goals
AVG NO ADR SPCS	average number of address spaces for the interval for the specific service class and period If the value in this field is zero and no other data is printed in this row, no address spaces were active, and no other information for this period is printed. If this number is zero for all periods within a service class, no information about that service class is printed. The average number is based on the MVS WLM sampling rate, which is currently 250 milliseconds.
TRANSACTIONS – PER/SEC	number of transactions executed for a period divided by the number of seconds in the workload interval If this field contains a value of 0.0, it could indicate that the service class includes Started Tasks or batch jobs, but there are not enough significant digits to display the true number. This field is blank only if no transactions executed.
TRANSACTIONS – TOTAL	total number of transactions that executed during the workload interval

Table 72 Field descriptions for the Extractor Summary Report (part 4 of 4)

Field	Description
RESPONSE TIME – GOAL	first column in field—the goal for the response time, in seconds; second column in field—the percentage of jobs that should meet or exceed that goal **Note**: For an average response time goal, only the first column (response time) contains information. For a percentage response time goal, both columns (response time and the percentage) contain information.
RESPONSE TIME – AVERAGE	average response time for the goal The value in this field is displayed as *mmmm:ss.hh*, where *mmmm* is minutes, *ss* is seconds, and *hh* is hundredths of seconds. This value can be quite large for a particular service class if the transaction length for a Started Task is the length of the Started Task. Asterisks (****) appear in this field if the number is too large to be printed.
EXECUTION VELOCITY – GOAL	goal for the execution velocity of a service class that has been defined with an execution velocity goal; for service classes with goals other than execution velocity goals, this field is blank
EXECUTION VELOCITY – ACTUAL	actual execution velocity for the service class
CPU USING	number of samples during which this service class was using CPU
EXE DELAY	total number of samples during which an execution delay was detected
TOTAL SERVICE UNITS – PER/SEC	number of service units that were used per second; includes all service units collected by WLM (TCB, IOC, central storage, and SRB)
TOTAL SERVICE UNITS – SU/1000	total number of service units that were used by the service class during the workload interval, divided by 1000

FICON Director Activity Report

The FICON Director Activity Report provides configuration and activity information for all FICON Directors (switches) to which the system was connected when the Extractor recorded the data. You can use this report to see how each FICON switch was configured and how it was being utilized

The FICON Director Activity Report is produced by using the FICONSW Analyzer control statement (see “FICONSW” on page 258). This report processes SMF 74-7 records that are created by the FCSW sampler. That sampler is started by the FICONSW Extractor control statement (see “FICONSW” on page 158).

This report consists of two sections:

- **IODF Section**

  The Input/Output Definition File (IODF) is a VSAM linear data set that contains I/O definition information. This section of the report identifies the number and name of the IODF that was in use at the time the data was collected by the Extractor.

- **FICON Director Section**

  This section of the report displays information about each FICON switch that was detected. This section also displays information about all of the ports that are assigned to each switch and how each port was being utilized during the recording period.

NOTE

The information shown for each port includes all of the I/O that is directed through this port, regardless of which system requests the I/O. That is, the measurements shown are not specific or limited to the system on which the Extractor gathered and recorded the measurements.

An example of the FICON Director Activity Report is shown in Figure 65.

Figure 65 FICON Director Activity Report (part 1 of 2)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    FICON DIRECTOR ACTIVITY REPORT                        RPTSEQ      3 PAGE       5
BMC SOFTWARE, INC.                                           BMC ENGINEERING                               REPORT DATE: DD MMM YY 10.52
ACTL 25 OCT YY 16.25.00  25 OCT YY 17.24.28                                                                SYSTEM ID: SJSC     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-7/24/12/0.99

----------------------------------------------------------- IODF SECTION -----------------------------------------------------------
IODF SUFFIX: 03   CREATED: 09/29/YYYY - 16.21.34   IODF DSN: BMCS.IODF03
ACTIVATED: IPL        CHANGED THIS REPORT PERIOD: NO

------------------------------------------------------ FICON DIRECTOR SECTION ------------------------------------------------------

SWITCH DEVICE: 0105   SWITCH ID: 61   TYPE: 006140   MODEL: 001   MFR: MCD   PLANT: 01   SERIAL: 00000131278E

 PORT     -CONNECTION-     AVG. FRAME     --AVG. FRAME SIZE--     TRANSFER RATE(MB/SEC)     ERROR
 ADDR      TYPE    ID      PACING(MS)       READ       WRITE       READ         WRITE        COUNT
 --       ------  ----     ----------     -------    -------     --------     --------       -----

 04       CU      ----              0        644       1,546        0.118        0.443          0
 05       ------  ----              0      1,604         181        0.187        0.009          0
 06       ------  ----              0      1,515          53        0.019            0          0
 07       ------  ----              0      1,518       1,041        0.119        0.067          0
 08       ------  ----              0          0           0            0            0          0
 09       ------  ----              0      1,497         553        0.117        0.030          0
 0A       ------  ----              0         68       1,911            0        0.011          0
 0B       ------  ----              0          0           0            0            0          0
 0C       CHP-H    E8               0      1,645          73        2.162        0.026          0
 0D       CHP-H    E9               0      1,961          76        2.314        0.007          0
 0E       CHP-H    EA               0      1,661          68        3.945        0.044          0
 0F       CU      ----              0         68       1,645        0.024        2.162          0
 10       CU      ----              0         76       1,961        0.007        2.314          0
 11       CU      ----              0         68       1,661        0.044        3.945          0
 12       CHP-H    E2               0      1,798       1,202        1.036        0.198          0
 13       CHP-H    E4               0      1,799       1,220        1.037        0.204          0
 14       CU      ----              0      1,202       1,798        0.198        1.036          0
```

(continued on next page)

Figure 65 FICON Director Activity Report (part 2 of 2)

```
15      CU      ----          0     1,220     1,799     0.204     1.037     0
16      CHP-H     F3          0       465     1,002     1.480     3.984     0
17      CHP-H     F4          0       455     1,002     1.427     3.944     0
18      CU      ----          0     1,002       465     3.984     1.480     0
19      CU      ----          0     1,002       455     3.944     1.427     0
1A      CHP-H     F0          0       263       989     1.075     3.362     0
1B      CU      ----          0       997       264     3.356     1.073     0
1C      CHP-H     F1          0       258     1,002     1.052     3.409     0
1D      CU      ----          0     1,002       258     3.409     1.052     0
1E      CHP       F0          0     1,060       640     0.816     0.330     0
1F      CU      ----          0       656     1,079     0.325     0.814     0
20      CHP       F2          0       765       924     0.828     0.892     0
21      CHP       F3          0         0         0         0         0     0
22      CU      ----          0     1,474       473     6.771     0.904     0
23      CHP       F0          0        92     1,621     0.038     2.937     0
24      CHP       E5          0         0         0         0         0     0
25      CHP       F2          0        92     1,621     0.038     2.942     0
26      CHP       F4          0        71       171         0         0     0
```

FICON Director Activity Report field descriptions

Table 74 describes each field in the FICON Director Activity Report.

Table 73 Field descriptions for the FICON Director Activity Report (part 1 of 2)

Field	Description
IODF Section	
IODF SUFFIX	two-character hexadecimal suffix for the IODF data set name
CREATED	date and time that this IODF data set was created
IODF DSN	data set name of the IODF
ACTIVATED	text indicating how this IODF was activated
CHANGED THIS REPORT PERIOD	text indicating whether the IODF was changed during this reporting period
FICON Director Section	
SWITCH DEVICE	four-digit hexadecimal address of the FICON switch
SWITCH ID	two-digit hexadecimal switch ID that is associated with this FICON switch **Note**: The SWITCH ID of cascaded switches is shown as '**'.
TYPE	device type of the FICON switch
MODEL	model number of the FICON switch
MFR	manufacturer of the FICON switch
PLANT	plant number where the FICON switch was manufactured
SERIAL	serial number of the FICON switch
PORT ADDRESS	hexadecimal address of the port on the FICON switch

Table 73 Field descriptions for the FICON Director Activity Report (part 2 of 2)

Field	Description
CONNECTION TYPE	identifies the type of connection; the values that you can see are as follows: CHP—denotes a channel path CHP-H—denotes a channel path of the system on which the Extractor gathered and recorded the data CU—denotes a control unit SWITCH—denotes a switch If the unit is not unique, dashes are displayed. For example, for CTC channels there might be a CU and a CHP connected to the same port.
CONNECTION ID	hexadecimal identifier of the connector; dashes are shown if the connection type is SWITCH, CU, or is not unique
AVG. FRAME PACING (MS)	average time (in milliseconds) that a frame had to wait before it could be transmitted during the report period **Note**: While the unit of time on the report is milliseconds, it is displayed in four decimal places, the equivalent of tenths of microseconds.
AVG. FRAME SIZE	average frame sizes (in bytes) used in reads and writes during this report period
TRANSFER RATE (MB/SEC)	rate (in million bytes per second) that data was read and written during the report period
ERROR COUNT	number of errors that were encountered during the report period

Graphics Trace Detail Report

The Graphics Trace Detail Report is produced by using the GRAPH TYPE=TRACE Analyzer control statement (see “GRAPH” on page 259). You can specify up to 12 measures. If you specify more than 12, only the first 12 measures are used. Data is gathered for the Graphics Trace Detail Report by using the TRACE76 Extractor control statement (see “TRACE76” on page 191).

Each measure selected is represented by a column in the report. The measures are printed in the same sequence as they appear in the GRAPH command.

One line of values is produced for each date and time interval specified on the GRAPH command. If the default time interval is used, one line is produced for each sample cycle section encountered in the input data record.

If the TRCETYPE parameter is not used in parallel with the TYPE=TRACE parameter, a separate report is generated with all six characteristics printed in the body of the report for each trace measure that does not have a corresponding characteristic value. Consequently, you can generate 12 separate reports with one graph control statement.

An example of the Graphics Trace Detail Report is shown in Figure 66.

Figure 66 Graphics Trace Detail Report

```
 PRODUCED BY CMF ANALYZER (v.r.mm)                   GRAPHICS TRACE DETAIL REPORT                      RPTSEQ     3 PAGE      5
 BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 10.24
 ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


   CPU ENVIRONMENT REQUESTED:  CPU  0
   DATE     TIME    CCVUTILP CCVUTILP CCVUTILP CCVUTILP CCVUTILP CCVUTILP
                       MIN      AVG      MAX      END      STD      DIF
10JUN 03  09.00.00      0.00     8.65    38.00     3.00     9.23     0.00
          09.01.00      0.00     3.92    29.00     4.00     4.18     0.00
          09.02.00      0.00     7.43    38.00     4.00     8.39     0.00
          09.03.00      0.00     8.75    40.00    11.00     8.49     0.00
          09.04.00      2.00     8.43    38.00    14.00     6.44     0.00
          09.05.00      2.00     6.82    35.00     8.00     5.42     0.00
          09.06.00      0.00     6.32    38.00     8.00     6.55     0.00
          09.07.00      0.00     6.05    38.00    12.00     8.12     0.00
          09.08.00      0.00     4.42    21.00     3.00     4.79     0.00
          09.09.00      0.00     4.08    12.00    10.00     2.92     0.00
          09.10.00      0.00     3.98    28.00     2.00     4.21     0.00
          09.11.00      0.00     4.78    35.00    13.00     6.81     0.00
          09.12.00      0.00     4.72    51.00     3.00     7.86     0.00
          09.13.00      0.00     5.03    14.00     3.00     3.84     0.00
          09.14.00      0.00     6.26    67.00    67.00     9.49     0.00
          09.15.00      0.00     7.20    36.00     2.00     7.82     0.00
          09.16.00      0.00     9.27    41.00    10.00    10.37     0.00
          09.17.00      0.00     4.88    50.00     4.00     7.52     0.00
          09.18.00      0.00     5.33    37.00    37.00     7.14     0.00
          09.19.00      0.00     3.60    27.00    11.00     4.01     0.00
          09.20.00      0.00     4.22    24.00     2.00     3.93     0.00
          09.21.00      0.00     4.58    45.00     8.00     6.71     0.00
          09.22.00      0.00     3.77    34.00     0.00     5.22     0.00
          09.23.00      0.00     3.50    15.00     4.00     3.11     0.00
          09.24.00      0.00     5.28    38.00    12.00     6.34     0.00
          09.25.00      0.00     5.50    34.00     3.00     5.67     0.00
          09.26.00      0.00     4.93    35.00     4.00     5.00     0.00
          09.27.00      0.00     6.83    96.00     8.00    12.92     0.00
          09.28.00      0.00     5.20    36.00    14.00     6.61     0.00
          09.29.00      0.00     7.23   100.00   100.00    13.45     0.00
          09.30.00      0.00    12.62   100.00    12.00    17.92     0.00
          09.31.00      0.00     7.37    66.00    13.00    11.14     0.00
          09.32.00      0.00     8.72    56.00     3.00    11.48     0.00
          09.33.00      0.00     3.62    42.00    42.00     5.73     0.00
          09.34.00      0.00     5.52    60.00    10.00     9.04     0.00
          09.35.00      0.00     4.72    23.00    12.00     4.14     0.00
          09.36.00      0.00     3.33    14.00    12.00     3.16     0.00
          09.37.00      0.00     6.15    36.00     4.00     6.36     0.00
          09.38.00      0.00     5.22    32.00     3.00     4.98     0.00
          09.39.00      0.00     4.03    41.00     2.00     7.00     0.00
          09.40.00      0.00     4.58    30.00     4.00     5.82     0.00
```

Graphics Trace Detail Report field description

Table 74 describes the field in the Graphics Trace Detail Report.

Table 74 Field description for the Graphics Trace Detail Report

Field	Description
CPU ENVIRONMENT REQUESTED	CPU selected in GRAPH control statement

Measure values for the Graphics Trace Detail Report are described in Appendix C, “Measure and trace values.”

HFS Statistics Report

The HFS Statistics Report provides information about caching by HFS buffer pools and mounted file systems.

The HFS Statistics Report is produced by using the HFS Analyzer control statement. The data for this report is obtained by using the HFS Extractor control statement.

This report consists of three sections:

- **Global Statistics**

  This section displays storage that is used by all buffer pools and overall caching of metadata and the first page of files.

- **Buffer Pool Statistics**

  This section displays storage that is used and I/O activity between buffers and DASDs for each buffer pool.

- **File System Statistics**

  This section displays storage that is used, I/O requests of files, caching of metadata, index, and the first page of files for each mounted file system.

NOTE

This section is formatted only if the FSNLIST= parameter is specified on the HFS Extractor control statement.

The report might issue the error message `ERROR: bad-data UNAVAILABLE RC=cccccccc-rrrrrrrr,` where *bad-data* is BUFFER LIMITS, BUFFER STATS, or STATISTICS. The return code *cccccccc* and the last four digits of reason code *rrrrrrrr* are documented in the *z/OS UNIX System Services Messages and Codes* manual.

An example of the HFS Statistics Report is shown in Figure 67.

Figure 67 HFS Statistics Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        HFS STATISTICS REPORT                              RPTSEQ     3 PAGE      5
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                                REPORT DATE: DD MMM YY11.35
REQD 27 JUL 06 09.45.00  27 JUL 06 23.59.59                  HOUSTON, TX.                                   SYSTEM ID: SJSD     Z v.rr.n
ACTL 27 JUL 06 09.45.00  27 JUL 06 11.15.00                                                                 REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-6/6/0/1.5
-------------------------------------------------------- GLOBAL STATISTICS -------------------------------------------------------

---------- STORAGE (MB) -----------   ----------------------- I/O ACTIVITY ------------------------
        -- LIMITS --   -- IN USE --          ----- FILE 1ST PAGE ------   -------- METADATA --------
                                              RATE    COUNT   HIT RATIO    RATE    COUNT   HIT RATIO
TOTAL   MAX    2,011          6.250   CACHE      0        1         100   1.068    5,768        99.8
FIXED   MIN        0              0   DASD       0        0               0.002       10

------------------------------------------------------ BUFFER POOL STATISTICS ----------------------------------------------------

                  -------------- STORAGE --------------   ------------- I/O ACTIVITY --------------
BUFFER     DATA   BUFFERS     PAGES     BYTES   % FIXED   --- TOTAL ----   --- FIXED ----   % FIXED
  SIZE   SPACES                                             RATE    COUNT    RATE    COUNT
     1        1       784       784    3,136K         0   0.291    1,571       0        0         0
     4        1         4        16       64K         0       0        0       0        0         0
    16        1        14       224      896K         0       0        0       0        0         0
    64        1         9       576    2,304K         0       0        0       0        0         0

------------------------------------------------------ FILE SYSTEM STATISTICS ----------------------------------------------------
FILE SYSTEM - BBO.V60.SJSE.BASE.CONFHFS                       MOUNTED: 07/27/2006 04:27:29   DURATION:     01.30.00
STORAGE     - ALLOCATED:    176M   USED:    137M   ATTR DIR: 11,740K   CACHED:  3,616K

------- DATA I/O --------            -1ST PAGE I/O-   -METADATA I/O-    -- INDEX I/O -   -- INDEX EVENTS --
            RATE    COUNT              RATE    COUNT    RATE    COUNT    RATE    COUNT                  COUNT
SEQUENTIAL     0        0  CACHE          0        0   2.086   11,265   2.137   11,543   NEW LEVELS       0
RANDOM         0        0  DASD           0        0       0        1       0        1   SPLITS           0
                           HIT RATIO    0               100              100             JOINS            0


FILE SYSTEM - USS.BBPLEX01.ROOT                               MOUNTED: 07/27/2006 00:14:09   DURATION:     01.30.00
STORAGE     - ALLOCATED:  5,040K   USED:  4,916K   ATTR DIR:     52K   CACHED:      0K

------- DATA I/O --------            -1ST PAGE I/O-   -METADATA I/O-    -- INDEX I/O -   -- INDEX EVENTS --
            RATE    COUNT              RATE    COUNT    RATE    COUNT    RATE    COUNT                  COUNT
SEQUENTIAL     0        0  CACHE          0        0   2.697   14,566   2.702   14,593   NEW LEVELS       0
RANDOM         0        0  DASD           0        0   0.002        9       0        0   SPLITS           0
                           HIT RATIO    0              99.9              100             JOINS            0


FILE SYSTEM - USS.SJSE.VAR                                    MOUNTED: 07/27/2006 00:14:10   DURATION:     01.30.00
STORAGE     - ALLOCATED:  6,528K   USED:  1,924K   ATTR DIR:     80K   CACHED:      0

------- DATA I/O --------            -1ST PAGE I/O-   -METADATA I/O-    -- INDEX I/O -   -- INDEX EVENTS --
            RATE    COUNT              RATE    COUNT    RATE    COUNT    RATE    COUNT                  COUNT
SEQUENTIAL     0        0  CACHE          0        0   1.977   10,676   1.977   10,676   NEW LEVELS       0
RANDOM         0        0  DASD           0        0       0        0       0        0   SPLITS           0
                           HIT RATIO    0               100              100             JOINS            0
```

HFS Statistics Report field descriptions

Table 75 describes each field in the Global Statistics section of the HFS Statistics Report.

Table 75 Field descriptions for the Global Statistics section (part 1 of 2)

Field	Description	
STORAGE (MB)	virtual storage (in megabytes) of HFS buffers at the end of the duration of the report	
TOTAL	**LIMITS - MAX**	maximum amount of virtual storage that HFS buffers should use; this value is specified by the VIRTUAL(max) parameter in BPXPRM*xx* member or the USS CONFIGHFS command
	IN USE	total amount of virtual storage in use by all HFS buffer pools
FIXED	**LIMITS - MIN**	minimum amount of fixed storage for HFS buffers; this value is specified by the FIXED(min) parameter in BPXPRM*xx* member or the USS CONFIGHFS command
	IN USE	total amount of fixed storage in use by all HFS buffer pools
I/O ACTIVITY	I/O activity of metadata and the first page of files	
FILE 1ST PAGE	**CACHE - RATE**	average number of times per second that the first page of a data file was requested and found in virtual storage (cache)
	CACHE - COUNT	number of times that the first page of a data file was requested and found in cache
	HIT RATIO	percentage of times that the first page of a data file was requested and found in cache
	DASD - RATE	average number of times per second that the first page of a data file was requested and not found in cache, and an I/O was necessary
	DASD - COUNT	number of times that the first page of a data file was requested and not found in cache, and an I/O was necessary

Table 75 Field descriptions for the Global Statistics section (part 2 of 2)

Field	Description	
METADATA	**CACHE - RATE**	average number of times per second that metadata was found in cache during file lookup
	CACHE - COUNT	number of times that metadata was found in cache during file lookup
	HIT RATIO	percentage of times that metadata was found in cache during file lookup
	DASD - RATE	average number of times per second that metadata was not found in cache during file lookup, and an index call was necessary that can result in I/O
	DASD - COUNT	number of times that metadata was not found in cache during file lookup, and an index call was necessary that can result in I/O

Table 76 describes each field in the Buffer Pool Statistics section of the HFS Statistics Report.

Table 76 Field descriptions for the Buffer Pool Statistics section

Field	Description	
BUFFER SIZE	size (in pages) of each buffer in the buffer pool	
DATA SPACES	number of data spaces allocated for the buffer pool	
STORAGE	**BUFFERS PAGES BYTES**	total amount of virtual storage in use by the HFS buffer pool in units of buffer, page, and byte
	% FIXED	percentage of buffer pool in permanently fixed storage
I/O ACTIVITY	I/O activity between buffer pools and DASDs	
	TOTAL	rate per second and count of all I/O requests
	FIXED	rate per second and count of I/O requests where a buffer was already fixed prior to I/O
	% FIXED	percentage of I/O requests where a buffer was already fixed prior to I/O

Table 77 describes each field in the File System Statistics section of the HFS Statistics Report.

Table 77 Field descriptions for the File System Statistics Section (part 1 of 2)

Field	Description	
FILE SYSTEM	name of the MVS HFS data set containing the file system	
	MOUNTED	date and time (*mm/dd/yyyy hh:mm:ss*) when the file system was last mounted If mount time changed (that is, the file system was remounted), an asterisk is displayed after the time. If the file system was never mounted, three dashes are printed.
	DURATION	sum of recording intervals (*ddd.hh.mm.ss*) during which statistics for the file system were collected; this is the divisor in the calculation of all **Rate** columns
STORAGE	**ALLOCATED**	amount of storage on DASDs allocated to the file system
	USED	amount of storage on DASDs actually used by the file system
	ATTR DIR	amount of storage on DASDs used by the file system for attribute directory
	CACHED	amount of storage in HFS buffer pools cached for this file system
DATA I/O	rate per second and count of sequential and random I/O requests of files	
1ST PAGE I/O	**CACHE**	rate per second and count of requests for the first page of a file where it was found in cache
	DASD	rate per second and count of requests for the first page of a file where it was not found in cache and I/O was necessary
	HIT RATIO	percentage of requests for the first page of a file where it was found in cache
METADATA I/O	**CACHE**	rate per second and count of requests for metadata where it was found in cache during file lookup
	DASD	rate per second and count of requests for metadata where it was not found in cache during file lookup
	HIT RATIO	percentage of requests for metadata where it was found in cache during file lookup

Table 77 Field descriptions for the File System Statistics Section (part 2 of 2)

Field	Description	
INDEX I/O	**CACHE**	rate per second and count of requests for index page where it was found in cache
	DASD	rate per second and count of requests for index page where it was not found in cache
	HIT RATIO	percentage of requests for index page where it was found in cache
INDEX EVENTS	**NEW LEVELS**	rate per second and count of requests for index page where it was found in cache
	SPLITS	rate per second and count of requests for index page where it was not found in cache
	JOINS	percentage of requests for index page where it was found in cache

HTTP Server Report

The HTTP Server Report provides information about the activities of HTTP servers. The information can be used to analyze problem servers and to view performance data.

The report consists of two parts:

- **HTTP Server Summary Report**

  This part summarizes the characteristics and activity for each server.

- **HTTP Server Details Report**

  This part consists of these sections:

  — Server Characteristics (obtained from SMF type 103-1 record)

  — Server Activity (obtained from SMF type 103-2 record)

There is one Details report per server. Data in the subtype 1 records is reported in the Server Characteristics only if there is at least one corresponding subtype 1 record.

NOTE

In this report, all rates are per second and all averages are based on counts divided by the number of samples taken during the report period.

Figure 68 HTTP Server Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      HTTP SERVER SUMMARY REPORT                       RPTSEQ      4 PAGE       6
BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 9.43
                                                         WORLDWIDE HEADQUARTERS                         SYSTEM ID: Z0        * N/A *

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  103-1/7/0/0  103-2/29/27/4.44
-----------------------------------------------------------------------------------------------------------------------------------
                                  AVAILABLE     REQUEST  RESPONSE  THROUGHPUT RATE    THREADS     CACHE SIZE    CACHE FILES  TIMEOUTS
SERVER NAME                       HHH.MM.SS        RATE      RATE        IN      OUT    MAX  USED     MAX  USED     MAX  USED

Z0EIP.PDL.POK.IBM.COM               1.20.33        4.21      4.21    860.70  19,151    200 64.67  5,120     0      NO     0          0
Z0EIP.PDL.POK.IBM.COM               0.17.43       20.59     20.59     4,087  103.3K    200 48.50  5,120     0      NO     0          0
Z0EIP.PDL.POK.IBM.COM               0.28.11       16.25     16.25     3,173  43,174    200 48.50  5,120     0      NO     0          0
Z0EIP.PDL.POK.IBM.COM               0.20.26       54.67     54.67    10,217  113.7K    200 48.50  5,120     0      NO     0          0
Z0EIP.PDL.POK.IBM.COM               0.48.41       11.23     11.23     2,003  21,195    200 24.25  5,120     0      NO     0          0
Z0EIP.PDL.POK.IBM.COM               1.25.46       16.00     16.00     2,948  32,755    200 16.17  5,120     0      NO     0          0
Z0EIP.PDL.POK.IBM.COM               0.00.00           0         0         0       0    200     0  5,120     0      NO     0          0
```

HTTP Server Summary Report field descriptions

Table 78 describes each field in the HTTP Server Summary Report.

Table 78 Field descriptions for the HTTP Server Summary Report

Field	Description
SERVER NAME	server name
AVAILABLE	total time (*hhh.mm.ss*) that the server was available during the interval
REQUEST RATE	number of requests that the HTTP server has successfully served per second
RESPONSE RATE	number of successful responses sent per second
THROUGHPUT RATE	number of bytes received or sent by this server per second
THREADS	**MAX** maximum number of threads that the server can have in the thread pool (or NO if no limit has been specified) **USED** number of currently active threads of the server This figure is an average for the report period.
CACHE SIZE	**MAX** maximum cache size (KB) for this server **USED** average cache size used by this server This figure is an average for the report period.
CACHE FILES	**MAX** maximum number of files allowed in the cache of this server **USED** average number of files cached for this server This figure is an average for the report period.
TIMEOUTS	number of timeouts on the server

Figure 69 HTTP Server Detail Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      HTTP SERVER DETAILS REPORT                      RPTSEQ     5 PAGE      7
BMC SOFTWARE, INC.                                             XYZ COMPANY                             REPORT DATE: DD MMM YY 9.43
                                                        WORLDWIDE HEADQUARTERS                         SYSTEM ID: Z0       * N/A *

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  103-1/7/0/0  103-2/29/27/4.44
-------------------------------------------------- SERVER CHARACTERISTICS SECTION --------------------------------------------------

NAME:       Z0EIP.PDL.POK.IBM.COM                SERVER ROOT IN HFS: /usr/lpp/internet/server_root
IP-ADDRESS: 9.12.17.38                           STARTUP:            06/10/YYYY-08.46.56
PORT:       80                                   SECURITY TYPE:      1
TYPE:       PROXY                                SSL-PORT:           443
APPL-LVL:   V5R3M0

--------- FLAGS ---------     ------- RESOURCES ------                        ------- CACHE -------     -- TIMEOUT THRESHOLDS --
DNS LOOKUP             NO     MAX BUFFER       102,400                        CACHE               NO     INPUT                 330
ACL SETTINGS           NO     MAX THREADS          200                        MAX SIZE         5,120     OUTPUT              3,600
META FILE             YES                                                     MAX FILES           NO     SCRIPT                600
DIRECTORY ACCESS       NO     -- GARBAGE COLLECTION --                        LIMIT 1            200     IDLE THREADS            0
SERVER IMBEDS HTML    YES     ENABLED               NO                        LIMIT 2          4,000     CACHE LOCK          1,200
NORMAL MODE           YES     INTERVAL          10,800                        TIME MARGIN        120
GMT                    NO     MEMORY USE           500                        KEEP EXPIRED       YES
PROXY                 YES                                                     CONNECT             NO

-------------------------------------------------------- SERVER ACTIVITY SECTION ---------------------------------------------------

------------- ACTIVITY -----------     - REQUEST TYPES -     ------ THREADS ------     -- CACHE USAGE --  ---- MISCELLANEOUS ---
                      COUNT    RATE     GET        20,347     MAX                200     KBYTES READ     0  TIMEOUTS              0
REQUESTS             20,353    4.21     POST            0     USED             64.67     HITS            0  CONNECTIONS       4,561
REQUESTS DISCARDED        0       0     CGI         2,268     NON-SSL WAIT     16.17     IN USE          0
REQUESTS IN ERROR         0       0     GWAPI       6,776     SSL WAIT         80.83     FILES           0
RESPONSES            20,353    4.21                           ASYNC I/O WAIT       0
RESPONSES DISCARDED       0       0                           MSG QUEUE WAIT       0

----------- THROUGHPUT -----------     ------------- RESPONSE TIMES -------------     ------------- ERROR STATISTICS ----------
                      BYTES    RATE                         MIN      MAX      AVG     200 (OK)                            20,353
IN                4,159,740  860.70     DNS LOOKUP            0        0        0     302 (MOVED TEMPORARILY)                  0
OUT               92,558.4K  19,151     SERVICE PLUGIN    0.002  230.629   39.097     401 (UNAUTHORIZED)                       0
UNKNOWN                   0       0     CGI               1.304   56.539   14.319     403 (FORBIDDEN)                          0
                                        SSL HANDSHAKE         0        0        0     404 (NOT FOUND)                          0
                                        PROXY                 0        0        0     407 (PROXY UNAUTHORIZED)                 0
                                                                                      500 (INTERNAL SERVER ERROR)              0
```

HTTP Server Detail Report field descriptions

Table 79 describes each field in the HTTP Server Detail Report.

Table 79 Field descriptions for the HTTP Server Detail Report (part 1 of 3)

Field	Description
Server Characteristics - Configuration Data	
NAME	server name
IP-ADDR	IP address of the host on which this HTTP server is executing
PORT	port number to which this HTTP server listens
TYPE	server role; the choices are as follows: HTTP—simple or normal HTTP server PROXY—proxy server CACHING—caching server CACHING PROXY—caching proxy UNKNOWN—unknown server role
APPL-LVL	HTTP server version and release levels
SERVER ROOT IN HFS	directory path and filename for servomotor
STARTUP	server startup date/time
SECURITY TYPE	security type
SSL-PORT	security (SSL) port
Server Characteristics - Flags	
DNS LOOKUP	DNS lookup flag
ACL SETTINGS	ACL settings
META FILE	meta file flag
DIRECTORY ACCESS	directory access flag
SERVER IMBEDS HTML	server embeds HTML flag
NORMAL MODE	normal mode flag
GMT	GMT flag
PROXY	proxy flag
Server Characteristics - Resources	
MAX BUFFER	maximum size of content buffer
MAX THREADS	maximum number of threads that the server can have in the thread pool
Server Characteristics - Garbage Collection	
ENABLED	indication whether garbage collection is enabled
INTERVAL	garbage collection interval in seconds
MEMORY USE	garbage collection memory usage
Server Characteristics - Cache	
CACHE	cache flag
MAX SIZE	maximum cache size (KB)

Table 79 Field descriptions for the HTTP Server Detail Report (part 2 of 3)

Field	Description
MAX FILES	maximum number of files in cache; NO indicates no defined maximum
LIMIT 1	cache limit 1
LIMIT 2	cache limit 2
TIME MARGIN	cache time margin (seconds)
KEEP EXPIRED	keep expired flag
CONNECT	cache connect flag
Server Characteristics - Timeout Thresholds (in seconds)	
INPUT	input timeout
OUTPUT	output timeout
SCRIPT	script timeout
IDLE THREADS	timeout for idle threads
CACHE LOCK	cache lock timeout
Server Activity - Requests The fields are given as COUNT and RATE (COUNT per second).	
REQUESTS	requests that were successfully processed by this server
REQUESTS DISCARDED	requests that were discarded as invalid by this server
REQUESTS IN ERROR	requests that this server responded to with an error
RESPONSES	number of responses that were successfully sent by this server
RESPONSES DISCARDED	responses that this server was unable to send back to the client
Server Activity - Request Types	
GET	number of GET requests that were received by this server
POST	number of POST requests that were received by this server
CGI	number of CGI requests that were received by this server
GWAPI	number of GWAPI requests that were received by this server
Server Activity - Threads	
MAX	maximum number of threads as specified in the HTTP server configuration file on the MaxActiveThreads directive
USED	average number of threads used; this figure is an average for the report period
NON-SSL WAITING	number of non Secure Sockets Layer (SSL) threads available for use If this value is 0, all non SSL threads are allocated.
SSL WAITING	number of Secure Sockets Layer (SSL) threads available for use If this value is 0, all SSL threads are allocated.
ASYNC I/O WAITING	If this server is running in Scalable Server mode, the number of asynchronous I/O threads available for use If this value is 0, all asynchronous I/O threads are allocated.

Table 79 Field descriptions for the HTTP Server Detail Report (part 3 of 3)

Field	Description
MSG QUEUE WAITING	if this server is running in Scalable Server mode, the number of message queue threads available for use If this value is 0, all message queue threads are allocated.
Server Activity - Cache Usage	
KBYTES READ	number of kilobytes read from the cache of this server
HITS	number of requests for files stored in the cache of this server
IN USE	average number of kilobytes (KB) of memory used by the cache for this server This figure is an average for the report period.
FILES	average number of files in the cache of this server This value is an average for the report period.
Server Activity - Miscellaneous	
TIMEOUTS	number of timeouts on the server This value is not affected by any configuration changes to this server.
CONNECTIONS	number of connections that this server has provided
Server Activity - Throughput The fields are given as BYTES and RATE (BYTES per second).	
IN	number of bytes sent to this server through requests
OUT	number of bytes sent by this server through responses
UNKNOWN	bytes that are not identified as part of a request
Server Activity - Response Times (in seconds) The values are minimum, maximum and average response times.	
DNS LOOKUP	time it takes to complete the search for a domain name in the Domain Name Server (DNS)
SERVICE PLUGINS	time it takes to complete customized application functions
CGI	time it takes to complete Common Gateway Interface (CGI) programs
SSL HANDSHAKE	time it takes to complete the exchange of security information between the HTTP server and browser
PROXY RESPONSE	if configured as a Proxy Web server, time it takes to complete a transaction between a browser, this proxy server, and the destination server
Server Activity - Error Statistics This value represents the number of responses with a specific error code.	
ERROR	Code—meaning 200—OK 302—moved temporarily 401—unauthorized 403—forbidden 404—not found 407—proxy unauthorized 500—internal server error

I/O Queuing Activity Report

The I/O Queuing Activity Report displays I/O configuration and activity data. It provides information about

- the IODF (input/output definition file)

  This is a VSAM linear data set that contains I/O definition information.

- I/O processor utilization

  One or more I/O processors are present for the channel subsystem. They manage and schedule I/O requests to devices within LCUs.

- for each Logical Control Unit or LCU, the mapping and activity of its channel paths and physical control units

  For dynamically managed channel paths, additional information is available.

Use the IOQ (see “IOQ” on page 272) Analyzer control statement to produce the I/O Queuing Activity Report.

The data for this report is obtained by using the IOQ (see “IOQ” on page 167) and DEVICE (see “DEVICE” on page 145) Extractor control statements.

There are three sections to the I/O Queuing Activity Report:

- IODF Section
- I/O Processor Utilization Section
- LCU Configuration and Activity Section

Figure 70 on page 465 shows an example of the I/O Queuing Activity Report.

Figure 70 I/O Queuing Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     I/O QUEUING ACTIVITY REPORT                         RPTSEQ     4 PAGE      6
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                              REPORT DATE: DD MMM YY 17.48
ACTL 06 JUL YY 17.00.00  06 JUL YY 17.45.00                  HOUSTON, TX.                                 SYSTEM ID: SJSC     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-1/21/536/0.75  78-3/3/897/0.75
---------------------------------------------------------- IODF SECTION ----------------------------------------------------------

IODF = 47   CREATED = 05/14/YYYY-09.45.23   ACTION = POR

----------------------------------------------------- I/O PROCESSORS SECTION -----------------------------------------------------

        -INITIATIVE QUEUE-     ------ IOP UTILIZATION -------    -- % I/O REQUESTS RETRIED --    -------- RETRIES / SSCH ---------
         ACTIVITY    AVG Q     % IOP   I/O START    INTERRUPT          CP    DP    CU    DV              CP     DP     CU     DV
IOP        RATE      LNGTH      BUSY     RATE          RATE      ALL  BUSY  BUSY  BUSY  BUSY      ALL   BUSY   BUSY   BUSY   BUSY
---     ---------  -------     -----   -------       -------    ----  ----  ----  ----  ----     -----  -----  -----  -----  -----
 00      2,606.77     0.34     36.98   2,606.6       2,719.3    86.8  83.7   2.4   0.5   0.2      6.58   6.34   0.18   0.04   0.01
 01      2,702.35     0.13     55.24   2,702.3       2,919.0    88.3  83.8   4.4     0   0.1      7.57   7.19   0.38      0   0.01
 02      1,578.16     0.01     10.54   1,578.1       1,660.8    38.5  33.5   4.4   0.1   0.5      0.63   0.55   0.07      0   0.01
 03      1,366.23     0.01     21.03   1,366.1       1,820.1    70.6  62.9   7.3     0   0.3      2.40   2.14   0.25      0   0.01
 04      2,690.40     0.06     47.24   2,690.3       2,861.1    81.3  74.5   6.5     0   0.3      4.36   3.99   0.35      0   0.01
 05      2,592.73     0.14     56.57   2,592.7       3,057.9    86.0  79.6   6.4     0   0.1      6.17   5.70   0.46      0   0.01

SUMMARY: 13,536.6     0.13     37.93    13,536        15,038    83.7  78.5   4.9   0.1   0.2      5.14   4.82   0.30   0.01   0.01

  ------------------------------------------------- LOGICAL CONTROL UNITS SECTION ------------------------------------------

                                                                             AVG    AVG                          AVG
                           --- DCM ---  CHAN  PREF  CHPID    % DP    % CU    CUB    CMR  CONTNTN  DELAY Q    CSS
 LCU     CONTROL UNITS     MIN MAX DEF  PATH  PATH  TAKEN    BUSY    BUSY    DLY    DLY    RATE   LENGTH     DLY
----  ------------------   --- --- ---   --    --  -------  -----   -----  -----  -----  -------  -------  -----
0000  0101                               70    PF    0.017      0       0      0      0
                                        ****         0.117      0       0      0      0        0        0  29.52

0001  0102                               78    NP        0      0       0      0      0
                                        ****             0      0       0      0      0        0        0      0

0002  0105                               E8    NS    0.025      0       0      0  29.63
                                        ****         0.025      0       0      0  29.63        0        0  29.74

0003  0106                               EC    PF    0.025      0       0      0  21.27
                                        ****         0.025      0       0      0  21.27        0        0  21.37

0005  01B0                               4E    NP    0.002      0   50.00      0      0
                                        ****         0.002      0   50.00      0      0    0.003        0   3.16

000A  03C2                               09    NS        0      0       0      0      0
                                        ****             0      0       0      0      0        0        0      0

000B  0570                               05    PF        0      0       0      0      0
                                        ****             0      0       0      0      0        0        0      0

000C  0590                               06    NP        0      0       0      0      0
                                        ****             0      0       0      0      0        0        0      0

0010  06C0                               B1    NS    0.008  20.00       0      0      0
                                         AF    NS    0.006  10.00       0      0
                                        ****         0.014  15.00       0      0      0    0.002        0   1.69
```

Table 80 describes each field in the I/O Queuing Activity Report.

Table 80 Field descriptions for the I/O Queuing Activity Report (part 1 of 3)

Field	Description
IODF SECTION	
IODF	two-character hexadecimal suffix for the IODF data set name
CREATED	date and time this IODF data set was created
ACTION	text indicating how this IODF data set was activated
I/O PROCESSORS SECTION	
IOP	I/O processor (IOP) numeric identifier; the IOP data section is sorted according ascending IOP numbers
Initiative Queue	for each IOP, there is one initiative queue
ACTIVITY RATE	rate, per second, at which the channel subsystem places requests on the I/O processor's initiative queue **Note**: This rate might be greater than the actual I/O rate because of requeued I/O requests.
AVG Q LNGTH	average number of entries present on the initiative queue
IOP Utilization	
% IOP BUSY	ratio of the number of times that the IOP was found busy and the total number of samples taken
I/O START RATE	rate, per second, at which I/O requests are initiated on this IOP
INTERRUPT RATE	rate, per second, at which interrupts are processed on this IOP
% I/O Requests Retried	
ALL	percentage of I/O requests that were retried for any reason; includes all initial I/O requests and all requests that were retried
CP BUSY	percentage of I/O requests that were retried because of CHANNEL busy conditions; includes all initial I/O requests and all requests that were retried
DP BUSY	percentage of I/O requests that were retried because of DIRECTOR PORT busy conditions; includes all initial I/O requests and all requests that were retried
CU BUSY	percentage of I/O requests that were retried because of CONTROL UNIT busy condition; includes all initial I/O requests and all requests that were retried
DV BUSY	percentage of I/O requests that were retried because of DEVICE busy conditions; includes all initial I/O requests and all requests that were retried

Table 80 Field descriptions for the I/O Queuing Activity Report (part 2 of 3)

Field	Description
Retries / SSCH	
ALL	ratio of I/O requests that were retried for any reason, versus the number of initial I/O requests
CP BUSY	ratio of I/O requests that were retried because of CHANNEL busy conditions, versus the number of initial I/O requests
DP BUSY	ratio of I/O requests that were retried because of DIRECTOR PORT busy conditions, versus the number of initial I/O requests
CU BUSY	ratio of I/O requests that were retried because of CONTROL UNIT busy conditions, versus the number of initial I/O requests
DV BUSY	ratio of I/O requests that were retried because of DEVICE busy conditions, versus the number of initial I/O requests
LOGICAL CONTROL UNITS SECTION	
LCU	Logical Control Unit identifier; a four-digit hexadecimal number; an LCU is a logical entity representing one or more physical control units
CONTROL UNITS	four-digit hexadecimal number of each physical control unit pertaining to this LCU; an LCU can have up to four physical control units
DCM **MIN** **MAX** **DEF**	minimum and maximum number of DCM-managed channels for this LCU for the reporting period Also, the number of installation-defined managed channels for the LCU. This data is available only for LCUs with DCM managed channel paths. Additionally, for all DCM-managed channel paths, the I/O activity rate, the director port contention, and the control unit contention are computed and shown. Note that these values can include managed channels that were only partially online.
CHAN PATH	two-digit hexadecimal number of each channel path that is assigned to this LCU An LCU can have up to four channel paths. Four asterisks (****) in this column indicate a summary line for all channel paths connected to this LCU.
PREF PATH	path attribute; displayed for subsystems that support Preferred Pathing (for example, the 1750) The path attribute values that can appear are preferred (PF), nonpreferred (NP), and not specified (NS).
CHPID TAKEN	rate, per second, of I/O requests handled by each CHPID during the reporting period If a channel path was offline or moved offline or online during the reporting period, this column could contain the following: OFFLINENOW OFFLINENOW ONLINE
% DP BUSY	percentage of I/O requests that were deferred because a director port was busy; this figure is a measure of director port contention
% CU BUSY	percentage of I/O requests that were deferred because the control unit was busy; this figure is a measure of control unit contention along each channel path
AVG CUB DLY	average number of milliseconds that the I/O request for the channel path waited because the control unit was busy

Table 80 Field descriptions for the I/O Queuing Activity Report (part 3 of 3)

Field	Description
AVG CMR DLY	average number of milliseconds for a start or resume function needed before the first command of the channel program sent to the device is marked Accepted by the device **Note**: CMR—Initial-command-response.
CONTNTN RATE	rate, per second, at which delayed I/O requests are placed on the Control Unit Header (CU-HDR) by the IOP
DELAY Q LENGTH	average number of delayed I/O requests present on the CU-HDR
AVG CSS DLY	average number of milliseconds for the start or resume function needed from the acceptance by the subchannel until the channel subsystem attempts to start the function; this figure is a measure of channel subsystem contention along each LCU
The following fields are present only if SMF- type 74-1 records are available:	
AVG SRVC TIME (MS)	average number of milliseconds it took to complete an I/O request for all devices attached to this LCU that were monitored by the device activity sampler
AVG IOSQ TIME (MS)	average number of milliseconds that I/O requests had to wait before a SSCH instruction was issued to queue the request, for all devices attached to this LCU that were monitored by the device activity sampler

Interval Bar Graph

The Interval Bar Graph (or histogram) is produced by using the GRAPH TYPE=PLOT Analyzer control statement (see "GRAPH" on page 259). The data in this graph can be used to pinpoint peaks and valleys in the performance of your system over time.

The Extractor control statements that are required for the Interval Bar Graph are dependent upon the information to be graphed. (See "Numeric list of record types" on page 47 for information about the record types generated by each Extractor control statement and a description of the specific type of data collected by each Extractor control statement.)

The Interval Bar Graph is scaled to the largest LIMIT value that is specified in the Analyzer GRAPH statement. The default limit is 100. Up to 16 measures can be specified; each measure is represented on the graph by a different character.

One line on the graph is produced for each interval that is specified in the GRAPH statement. When more than one measure is represented, the character representing the smallest measure value during a single interval is displayed first on the line. If the interval default is used, one line is produced for each record interval encountered in the data.

The Interval Bar Graph is read from left to right. The value of each measure on the graph is determined by locating the rightmost character representing that measure; the value shown at that point is the value of the measure. For measures that have exactly the same value during a single interval, the character O is issued to indicate an overlap.

An example of the Interval Bar Graph is shown in Figure 71.

Figure 71 Interval Bar Graph

```
PRODUCED BY CMF ANALYZER (v.r.mm)                          INTERVAL BAR GRAPH                                  RPTSEQ    23 PAGE     58
BMC SOFTWARE, INC.                                             XYZ COMPANY                                     REPORT DATE: DD MMM YY 9.33
ACTL 10 JUN YY 06.45.01  10 JUN YY 07.45.00              WORLDWIDE HEADQUARTERS                                SYSTEM ID: SYSB     Z v.rr.n


                * REPRESENTS PAGESEC
                + REPRESENTS PAGETIMEC
                = REPRESENTS CPU
                O REPRESENTS OVERLAP
                  CPU ENVIRONMENT REQUESTED:  ALL CPU'S

  DATE     TIME                                               GRAPHIC DISPLAY OF INTERVAL TOTAL TIME

                               0         10        20        30        40        50        60        70        80        90        100
                               +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
23MAY YY 17.56.53              O=======================================.         .         .         .         .         .         .
         18.00.00              +*======================================.         .         .         .         .         .         .
         18.30.00              +*=======================================         .         .         .         .         .         .
         18.45.00              O========================================         .         .         .         .         .         .
         19.00.00              O=======================================.         .         .         .         .         .         .
                               +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
```

Interval Bar Graph field descriptions

Table 81 describes each field in the Interval Bar Graph.

Table 81 Field descriptions for the Interval Bar Graph Report

Field	Description
c REPRESENTS xxxxxxx	character, *c*, used to represent various measures, *xxxxxxx*, selected in GRAPH statement
DATE TIME	date and time of intervals at which samplings of the measures were made
GRAPHIC DISPLAY OF INTERVAL TOTAL TIME	axis representing values for each measure, according to the far right character The values are not a continuum-type scale. They can be percentages, numerical count, and so on, depending on the measure being scaled. See Appendix C, “Measure and trace values,” for a description of each measure being reported.

Kiviat Graph

The Kiviat Graph is produced by using the GRAPH TYPE=KIVIAT Analyzer control statement (see “GRAPH” on page 259). The shape inside the circle represents the overall performance of the measures specified in the GRAPH statement.

The Extractor control statements that are required for the Kiviat Graph are dependent upon the information to be graphed. (See “Record types” on page 45 for information about the record types generated by each Extractor statement and a description of the specific type of data collected by each Extractor statement.)

A sample of the Kiviat graph is shown in Figure 72 on page 471. Each axis is labeled with the name specified on the GRAPH control statement. The actual value observed for each measure is printed as a number. A description of the value and a limit value (which corresponds to the circle for each measure) is also printed. If the graphed value exceeds the limit, the value is plotted on the circumference of the circle.

Figure 72 Kiviat Graph

```
PRODUCED BY CMF ANALYZER (v.r.mm)                            KIVIAT GRAPH                               RPTSEQ    24 PAGE     59
BMC SOFTWARE, INC.                                            XYZ COMPANY                               REPORT DATE: DD MMM YY 9.33
ACTL 10 JUN YY 06.45.01  10 JUN YY 07.45.00             WORLDWIDE HEADQUARTERS                          SYSTEM ID: SYSB     Z v.rr.n


                  CPU ENVIRONMENT REQUESTED:  ALL CPU'S
                                                           CPU
                                           VAL =      39.12LIM =     100.00
                                                   ********.********
                                               ****                 ****
                                          *****                         *****
                                     *****                                   ****
                                   **                                            **
                                 **                        .                       **
                                *                                                    **
                              **                                                       **
          SQA-AVG           **                                                           **   BTCH-MAX
          VAL =      34.49 .                                                               .  VAL =       5.00
          LIM =     100.00*                                .                                * LIM =     100.00
                         *                                                                   *
                        *        .                                                            *
                       *                                                             .         *
                      *                                                                         *
                     *                                    *O                                     *
                    *                   .               ****                  .                   *
                    *                                 ******                                      *
                   *                                *********                                      *
                   *                          .   ***********          .                           *
                  *                             O**********.*                                       *
                  *                                *************                                    *
                  *                            ******.********    .                                 *
                  *                            ***************                                       *
                 *                            **************O                                        *
                 *                            *************+O                                        *
                 *                           ******.**********     .                                 *
                 *                         . O*****************            .                        *
                  *                .         ******************                    .                *
                  *        .                 **********.*******.                           .        *
 CSA-AVG          .                          ********************                                   *
 VAL =      35.24  *                          *******************                                  .    BTCH-AVG
 LIM =     100.00  *                          *******************                                  *    VAL =       2.77
                    *                         ********************                                *     LIM =     100.00
                    *                         ******.*************.*                              *
                     *                        ***********************                            *
                      *                       ***********************                           *
                       *                       ***********************                         *
                        *                      *.*********************.                       *
                         *                     *************************                     *
                          *                    O************************                    *
                           *                                ************O                  *
                            **                                                           **
                              **            .                            .             **
                                **                                                    *
                                  **                                                **
                                                                                  **
                        STC-AVG         *.                                   .***
                        VAL =      67.87  *****                         *****    STC-MAX
                        LIM =     100.00       ****                 ****         VAL =      70.00
                                                   *****************             LIM =     100.00
```

Kiviat Graph field description

Table 82 describes the field in the Kiviat Graph.

Table 82 Field description for the Kiviat Graph

Field	Description
CPU ENVIRONMENT REQUESTED	CPU selected in GRAPH control statement

Measure values for the Kiviat Graph are described in Appendix C, “Measure and trace values.”

Link Pack Area Report

The Link Pack Area report provides information about the use of the modified, fixed, and pageable link pack areas. This report can be used as a guide for determining where a link pack area module should reside and for packing the PLPA and MLPA to reduce page faults.

The Link Pack Area Report is produced by using the LINKPACK Analyzer control statement (see “LINKPACK” on page 273). The data for this report is obtained by using the LINKMAP (see “LINKMAP” on page 169) Extractor control statement and defining GBLS=YES to the REPORT Extractor control statement (see “REPORT” on page 176).

The Link Pack Area Report is divided into four sections with an optional fifth section:

- **Performance Statistics**

  This section lists module counts for the MLPA, FLPA, and PLPA, in both the lower and extended areas. It also shows the user-supplied threshold rank for the run.

- **Link Pack Area Module Mapping**

  This section lists LPA modules sorted by page address. The action flag included in this section marks LPA pages, which contain modules with different levels of activity. Such modules are often candidates for inclusion in the packlist.

- **Portion of the Link Pack Areas Selected for Potential Substitution into the Packlist**

  This section lists all modules with a busy rank greater than the threshold rank, sorted in descending order by busy rank.

- **Time in Use**

  This section lists all modules with a busy rank greater than zero (one section each for the MLPA, FLPA, and PLPA). You can use these reports to make decisions regarding packlist inclusion and FLPA module placement.

- **Map Error Log**

  This report is generated only when errors occur. Any of the following errors might be reported:

  — **module name mismatch**: More than one entry was found for the same location, but the module names were not the same.

  — **module length mismatch:** More than one entry was found for the same location, but the module length was not the same.

  — **address mismatch:** More than one entry was found for the same location, but the entry points were not the same.

  — **lower flag mismatch:** More than one entry was found for the same location, but the characteristics of the entry were not the same. (State indicators.)

  — **extended flag mismatch:** More than one entry was found for the same location, but the extended characteristics of the entry were not the same. (Location and type indicators.)

  — **unknown:** A mismatch between at least two entries was found, but it could not be resolved into one of the previous categories.

If the LPA configuration was changed during the interval covered by the input data, two or more Link Pack Area reports are produced.

NOTE

This report requires the *xxx*-16 records written at CMF MONITOR Extractor initialization, where *xxx* is the SMF ID for the CMF MONITOR user records (the default is 240). To obtain the desired report, ensure that the input data includes the records written when extraction began. It is not necessary to include the reporting period of these records in the DATETIME or SHIFT ranges.

An example of the Link Pack Area Report is shown in Figure 73.

Figure 73 Link Pack Area Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                          LINKPACK AREA REPORT                          RPTSEQ    22 PAGE    164
BMC SOFTWARE, INC.                                              XYZ COMPANY                              REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 15.59.59               WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-11/224/24.9K/6.99  240-16/5/0/0
-----------------------------------------------------  PERFORMANCE STATISTICS  ----------------------------------------------------
 MODULE COUNT:     PLPA    FLPA      MLPA      CSA                             THRESHOLD RANK
    BELOW 16M     (339)    (0)       (1)      (24)                                         (0)
    ABOVE 16M     (873)    (2)       (6)      (36)

---------------------------------------------  LINKPACK AREA MODULE MAPPING SECTION  ---------------------------------------------
    ACTION  BUSY  MODULE  MODULE    PAGE      ACTION  BUSY  MODULE  MODULE    PAGE      ACTION  BUSY  MODULE  MODULE    PAGE
     FLAG   RANK   NAME   LENGTH   ADDRESS     FLAG   RANK   NAME   LENGTH   ADDRESS     FLAG   RANK   NAME   LENGTH   ADDRESS
               0 ISPKEY     1C68   54FE000               0 FLMXI     731C0   54D5000               0 FLMXI     731C0   54A6000
               0 IKJIFR00    C90   54FD000               0 FLMXI     731C0   54D4000               0 FLMXI     731C0   54A5000
               0 ISPISM     2358   54FD000               0 FLMXI     731C0   54D3000               0 FLMXI     731C0   54A4000
               0 ISPISM     2358   54FC000               0 FLMXI     731C0   54D2000               0 FLMXI     731C0   54A3000
               0 ISPISM     2358   54FB000               0 FLMXI     731C0   54D1000               0 FLMXI     731C0   54A2000
               0 IEAETEQ     428   54FA000               0 FLMXI     731C0   54D0000               0 FLMXI     731C0   54A1000
               0 ISPDTWIN   4BD8   54FA000               0 FLMXI     731C0   54CF000               0 FLMXI     731C0   54A0000
               0 ISPDTWIN   4BD8   54F9000               0 FLMXI     731C0   54CE000               0 FLMXI     731C0   549F000
               0 ISPDTWIN   4BD8   54F8000               0 FLMXI     731C0   54CD000               0 FLMXI     731C0   549E000
               0 ISPDTWIN   4BD8   54F7000               0 FLMXI     731C0   54CC000               0 FLMXI     731C0   549D000
               0 ISPDTWIN   4BD8   54F6000               0 FLMXI     731C0   54CB000               0 FLMXI     731C0   549C000
               0 IGWAMCS3    9D8   54F5000               0 FLMXI     731C0   54CA000               0 FLMXI     731C0   549B000

----------------------  PORTION OF LINKPACK AREAS SELECTED FOR POTENTIAL SUBSTITUTION INTO THE PACKLIST  ----------------------
                                              (MODULES WHOSE RANK EXCEED THE THRESHOLD RANK)
    MODULE    MODULE     BUSY          MODULE    MODULE     BUSY          MODULE    MODULE     BUSY          MODULE    MODULE     BUSY
     NAME     LENGTH     RANK           NAME     LENGTH     RANK           NAME     LENGTH     RANK           NAME     LENGTH     RANK
   ECNDLL     462C48       34        ACYAPCNP     7260         3        IKJEFT04     AFA8         2        IGC0002F     4120         1
   CSVEXPR      73F0       20        BPXINLPA    8E338         3        IKTLTERM     17D0         2        IGC00020     F510         1
   EUVPDLL    8CA658       18        EUVFDLL     F1638         3        IRRMNGR     26FA0         2        IGC00030     1008         1
   IEFW21SD    965D0       17        IGC0003E     B518         3        ISGLRELS     11D8         2        IGC0005E    135C8         1
   PSCLBSXP     1CF80      17        IGFDI0       A150         3        ISRPLEX      E748         2        IGC0009C     77E8         1
   ISPSUBS      DC8B0      14        IGG019BO      3A0         3        ISRPX        A0F8         2        IGG019BB      740         1
   IEFJRASP     1830       13        IGWBBMF1    641D0         3        IWMI2LPA    19750         2        IGG019FL      370         1
-----------------------------------------  MLPA MODULES BY TIME IN USE  -------------------------------------------------------
 % MLPA   MODULE    MODULE     LOAD     PAGE    FRAME  % MLPA      % MLPA   MODULE    MODULE     LOAD     PAGE    FRAME  % MLPA
ACTIVITY   NAME     LENGTH    ADDRESS  FIXED    USAGE   PAGE       ACTIVITY  NAME     LENGTH    ADDRESS  FIXED    USAGE   PAGE
                                        PCT            FAULTS                                             PCT            FAULTS
100.00   XDC31      48B60    5BEEB18    0.0      4.0    0.0
-----------------------------------------  PLPA MODULES BY TIME IN USE  -------------------------------------------------------
 % PLPA   MODULE    MODULE     LOAD     PAGE    FRAME  % PLPA      % PLPA   MODULE    MODULE     LOAD     PAGE    FRAME  % PLPA
ACTIVITY   NAME     LENGTH    ADDRESS  FIXED    USAGE   PAGE       ACTIVITY  NAME     LENGTH    ADDRESS  FIXED    USAGE   PAGE
                                        PCT            FAULTS                                             PCT            FAULTS
  8.21   ECNDLL    462C48    2B5C000    0.0      7.0   16.7          0.72   IXGINLPA   E66E0    4CDA000    0.0      3.0    1.5
  4.83   CSVEXPR     73F0    2AE7000    0.0      1.0    0.0          0.48   EAGRTLIB   42508    5AF9000    0.0      1.0    0.0
  4.35   EUVPDLL   8CA658    30CE000    0.0     15.0    8.8          0.48   GDEIGCS    28980    3D5E000    0.0      2.0    1.0
  4.11   IEFW21SD   965D0    4141000    0.0     13.0    1.5          0.48   ICEMAN     11200     C7B000    0.0      2.0    1.0
  4.11   PSCLBSXP   1CF80     C4B000    0.0      1.0    8.3          0.48   IDA0192A   6F070    3F5D000    0.0      1.0    1.0
  3.38   ISPSUBS    DC8B0    564B000    0.0     12.0    1.5          0.48   IEEMB887    1C70    4081000    0.0      1.0    1.0
  3.14   IEFJRASP    1830    410C000    0.0      1.0    0.0          0.48   IEWLDR00   1F7D8    436B000    0.0      2.0    0.0
```

If an action flag (shown as ****) is displayed in this report, it might not require action. Sometimes the system automatically assigns smaller modules to whatever space is available in the PLPA page. In such situations, it might not be possible or desirable to make changes.

In the PLPA Modules by Time in Use section, the PAGE FIXED PCT value is calculated as

$$\text{Page Fixed Pct} = \frac{\text{\# of times at least one page was fixed for a module}}{\text{\# of times the module was active}}$$

In this same section, the PERCENT PAGE FAULTS value is calculated as

$$\text{Percent Page Faults} = \frac{\text{page fault observations for the module in LPA}}{\text{page fault observations for all modules in LPA}}$$

Link Pack Area Report field descriptions

Table 83 describes each field in Link Pack Area Report.

Table 83 Field descriptions for the Link Pack Area Report (part 1 of 2)

Field	Description
MODULE COUNT	number of modules in PLPA, FLPA, MLPA, or CSA
PAGE FIXED PCT	this field is blank
THRESHOLD RANK	user-supplied value with two functions: ■ to determine the level of activity a module must reach before being included in the substitution section ■ to describe the allowable tolerance of differences in the busy ranks of modules in the same page
ACTION FLAG	four asterisks (****) are printed whenever a busy rank for a module sharing a page with another module has a difference greater than the user-supplied threshold value
BUSY RANK	number of references to a module
MODULE NAME	name of the link pack area module
MODULE LENGTH	length of the link pack area module in bytes
PAGE ADDRESS	hexadecimal address of the page that this module occupies in the link pack area
% MLPA ACTIVITY	percentage of MLPA activity that was caused by this module
LOAD ADDRESS	load address of this module
PAGE FIXED PCT	percentage of times (frequency) that this module was observed in fixed pages
FRAME USAGE	number of pages in the module where BUSY RANK was greater than zero
% MLPA PAGE FAULTS	percentage of MLPA page faults caused by this module
% PLPA ACTIVITY	percentage of PLPA activity that was caused by this module

Table 83 Field descriptions for the Link Pack Area Report (part 2 of 2)

Field	Description
% PLPA PAGE FAULTS	percentage of PLPA page faults that were caused by this module
% FLPA ACTIVITY	percentage of FLPA activity that was caused by this module
% FLPA PAGE FAULTS	percentage of FLPA page faults that were caused by this module

Logical Partition Report

The Logical Partition Report shows the activity of logical partition (LPAR) management time incurred in a Processor Resource / Systems Manager (PR / SM) environment.

This report is produced by using the PRSM Analyzer control statement. The data for this report is obtained by using the CPU Extractor control statement (see “CPU” on page 136).

Figure 74 is an example of the LPAR Report.

Figure 74 Logical Partition Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                       LOGICAL PARTITION REPORT                                 RPTSEQ     3 PAGE      5
BMC SOFTWARE, INC.                                             XYZ COMPANY                                       REPORT DATE: DD MMM YY 16.36
ACTL 10 JUN YY 09.00.00  10 JUN YY 12.45.00              WORLDWIDE HEADQUARTERS                                  SYSTEM ID: SJSC     02.09.00

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  70-1/15/6,726/3.75
CPU FAMILY: 9672  VERSION: A5  MODEL: Z57
HOME PARTITION = SJSC             NUMBER OF PHYSICAL PROCESSORS =    7       DISPATCH INTERVAL = DYNAMIC
DURATION = DAYS:   0   HOURS:  3   MINUTES: 45   SECONDS:  0

                                                                  ------- LOGICAL PROCESSORS -------     --- PHYSICAL PROCESSORS ----
                                                                        ---- PERCENT DISPATCHED ----     ---- PERCENT DISPATCHED ----
                                       WAIT      WEIGHTING                          PARTITION                        PARTITION
  NAME                  STATUS       COMPLETION    FACTOR         NUM  EFFECTIVE  OVERHEAD   TOTAL      EFFECTIVE  OVERHEAD   TOTAL
--------            --------------   ----------  ----------       ----  ---------  --------- -------     ---------  --------- -------

CF01                ACTIVE           N/A         DED                 2      99.92       0.01   99.93         28.55       0.00   28.55

SJSB                ACTIVE           NO                    20        3      19.23       0.55   19.78          8.24       0.24    8.48

SJSC                ACTIVE           NO                    15        4      10.80       0.64   11.43          6.17       0.36    6.53

SJSG                ACTIVE           NO                     3        2       3.98       0.79    4.77          1.14       0.23    1.36

SJSH                ACTIVE           NO                     8        2       4.72       0.68    5.40          1.35       0.20    1.54

SYSO                ACTIVE           NO                    15        3      41.53       0.20   41.73         17.80       0.09   17.89

VM5                 ACTIVE           NO                    10        2      20.54       1.74   22.28          5.87       0.50    6.37

          PARTITION UNATTRIBUTABLE                                                                                   2.01    2.01

          TOTAL                                                                                            69.11       3.61   72.73
```

Logical Partition Report field descriptions

Table 84 describes each field in the LPAR Report.

Table 84 Field descriptions for the Logical Partition Report (part 1 of 3)

Field	Description
CPU FAMILY VERSION MODEL	CPU family, version, and model
HOME PARTITION	home partition where CMF MONITOR executed and collected the data used as the basis for this report This home partition is highlighted throughout the report.
NUMBER OF PHYSICAL PROCESSORS	average number of physical processors available to the complex
DISPATCH INTERVAL	average time slice, which is the time that each partition accesses the physical system before another partition is allowed to access it; this value might be in milliseconds or it might be the word DYNAMIC If DYNAMIC, it means that the system determines the most effective dispatch interval based on system performance.
DURATION	total length of all measurement intervals being reported
NAME	name of the partition that is displayed on this row; the home partition is highlighted **Note**: There is only one row for each partition unless you have specified a criterion for splitting rows based on changes to the PR/SM system and a change has occurred.
SHARE %	displayed when you specify a DETAIL= operand in the PRSM control statement that causes new rows to generate based on changes to the PR/SM system A value is displayed in this column only if a change has occurred in the area and a split on this value has been specifically requested. The value displayed represents the percent of data record intervals that are contributing to the values and calculations of this row. **Note**: This value might not be precisely the time interval percentage of the duration of the change, since the smallest time unit available to record the duration of the change is the Extractor record interval. For example, if you change the weighting factor for five minutes and then change it back, and the Extractor record interval is one hour, an SMF type 70 record would indicate that a change had occurred to the weighting factor during an hour of data collection. There is no way to determine that the change lasted only five minutes.

Table 84 Field descriptions for the Logical Partition Report (part 2 of 3)

Field	Description
STATUS	two status flags are reported here: ■ whether the partition is active (ACTIVE) or deactivated (DEACT) ■ whether the partition is CAPPED Capping limits the maximum share of the system complex to the weighting factor. Even if other partitions are idle, the system does not provide more system access to a partition that has reached its weighting factor.
WAIT COMPLETION	wait completion attribute of nondedicated logical processors assigned to a partition; possible values for this field are YES — processor dispatched to a partition remains dispatched until the time slice expires NO — processor dispatched to a partition is returned to PR/SM and becomes available to other partitions as soon as the partition enters an enabled wait state MIX — partition has a mixture of nondedicated logical processors with wait completion attributes of YES and NO N/A — partition has only dedicated processors
WEIGHTING FACTOR	dispatching weight assigned to a partition; possible values for this field are *nnn* — all logical processors are not dedicated and have the same weight of *nnn* MIX — all logical processors are not dedicated, and they do not have the same weights DED — all logical processors are dedicated DNE — partition has a mixture of dedicated and nondedicated processors; nondedicated processors have equal weights DNN — partition has a mixture of dedicated and nondedicated processors; nondedicated processors do not have equal weights
LOGICAL PROCESSORS – NUM	average number of logical processors that are assigned to the partition
LOGICAL PROCESSORS – EFFECTIVE	percent of time that this partition was processing applications or waiting, if wait assist is enabled This value does not include time spent for partition management. It is affected by the number of logical processors assigned to the partition.
LOGICAL PROCESSORS – PARTITION OVERHEAD	percent of time that this partition was performing partition management This value is affected by the number of logical processors assigned to the the partition.

Table 84 Field descriptions for the Logical Partition Report (part 3 of 3)

Field	Description
LOGICAL PROCESSORS – TOTAL	sum of effective and partition overhead This value is affected by the number of logical processors assigned to the partition.
PHYSICAL PROCESSORS – EFFECTIVE	Percent of time that this partition was processing applications or waiting, if wait assist is enabled This value does not include time spent for partition management. It is affected by the number physical processors available.
PHYSICAL PROCESSORS – PARTITION OVERHEAD	percent of time that this partition was performing partition management This value is affected by the number of physical processors available.
PHYSICAL PROCESSORS – TOTAL	sum of effective and partition overhead This value is affected by the number of physical processors available.
PARTITION UNATTRIBUTABLE	percent of time spent by the system in partition management that could not be assigned to a specific partition
TOTAL	overall sum of effective, partition overhead, and effective and partition overhead values for both logical and physical processors that are running on your system

LOTUS DOMINO Server Report

The LOTUS DOMINO Server Report provides information about the activities of LOTUS DOMINO servers. The information can be used to analyze problem servers and to view performance data.

The report consists of two parts:

- **LOTUS DOMINO Server Summary Report**

  The summary contains one line for each server.

- **LOTUS DOMINO Server Details Report**

  This part consists of these sections:

  — Server Activity (obtained from SMF type 108-1 and 108-3 records)
  — Transaction Activity (obtained from SMF type 108-1 and 108-3 records)
  — Port Activity (obtained from SMF type 108-1 and 108-3 records)

Configuration data is reported together with performance data. Configuration data is not reported if at least one corresponding performance data record is not present.

NOTE

In this report, all rates are per second and all averages are based on counts divided by the number of samples taken during the report period.

Figure 75 LOTUS DOMINO Server Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  LOTUS DOMINO SERVER SUMMARY REPORT                     RPTSEQ      3 PAGE       5
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                             REPORT DATE: DD MMM YY 18.42
ACTL 10 JUN YY 23.45.00  10 JUN YY 23.45.00                   HOUSTON, TX.                                SYSTEM ID: ML96     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  108-1/192/96/24  108-3/192/96/24
-------------------------------------------------------------------------------------------------------------------------------------
                                 AVAILABLE  -- AVG. USERS -- -AVG-  TRANSACTION   ASYNC I/O RATE   --- MAIL RATE ---    - SMTP RATE -
SERVER NAME                      HHH.MM.SS  CONNECTED ACTIVE TASKS       RATE       READS  WRITES  DELIVERED      SENT   READS WRITES
-------------------------------- ---------  --------- ------ -----  -----------   ------- -------  --------- --------   ------ ------
D01MLC83/01/M/IBM                024.00.00        138      2   258         8.77     26.31   14.71       0.27     0.05  49,710   0.01
D01MLC96/01/M/IBM                024.00.00        608     13   751        38.22    110.68   66.37       1.16     0.23    0.29   0.08
```

LOTUS DOMINO Server Summary Report field descriptions

Table 85 describes each field in the LOTUS DOMINO Server Summary Report.

Table 85 Field descriptions for the LOTUS DOMINO Server Summary Report

Field	Description
SERVER NAME	server name
AVAILABLE	total time (*hhh.mm.ss*) that the server was available during the interval
AVG. USERS - CONNECTED	average number of connected users
AVG. USERS - ACTIVE	average number of active users
AVG TASKS	average number of tasks in use
TRANSACTION RATE	rate (per second) at which transactions were processed for the reporting period
ASYNC I/O RATE - READS	rate (per second) of asynchronous reads
ASYNC I/O RATE - WRITES	rate (per second) of asynchronous writes
MAIL RATE - DELIVERED	rate (per second) of Domino mail messages delivery to local users
MAIL RATE - SENT	rate (per second) of Domino mail messages that were sent to other servers
SMTP RATE - READS	rate (per second) of SMTP messages that were received from other servers
SMTP RATE - WRITES	rate (per second) of SMTP messages that were sent to other servers

Figure 76 LOTUS DOMINO Server Detail Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  LOTUS DOMINO SERVER DETAILS REPORT                   RPTSEQ     7 PAGE     15
BMC SOFTWARE, INC.                                              XYZ COMPANY                             REPORT DATE: DD MMM YY 9.43
ACTL 10 JUN YY 23.45.00  12 JUN YY 23.45.00              WORLDWIDE HEADQUARTERS                         SYSTEM ID: ML96     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  108-1/192/96/24  108-3/192/96/24
--------------------------------------------------- SERVER ACTIVITY SECTION --------------------------------------------------------
NAME: D01MLC83/01/M/IBM

--- USER ACTIVITY ---   ------ TASKS ------   ------------- MESSAGES -------------   --- ACCESS RATES ---   ---- DATABASE CACHE ----
LIMIT              0   HI WATER MARK   474   MAILBOXES                          2   AS I/O READ    26.31   STATUS                   OK
CONNECTED        138   AVG IN USE      258                    COUNT  RATE AVG.SIZE   AS I/O WRITE   14.71   MAX ENTRIES             768
ACTIVE             2   MAX UPDATES       0   MAIL DELIVERED 23,055  0.27        57   POP3 READ          0   HIGH WATER MARK         706
 WITHIN  1 MIN    18   MAX REPLICS       0   MAIL SENT       4,091  0.05        75   IMAP READ          0   AVG ENTRIES IN USE      297
 WITHIN  3 MIN    34   COUNT REPLICS 2,184   SMTP RECEIVED  4,295M 49.7K        19   HTTP READ          0   AVG INITIAL OPENS     1,330
 WITHIN  5 MIN    48                         SMTP SENT         843  0.01        13   HTTP WRITE         0   AVG REJECTIONS            0
 WITHIN 15 MIN    94                                                                                       AVG CACHE HITS          599
 WITHIN 30 MIN   136
                        - VIRTUAL THREADS -   - PHYSICAL THREADS -                   --- AVAILABILITY ---   --- NSF BUFFER POOL ----
                        HI WATER MARK   355   HI WATER MARK     19                   THRESHOLD          0   MAX         268,435,456
                        AVG IN USE      139   AVG IN USE         0                   INDEX             86   AVG IN USE  267,665,340
                                              TOTAL             42

------------------------------------------------- TRANSACTION ACTIVITY SECTION -----------------------------------------------------

MAXIMUM NUMBER OF CONCURRENT TRANSACTIONS: NO LIMIT

--------------- TOP 10 TRANSACTIONS, BY COUNT ------------------    ---------- TOP 10 TRANSACTIONS, BY TOTAL RESPONSE TIME ---------
                                     % OF   RATE  RESPONSE TIME                                          % OF   RATE  RESPONSE TIME
TYPE   NAME                    COUNT TOTAL  /SEC   AVG     TOTAL    TYPE   NAME                    COUNT TOTAL  /SEC   AVG     TOTAL
---- --------------------   ------- ----- ----- -----   --------    ---- --------------------   ------- ----- ----- -----   --------
     TOTAL                  757,565 100.0  8.77 14.02   10,621K          TOTAL                  757,565 100.0  8.77 14.02   10,621K

   1 OPEN_DB_RQST            99,708 13.16  1.15 1.750   174,529      142 START_SERVER_RQST       53,672  7.08  0.62 143.7  7,710.9K
  48 CLOSE_DB_RQST_ALT       91,795 12.12  1.06 0.947  86,910.3        8 UPDATE_NOTE_RQST_ALT    60,230  7.95  0.70 14.56   876,963
   6 OPEN_NOTE_RQST          74,326  9.81  0.86 5.290   393,187       35 WRITE_OBJECT_RQST       20,297  2.68  0.23 22.14   449,281
   8 UPDATE_NOTE_RQST_ALT    60,230  7.95  0.70 14.56   876,963        6 OPEN_NOTE_RQST          74,326  9.81  0.86 5.290   393,187
 142 START_SERVER_RQST       53,672  7.08  0.62 143.7  7,710.9K        1 OPEN_DB_RQST            99,708 13.16  1.15 1.750   174,529
  55 READ_ENTRIES_RQST       40,982  5.41  0.47 2.609   106,934      134 APPLY_FOLDER_REPL_OPS_R  7,685  1.01  0.09 22.67   174,188
 114 GET_REPLICA_MATCHES_RQS 28,791  3.80  0.33 0.002    68.740       34 READ_OBJECT_RQST        23,840  3.15  0.28 4.815   114,786
  23 DB_REPLINFO_GET_RQST    27,034  3.57  0.31 0.001    20.169       55 READ_ENTRIES_RQST       40,982  5.41  0.47 2.609   106,934
  34 READ_OBJECT_RQST        23,840  3.15  0.28 4.815   114,786       53 UPDATE_COLLECTION_RQST   8,296  1.10  0.10 10.50  87,073.0
  35 WRITE_OBJECT_RQST       20,297  2.68  0.23 22.14   449,281       48 CLOSE_DB_RQST_ALT       91,795 12.12  1.06 0.947  86,910.3

----------------------------------------------------- PORT ACTIVITY SECTION --------------------------------------------------------
MAXIMUM NUMBER OF CONCURRENT SESSIONS:  65,535
SESSION TIMEOUT:                        54,928

NAME     -- SESSIONS IN --  -- SESSIONS OUT --  ----- BYTES ----
             COUNT  RATE          COUNT  RATE       SENT  RECEIVED
TCPIP0A    58,634  0.68          5,996  0.07     3,541M    1,144M
```

LOTUS DOMINO Server Detail Report field descriptions

Table 86 describes each field in the LOTUS DOMINO Server Detail Report.

Table 86 Field descriptions for the LOTUS DOMINO Server Detail Report (part 1 of 2)

Field	Description
NAME	server name
	User Activity
LIMIT	maximum number of users that are allowed to access the server; a value of 0 means that there is no limit
CONNECTED	average number of connected users
ACTIVE	average number of active users
WITHIN n MIN	average number of connected users that have been active within the last 1, 3, 5, 15, and 30 minutes
	Tasks
HI WATER MARK	maximum number of tasks in use
AVG IN USE	average number of tasks in use
MAX UPDATES	maximum number of concurrent update tasks
MAX REPLICS	maximum number of concurrent replication tasks
COUNT REPLICS	average number of replications initiated by server
	Messages
MAILBOXES	number of mail boxes
MAIL DELIVERED	number, rate per second, and average size of Domino mail messages delivered to local users
MAIL SENT	number, rate per second, and average size of Domino mail messages sent to other servers
SMTP RECEIVED	number, rate per second, and average size of SMTP mail messages received from other servers
SMTP SENT	number, rate per second, and average size of SMTP mail messages sent to other servers
	Access Rates
AS I/O READ	rate (per second) of asynchronous I/O reads
AS I/O WRITE	rate (per second) of asynchronous I/O writes
POP3 READ	rate (per second) of POP3 reads
IMAP READ	rate (per second) of IMAP reads
HTTP READ	rate (per second) of HTTP reads
HTTP WRITE	rate (per second) of HTTP writes
	Database Cache
STATUS	status of the database cache: either OK or ? (unknown)
MAX ENTRIES	maximum number of database entries allowed in cache at any one time
HIGH WATER MARK	high water mark of database entries in cache

Table 86 Field descriptions for the LOTUS DOMINO Server Detail Report (part 2 of 2)

Field	Description
AVG ENTRIES IN USE	average number of database entries in cache
AVG INITIAL OPENS	average number of initial database opens
AVG REJECTIONS	average number of overcrowding rejections
AVG CACHE HITS	average number of database cache hits
Virtual Threads	
HI WATER MARK	maximum number of virtual thread pool threads
AVG IN USE	average number of virtual thread pool threads in use
Physical Threads	
HI WATER MARK	maximum number of physical thread pool threads
AVG IN USE	average number of physical thread pool threads in use
TOTAL	total number of physical thread pool threads used
Availability	
THRESHOLD	server availability threshold
INDEX	server availability index
NSF Buffer Pool	
MAX	maximum size (in bytes) of the NSF (Notes Storage Facility) buffer pool
AVG IN USE	average number of bytes of the NSF buffer pool in use
Transaction Activity	
MAXIMUM NUMBER OF CONCURRENT TRANS	limit for the number of concurrent transactions in a server
Top-10 List of Transaction Types -- Sorted by COUNT and by R/T TOTAL.	
TYPE	transaction type
NAME	transaction name
COUNT	number of transactions processed
% OF TOTAL	percentage of all transactions
RATE/SEC	rate (per second) of processed transactions
RESPONSE TIME - AVG	average response time (in seconds) of all completed transactions
RESPONSE TIME - TOTAL	total response time (in seconds) of all completed transactions
Port Activity	
MAXIMUM NUMBER OF CONCURRENT SESSIONS	maximum number of sessions that can run concurrently on the server
SESSION TIMEOUT	time limit (minutes) after which idle connections are terminated
NAME	port name
SESSIONS IN	count and rate (per second) of incoming sessions established during the interval
SESSIONS OUT	count and rate (per second) of outgoing sessions established during the interval
BYTES - SENT	number of bytes sent to the network
BYTES - RECEIVED	number of bytes received from the network

LOTUS DOMINO Database Activity Report

The LOTUS DOMINO Database Activity Report displays information, such as the number of documents created / deleted and database operations, for each database of the LOTUS DOMINO application server.

Use the DOMINO (see “DOMINO” on page 246) Analyzer control statement (specifying either TYPE=DATABASE or TYPE=ALL) to produce the LOTUS DOMINO Database Activity Report.

The data for this report is obtained from the SMF type 108-6 records that are created by the LOTUS DOMINO server.

Figure 77 shows an example of the LOTUS DOMINO Database Activity Report.

Figure 77 LOTUS DOMINO Database Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)              LOTUS DOMINO DATABASE ACTIVITY REPORT                  RPTSEQ     7 PAGE     89
BMC SOFTWARE, INC.                                        BMC SOFTWARE, INC.                           REPORT DATE: DD MMM YY 20.21
ACTL 31 MAY YY 23.45.00  01 JUN YY 23.45.00                  HOUSTON, TX.                              SYSTEM ID: ML96     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  108-1/192/96/24  108-2/205/96/24  108-3/192/96/24  108-6/198/96/24
-------------------------------------------------------------------------------------------------------------------------------------
```

DATABASE NAME	NO. OF OPERATIONS	NO. OF REPLICATIONS	-- DOCUMENTS --- CREATED	DELETED	TOTAL DB ACTIVITY
/d01mlc96/mail155/hilligra.nsf	0	0	121	0	121
/d01mlc96/mail39/richking.nsf	0	0	120	0	120
/d01mlc83/mail19/rekowal.nsf	0	0	119	0	119
/d01mlc96/mail36/markan.nsf	0	0	117	0	117
/d01mlc96/mail83/dnewlan.nsf	0	0	114	3	117
/d01mlc83/mail4/tbradley.nsf	0	45	71	0	116
/d01mlc83/mail46/mhoffsta.nsf	0	47	69	0	116
/d01mlc83/mail65/krtilley.nsf	0	0	116	0	116
/d01mlc96/mail81/croushor.nsf	0	0	116	0	116
/d01mlc83/mail15/drhoderi.nsf	0	0	115	0	115
/d01mlc83/mail56/fraserf.nsf	0	0	115	0	115
/d01mlc96/mail106/sadana.nsf	0	0	115	0	115
/d01mlc96/mail115/cusimano.nsf	0	0	115	0	115
/d01mlc96/mail27/rolfschm.nsf	0	0	115	0	115
/d01mlc96/mail72/elainew.nsf	0	0	115	0	115
/d01mlc96/mail85/ambrusod.nsf	2	0	113	0	115
/d01mlc83/mail60/rayricci.nsf	0	41	73	0	114
/d01mlc96/mail6/luannes.nsf	0	0	114	0	114
/d01mlc96/mail88/tmacey.nsf	0	0	114	0	114
/d01mlc83/mail34/rockyt.nsf	0	0	113	0	113
/d01mlc83/mail61/beattie.nsf	0	43	70	0	113
/d01mlc96/mail18/voltman.nsf	0	0	113	0	113
/d01mlc96/mail20/racfscr.nsf	0	0	88	25	113
/d01mlc83/mail22/cutler.nsf	0	0	112	0	112
/d01mlc83/mail37/jtison.nsf	0	0	112	0	112
/d01mlc96/mail55/efernand.nsf	0	0	112	0	112
/d01mlc96/mail67/muriels.nsf	0	0	112	0	112
/d01mlc96/mail70/lorettat.nsf	0	0	111	1	112

LOTUS DOMINO Database Activity Report field descriptions

Table 87 describes each field in the LOTUS DOMINO Database Activity Report.

Table 87 Field descriptions for the LOTUS DOMINO Database Activity Report

Field	Description
DATABASE NAME	name of the LOTUS DOMINO server
NO. OF OPERATIONS	number of indexing operations performed during the reporting period
NO. OF REPLICATIONS	number of database replications performed during the reporting period
DOCUMENTS CREATED	number of LOTUS DOMINO documents created in the database
DOCUMENTS DELETED	number of LOTUS DOMINO documents deleted from the database
TOTAL DB ACTIVITY	sum of the four preceding columns - total activity versus the database

LOTUS DOMINO User Activity Report

LOTUS DOMINO User Activity Report displays information, such as the amount of CPU time and bytes read and written for each user of the LOTUS DOMINO application server.

Use the DOMINO (see “DOMINO” on page 246) Analyzer control statement (specifying either TYPE=USER or TYPE=ALL) to produce the LOTUS DOMINO User Activity Report.

The data for this report is obtained from the SMF type 108-2 records that are created by the LOTUS DOMINO server.

Figure 78 on page 487 shows an example of the LOTUS DOMINO Database Activity Report.

Figure 78 LOTUS DOMINO User Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  LOTUS DOMINO USER ACTIVITY REPORT                    RPTSEQ     6 PAGE     11
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                            REPORT DATE: DD MMM YY 20.21
ACTL 31 MAY YY 23.45.00  01 JUN YY 23.45.00                   HOUSTON, TX.                               SYSTEM ID: ML96     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  108-1/192/96/24  108-2/205/96/24  108-3/192/96/24  108-6/198/96/24
-----------------------------------------------------------------------------------------------------------------------------------
                                  CONN                                                           - CPU -    -- NO. OF BYTES -
SERVER NAME                       TYPE   -- IP ADDRESS --   --------- DOMINO USER NAME ---------   (MS)       READ   WRITTEN

D01MLC96/01/M/BMC                 NRPC   9.117.165.174      CN=Abba King/OU=Southbury/O=BMC       575,177    1,311K   11,347K
D01MLC96/01/M/BMC                 SMTP   9.99.140.22                                              570,054   41,366K    1,120K
D01MLC83/01/M/BMC                 NRPC   9.242.197.238      CN=Rosemarie Okie/OU=Somers/O=BMC     457,044   454,364    149.1M
D01MLC96/01/M/BMC                 NRPC   9.117.73.244       CN=Lynne Swamp/OU=Houston/O=I         403,418    1,775K   43,927K
D01MLC96/01/M/BMC                 NRPC   9.45.56.141        CN=AUTO TEST1/OU=Houston/O=BMC        267,868   219,600    7,306K
D01MLC96/01/M/BMC                 NRPC   9.117.127.125      CN=D01HUB02/OU=01/OU=H/O=BMC          245,382    172.5M   74,445K
D01MLC83/01/M/BMC                 SMTP   9.37.3.210                                               237,674   15,774K   564,590
D01MLC83/01/M/BMC                 NRPC   9.117.127.108      CN=D01HUB01/OU=01/OU=H/O=BMC          234,339      370M    8,099K
D01MLC83/01/M/BMC                 SMTP   9.99.140.24                                              216,840   11,414K   572,656
D01MLC83/01/M/BMC                 SMTP   9.37.3.208                                               177,556   22,672K   319,250
D01MLC83/01/M/BMC                 SMTP   9.99.140.22                                              136,483   13,366K   244,858
D01MLC96/01/M/BMC                 NRPC   9.117.127.108      CN=D01HUB01/OU=01/OU=H/O=BMC          136,002    424.2M    1,873K
D01MLC83/01/M/BMC                 SMTP   9.117.200.23                                             132,549    7,441K   226,140
D01MLC96/01/M/BMC                 NRPC   9.14.6.41          CN=Robbie Williams/OU=Houston/        124,549    6,655K      111M
D01MLC96/01/M/BMC                 NRPC   9.117.182.201      CN=Krispy Kreme/OU=Houston/O=I        113,657   429,098   66,914K
D01MLC96/01/M/BMC                 NRPC   9.14.6.44          CN=Kim Lee/OU=Sterling Forest/O=BMC   108,708    5,970K   85,570K
D01MLC96/01/M/BMC                 NRPC   9.14.6.45          CN=Heather Hendrson/OU=Endicott/O=IB   95,690    5,078K   65,867K
D01MLC96/01/M/BMC                 NRPC   9.117.76.149       CN=Jim Jones/OU=Houston/O=BMC          68,509   623,342   15,363K
D01MLC96/01/M/BMC                 NRPC   9.117.139.181      CN=George Dillman/OU=Houston/O=        62,743   820,998   36,619K
D01MLC83/01/M/BMC                 SMTP   9.117.200.21                                              59,212    1,582K   192,485
D01MLC96/01/M/BMC                 NRPC   9.38.93.53         CN=Marlene Fine/OU=Rochester/OU=Cont   54,308   538,544   23,802K
D01MLC96/01/M/BMC                 NRPC   9.117.109.220      CN=Joseph Jokers/OU=Houston/O=I        53,708   471,056   33,467K
D01MLC96/01/M/BMC                NRPC   9.14.6.43          CN=John Johnson/OU=Southbury/O=BMC     52,926   5,830K   46,427K
```

LOTUS DOMINO Database Activity Report field descriptions

Table 88 describes each field in the LOTUS DOMINO User Activity Report.

Table 88 Field descriptions for the LOTUS DOMINO User Activity Report

Field	Description
SERVER NAME	name of the LOTUS DOMINO server
CONN TYPE	type of connection
IP ADDRESS	IP address
DOMINO USER NAME	user name
CPU (MS)	number of milliseconds or CPU time consumed by the user
NO. OF BYTES READ	number of bytes read over the connection by the user
NO. OF BYTES WRITTEN	number of bytes written over the connection by the user

OMVS Kernel Activity Report

The OMVS Kernel Activity report displays information about resource consumption by OMVS address spaces.

The OMVS Kernel Activity report is produced by using the OMVS Analyzer control statement (see "OMVS" on page 276).

Figure 79 is an example of the OMVS Kernel Activity Report.

Figure 79 OMVS Kernel Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    OMVS KERNEL ACTIVITY REPORT                    RPTSEQ    23 PAGE    271
BMC SOFTWARE, INC.                                           XYZ COMPANY                            REPORT DATE: DD MMM YY 13.40
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                       SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-3/28/25.1K/7
---------------------------------------------------  OMVS SYSTEM CALL ACTIVITY  ---------------------------------------------------
                                             SYSCALLS                    CPU TIME (HUNDREDTHS)
                                            PER SECOND                        PER SECOND
                                            ----------                        ----------
                      MINIMUM -               1.000                             0.000
                      AVERAGE -               7.273                             0.009
                      MAXIMUM -             115.000                             1.000
-----------------------------------------------------  OMVS PROCESS ACTIVITY  -----------------------------------------------------
             MAX ALLOWABLE PROCESSES:     200         MAX ALLOWABLE USERS:         200         MAX PROCESSES PER USER:      100
             ---------- PROCESSES -----------         ------------ USERS ------------         ------- PROCESSES / USER -------
              CURRENT #         OVERRUNS/SEC           CURRENT #         OVERRUNS/SEC           CURRENT #         OVERRUNS/SEC
  MINIMUM -         14              0.000                      0             0.000                                    0.000
  AVERAGE -       17.0              0.000                    0.0             0.000                    0.0             0.000
  MAXIMUM -         19              0.000                      0             0.000                                    0.000
------------------------------------------------  OMVS INTER-PROCESS COMMUNICATION  ------------------------------------------------
             MAX MESSAGE QUEUE IDS:   500  MAX SEMPHORE IDS:        500  MAX SHRD MEMORY IDS:     500  MAX SHRD MEMORY PAGES:  262K
             ----- MESSAGE QUEUE IDS ----  ------- SEMAPHORE IDS ------  ----- SHARED MEMORY IDS ----  ---- SHARED MEMORY PAGES ---
              CURRENT #     OVERRUNS/SEC   CURRENT #      OVERRUNS/SEC   CURRENT #      OVERRUNS/SEC   CURRENT #      OVERRUNS/SEC
  MINIMUM -        0.0          0.000            0.0          0.000            0.0          0.000            0.0          0.000
  AVERAGE -        0.0          0.000            0.0          0.000            0.0          0.000            0.0          0.000
  MAXIMUM -        0.0          0.000            0.0          0.000            0.0          0.000            0.0          0.000
--------------------------------------------------------  OMVS MEMORY MAP  --------------------------------------------------------
             MAX MEM MAP STORAGE PAGES: 40960         MAX SHARED STORAGE PAGES:    131K
             --- MEMORY MAP STORAGE PAGES ---         ----- SHARED STORAGE PAGES -----
              CURRENT #         OVERRUNS/SEC           CURRENT #         OVERRUNS/SEC
  MINIMUM -        0.0              0.000                    0.0             0.000
  AVERAGE -        0.0              0.000                    0.0             0.000
  MAXIMUM -        0.0              0.000                    0.0             0.000
```

OMVS Kernel Activity Report field descriptions

Table 89 describes each field in the OMVS Kernel Activity Report.

Table 89 Field descriptions for the OMVS Kernel Activity Report (part 1 of 2)

Field	Description
SYSCALLS PER SECOND – MINIMUM	minimum number of OMVS system calls invoked per second during any single cycle of the interval
SYSCALLS PER SECOND – AVERAGE	average number of OMVS system calls invoked per second during the interval
SYSCALLS PER SECOND – MAXIMUM	maximum number of OMVS system calls invoked per second during any single cycle of the interval
CPU TIME PER SECOND – MINIMUM	minimum CPU time per second, reported in hundredths, spent processing system calls or Syscalls in the OMVS kernel address space during any single cycle during the interval
CPU TIME PER SECOND – AVERAGE	average CPU time per second, reported in hundredths, spent processing system calls or Syscalls in the OMVS kernel address space during the interval
CPU TIME PER SECOND – MAXIMUM	maximum CPU time per second, reported in hundredths, spent processing system calls or Syscalls in the OMVS kernel address space during any single cycle during the interval
MAX ALLOWABLE PROCESSES	maximum number of OMVS processes allowed, defined as a constant in SYS1.PARMLIB member BPXPRM*xx*
MAX ALLOWABLE USERS	maximum number of OMVS users allowed, defined as a constant in SYS1.PARMLIB member BPXPRM*xx*
MAX PROCESSES PER USER	maximum number of OMVS processes per user allowed, defined as a constant in SYS1.PARMLIB member BPXPRM*xx*
PROCESSES: NUMBER OF PROCESSES – MINIMUM	minimum number of OMVS processes during any single cycle of the interval
PROCESSES: NUMBER OF PROCESSES – AVERAGE	average number of OMVS processes during the interval
PROCESSES: NUMBER OF PROCESSES – MAXIMUM	maximum number of OMVS processes during any single cycle of the interval
PROCESSES: OVERRUNS PER SECOND – MINIMUM	minimum number of times Fork / Dub failed because the maximum number of processes was exceeded during any single cycle of the interval
PROCESSES: OVERRUNS PER SECOND – AVERAGE	average number of times Fork / Dub failed because the maximum number of processes was exceeded during any single cycle of the interval
PROCESSES: OVERRUNS PER SECOND – MAXIMUM	maximum number of times Fork / Dub failed because the maximum number of processes was exceeded during any single cycle of the interval
USERS: NUMBER OF USERS – MINIMUM	minimum number of OMVS users during any single cycle of the interval
USERS: NUMBER OF USERS – AVERAGE	average number of OMVS users during the interval
USERS: NUMBER OF USERS – MAXIMUM	maximum number of OMVS users during any single cycle of the interval

Table 89 Field descriptions for the OMVS Kernel Activity Report (part 2 of 2)

Field	Description
USERS: OVERRUNS PER SECOND – MINIMUM	minimum number of times Fork / Dub failed because the maximum number of users was exceeded during any single cycle of the interval
USERS: OVERRUNS PER SECOND – AVERAGE	average number of times Fork / Dub failed because the maximum number of users was exceeded during any single cycle of the interval
USERS: OVERRUNS PER SECOND – MAXIMUM	maximum number of times Fork / Dub failed because the maximum number of users was exceeded during any single cycle of the interval
PROCESSES/USER: NUMBER OF PROC/USER – MINIMUM	minimum number of OMVS processes per user during any single cycle of the interval
PROCESSES/USER: NUMBER OF PROC/USER – AVERAGE	average number of OMVS processes per user during the interval
PROCESSES/USER: NUMBER OF PROC/USER – MAXIMUM	maximum number of OMVS processes per user during any single cycle of the interval
PROCESSES/USER: OVERRUNS PER SECOND – MINIMUM	minimum number of times Fork / Dub failed because the maximum number of processes per user was exceeded during any single cycle of the interval
PROCESSES/USER: OVERRUNS PER SECOND – AVERAGE	average number of times Fork / Dub failed because the maximum number of processes per user was exceeded during the interval
PROCESSES/USER: OVERRUNS PER SECOND – MAXIMUM	maximum number of times Fork / Dub failed because the maximum number of processes per user was exceeded during any single cycle of the interval

Performance Summary Report

The Performance Summary Report summarizes key measures in the system so that potential trouble spots can be located.

The Performance Summary Report is produced by using the PERFSUM (see "PERFSUM" on page 279) and PERFORM (see "PERFORM" on page 277) Analyzer control statements. The data for this report is obtained by using the following Extractor control statements:

- "ASMDATA" on page 127
- "CHANNEL" on page 134
- "CPU" on page 136
- "DEVICE" on page 145
- "ENQUEUE" on page 152
- "EXTSUM" on page 154
- "IOQ" on page 167
- "PAGING" on page 173
- "TSODATA" on page 194
- "WORKLOAD" on page 203

This report (see Figure 80 on page 493) is divided into the following sections:

- **Bottleneck Detection**

  This section lists and describes bottlenecks detected in key system areas. The types of bottlenecks can be defined by you in the PERFSUM control statement (see "PERFSUM" on page 279). If you do not define bottleneck types in the PERFSUM control statement, defaults are used.

  Three dashes (---) in the DETECTED column indicate that the data necessary to determine if this bottleneck existed is not available. This situation is most likely due to missing record types on which this determination is based.

- **Performance Measures Summary**

  This section lists busy times and queue depths for key areas of the system.

- **Job Class Activity**

  This section lists active and completed jobs for the extracted time frame.

- **Policy Activations**

  This section lists the service policies that were active during the time frame for this report. The list includes the policy name, its description, and the date and time it was activated.

- **Workload Activity**

  This section lists service classes and report classes, displaying their performance index, response time, transactions, and service units.

NOTE

The structure of the summary period and workload description is defined according to the following order of precedence:

1. the PERFORM Analyzer control statement (see “PERFORM” on page 277)

2. the presence of SMF type *xxx*-07 user records in the Analyzer output

   This record is written at CMF MONITOR initialization time when the PERFORM= operand is coded on the EXTSUM Extractor control statement.

3. the summary periods are reported as ALL, and the workload description field is either blank or contains the name of the job class, if applicable

The Performance Summary Report is shown in Figure 80.

Figure 80 Performance Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      PERFORMANCE SUMMARY REPORT                          RPTSEQ     5 PAGE     15
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                              REPORT DATE: DD MMM YY 13.19
ACTL 28 JAN YY 10.00.00  28 JAN YY 10.30.00                   HOUSTON, TX.                                 SYSTEM ID: SJSD     Z v.rr.n

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/2/878/0.5  240-3/2/352/0.5  240-5/18/876/0.5  240-6/2/1,754/0.5
240-20/2/1,756/0.5  70-1/2/878/0.5  71-1/2/352/0.5  72-3/98/0/0.5  73-1/2/0/0.5  75-1/22/878/0.5  77-1/2/0/0.5  78-2/2/350/0.5
78-3/2/594/0.5  74-1/18/876/0.5
-------------------------------------------------------  BOTTLENECK DETECTION  ------------------------------------------------------
 DETECTED     BOTTLENECK DESCRIPTION
   YES        ENQUEUE CONFLICT                      --  (CUMULATIVE DELAY GT  10 PERCENTAGE OF TOTAL TIME)
   NO         CHANNEL PATH OVERLOAD                 --  (BUSIEST CHANNEL GT  50 PERCENT UTILIZATION)
   NO         CPU OVER/UNDERLOADING                 --  (AVERAGE CPU UTILIZATION GT  95 PERCENT OR CPU UTILIZATION LT  20 PERCENT)
   NO         PAGING OVERLOAD                       --  (AVERAGE PAGING RATE GT     50 PAGES PER SECOND)
   NO         TSO RESPONSE PROBLEM                  --  (AVERAGE RESPONSE GT  15 SECONDS)
   YES        SUPERVISOR OVERUTILIZATION            --  (SUPERVISOR STATE MODE GT  50 PERCENT OF CPU BUSY)
   NO         AUXILIARY STORAGE OVERUTILIZATION     --  (SLOTS IN USE GT  90 PERCENT OF SLOTS AVAILABLE)

---------------------------------------------  PERFORMANCE MEASURES SUMMARY SECTION  ---------------------------------------------
 CPU UTILIZATION % BY PROCESSOR TYPE                                AVERAGE NUMBER OF TSO USERS                              =   23.9
                         CP      ZAAP     ZIIP                      PAGING RATE (PAGES PER SECOND)                           =     .1
   LPAR                 44.1      0.2     25.4                      SWAP RATE (SWAPS PER MINUTE)                             =     .0
   MVS                  48.1      0.1     27.8                      AVERAGE PRIVATE AREA WORKING SET (PAGES)                 =     .2
 AVERAGE CPU QUEUE DEPTH                             =       3.1    AVERAGE JES SPOOL SPACE USED                             =   41.4
 TIME CPU QUEUE EXISTED (PERCENT OF SAMPLES)         =       9.1    3590 MOUNTS     =     3  MAX CONCURRENTLY ALLOCATED      =      1
 TOTAL I/O INTERRUPT RATE FOR ALL CPU'S              =     742.7
 SSCH INSTRUCTIONS PER SECOND                        =     704.7
 AVERAGE TSO TRANSACTION RATE (TRANS PER MINUTE)     =      55.9

--------------------------------------------------------  JOB CLASS ACTIVITY  -----------------------------------------------------
                                 JOB      JOB CLASS            ACTIVE  JOBS               COMPLETED
                                CLASS    DESCRIPTION        AVG     END     %CPU           JOBS    STEPS
                                OTHER    ALL_OTHERS         2.3       2    38.8             2        6

                                BATCH    ALL_BATCH          2.3       2    38.8             2        6
                                TSO      ALL_TSO           17.9      17     1.0             4        4
                                STC      ALL_STC           49.8      67    60.2             2        2
                                TOTAL                      70.0      86   100.0             8       12

------------------------------------------------------- POLICY ACTIVATIONS -------------------------------------------------------
 POLICY:BBPLEX01 BMC Software Service Definition   ACTIVATED 25 JAN YYYY AT 10:21
------------------------------------------------------- WORKLOAD ACTIVITY --------------------------------------------------------
SERVICE                                      PERF   RESPONSE TIME      TRANSACTIONS    EXEC VEL    SERVICE UNITS      SERVICE
 CLASS    DESCRIPTION                  PER   INDX  ACTUAL   GOAL       /SEC  TOTAL    ACTL GOAL  PER SEC    TOTAL      CLASS
BATCH                                   1    0.68   8.71 M             0.00      1     15%  10%     65     116,743   BATCH
                                        2    0.01  17.30 M             0.00      1     72%   1%   3372   6,070,415
CICS                                    1    0.01   0.01 S  0.50 S     7.46  13106                                   CICS
STC                                     1          89.22 M             0.00      5    11%         7430      13,375K  STC
                                        2                                              5%           14      25,381
TSO                                     1    6.27   3.86 S             1.04   1822     14%  90%     79     142,662   TSO
                                        2    3.56   1.48 S             0.01     19     22%  80%     27      49,435
                                        3    2.52   9.76 S             0.01      9     28%  70%     15      26,366
                                        4    7.13                                       3%  20%     10      17,824
BATNRM                                  1    0.68   8.71 M             0.00      1     15%  10%     65     116,743   BATNRM
                                        2    0.01  17.30 M             0.00      1     72%   1%   3372   6,070,415
CICSHOT                                 1    0.02   0.01 S  0.50 S     2.71   4770                                   CICSHOT
CICSNRM                                 1    0.01   0.00 S  0.50 S     4.74   8336                                   CICSNRM
GRS                                     1   59.53                                       2%  99%     44      79,047   GRS
STCNRM                                  1   19.59  85.63 M             0.00      4      3%  60%   1226   2,206,047   STCNRM
STCPAS                                  1    6.33                                       6%  40%     83     150,243   STCPAS
STCPROD                                 1    0.40                                     100%  40%      0         335   STCPROD
STCSLOW                                 2    0.18                                       5%   1%     14      25,381   STCSLOW
SYSSTC                                  1           1.73 H             0.00      1     29%        5211   9,380,467   SYSSTC
SYSTEM                                  1                                              13%         866   1,558,699   SYSTEM
TSONRM                                  1    6.27   3.86 S             1.04   1822     14%  90%     79     142,662   TSONRM
                                        2    3.56   1.48 S             0.01     19     22%  80%     27      49,435
                                        3    2.52   9.76 S             0.01      9     28%  70%     15      26,366
                                        4    7.13                                       3%  20%     10      17,824
```

Bottleneck Detection Section field descriptions

Table 90 describes each field in the Bottleneck Detection section of the Performance Summary Report.

Table 90 Field descriptions for the Bottleneck Detection section

Field	Description
DETECTED	indication of whether the listed bottleneck existed during the measurement interval
BOTTLENECK DESCRIPTION	general description of the bottleneck (as defined in the PERFSUM control statement), followed by data on the severity of the bottleneck
ENQUEUE CONFLICT	average percent of time that an enqueue conflict existed
CHANNEL PATH OVERLOAD	average use of the busiest channel path
CPU OVER/UNDERLOADING	average CPU busy time
PAGING OVERLOAD	average pages per second
TSO RESPONSE PROBLEM	average TSO response time This field measures commands typed at the READY prompt and TSO-in-batch jobs. It ignores commands issued within ISPF (except EDIT and TEST).
SUPERVISOR OVERUTILIZATION	average amount of CPU busy time that was spent in supervisor state
AUXILIARY STORAGE OVERUTILIZATION	average ASM slots in use

The total reporting period is the time during which the bottleneck existed. This period is defined by the DATETIME, CYCLE, and PERIOD statements, or the actual input data.

Performance Measures Summary field descriptions

Table 91 describes each field in the Performance Measures Summary Section of the Performance Summary Report.

Table 91 Field descriptions for the Performance Measures Summary Section (part 1 of 2)

Field	Description
CPU UTILIZATION % BY PROCESSOR TYPE	LPAR CPU busy percentage of standard CPs, zAAPs, and zIIPs; the formula is percentage of busy time / online time MVS CPU busy percentage of standard CPs, zAAPs, and zIIPs; the formula is **LPAR mode:** (online time - wait time) /online time **Basic mode or under VM:** (interval - wait time) / interval
AVERAGE CPU QUEUE DEPTH	average number of dispatchable address spaces waiting for CPU
TIME CPU QUEUE EXISTED (PERCENT OF SAMPLES)	percentage of time that address spaces were found dispatchable and waiting for CPU
TOTAL I/O INTERRUPT RATE FOR ALL CPU'S	rate per second at which I/O interrupts were handled
SSCH INSTRUCTIONS PER SECOND	average number of SSCH instructions per second This value is based on the SMF type 74 device records and is limited by the devices sampled.
AVERAGE TSO TRANSACTION RATE (TRANS PER MINUTE)	average number of TSO transactions per minute This field measures commands typed at the READY prompt and TSO-in-batch jobs. It ignores commands issued within ISPF (except EDIT and TEST).
AVERAGE NUMBER OF TSO USERS	average number of TSO users logged on simultaneously
PAGING RATE (PAGES PER SECOND)	total paging rate, in pages per second, between auxiliary storage and central storage, including VIO, swap, hiperspace, and block paging
SWAP RATE (SWAPS PER MINUTE)	average number of swaps to auxiliary storage per minute
AVERAGE PRIVATE AREA WORKING SET (PAGES)	average working set size for storage obtained in private area in pages
AVERAGE JES SPOOL SPACE USED	average percentage of JES spool space used **Note**: If the Extractor EXTSUM statement does not have the parameter JES=YES specified, --- (dashes) appear in this field.
MOUNTS, MAX CONCURRENTLY ALLOCATED	for each type of tape drives observed (3490, 3590, ...), the number of mounts that occurred and the maximum number of tape drives allocated simultaneously
COUNT OF REEL VOLUMES MOUNTED	total number of reel mounts that occurred, including all mounts and any remounts of the same tape volume

Table 91 Field descriptions for the Performance Measures Summary Section (part 2 of 2)

Field	Description
COUNT OF CARTRIDGE VOLUMES MOUNTED	total number of cartridge mounts that occurred, including all mounts and any remounts of the same tape volume
MAXIMUM 3420 DRIVES CONCURRENTLY ALLOCATED	maximum number of 3420-type tape drives allocated simultaneously
MAXIMUM 3480 DRIVES CONCURRENTLY ALLOCATED	maximum number of 3480-type tape drives allocated simultaneously
MAXIMUM 3490 DRIVES CONCURRENTLY ALLOCATED	maximum number of 3490-type tape drives allocated simultaneously

Job Class Activity field descriptions

Table 92 describes each field in the Job Class Activity section of the Performance Summary Report.

Table 92 Field descriptions for the Job Class Activity section

Field	Description	
JOB CLASS	job class as specified on the Extractor EXTSUM statement	
JOB CLASS DESCRIPTION	job class description as specified on the Extractor EXTSUM statement	
ACTIVE JOBS	**AVG**	average number of jobs active in a given class during the Extractor interval
	END	number of jobs active in a given job class at the end of the Extractor interval
	% CPU	percentage of busy time percent used by a given job class during the Extractor interval
COMPLETED	**JOBS**	total number of jobs completed for a given job class during the Extractor interval
	STEPS	total number of job steps completed for a given job class during the Extractor interval

Workload Activity field descriptions

Table 93 describes each field in the Workload Activity section of the Performance Summary Report.

Table 93 Field descriptions for the Workload Activity section

Field	Description	
POLICY	service class policy name and description For more information about service policies, see Appendix B, "Workload measurement."	
SERVICE CLASS	name of the service class For more information about service classes, see Appendix B, "Workload measurement."	
DESCRIPTION	description of the service class	
PER	period number of the service class	
PERF INDX	performance index For more information about the interpretation of a performance index value, see Appendix B, "Workload measurement." This value is not shown if the goal is discretionary or system.	
RESPONSE TIME	applicable only to average response time goal and response time with percentile goal	
	ACTUAL	actual response time
	GOAL	expected response time
TRANSACTIONS	**/SEC**	average number of transactions completed per second
	TOTAL	total number of transactions completed in the report interval
EXEC VEL	**ACTL**	actual execution velocity
	GOAL	for velocity goal, the expected execution velocity
SERVICE UNITS	**PER SEC**	rate at which service units were used per second
	TOTAL	number of service units consumed in the report interval

Pie Chart

The pie (or multiplot) chart is produced by using the GRAPH TYPE=PIE Analyzer control statement. The measures specified in the GRAPH statement for this chart should be mutually exclusive (for example, CPU busy time and CPU wait time) to show an accurate representation of activity. The pie chart is most useful when measures that add up to 100% are specified (for example, PPB and SUP); however, any measures that are valid for the GRAPH statement can be charted. Up to 16 measures can be specified.

The Extractor statements required for the pie chart are dependent upon the information to be charted. (See “Record types” on page 45 for information about the record types generated by each Extractor statement and a description of the specific type of data collected by each Extractor statement.)

Each axis on the chart is labeled with a letter. The letter corresponds to a measure specified in the control statement. A key that shows the correspondence of letters to measures is printed at the upper right of the report page, outside the chart. The value listed for each letter represents a percentage of the sum of all items in that lettered category.

The charted value occupies a segment of the circle. The segment is read clockwise from the axis that describes the measure.

A sample of the Pie Graph is shown in Figure 81; explanations of the measures follow.

Figure 81 Pie Graph

```
PRODUCED BY CMF ANALYZER (v.r.mm)                          PIE GRAPH                         RPTSEQ    25 PAGE     60
BMC SOFTWARE, INC.                                        XYZ COMPANY                        REPORT DATE: DD MMM YY 9.33
ACTL 10 JUN YY 06.45.01  10 JUN YY 07.45.00         WORLDWIDE HEADQUARTERS                   SYSTEM ID: SYSB     Z v.rr.n


                  CPU ENVIRONMENT REQUESTED:  CPU 0
```

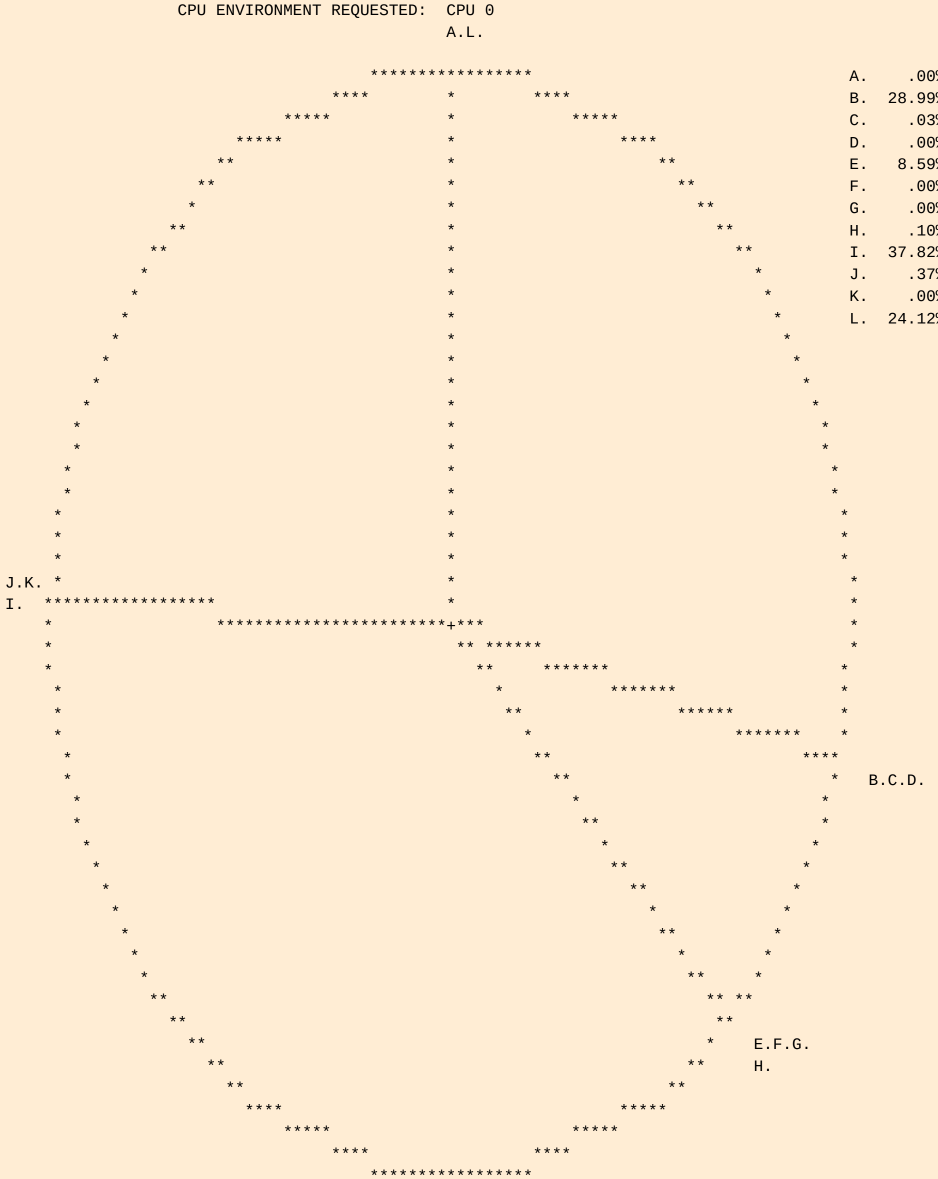


```
A.     .00% SWP-AXS
B.  28.99% SWP-DW
C.     .03% SWP-ENQ
D.     .00% SWP-EXC
E.   8.59% SWP-LW
F.     .00% SWP-REQ
G.     .00% SWP-RSS
H.     .10% SWP-UNI
I.  37.82% SWP-TOT
J.     .37% SWP-TXS
K.     .00% SWP-LTT
L.  24.12% SWP-TWT
```

Pie Graph field descriptions

Table 94 describes each field in the Pie Graph.

Table 94 Field descriptions for the Pie Graph

Field	Description
CPU ENVIRONMENT REQUESTED	CPUs for which measures are given, according to the GRAPH statement
A through L	breakdown by percentages of the event counts selected in the statement For information about the measures that are charted, see Appendix C, “Measure and trace values.”

Processor Concurrency Report

The Processor Concurrency Report graphically displays the number of processors (CPUs) within the configuration and the percentage of time that these CPUs were concurrently busy.

The Processor Concurrency Report is produced by using the CPUCON Analyzer control statement (see “CPUCON” on page 232). The data for this report is obtained by using the CPU (see “CPU” on page 136) Extractor control statement.

The vertical axis of the graph represents the percentage of time that zero or more CPUs were concurrently busy; the horizontal axis represents the number of CPUs that were concurrently busy. The number of CPUs concurrently busy is represented by columns of asterisks. For example, if two processors in a system were concurrently busy 75% of the time, a column of asterisks would be printed above the number 2 on the horizontal axis and plotted up to the number 75 on the vertical axis.

A sample of the Processor Concurrency Report is shown in Figure 82.

Figure 82 Processor Concurrency Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    PROCESSOR CONCURRENCY REPORT                   RPTSEQ    11 PAGE     33
BMC SOFTWARE, INC.                                            XYZ COMPANY                           REPORT DATE: DD MMM YY 9.33
ACTL 10 JUN YY 06.45.01  10 JUN YY 07.45.00             WORLDWIDE HEADQUARTERS                      SYSTEM ID: SYSB     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/4/709/1.00

-----------------------------------------------------  CPU OPERATION PATTERN  -----------------------------------------------------
                                              NUMBER OF CPU'S IN THE CONFIGURATION ( 3)
                                      ALL CPU'S CONCURRENTLY BUSY ( 57.1 PERCENT OF TOTAL TIME)

    CPU MEASURES   - AVG %  9   18   16   57

                    100 +                                                                                       +
                        I                                                                                       I
                        I                                                                                       I
                        I                                                                                       I
                        I                                                                                       I
                        I                                                                                       I
                        I                                                                                       I
                        I                                                                                       I
                    75  +                                                                                       +
                        I                                                                                       I
                S       I                                                                                       I
                A       I                                                                                       I
                M       I                                                                                       I
                P       I                                                                                       I
                L       I                ***                                                                    I
                E       I                * *                                                                    I
                    50  +                * *                                                                    +
                P       I                * *                                                                    I
                E       I                * *                                                                    I
                R       I                * *                                                                    I
                C       I                * *                                                                    I
                E       I                * *                                                                    I
                N       I                * *                                                                    I
                T       I                * *                                                                    I
                    25  +                * *                                                                    +
                        I                * *                                                                    I
                        I                * *                                                                    I
                        I      ***  ***  * *                                                                    I
                        I      * *  * *  * *                                                                    I
                        I      * *  * *  * *                                                                    I
                        I ***  * *  * *  * *                                                                    I
                        I * *  * *  * *  * *                                                                    I
                        ---+----+----+----+----+----+----+----+----+----+----+----+----+----+----+----+----+----

                           0    1    2    3    4    5    6    7    8    9   10   11   12   13   14   15   16


                                                NUMBER OF CPU'S CONCURRENTLY BUSY
```

Processor Concurrency Report field descriptions

Table 95 describes each field in the Processor Concurrency Report.

Table 95 Field descriptions for the Processor Concurrency Report

Field	Description
NUMBER OF CPU'S IN THE CONFIGURATION	total number of CPUs measured
ALL CPU'S CONCURRENTLY BUSY	total percentage of time that all CPUs measured were concurrently busy
CPU MEASURES	average percentages of time that all CPUs were not busy and multiple CPU busy percentages; these figures correspond to columns below
SAMPLE PERCENT	axis representing percentage of time that zero or more CPUs were concurrently busy
NUMBER OF CPU'S CONCURRENTLY BUSY	axis representing number of CPUs

Profile Bar Graph

The Profile Bar Graph (or histogram) is produced by using the GRAPH TYPE=PROFILE Analyzer control statement. The data in this graph can be used to locate peaks and valleys in the performance of the system over time of day.

The Profile Bar Graph is scaled to the largest LIMIT= value that is specified in the GRAPH statement. The default limit is 100 (coded `LIMIT=10000`). Up to 16 measures can be specified; each measure is represented on the graph by a different character. When more than one measure is represented, the one with the smallest value is displayed first on the line.

The Extractor statements required for the Profile Bar Graph are dependent upon the information to be graphed. (See "Record types" on page 45 for information about the record types generated by each Extractor statement and a description of the specific type of data collected by each Extractor statement.)

One line on the graph is produced for each interval specified in the GRAPH command. If the interval default is used, one line is produced for each record interval encountered in the data.

The Profile Bar Graph is read from left to right. The value of each measure on the graph is determined by locating the rightmost character representing that measure; the value shown at that point is the value of the measure. For measures that have exactly the same value on the line, the character O is printed to indicate an overlap.

A sample of the Profile Bar Graph is shown in Figure 83. In this figure, one measure, CPU, was selected.

Figure 83 Profile Bar Graph

```
PRODUCED BY CMF ANALYZER (v.r.mm)                          PROFILE BAR GRAPH                                RPTSEQ    23 PAGE     58
BMC SOFTWARE, INC.                                            XYZ COMPANY                                   REPORT DATE: DD MMM YY 9.33
ACTL 10 JUN YY 06.45.01  10 JUN YY 07.45.00             WORLDWIDE HEADQUARTERS                              SYSTEM ID: SYSB     Z v.rr.n


                * REPRESENTS CPU
                  CPU ENVIRONMENT REQUESTED:  ALL CPU'S
  DATE     TIME        CPU                                 GRAPHIC DISPLAY OF INTERVAL TOTAL TIME
                               0        10        20        30        40        50        60        70        80        90       100
                               +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
         07.45.00          56  .********************************************************   .         .         .         .         .
         08.00.00          60  .************************************************************         .         .         .         .
         08.15.00          54  .******************************************************     .         .         .         .         .
         08.30.00          55  .*******************************************************    .         .         .         .         .
         08.45.01          54  .******************************************************     .         .         .         .         .
         09.00.00          52  .****************************************************       .         .         .         .         .
         09.15.00          50  .**************************************************         .         .         .         .         .
                               +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
```

The Profile Bar Graph is most effective when used over time to compare one time frame to another.

Profile Bar Graph field descriptions

Table 96 describes each field in the Profile Bar Graph.

Table 96 Field descriptions for the Profile Bar Graph

Field	Description
***n* REPRESENTS**	character *n* used to represent each measure selected in GRAPH statement
DATE	date when samplings were made
TIME	time frames at which samplings were made
GRAPHIC DISPLAY OF INTERVAL TOTAL TIME	axis representing values for each measure, according to rightmost character Values can be percentages, numerical count, and so on, depending on the measure being scaled.

Report Table of Contents

The Report Table of Contents always precedes the CMF MONITOR logs and selected reports. It lists the CMF MONITOR reports by sequence, page number, and report title, and, if specified, by subtitle.

The Report Table of Contents is generated only if the //RPTCONTS DD JCL control statement is defined in the Analyzer JCL. (See “Defining Analyzer JCL manually” on page 84.) The //RPTCONTS DD statement should precede the //SYSPRINT DD statement in the Analyzer JCL.

No Analyzer or Extractor control statements are required to produce the Report Table of Contents. It is produced automatically as part of the preliminary reporting information that CMF MONITOR supplies with each batch report job (see “Preliminary reporting information” on page 330 for more information).

An example of the Report Table of Contents is shown in Figure 84.

Figure 84 Report Table of Contents

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    REPORT TABLE OF CONTENTS                                     CONTENTS    PAGE        i
BMC SOFTWARE, INC.                                                                                                REPORT DATE: DD MMM YY
                                                                                                                  REPORT TIME:    11.23.55


 RPT SEQ NO   PAGE NO      REPORT TITLE                                                  REPORT SUBTITLE

     1           1      CONTROL CARD LOG
     2           2      COLLECTION PHASE LOG
     3           6      AUXILIARY STORAGE REPORT
     4           8      CPU UTILIZATION REPORT
     5          11      VIRTUAL STORAGE ACTIVITY REPORT
     6          17      REPORT PHASE LOG
```

Report Table of Contents field descriptions

Descriptions of the fields in the Report Table of Contents are listed in Table 97.

Table 97 Field descriptions for the Report Table of Contents

Field	Description
RPT SEQ NO	sequence number of when the report is displayed in the output
PAGE NO	page number within the report output where the report begins
REPORT TITLE	report title
REPORT SUBTITLE	report subtitle, if SUBTITLE control statement is chosen (see “SUBTITLE” on page 309 for more information)

Shared Device Activity Report

The Shared Device Activity report displays shared DASD and tape activity (see Figure 85).

The Shared Device Activity Report is produced by using the SHARDEV Analyzer control statement (see “SHARDEV” on page 297). The data for this report is obtained by using the DEVICE Extractor control statement (see “DEVICE” on page 145).

Figure 85 Shared Device Activity Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     SHARED DEVICE ACTIVITY REPORT                    RPTSEQ     3 PAGE     72
BMC SOFTWARE, INC.                                         BMC SOFTWARE, INC.                          REPORT DATE: DD MMM YY 15.32
REQD 06 JUL YY 13.30.00  06 JUL YY 15.29.59                   HOUSTON, TX.                             SYSTEM ID: **ALL**  COMB-MVS
ACTL 06 JUL YY 13.30.00  06 JUL YY 15.30.00                                                            REPORT CYCLE: CYCLE099

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  74-1/112/7,125/4
                                 ---------------- AVERAGE TIME IN MSEC ----------------   *** PERCENT OF TOTAL TIME ***
VOLSER/DEVTYPE /PAV MX                              INIT    DEV                                      DEV                         AVG
DEV  IODF     SYSNAME/    SSCH     TOTAL           COMND   BUSY                               DEV    IN     DEV  DEV  MOUNT    DSETS
NUM  SFX  LCU  SYSID     PER SEC  SERVICE    IOSQ    RESP  DELAY    PEND    DISC    CONN    CONN  USE    RESV ALLOC  PEND     ALLOC
---- --- ---- -------- --------  ------- -------- ------ ------ ------- ------- -------   ----- ----- ----- ----- -----     -----
PAGC31 3380K
8320  47 009B SJSC       0.287      3.7      1.0    0.0    0.0     1.2     1.0     0.5     0.0   0.0   0.0 100.0   0.0       4.0
8320  47 0067 SJSE       0.024      7.6      0.0    0.0    0.0     7.4     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.312      4.0      0.9    0.0    0.0     1.7     1.0     0.4     0.0   0.0   0.0 100.0   0.0       4.0

PAGC32 33903         3
4103  47 0078 SJSC       9.142      0.8      0.0    0.1    0.0     0.3     0.0     0.5     0.2   0.2   0.0 100.0   0.0       2.0
4103  47 0045 SJSE       0.024      0.4      0.0    0.1    0.0     0.2     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    9.166      0.8      0.0    0.1    0.0     0.3     0.0     0.5     0.2   0.2   0.0 100.0   0.0       2.0

PAGC33 33903
85CE  47 009D SJSC       0.386      1.1      0.0    0.0    0.0     0.7     0.1     0.4     0.0   0.0   0.0 100.0   0.0       3.0
85CE  47 0069 SJSE       0.024      1.8      0.0    0.0    0.0     1.6     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.410      1.2      0.0    0.0    0.0     0.8     0.0     0.4     0.0   0.0   0.0 100.0   0.0       3.0

PAGC34 3380K
833A  47 009B SJSC       0.385      1.5      0.0    0.0    0.0     1.1     0.1     0.4     0.0   0.0   0.0 100.0   0.0       3.0
833A  47 0067 SJSE       0.024      8.8      0.0    0.0    0.0     8.6     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.409      1.9      0.0    0.0    0.0     1.5     0.1     0.3     0.0   0.0   0.0 100.0   0.0       3.0

PAGC35 33903
874D  47 009F SJSC       0.262      1.2      0.0    0.0    0.0     0.7     0.1     0.4     0.0   0.0   0.0 100.0   0.0       2.0
874D  47 006B SJSE       0.024      1.8      0.0    0.0    0.0     1.6     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.286      1.2      0.0    0.0    0.0     0.8     0.0     0.4     0.0   0.0   0.0 100.0   0.0       2.0

PAGD27 3380K
8321  47 009B SJSC       0.012     11.7      0.0    0.0    0.0    11.4     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
8321  47 0067 SJSE       0.024      8.4      0.0    0.0    0.0     8.2     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.036      9.5      0.0    0.0    0.0     9.3     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0

PAGE25 3380K
8322  47 009B SJSC       0.012      9.0      0.0    0.0    0.0     8.7     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
8322  47 0067 SJSE       0.024     31.5     22.9    0.0    0.0     8.3     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.037     23.9     15.2    0.0    0.0     8.5     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0

PAGF49 3380K
8307  47 009B SJSC       0.004     13.4      0.0    0.0    0.0    13.1     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
8307  47 0067 SJSE       0.019      4.6      0.0    0.0    0.0     4.4     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0
SUMMARY  (  2 ACTIVE)    0.024      6.3      0.0    0.0    0.0     6.1     0.0     0.2     0.0   0.0   0.0 100.0   0.0       0.0

                                                                  SUMMARY
                                              1015 DASD AND   107 TAPE DEVICES WERE FOUND
                                                                        # OF DEVICES
                                                        SYSTEM         DASD     TAPE
                                                        SJSC           1015      104
                                                        SJSE           1015      104
```

Shared Device Activity Report field descriptions

Table 98 describes each field in the Shared Device Activity Report.

Table 98 Field descriptions for the Shared Device Activity Report (part 1 of 3)

Field	Description
VOLSER	volser from the UCB of the last volume mounted on this device
DEV TYPE	type of the physical I/O device on which the volume is mounted
PAV MX	number of exposures (base and aliases) of a Parallel Access Volume (PAV) at the end of the report duration; applicable only to DASDs An asterisk (*) following this number indicates that the number of exposures changed during the report duration.
DEV NUM	four-digit hexadecimal number that identifies this device An asterisk (*) following this number indicates that the online/offline status of the device was changed during the report interval.
IODF SFX	IODF suffix in effect for the system
LCU	four-digit hexadecimal identifier of the Logical Control Unit
SYSNAME/SYSID	name of the system by either SYSNAME or SYSID See “SYSPLEX” on page 310 for more information about SYSNAME and SYSID.
SSCH PER SEC	rate per second at which SSCH instructions issued for this device completed successfully OFFLINE is printed under this column if the device was offline during the entire report interval. If the online/offline status of the device was changed during a recording interval, it is considered offline and its data is skipped for that particular interval. The device SUMMARY line contains the sum of SSCH/sec for all systems reported.
TOTAL SERVICE	average number of milliseconds the device required to service an I/O request (IOSQ + PEND + DISC + CONN) The device SUMMARY line contains the sum of the weighted IOSQ, PEND, DISC, and CONN.
IOSQ	average number of milliseconds of delay an I/O request encountered because the device was busy performing an I/O from the local system The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the I/Os by that system over the total I/Os by all reported systems.

Table 98 Field descriptions for the Shared Device Activity Report (part 2 of 3)

Field	Description
INIT COMND RESP	average number of milliseconds of delay that an I/O request encountered beginning from when the first command of the channel program is sent to the device until the device indicates that it has accepted the command This delay is part of pending time. Delay is available only on z990 or later processors.
DEV BUSY DELAY	average number of milliseconds of delay that an I/O request encountered because the device was busy due to I/O from another system; this delay is part of pending time The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the I/Os by that system over the total I/Os by all reported systems.
PEND	average number of milliseconds that an I/O request must wait for hardware, such as an available channel path or control unit, as well as time between the SSCH pending at the channel and the device active on the subchannel; also includes delays caused by another processor reserving this device in a shared DASD environment The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the I/Os by that system over the total I/Os by all reported systems.
DISC	average number of milliseconds during which the device was processing an SSCH instruction but not transferring data The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the I/Os by that system over the total I/Os by all reported systems.
CONN	average number of milliseconds during which the device was processing an SSCH instruction and transferring data The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the I/Os by that system over the total I/Os by all reported systems.
% DEV CONN	percentage of time during the measurement interval that the device was connected to a channel path The device SUMMARY line contains the sum of % DEV CONN for each system.
% DEV IN USE	percentage of time during the measurement interval that the device was in use; includes device connect and disconnect time The device SUMMARY line contains the sum of % DEV IN USE for each system.
% DEV RESV	percentage of time during the measurement interval that this device was reserved by the processor on which the CMF MONITOR Extractor was executing The device SUMMARY line contains the sum of % DEV RESV for each system.

Table 98 Field descriptions for the Shared Device Activity Report (part 3 of 3)

Field	Description
% DEVICE ALLOC	percentage of time during the measurement interval that this device was allocated to one or more data sets; DASD devices always show 100% allocation The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the I/Os by that system over the total I/Os by all reported systems.
% MOUNT PEND	percentage of time during the measurement interval that this device had an outstanding mount pending request The device SUMMARY line contains the sum of % MOUNT PEND for each system.
AVG DSETS ALLOC	average number of data sets allocated on this DASD The device SUMMARY line contains the sum of AVG DSETS ALLOC for each system.
NUM OF MOUNTS	total number of mounts for the tape device during the reporting interval An asterisk (*) to the left of this value indicates that a mount pending condition existed at the start of the recording interval. An asterisk (*) to the right of the value indicates that a mount pending condition existed at the end of the recording interval. The device SUMMARY line contains the sum of NUM OF MOUNTS for each system.
AVG MOUNT TIME	average mount pending time for the tape device The field is reported as *hh:mm:ss*; the maximum time reported in this field is 99:59:59. The device SUMMARY line contains the weighted average of the values of the reported systems. The weighting factor for a system is the ratio of the number of mounts on that system over the total number of mounts on all reported systems.
TIME DEVICE ALLOC	total time that the tape device was allocated during the reporting interval The field is reported as *hh:mm:ss*; the maximum time reported in this field is 99:59:59. The device SUMMARY line contains the sum of TIME DEVICE ALLOC for each system.
SUMMARY	total number of DASD and tape devices found on all systems
# OF DEVICES – DASD	total number of DASD devices used on the system
# OF DEVICES – TAPE	total number of tape devices used on the system

Storage Management Report

This report provides detailed information about paging activity, frame counts, and storage movement. The example in Figure 86 on page 510 illustrates the maximum amount of data available. When data collected on systems that do not make certain values available is used, the display fields associated with those values are not displayed. Notes in the example indicate the affected fields, columns, rows, or report sections.

The Storage Management Report is produced by using the STORAGE Analyzer control statement (see "STORAGE" on page 308). The data for this report is obtained by using the PAGING (see "PAGING" on page 173) Extractor control statement.

The Storage Management Report consists of these sections:

- **Detail Paging Activity**

  This section displays paging rates and activity percents viewed by page origin.

- **Expanded Storage Movement**

  This section displays special measurements pertaining to expanded storage (Migration age and High UIC) and storage movement between mediums in terms of rate, percent of total, and frame count measures.

- **Page Frame Counts**

  This section displays frame count measures in four categories:

  — CENTRAL STORAGE
  — FIXED CENTRAL STORAGE
  — EXPANDED STORAGE
  — SHARED PAGE GROUPS

A sample of the Storage Management Report is shown in Figure 86 on page 510.

Figure 86 Storage Management Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    STORAGE MANAGEMENT REPORT                        RPTSEQ     4 PAGE      7
BMC SOFTWARE, INC.                                          XYZ COMPANY                               REPORT DATE: DD MMM YY 10.21
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00           WORLDWIDE HEADQUARTERS                          SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  71-1/30/450/7.5
MODE = ESAME --------------------------------------- DETAIL PAGING ACTIVITY ----------------------------------------------------
                                                                                                                     PAGE
                             R A T E S   (PAGES PER SECOND)                P E R C E N T   O F   T O T A L          MOVEMENT
                          ------------------------------------          ----------------------------------     ---------------
CENTRAL STORAGE             PAGE-IN      PAGE-OUT      RECLAIMS             PAGE-IN      PAGE-OUT      RECLAIMS     PG/SEC    26.99
  LPA                          0.60                        0.00               100.0                         0.0     TIME %     0.0
  LPA,BLOCK                    0.05                                           100.0
  CSA                          4.93          0.35          0.00                93.4           6.6           0.0
  CSA,BLOCK                    0.02                                           100.0
SYSTEM AREA TOTAL              5.60          0.35          0.00                94.2           5.8           0.0
                                                                                                                 BLOCK PAGING
  NON-VIO, NON-SWAP            0.00         17.88          0.00             *ERROR*       *ERROR*           0.0    ---------------
  NON-VIO, SWAP                2.57          2.56                              50.0          50.0                  PG/BLK     4.48
  NON-VIO,NON-SWAP,BLOCK      11.18                                           100.0
  VIO                          0.00          0.00          0.00                 0.0           4.8          95.2
  HIPERSPACE                   0.00          0.00                               0.0           0.0
  SHARED                       0.03          0.02                              64.8          35.2
ADDRESS SPACE TOTAL           13.75         20.46          0.00                40.2          59.8           0.0

TOTAL                         19.36         20.81          0.00                48.2          51.8           0.0

PAGE FAULTS/SECOND             8.08

---------------------------------------------------- EXPANDED STORAGE MOVEMENT ---------------------------------------------------

MIGRATION AGE   MIN       0    MAX       0    AVG      0.0          HIGH UIC       MIN      154    MAX    2,540    AVG   2089.5
FREED WITHOUT MIGRATION:  TOTAL        0   PERCENT OF TOTAL      0.0

                   R A T E S   (PAGES PER SECOND)         P E R C E N T   O F   T O T A L       EXPANDED STORAGE FRAME COUNTS
                  ----------------------------------     --------------------------------    -----------------------------------
                   FROM         TO           TO            FROM          TO          TO
                  CENTRAL      CENTRAL     AUXILIARY      CENTRAL      CENTRAL     AUXILIARY      MINIMUM      MAXIMUM      AVERAGE
  VIO               0.00         0.00         0.00           0.0          0.0         0.0            0            0            0
  HIPERSPACE        0.00         0.00         0.00           0.0          0.0         0.0            0            0            0
  SHARED            0.00         0.00                        0.0          0.0                        0            0            0
TOTAL               0.00         0.00         0.00                                                   0            0            0

------------------------------------------------------- PAGE FRAME COUNTS -------------------------------------------------------

CENTRAL STORAGE            MINIMUM       MAXIMUM       AVERAGE    *   EXPANDED STORAGE          MINIMUM       MAXIMUM       AVERAGE
  SQA                        5,886         7,540         7,076    *     SQA                           0             0             0
  LPA                            0         4,187           886    *     LPA                           0             0             0
  CSA                        8,914        17,072        14,445    *     CSA                           0             0             0
  LSQA                      15,440        27,201        23,321    *     LSQA                          0             0             0
  ADDRESS SPACE             64,455       165,791       137,123    *     ADDRESS SPACE                 0             0             0
  AVAILABLE                     22       103,336        20,906    *     AVAILABLE                     0             0             0
TOTAL                      195,063       195,063       195,063    *
                                                                  *     ONLINE                        0
                                                                  *     INSTALLED                     0

BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  71-1/30/450/7.5
FIXED CENTRAL STORAGE      MINIMUM       MAXIMUM       AVERAGE    *   SHARED PAGE GROUPS        MINIMUM       MAXIMUM       AVERAGE
  NUCLEUS                    1,905                                *     CENTRAL STORAGE              34           438           170
  SQA                        5,375         7,531         6,986    *     EXPANDED STORAGE              0             0             0
  LPA                            0            94            12    *     FIXED TOTAL                   1             2             1
  CSA                          196           570           467    *     FIXED BELOW 16M               0             0             0
  LSQA                       7,413        15,605        12,807    *     AUXILIARY SLOTS               0           389           317
  ADDRESS SPACE              3,258         4,842         3,986
  BELOW 16 MEGABYTES           247           479           416
TOTAL                       19,623        29,697        26,165
```

Detail Paging Activity Section field descriptions

Table 99 describes each field in the Detail Paging Activity section of the Storage Management Report.

NOTE

A value of *ERROR* in any of the fields of this report indicates that invalid data was encountered. One possible cause of invalid data is that expanded storage is offline; however, other causes are also possible.

Table 99 Field descriptions for the Detail Paging Activity section (part 1 of 2)

Field	Description
MODE	architecture level of the operating system; ESAME indicates 64-bit real support
RATES – LPA	paging rate of non swap, non block pages transferred from auxiliary storage to pageable LPA central storage
RATES – LPA,BLOCK	paging rate of non swap, block pages transferred from auxiliary storage to pageable LPA central storage
RATES – CSA	paging rate of non swap, non block pages transferred between auxiliary storage and pageable CSA central storage
RATES – CSA,BLOCK	paging rate of non swap, block pages transferred from auxiliary storage to pageable CSA central storage
RATES – SYSTEM AREA TOTAL	paging rate of pages transferred between auxiliary storage and pageable LPA and CSA central storage
RATES – NON-VIO, NON-SWAP	paging rate of non VIO, non swap pages transferred between auxiliary and central storage
RATES – NON-VIO, SWAP	paging rate of non VIO, swap pages transferred between auxiliary and central storage
RATES – NON-VIO, NON-SWAP, BLOCK	paging rate of non VIO, non swap, and block pages transferred between auxiliary to central storage
RATES – VIO	paging rate of all VIO pages transferred between auxiliary and central storage
RATES – HIPERSPACE	paging rate of hiperspace pages transferred between auxiliary and central storage
RATES – SHARED	paging rate of shared pages transferred between auxiliary and central storage This field is only displayed in MVS 5.2 and later systems.
RATES – ADDRESS SPACE TOTAL	paging rate of all VIO, non VIO, hiperspace, and shared pages transferred between auxiliary and central storage
RATES – PAGE-IN	paging rate of page-ins transferred from auxiliary to central storage
RATES – PAGE-OUT	paging rate of page-outs transferred from central storage to auxiliary storage
RATES – RECLAIMS	paging rate of pages reused without reading from external storage
PERCENT OF TOTAL – PAGE-IN	percentage of the total of the three rates (page-ins, page-outs, and reclaims) for this row that was the page-in rate
PERCENT OF TOTAL – PAGE-OUT	percentage of the total of the three rates (page-ins, page-outs, and reclaims) for this row that was the page-out rate

Table 99 Field descriptions for the Detail Paging Activity section (part 2 of 2)

Field	Description
PERCENT OF TOTAL – RECLAIMS	percentage of the total of the three rates (page-ins, page-outs, and reclaims) for this row that was the reclaim rate
PAGE MOVEMENT PG/SEC	rate of page movement between below the 16-MB line and above the 16-MB line in central storage
PAGE MOVEMENT TIME %	percentage of time spent in PREF STEAL during the reporting interval; this field is reported as a system area total value
BLOCK PAGING PG/BLK	average number of pages per block
TOTAL	paging rate of all pages transferred
PAGE FAULTS/SEC	rate of page faults per second; the rate of pages read from DASD only

Expanded Storage Movement section field descriptions

Table 100 describes each field in the Expanded Storage Movement section of the Storage Management Report.

Table 100 Field descriptions for the Expanded Storage Movement section (part 1 of 2)

Field	Description
MIGRATION AGE – MIN	minimum age in seconds of a page before it was migrated from expanded storage
MIGRATION AGE – MAX	maximum age in seconds of a page before it was migrated from expanded storage
MIGRATION AGE – AVG	average age in seconds of a page before it was migrated from expanded storage
HIGH UIC – MIN	minimum high Unreferenced Interval Count (UIC) in seconds (time that a page of central storage frames has not been referenced)
HIGH UIC – MAX	maximum high Unreferenced Interval Count (UIC) in seconds (time that a page of central storage frames has not been referenced)
HIGH UIC – AVG	average high Unreferenced Interval Count (UIC) in seconds (time that a page of central storage frames has not been referenced)
RATES – FROM CENTRAL	pages moved from central storage to expanded storage **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
RATES – TO CENTRAL	pages moved from expanded storage to central storage **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
RATES – TO AUXILIARY	pages moved from expanded storage to auxiliary storage **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.

Table 100 Field descriptions for the Expanded Storage Movement section (part 2 of 2)

Field	Description
PERCENT OF TOTAL – FROM CENTRAL	percentage of the total page movements for this row's area that were movements from central storage to expanded storage **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
PERCENT OF TOTAL – TO CENTRAL	percentage of the total page movements for this row's area that were movements from expanded storage to central storage **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
PERCENT OF TOTAL – TO AUXILIARY	percentage of the total page movements for this row's area that were movements from expanded storage to auxiliary storage **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
EXPANDED STORAGE FRAME COUNTS – MINIMUM	minimum number of VIO, HIPERSPACE, or SHARED page frames in use in expanded storage A VIO page frame is allocated to a VIO address space, a HIPERSPACE frame is allocated to a HIPERSPACE, and a SHARED page frame is allocated to a shared address space. **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
EXPANDED STORAGE FRAME COUNTS – MAXIMUM	maximum number of VIO, HIPERSPACE, or SHARED page frames in use in expanded storage A VIO page frame is allocated to a VIO address space, a HIPERSPACE frame is allocated to a HIPERSPACE, and a SHARED page is allocated to a shared address space. **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
EXPANDED STORAGE FRAME COUNTS – AVERAGE	average number of VIO, HIPERSPACE, or SHARED page frames in use in expanded storage A VIO page frame is allocated to a VIO address space, a HIPERSPACE frame is allocated to a HIPERSPACE, and a SHARED page is allocated to a shared address space. **Note**: The VIO, HIPERSPACE, and SHARED values are subsets of the TOTAL values.
EXPANDED STORAGE FRAME COUNTS – ONLINE	number of expanded storage page frames online
EXPANDED STORAGE FRAME COUNTS – INSTALLED	number of expanded storage page frames installed

Page Frame Counts section field descriptions

Table 101 describes each field in the Page Frame Counts section of the Storage Management Report.

Table 101 Field descriptions for the Page Frame Counts section

Field	Description
CENTRAL STORAGE – MINIMUM	minimum number of page frame counts in SQA, LPA, CSA, LSQA, address space, available, or total
CENTRAL STORAGE – MAXIMUM	maximum number of page frame counts in SQA, LPA, CSA, LSQA, address space, available, or total
CENTRAL STORAGE – AVERAGE	average number of page frame counts in SQA, LPA, CSA, LSQA, address space, available, or total
EXPANDED STORAGE – MINIMUM	minimum number of page frames backing SQA, LPA, CSA, LSQA, address space, available, or total page frames in expanded storage
EXPANDED STORAGE – MAXIMUM	maximum number of page frames backing SQA, LPA, CSA, LSQA, address space, available, or total page frames in expanded storage
EXPANDED STORAGE – AVERAGE	average number of page frames backing SQA, LPA, CSA, LSQA, address space, available, or total page frames in expanded storage
EXPANDED STORAGE – ONLINE	number of expanded storage page frames online
EXPANDED STORAGE – INSTALLED	number of expanded storage page frames installed
FIXED CENTRAL STORAGE – MINIMUM	minimum fixed central storage page frame count
FIXED CENTRAL STORAGE – MAXIMUM	maximum fixed central storage page frame count
FIXED CENTRAL STORAGE – AVERAGE	average fixed central storage page frame count
SHARED PAGE GROUPS – MINIMUM	minimum number of shared page group page-ins and page-outs in central storage
SHARED PAGE GROUPS – MAXIMUM	maximum number of shared page group page-ins and page-outs in central storage
SHARED PAGE GROUPS – AVERAGE	average number of shared page group page-ins and page-outs in central storage

System Resources Manager Report

The System Resources Manager Report provides general information about the System Resource Manager (SRM) in five sections.

The System Resources Manager Report is produced by using the SRM Analyzer control statement (see “SRM” on page 307). The data for this report is obtained by using the CPU (see “CPU” on page 136) and PAGING (see “PAGING” on page 173) Extractor control statements.

- **Queue Measures**

  This section shows queuing and swapping data for address spaces. This section is divided into two parts, physical queues and logical queues.

- **Paging Activity**

  This section shows paging activity by average pages per second and average percentage. These measures are further divided into system area and private area measures.

  Reclaims are shown under Paging Activity. A reclaim occurs when a page frame is stolen from a private or system pageable area but is retrieved for reuse before being reallocated. The request for a page is satisfied without starting a page-in.

- **Expanded Storage Measures**

  This section shows paging activity to the expanded storage and statistical data on available frames and migration age.

- **SRM Data**

  This section contains counts of SRM related events, such as SQA and available frame queue shortages. It also displays percentages of threshold imbalance as defined in the following control blocks:

  — CPU management control table (CCT)

  — I/O management control table (ICT)

  — storage management control table (MCT)

- **Swapping Measures**

  This section shows a count of reasons for swapping by destination of the swap, rate per minute, and percentage value. This section of the report prints on two pages, separate from the other four report sections. (Two pages are required for the Swapping Measures section because of the volume of information presented.)

A sample of the System Resource Manager Report is shown in Figure 87, except for the Swapping Measures section. The Swapping Measures section, which prints on two pages, is shown in Figure 88 on page 523.

Figure 87 SRM Report (except Swapping Measures section)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  SYSTEM RESOURCES MANAGER REPORT                    RPTSEQ     4 PAGE      9
BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 11.55
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/30/5,393/7.5  240-3/30/450/7.5  70-1/30/5,393/7.5  71-1/30/450/7.5
-------------------------------------------------------  QUEUE MEASURES  ------------------------------
                           PHYSICAL QUEUES                                        LOGICAL QUEUES                ADDRESS SPACES
    MEASURE      IN & READY      IN QUEUE  OUT & READY  WAIT QUEUE           OUT & READY    WAIT QUEUE             IN AND NONSWAP
    PCT QUEUED      100.0          100.0        .3          57.7                   .4        100.0                          N/A
    AVERAGE           1.8           97.8        .0          34.5                   .0         73.4                         95.2
    MINIMUM           1.0           67.0        .0            .0                   .0          5.0                         66.0
    MAXIMUM          81.0          158.0       1.0         117.0                  2.0        118.0                        132.0
    STD DEV           0.5           10.3       0.5          21.7                  0.5         19.5                         10.2
-------------------------------------------------------  PAGING ACTIVITY  ----------------------------------------------------
ACTIVITY TYPE                  SYSTEM AREA     PRIVATE AREA      *      ACTIVITY TYPE                          RCL/MIN      PCT
                              PG/SEC     PCT    PG/SEC     PCT      *
DEMAND PAGING PAGE IN          5.60    94.18    11.16    32.66      *
DEMAND PAGING PAGE OUT          .35     5.82    17.88    52.33      *      COMMON AREA RECLAIMS                       .00       .00
VIO PAGE IN                                       .00      .00      *      PRIVATE AREA RECLAIMS (NON-VIO)            .00       .00
VIO PAGE OUT                                      .00      .00      *      VIO RECLAIMS                               .04    100.00
SWAP PAGE IN                                     2.57     7.51      *
SWAP PAGE OUT                                    2.56     7.50      *      TOTAL RECLAIMS                             .04    100.00
TOTAL PAGING                   5.95   100.00    34.17   100.00      *      ACTIVITY TYPE                      PAGES MOVED    PG/SEC
                                                                    *      PAGE MOVEMENT                           728729     26.99
-------------------------------------------------  EXPANDED STORAGE MEASURES  -------------------------------------------------
                    EXPANDED STORAGE ACTIVITY                       *                 EXPANDED STORAGE CONFIGURATION
PAGES TO EXPANDED STORAGE                          .00              *      INSTALL PAGE FRAMES                          0
MIGRATION RATE TO AUX. STORAGE                     .00              *      ONLINE PAGE FRAMES                           0
                  AVAILABLE         SYSTEM            MIGRATION     *
MEASURE           ES FRAMES        HIGH UIC              AGE        *
AVERAGE               0.0           2089.5               0.0        *
MINIMUM                 0              154                 0        *
MAXIMUM                 0             2540                 0        *
-------------------------------------------------------     SRM DATA    -----------------------------------------------------------
KEY SRM MEASURES                                        VALUE       *      KEY SRM MEASURES                                       PCT
SQA LOW EVENT COUNT                                         0       *      CPU OVERLOAD PERCENTAGE                                .00
AVAILABLE QUEUE LOW EVENT COUNT                          6334       *      I/O OVERLOAD PERCENTAGE                                .00
ENQUEUE HOLD SWAPOUT EVENT COUNT                            0       *      I/O UNDERLOAD PERCENTAGE                               .00
MAX PAGE-INS/SEC                                        392.2       *      CPU IMBALANCE PERCENTAGE                               .00
MAX SWAPS/MIN                                             230       *      STORAGE IMBALANCE PERCENTAGE                           .00
AVG PAGES/SWAP OUT                                        268       *      TERMINAL I/O SWAPS PERCENTAGE DUE TO OUTPUT           4.92
AVG PAGES/SWAP IN                                         268       *
SRM OPTIONS MEMBER FROM PARMLIB                      IEAOPTE1       *
```

Queue Measures section field descriptions

Table 102 describes each field in the Queue Measures section of the System Resources Manager Report.

Table 102 Field descriptions for the Queue Measures section

Field	Description	
MEASURE	measure being reported, which can be	
	PCT QUEUED	percentage of time the queue existed
	AVERAGE	average number of address spaces queued over time
	MINIMUM	minimum number of address spaces queued over time
	MAXIMUM	maximum number of address spaces queued over time
	STD DEV	standard deviation of the average
PHYSICAL QUEUES	**IN & READY**	number of address spaces swapped in with dispatchable work to do
	IN QUEUE	number of address spaces swapped in, but were nondispatchable
	OUT & READY	number of address spaces swapped out, but were ready to be swapped in
	WAIT QUEUE	number of address spaces swapped out and waiting for an event to finish (such as a WTOR or mount)
LOGICAL QUEUES	**OUT & READY**	address spaces on the SRM out queue that are physically in central storage but are logically swapped out of central storage and ready to execute
	WAIT QUEUE	address spaces on the SRM wait queue that are physically in central storage but are logically swapped out to central storage and not ready
ADDRESS SPACES IN AND NONSWAP	number of swapped in address spaces marked not swappable	

Paging Activity section field descriptions

Table 103 describes each column field in the Paging Activity section of the System Resources Manager Report.

Table 103 Field descriptions for the Paging Activity section (part 1 of 2)

Field	Description	
SYSTEM AREA	**PG/SEC**	rate per second at which pages were moved in the system area for the ACTIVITY TYPE
	PCT	percentage of total pages moved in the system area for the ACTIVITY TYPE
PRIVATE AREA	**PG/SEC**	rate per second at which pages were moved in the private for the ACTIVITY TYPE
	PCT	percentage of total pages moved in the private for the ACTIVITY TYPE
ACTIVITY TYPE	**DEMAND PAGING PAGE IN**	page-ins triggered to satisfy address space requests for pages that were not in central storage
	DEMAND PAGING PAGE OUT	central storage page in an address space was paged out so that a page-in request could be satisfied
	VIO PAGE IN	VIO pages brought into central storage
	VIO PAGE OUT	VIO pages paged out of central storage
	SWAP PAGE IN	page-ins into central storage because an address space was swapped in
	SWAP PAGE OUT	page-outs that occurred because an address space was swapped out
	TOTAL PAGING	total paging activity
	COMMON AREA RECLAIMS	COMMON AREA RECLAIMS
	RCL/MIN	pages reclaimed per minute in CSA and PLPA
	PCT	percentage of total reclaims that were common area reclaims
	PRIVATE AREA RECLAIMS (NON-VIO)	PRIVATE AREA RECLAIMS (NON-VIO)
	RCL/MIN	pages reclaimed per minute in the private area
	PCT	percentage of total reclaims that handled in the address space

Table 103 Field descriptions for the Paging Activity section (part 2 of 2)

Field	Description	
ACTIVITY TYPE (*continued*)	**VIO RECLAIMS**	
	RCL/MIN	pages reclaimed per minute due to a VIO page request that was satisfied without a page I/O (handled by an explicit VIO reclaim interface)
	PCT	percentage of total reclaims that were VIO reclaims
	TOTAL RECLAIMS	
	RCL/MIN	total reclaims per minutes
	PCT	total percentage of all reclaims (by definition, 100%)
PAGE MOVEMENT	**PAGES MOVED**	total number of pages that were moved between below the 16-MB line and above the 16-MB line in central storage
	PG/SEC	rate per second at which pages were moved between below the 16-MB line and above the 16-MB line in central storage

Expanded Storage Measures section field descriptions

Each field in the Expanded Storage Measures section of the System Resources Manager Report is described in Table 104.

Table 104 Field descriptions for the Expanded Storage Measures section (part 1 of 2)

Field	Description	
EXPANDED STORAGE ACTIVITY	**PAGES TO EXPANDED STORAGE**	total rate per minute of pages sent to expanded storage for paging and swapping requests
	MIGRATION RATE TO AUX. STORAGE	total rate per minute of pages migrated from expanded storage to auxiliary storage
	AVAILABLE ES FRAMES	
	AVERAGE	average number of page frames actually available in expanded storage for the reporting interval
	MINIMUM	minimum number of page frames actually available in extended storage for the reporting interval
	MAXIMUM	maximum number of page frames actually available in expanded storage for the reporting interval

Table 104 Field descriptions for the Expanded Storage Measures section (part 2 of 2)

Field	Description	
EXPANDED STORAGE ACTIVITY *(continued)*	**SYSTEM HIGH UIC**	
	AVERAGE	average system high Unreferenced Interval Count (UIC) in seconds (time that a central storage frame has not been referenced) for the reporting interval
	MINIMUM	minimum system high Unreferenced Interval Count (UIC) in seconds (time that a central storage frame has not been referenced) for the reporting interval
	MAXIMUM	maximum system high Unreferenced Interval Count (UIC) in seconds (time that a central storage frame has not been referenced) for the reporting interval
	MIGRATION AGE	
	AVERAGE	average length of time that a page remains unreferenced in expanded storage before being migrated to auxiliary storage for the reporting interval, in seconds
	MINIMUM	minimum length of time that a page remains unreferenced in expanded storage before being migrated to auxiliary storage for the reporting interval, in seconds
	MAXIMUM	maximum length of time that a page remains unreferenced in expanded storage before being migrated to auxiliary storage for the reporting interval, in seconds
EXPANDED STORAGE CONFIGURATION	**INSTALLED PAGE FRAMES**	total number of page frames installed in expanded storage
	ONLINE PAGE FRAMES	total number of page frames in expanded storage that are online to the system

Table 105 describes each field in the SRM Data section of the System Resources Manager Report.

Table 105 Field descriptions for the SRM Data section

Field	Description
SQA LOW EVENT COUNT	number of times that SQA low threshold was reached; an SQALOW SYSEVENT was issued
AVAILABLE QUEUE LOW EVENT COUNT	number of times that available frame queue dropped below the low threshold; an AVQLOW SYSEVENT was issued
ENQUEUE HOLD SWAPOUT EVENT COUNT	number of times that an owner of a resource in contention was swapped out of storage
MAX PAGE-INS/SEC	maximum number of page-ins per second; this includes demand page-ins, swap page-ins, and VIO page-ins
MAX SWAPS/MIN	maximum number of swaps per minute
AVG PAGES/SWAP OUT	average number of pages swapped out for each swap-out
AVG PAGES/SWAP IN	average number of pages swapped in for each swap-in
SRM OPTIONS MEMBER FROM PARMLIB	name of the member in SYS1.PARMLIB that was used for the SRM options
CPU OVERLOAD PERCENTAGE	percentage of time that the CPU was observed to be overloaded This figure is determined by measuring whether the **CCVUTILP** field of the CCT contains a value greater than 100.
I/O OVERLOAD PERCENTAGE	if I/O load balancing is active, percent of time that any logical path block was overutilized The LPB thresholds are user-specifiable in the IEAOPT member of SYS1.PARMLIB.
I/O UNDERLOAD PERCENTAGE	if I/O load balancing is active, percent of time that any logical path block was underutilized The LPB thresholds are user-specifiable in the IEAOPT member of SYS1.PARMLIB.
CPU IMBALANCE PERCENTAGE	if CPU load balancing is active, percent of time that the CPUs were out of balance
STORAGE IMBALANCE PERCENTAGE	if storage balancing is active, percent of time that storage utilization was out of balance
TERMINAL I/O SWAPS PERCENTAGE DUE TO OUTPUT	percentage of terminal swaps out of the total TERMINAL I/O WAITS

Swapping Measures section

In interpreting the Swapping Measures section of Figure 88, columns represent swap reasons and rows represent swap types. While swap reasons are mutually exclusive, swap types can overlap or otherwise interact:

- The row labeled LOG SWAP EFFECTIVE is a subset of LOG SWAP TOTAL.
- AUX STOR DIRECT and AUX STOR VIA TRANSITION are added together to produce AUX STOR TOTAL.
- LOG SWAP EFFECTIVE and EXPANDED STORAGE EFFECTIVE are added together to produce LOG SWAP +EXT STOR EFFECTIVE.
- The TOTAL row shows the sum of LOG SWAP EFFECTIVE, AUX STOR TOTAL, and EXTENDED STORAGE TOTAL.

The TOTAL values shown under each column heading represent each swap type percentage of that column reason. The totals for each row in the **TOTAL** column represent that swap type percentage of all swaps.

Figure 88 on page 523 shows the two pages that contain the Swapping Measures section of the SRM Report.

Figure 88 SRM Report, Swapping Measures section (part 1 of 2)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                      SYSTEM RESOURCES MANAGER REPORT                     RPTSEQ     4 PAGE     10
 BMC SOFTWARE, INC.                                               XYZ COMPANY                               REPORT DATE: DD MMM YY 11.55
 ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00             WORLDWIDE HEADQUARTERS                            SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/30/5,393/7.5  240-3/30/450/7.5  70-1/30/5,393/7.5  71-1/30/450/7.5
 ------------------------------------------------------ SWAPPING MEASURES  -------------------------------------------------------
                                                        CENTRAL
                        TERMINAL    LONG    AUX STOR  STORAGE   DETECTED                                     UNI-     TRANS
 SWAP TYPE   MEASURE    I/O WAIT    WAIT    SHORTAGE  SHORTAGE    WAIT    REQUEST   ENQUEUE   EXCHANGE  LATERAL   NONSWAP
 ---------- ----------  --------  --------  --------  --------  --------  --------  --------  --------  --------  --------
 LOG SWAP   COUNT          15351     11402         0         0      3575         0         0         2        50         0
 EFFECTIVE  RATE/MIN       34.11     25.34       .00       .00      7.94       .00       .00       .00       .11       .00
            % TOTAL        99.29     99.63       .00       .00     99.69       .00       .00    100.00    100.00       .00
 LOG SWAP   COUNT          15461     11444         0         0      3586         0         0         2        50         1
 TOTAL      RATE/MIN       34.36     25.43       .00       .00      7.97       .00       .00       .00       .11       .00
            % TOTAL       100.00    100.00       .00       .00    100.00       .00       .00    100.00    100.00      1.05
 AUX STOR   COUNT              0         0         0         0         0         0         0         0         0        94
 DIRECT     RATE/MIN         .00       .00       .00       .00       .00       .00       .00       .00       .00       .21
            % TOTAL          .00       .00       .00       .00       .00       .00       .00       .00       .00     98.95
 AUX STOR   COUNT            110        42         0         0        11         0         0         0         0         1
 VIA        RATE/MIN         .24       .09       .00       .00       .02       .00       .00       .00       .00       .00
 TRANSITION % TOTAL          .71       .37       .00       .00       .31       .00       .00       .00       .00      1.05
 AUX STOR   COUNT            110        42         0         0        11         0         0         0         0        95
 TOTAL      RATE/MIN         .24       .09       .00       .00       .02       .00       .00       .00       .00       .21
            % TOTAL          .71       .37       .00       .00       .31       .00       .00       .00       .00    100.00
 EXPANDED   COUNT              0         0         0         0         0         0         0         0         0         0
 STORAGE    RATE/MIN         .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 DIRECT     % TOTAL          .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 EXPANDED   COUNT              0         0         0         0         0         0         0         0         0         0
 STORAGE    RATE/MIN         .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 EFFECTIVE  % TOTAL          .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 EXPANDED   COUNT              0         0         0         0         0         0         0         0         0         0
 STORAGE    RATE/MIN         .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 TOTAL      % TOTAL          .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 MIGRATED   COUNT              0         0         0         0         0         0         0         0         0         0
 FROM       RATE/MIN         .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 EXP STOR   % TOTAL          .00       .00       .00       .00       .00       .00       .00       .00       .00       .00
 LOG SWAP   COUNT          15351     11402         0         0      3575         0         0         2        50         0
 +EXP STOR  RATE/MIN       34.11     25.34       .00       .00      7.94       .00       .00       .00       .11       .00
 EFFECTIVE  % TOTAL        99.29     99.63       .00       .00     99.69       .00       .00    100.00    100.00       .00
 TOTAL      COUNT          15461     11444         0         0      3586         0         0         2        50        95
            RATE/MIN       34.36     25.43       .00       .00      7.97       .00       .00       .00       .11       .21
            % GRND TOT     50.46     37.35       .00       .00     11.70       .00       .00       .01       .16       .31
```

(continued on next page)

Figure 88 SRM Report, Swapping Measures section (part 2 of 2)

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  SYSTEM RESOURCES MANAGER REPORT                    RPTSEQ     4 PAGE     11
BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 11.55
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/30/5,393/7.5  240-3/30/450/7.5  70-1/30/5,393/7.5  71-1/30/450/7.5
------------------------------------------------------  SWAPPING MEASURES  -----------------------------------------------------
                                                                  OMVS      OMVS
                        IMPROVE   IMPROVE   USER OUT    APPC      INPUT    OUTPUT
SWAP TYPE   MEASURE     CS USAGE  PAGE RT   TOO LONG    WAIT      WAIT      WAIT     TOTAL
---------- ----------   --------  --------  --------  --------  --------  --------  --------
LOG SWAP   COUNT               0         0         0         0         0         0     30380
EFFECTIVE  RATE/MIN          .00       .00       .00       .00       .00       .00     67.51
           % TOTAL           .00       .00       .00       .00       .00       .00     99.46
LOG SWAP   COUNT               0         0         0         0         0         0     30544
TOTAL      RATE/MIN          .00       .00       .00       .00       .00       .00     67.88
           % TOTAL           .00       .00       .00       .00       .00       .00     99.69
AUX STOR   COUNT               0         0         0         0         0         0        94
DIRECT     RATE/MIN          .00       .00       .00       .00       .00       .00       .21
           % TOTAL           .00       .00       .00       .00       .00       .00       .31
AUX STOR   COUNT               0         0         0         0         0         0       164
VIA        RATE/MIN          .00       .00       .00       .00       .00       .00       .36
TRANSITION % TOTAL           .00       .00       .00       .00       .00       .00       .54
AUX STOR   COUNT               0         0         0         0         0         0       258
TOTAL      RATE/MIN          .00       .00       .00       .00       .00       .00       .57
           % TOTAL           .00       .00       .00       .00       .00       .00       .84
EXPANDED   COUNT               0         0         0         0         0         0         0
STORAGE    RATE/MIN          .00       .00       .00       .00       .00       .00       .00
DIRECT     % TOTAL           .00       .00       .00       .00       .00       .00       .00
EXPANDED   COUNT               0         0         0         0         0         0         0
STORAGE    RATE/MIN          .00       .00       .00       .00       .00       .00       .00
EFFECTIVE  % TOTAL           .00       .00       .00       .00       .00       .00       .00
EXPANDED   COUNT               0         0         0         0         0         0         0
STORAGE    RATE/MIN          .00       .00       .00       .00       .00       .00       .00
TOTAL      % TOTAL           .00       .00       .00       .00       .00       .00       .00
MIGRATED   COUNT               0         0         0         0         0         0         0
FROM       RATE/MIN          .00       .00       .00       .00       .00       .00       .00
EXP STOR   % TOTAL           .00       .00       .00       .00       .00       .00       .00
LOG SWAP   COUNT               0         0         0         0         0         0     30380
+EXP STOR  RATE/MIN          .00       .00       .00       .00       .00       .00     67.51
EFFECTIVE  % TOTAL           .00       .00       .00       .00       .00       .00     99.16
TOTAL      COUNT               0         0         0         0         0         0     30638
           RATE/MIN          .00       .00       .00       .00       .00       .00     68.08
           % GRND TOT        .00       .00       .00       .00       .00       .00    100.00
```

Swapping Measures section column field descriptions

Table 106 describes each column field in the Swapping Measures section of the System Resources Manager Report.

Table 106 Column field descriptions for the Swapping Measures section

Field	Description
TERMINAL I/O WAIT	number of swap-outs that occurred because a terminal was waiting for input or output buffers, but was not a candidate for logical swapping
LONG WAIT	number of swap-outs that occurred because the user address space was issued a WAIT,LONG=YES
AUX STOR SHORTAGE	number of swap-outs that occurred because 70% of all local page slots are allocated
CENTRAL STORAGE SHORTAGE	number of swap-outs that occurred due to a shortage of central storage page frames
DETECTED WAIT	number of swap-outs that occurred because the user address space has been in a wait greater than the SRM-specified threshold of two seconds
REQUEST	number of swap-outs that occurred due to the user address space being set nonswappable by PPT or by the CONFIG STOR,OFFLINE command
ENQUEUE	number of swap-outs that occurred because a user address space was enqueued on a required system resource
EXCHANGE	number of swap-outs that occurred to allow a user with a greater workload level to be swapped in
UNILATERAL	number of swap-outs that occurred because the target multiprogramming level for the domain has been exceeded
TRANS NOSWAP	number of swap-outs that occurred due to a user issuing the TRANSWAP SYSEVENT
IMPROVE CS USAGE	number of swaps that occurred because the SRM is attempting to improve Central Storage (CS) usage
IMPROVE PAGE RT	number of swaps that occurred because the SRM is attempting to improve the paging rate of the entire system
USER OUT TOO LONG	number of swaps that occurred because a user is swapped-out too long
APPC WAIT	number of swaps that occurred because the SRM is attempting to improve Advanced Program-to-Program Communication (APPC / MVS)
OMVS INPUT WAIT	number of swap-outs that occurred because the OMVS address space was waiting for terminal input
OMVS OUTPUT WAIT	number of swap-outs that occurred because the OMVS address space was waiting for terminal output
IN-REAL	number of swap-outs that occurred due to recovery of real storage frames from a swappable address space
TOTAL	number of swap-outs for a specific placement

Swapping Measures section row field descriptions

Table 107 describes each row field in the Swapping Measures section of the System Resources Manager Report.

Table 107 Row field descriptions for the Swapping Measures section

Field	Description
LOGICAL SWAP EFFECTIVE	logical swap candidates that were logically swapped but not physically swapped
LOGICAL SWAP TOTAL	logical swap candidates that were logically swapped
AUX STORAGE DIRECT	swap requests directed to auxiliary storage that went directly to auxiliary storage
AUX STORAGE VIA TRANSITION	swap requests that were directed to auxiliary storage but either were logically swapped or were swapped to expanded storage before being placed on auxiliary storage
AUX STORAGE TOTAL	swap requests that were either directly or indirectly placed on auxiliary storage
EXPANDED STORAGE DIRECT	physical swaps that were directed to expanded storage
EXPANDED STORAGE EFFECTIVE	physical and logical swaps directed to expanded storage that were not migrated to auxiliary storage
EXPANDED STORAGE TOTAL	logical swap candidates that were directly or indirectly placed on the expanded storage
MIGRATED FROM EXP STORAGE	swap requests that migrated from the expanded storage to auxiliary storage
LOGICAL SWAP + EXP STORAGE EFFECTIVE	swap requests that remained logically swapped or were swapped to expanded storage; sum of Logical Swap Effective and Expanded Storage Effective
TOTAL	total swap requests

Tabular Subreport

The Tabular Subreport is produced by using the Analyzer GRAPH TYPE=TAB control statement. Up to 12 measures can be specified in the GRAPH statement (see "GRAPH" on page 259); if more than 12 are specified, only the first 12 measures are used.

Each measure that is selected is represented by a column on the report. The measures are printed in the same sequence as they appear in the GRAPH command, allowing you to group corresponding measures together.

The Extractor statements required for the Tabular Subreport are dependent upon the information to be graphed. (See "Record types" on page 45 for information about the record types generated by each Extractor statement and a description of the specific type of data collected by each Extractor statement.)

One line of values is produced for each date and time interval specified on the GRAPH command. If the default is used, one line is produced for each record interval encountered in the data.

A sample of the Tabular Subreport is shown in Figure 89.

Figure 89 Tabular Subreport

```
PRODUCED BY CMF ANALYZER (v.r.mm)                         TABULAR   SUBREPORT                                   RPTSEQ     5 PAGE      8
 BMC SOFTWARE, INC.                                            XYZ COMPANY                                       REPORT DATE: DD MMM YY 17.14
 ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00             WORLDWIDE HEADQUARTERS                                  SYSTEM ID: SJSE     Z v.rr.n


   CPU ENVIRONMENT REQUESTED:  ALL CPU'S
   DATE      TIME     BTCH-MAX BTCH-MIN  NTU-MAX  NTU-MIN  STC-MAX  STC-MIN TSO-COMM

 10JUN YY 09.00.00       1.00     0.00    36.00    34.00   110.00   107.00    61.00
          09.15.00       3.00     0.00    53.00    35.00   109.00   105.00   800.00
          09.30.00       2.00     2.00    67.00    53.00   114.00   108.00   804.00
          09.45.00       3.00     2.00    80.00    66.00   116.00   114.00   790.00
          10.00.00       7.00     2.00    91.00    80.00   120.00   111.00   370.00
          10.15.00       4.00     3.00    93.00    90.00   121.00   119.00   205.00
          10.30.00       8.00     3.00    97.00    92.00   122.00   117.00   148.00
          10.45.00       7.00     5.00    97.00    97.00   128.00   119.00    18.00
          11.00.00       7.00     6.00    98.00    97.00   126.00   123.00    20.00
          11.15.00       8.00     6.00    98.00    96.00   126.00   124.00    18.00
          11.30.00       9.00     6.00    99.00    97.00   128.00   122.00    46.00
          11.45.00      12.00     7.00    99.00    98.00   130.00   123.00     9.00
          12.00.00      11.00     7.00    98.00    47.00   129.00   125.00     0.00
          12.15.00       8.00     7.00    47.00    46.00   129.00   128.00     0.00
          12.30.00       8.00     7.00    47.00    46.00   129.00   128.00     0.00
          12.45.00      10.00     7.00    48.00    46.00   129.00   126.00     0.00
          13.00.00      10.00     8.00    49.00    47.00   127.00   124.00     0.00
          13.15.00      10.00     9.00    48.00    46.00   125.00   124.00     0.00
          13.30.00      10.00     9.00    48.00    45.00   126.00   124.00     0.00
          13.45.00      10.00     8.00    47.00    45.00   127.00   124.00     0.00
          14.00.00       9.00     8.00    47.00    46.00   127.00   126.00     0.00
          14.15.00      13.00     8.00    47.00    47.00   127.00   121.00     0.00
          14.30.00      13.00    10.00    47.00    44.00   125.00   122.00     0.00
          14.45.00      10.00     8.00    44.00    41.00   127.00   125.00     0.00
          15.00.00       8.00     7.00    41.00    33.00   128.00   127.00     0.00
          15.15.00       7.00     7.00    33.00    32.00   127.00   127.00     0.00
          15.30.00       7.00     7.00    32.00    32.00   127.00   127.00     0.00
          15.45.00       7.00     7.00    33.00    32.00   128.00   126.00     0.00
```

Tabular Subreport field descriptions

Table 108 describes each field in the Tabular Subreport.

Table 108 Field descriptions for the Tabular Subreport

Field	Description
CPU ENVIRONMENT REQUESTED	CPU number or ALL selected in the GRAPH control statement
DATE TIME	date and time intervals selected
nnn-nnn	remaining columns show values for specific measures See Appendix C, "Measure and trace values," for information about specific measures.

Trace Report

The Trace Report lists the data in the trace records generated by the Extractor. The data for this report is obtained from CMF type 240-18 user records that are generated by using the Extractor TRACE control statement (see "TRACE" on page 183). The report is produced by using the Analyzer TRACE control statement (see "TRACE" on page 312).

The Analyzer TRACE control statement causes the data in the Extractor trace records to be formatted. Unlike other Analyzer reporters, the trace formatter has no data collection phase. The trace records are formatted and written out as they are read. For this reason, a separate CMXTRACE sysout data set is dynamically allocated by the Analyzer. However, the ddname //CMXTRACE can be used to direct the output from the Trace Report.

The TITLE and LOCATION parameter values from the Analyzer HEADERS control statement do not appear on the Trace Report. The Analyzer SUBTITLE control statement parameter values are used.

The Trace Report terminates with one of the following messages:

- `total trace entries processed ===> x,xxx,xxx`

  This message shows how many trace entries were formatted.

- `no trace entries selected for formatting`

  Although trace records were present, none of the records found matched the criteria specified on the Analyzer TRACE control statement.

- `no trace records encountered`

  No CMF type 240-18 user records were encountered in the EXTDATA data set. If this message is issued unexpectedly, make sure that the record type and monitor ID specified for this Analyzer run match the corresponding parameters in the extraction run.

NOTE

The default value for CMF MONITOR user type records is 240; a different value could have been specified at the SMFRECID parameter of the Extractor report control statement.

A sample of the Trace Report is shown in Figure 90.

Figure 90 Trace Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)              C M F  T R A C E  R E P O R T                        MODE = ALL   PAGE     1
 BMC SOFTWARE, INC.                                                                                   REPORT DATE: DD MMM YY 16.29

        DATE         TIME       ID  MODE   ADDR                DATA
 DD MMM YY         12:01:02     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:02     20  SRM   0E90FD               C3D4C4E2 C1E4D7F9 07DC0014 7FFD8800    *CMDSAUP9........*
                   12:01:02     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:02     21  SRM   0E90FD               C3D4C4E2 C1E4D7F9 07DC0015 7FFD8800    *CMDSAUP9........*
                   12:01:02     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:02     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:02     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:02     20  SRM   0E90FD               D4C9D4C3 D6D54040 001A0014 7FFD8800    *MIMCON  ........*
                   12:01:02     21  SRM   0E90FD               D4C9D4C3 D6D54040 001A0015 7FFD8800    *MIMCON  ........*
                   12:01:03     20  SRM   0E90FD               E2D3E2F0 40404040 003C0014 7FFD8800    *SLS0    ........*
                   12:01:03     21  SRM   0E90FD               E2D3E2F0 40404040 003C0015 7FFD8800    *SLS0    ........*
                   12:01:03     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:03     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:03     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:03     20  SRM   0E90FD               D4C9D4C3 D6D54040 001A0014 7FFD8800    *MIMCON  ........*
                   12:01:03     21  SRM   0E90FD               D4C9D4C3 D6D54040 001A0015 7FFD8800    *MIMCON  ........*
                   12:01:03     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               C3C1E3C1 D3D6C740 00230014 7FFD8800    *CATALOG ........*
                   12:01:04     21  SRM   0E90FD               C3C1E3C1 D3D6C740 00230015 7FFD8800    *CATALOG ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*
                   12:01:04     20  SRM   0E90FD               D4C9D4C7 D9404040 00190014 7FFD8800    *MIMGR   ........*
                   12:01:04     21  SRM   0E90FD               D4C9D4C7 D9404040 00190015 7FFD8800    *MIMGR   ........*

                   TOTAL TRACE ENTRIES PROCESSED ===>        50
```

Trace Report field descriptions

Table 109 describes each field in the Trace Report.

Table 109 Field descriptions for the Trace Report

Field	Description
DATE	date entries being formatted were written by the Extractor Date is repeated if it changes or if a page eject occurs. This date is affected by the START and STOP parameters in the Analyzer TRACE control statement.
TIME	time the entry being formatted was traced by the Extractor; repeated for each entry This time is affected by the START and STOP parameters in the TRACE control statement, but is taken from the time stamp on the record, indicating when it was written.
ID	ID of the entry as specified in the Extractor TRACE control statement or on the CMFTRACE macro in the user-written SRB routine
MODE	mode in which this entry was traced: SRB — trace made from the user-written SRB routine SRM — trace resulted from a SYSEVENT ID specified in the Extractor TRACE control statement
ADDR	virtual address storage location from which the data was traced
DATA	actual data traced; can be from 1 to 112 characters long Each line for a trace entry formats up to 16 characters. Data is displayed in hexadecimal format, followed by an EBCDIC translation.

TSO Command Summary Report

The TSO Command Summary Report graphically summarizes TSO command usage. The frequency and average response time of each command are listed and plotted individually in this report. The data is obtained from type 240-20 user records, which are generated by the Extractor TSODATA control statement (see "TSODATA" on page 194). The TSO Command Summary Report is produced by using the Analyzer TSOPERF control statement (see "TSOPERF" on page 315).

A sample of the TSO Command Summary Report is shown in Figure 91 on page 533. In the TSODATA Extractor control statement at the CMDS parameter, either proper command names or TSO aliases (command short form) can be defined for sampling. For example, you can list either EDIT or E for the EDIT command.

The number of commands reported on is controlled by the LIMIT parameter of the TSOPERF control statement. When the number of commands is limited, the commands displayed are those most frequently used during the measurement interval. When the number of commands is not limited, all commands sampled are listed.

When a TSO user enters a command for which a response is never received, the Extractor stops response timing at five minutes. In this way, infinite response times are not collected.

TSO commands executed under SPF do not issue SYSEVENT ZERO; therefore, these commands are not timed under SPF. However, the EDIT and TEST subcommands under EDIT and TEST modes are counted and monitored. If FSE is being used and the subcommands are stacked, FSE replaces the last character of the subcommand with the character Z.

When executing programs in foreground, response times tend to be high. For this reason, use a value between 5 and 15 for the scale on this report so that the graphed data is more meaningful.

Figure 91 is an example of the TSO Command Summary Report.

Figure 91 TSO Command Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    TSO COMMAND SUMMARY REPORT                       RPTSEQ     9 PAGE     12
BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 11.55
ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00            WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/30/5,393/7.5  240-20/12/2,171/3
                                                    C O M M A N D  S U M M A R Y
          TSO          COMMAND       RESPONSE        0             USAGE         2,054       0.0          RESPONSE        11.22
        COMMAND         USAGE          TIME          +--+--+--+--+--+--+--+--+--+--+--+      +--+--+--+--+--+--+--+--+--+--+--+
        ALLOCATE       2,054           0.78          .*********************************      .*I                               .
        ATTRIB           100           0.12          .**                              .      . I                               .
        CONSOLE            3           0.23          .                                .      .*I                               .
        DEFINE             1          11.22          .                                .      .*I*******************************
        DSN               23           4.01          .                                .      .*I**********                     .
        DYNASTEP           1           0.11          .                                .      . I                               .
        END               13           4.18          .                                .      .*I**********                     .
        EXEC             385           0.49          .******                          .      .*I                               .
        EXECUTIL          11           0.12          .                                .      . I                               .
        FREE             337           0.11          .*****                           .      . I                               .
        IKJEFF76          64           0.31          .*                               .      .*I                               .
        IPCS               1           5.15          .                                .      .*I*************                  .
        ISPEXEC            1           0.20          .                                .      .*I                               .
        ISPF               8           0.70          .                                .      .*I                               .
        LISTALC            5           0.88          .                                .      .*I*                              .
        LOCATE             1           0.20          .                                .      .*I                               .
        PDF               73           2.12          .*                               .      .*I****                           .
        PMGLAUTH          12           1.80          .                                .      .*I***                            .
        PROFILE          172           0.10          .***                             .      . I                               .
        SYSPROG            1           0.44          .                                .      .*I                               .
        TAOINQ             1           0.52          .                                .      .*I                               .
        TEST              20 (1,862)   0.70          .                                .      .*I                               .
        TESTXA             2           0.08          .                                .      . I                               .
                                                     +--+--+--+--+--+--+--+--+--+--+--+      +--+--+--+--+--+--+--+--+--+--+--+
        TOTAL/AVG      3,289           0.68          0             USAGE         2,054       0.0          RESPONSE        11.22
```

TSO Command Summary Report field descriptions

Table 110 describes each field in the TSO Command Summary Report.

Table 110 Field descriptions for the TSO Command Summary Report

Field	Description
TSO COMMAND	all major TSO commands issued, or commands most frequently used during the measurement interval
COMMAND USAGE	total number of times command was used; for EDIT and TEST, the number of times a subcommand was issued is shown in parentheses
RESPONSE TIME	average command response time; defined internally as SYSEVENT ZERO (enter) to TPUT (screen); line speed and other external activities are not included; reported in seconds per command
USAGE	usage count for each command plotted on a self-scaling bar graph
RESPONSE	response time for each command plotted on a bar graph; the graph scales to the value specified, to the highest value encountered (if SCALE=999 is specified), or to the default of the TSOPERF statement
TOTAL/AVG	total count of all commands reported; average response time for all commands is reported

TSO Interval Summary Report

The TSO Interval Summary Report graphically summarizes TSO workload on an interval basis. The average number of TSO users and either the average response time or the average TSO overhead for the interval are plotted.

The data for this report is obtained from type 20 user records and type 1 CPU records. Type 20 records are generated by the Extractor TSODATA control statement (see "TSODATA" on page 194). Type 1 CPU records are generated by the Extractor CPU control statement (see "CPU" on page 136). This report is produced by using the Analyzer TSOPERF control statement (see "TSOPERF" on page 315).

NOTE

If no type 20 records are encountered, this report is not produced.

An example of the TSO Interval Summary Report is shown in Figure 92.

Figure 92 TSO Interval Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     TSO INTERVAL SUMMARY REPORT                      RPTSEQ    31 PAGE    340
 BMC SOFTWARE, INC.                                            XYZ COMPANY                              REPORT DATE: DD MMM YY 13.40
 ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00             WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-1/28/25.1K/7  240-20/28/25.2K/7
                                                     I N T E R V A L  S U M M A R Y
  INTERVAL   INTERVAL    AVERAGE   AVERAGE   % CPU     0.0       AVERAGE USERS        30.6      0.0          RESPONSE          2.07
    DATE       TIME       USERS    RESPONSE   TSO       +--+--+--+--+--+--+--+--+--+--+--+       +--+--+--+--+--+--+--+--+--+--+--+
 10 JUN 03   09:00:00      16.1      1.11      0.6      .*****************         I     .       .*****I************              .
             09:15:00      19.1      0.57      0.6      .*********************     I     .       .*****I***                       .
             09:30:00      20.7      0.25      0.4      .**********************    I     .       .**** I                          .
             09:45:00      18.8      0.14      0.3      .********************      I     .       .**   I                          .
             10:00:00      17.8      0.07      0.2      .*******************       I     .       .*    I                          .
             10:15:00      20.6      0.25      0.5      .**********************    I     .       .**** I                          .
             10:30:00      25.9      0.20      0.8      .**************************I*    .       .***  I                          .
             10:45:00      27.8      0.48      1.0      .**************************I***  .       .*****I**                        .
             11:00:00      28.6      0.77      0.4      .**************************I**** .       .*****I******                    .
             11:15:00      29.2      0.32      0.2      .**************************I**** .       .*****I                          .
             11:30:00      30.6      0.75      1.0      .**************************I******       .*****I******                    .
             11:45:00      30.0       ---      0.3      .**************************I*****.       .     I                          .
             12:00:00      30.0      0.28      0.7      .**************************I*****.       .**** I                          .
             12:15:00      30.0      0.22      0.6      .**************************I*****.       .**** I                          .
             12:30:00      29.0      0.75      1.1      .**************************I**** .       .*****I******                    .
             12:45:00      28.8      1.07      0.5      .**************************I**** .       .*****I***********               .
             13:00:00      29.0       ---      0.4      .**************************I**** .       .     I                          .
             13:15:00      28.6      0.76      1.0      .**************************I**** .       .*****I******                    .
             13:30:00      27.8      1.66      1.6      .**************************I***  .       .*****I********************      .
             13:45:00      27.7      0.63      1.6      .**************************I***  .       .*****I****                      .
             14:00:00      26.0      0.57      1.3      .**************************I*    .       .*****I***                       .
             14:15:00      25.1      0.31      0.7      .**************************I     .       .*****I                          .
             14:30:00      26.3      0.86      1.4      .**************************I*    .       .*****I********                  .
             14:45:00      25.2      0.18      0.8      .**************************I     .       .***  I                          .
             15:00:00      24.1      0.11      0.3      .**************************I     .       .**   I                          .
             15:15:00      21.7      2.07      0.3      .***********************   I     .       .*****I***************************
             15:30:00      20.6      0.09      0.1      .**********************    I     .       .*    I                          .
             15:45:00      21.2      0.19      0.5      .***********************   I     .       .***  I                          .
                                                        +--+--+--+--+--+--+--+--+--+--+--+       +--+--+--+--+--+--+--+--+--+--+--+
             AVERAGES      25.2      0.40      0.7     0.0       AVERAGE USERS        30.6      0.0          RESPONSE          2.07
```

Fields and data columns are described in the legend of the figure. Dashes (---) are printed in the report to indicate values that were not calculated because of missing record types or zero samples for an interval.

In Figure 92 on page 534, the two graphs at the right show that at the sample interval beginning at 0730 hours, the number of users was 113.7, which was above the average number of users of 101.0. The average response time was 1.13, which was below the average of 1.24 seconds.

NOTE

When recording at relatively short intervals, the average response time for an interval could be skewed because of a few large response times. For example, a foreground compilation run during a five-minute recording interval could cause an unusually high average response time to be reported.

TSO Interval Summary Report field descriptions

Table 111 describes each field in the TSO Interval Summary Report.

Table 111 Field descriptions for the TSO Interval Summary Report

Field	Description
INTERVAL DATE	start date of measurement interval
INTERVAL TIME	start time of measurement interval
AVERAGE RESPONSE	average response time for all commands issued during measurement interval **Note**: This field shows the average response time for all commands issued and not the average response time for TSO as a whole. For example, for a particular CLIST this field shows the average response time for just the commands issued within that CLIST, rather than the average response time for the execution of the entire CLIST.
% CPU TSO	percentage of time during measurement interval that the CPU was busy for a TSO user (You can plot this value on a self scaling bar graph by specifying PLOT=TSO in the Analyzer TSOPERF control statement. The graph scales to highest value encountered. A column of I's marks the average TSO / CPU line.)
AVERAGE USERS	average number of active TSO users during the measurement interval This number of active TSO users is plotted on a self scaling bar graph, which is scaled to highest value encountered. A column of I's marks average number of active TSO users.
RESPONSE	average response time for all commands issued plotted on a self scaling bar graph; PLOT=RESPONSE must be specified in the Analyzer TSOPERF control statement The graph is scaled to the highest value encountered. A column of I's marks average response time.

TSO User Summary Reports

The TSO User Summary Report provides information about TSO usage by TSO logon ID, displaying it in these categories:

- swapping
- paging
- SRM usage

The TSO User Summary Report is produced by using the TSOUSER Analyzer control statement (see “TSOUSER” on page 317). The data for this report is obtained by using the TSODATA Extractor control statement with the parameter USERS=YES, the default, specified (see “TSODATA” on page 194). Report data is obtained from type 21 user records, which are generated only if USERS=YES is defined in the Extractor TSODATA control statement.

NOTE

This report reflects TSO activities of users logging off during the report interval. Transactions executed prior to this report period are recorded during the interval in which the user logs off. Information about swapping, paging, and SRM service consumption is also recorded in that interval.

If PCF is installed and the command accounting option is used, PCF resets certain system data fields used in the CMF type 21 record, so that the type 21 record is invalid.

A sample of the TSO User Summary Report is shown in Figure 93.

Figure 93 TSO User Summary Report

```
PRODUCED BY CMF ANALYZER (v.r.mm)                        TSO USER SUMMARY REPORT                                RPTSEQ    11 PAGE     14
 BMC SOFTWARE, INC.                                              XYZ COMPANY                                     REPORT DATE: DD MMM YY 11.55
 ACTL 10 JUN YY 09.00.00  10 JUN YY 16.00.00              WORLDWIDE HEADQUARTERS                                 SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  240-21/12/0/3
                                         **** S W A P P I N G ****     ******** P A G I N G ********    ********** S R M ***********
   TSO USER    TOTAL   AVERAGE   TOTAL   AVERAGE  PAGES/   PAGES/       PAGES  PAGE-INS/  PAGE-OUTS/     SERVICE  ABSORBTION  AVERAGE
             SESSIONS  SESSION   TRANS    SWAPS  SWAP-IN  SWAP-OUT     STOLEN    TRANS       TRANS         RATE       RATE     TRANS
   BAOBKS4       1      02:44      116     1.0    392.0    392.0        111.0      0.1         0.3       8292.6    8338.3     116.0
   BCVKSR3       1      06:10      342     2.0    437.0    437.0       1269.0      0.6         0.7       8860.9     123.0     342.0
   BMCGPA3       1      03:07      465     0.0      0.0      0.0          0.0      0.0         0.0       1209.3    1556.5     465.0
   BMCHVM1       1      00:01        5     0.0      0.0      0.0          0.0      0.0         0.0       1607.9    1607.9       5.0
   BMCHVM2       1      00:07        5     0.0      0.0      0.0          0.0      0.0         0.0       4143.1    4143.1       5.0
   BMCPXK2       1      03:50      404     7.0    435.0    435.0        684.0      0.4         0.3         46.4     946.8     404.0
   BMCQC01       1      00:57        9     1.0    505.0    505.0        593.0      2.7         0.2       9565.2    9866.1       9.0
   BMVBHM6       1      00:18       27     0.0      0.0      0.0          0.0      0.0         0.0       5115.3    5115.3      27.0
   BMVDLL3       1      02:02       33     0.0      0.0      0.0          0.0      0.0         0.0       4883.1    4883.1      33.0
   BMVGKC2       2      01:27      684     0.5    282.0    282.0        101.5      0.1         0.1       4239.0    4285.3     342.0
   BMVJES1       1      04:12      173     0.0      0.0      0.0        815.0      0.3         0.6       9354.1     243.3     173.0
   BMVSAU2       1      00:15      144     0.0      0.0      0.0          0.0      0.0         0.0        902.8     928.5     144.0
   BOLGBG3       2      00:26       64     1.0    509.0    509.0       4539.5     10.5         7.3       6692.2    6698.3      32.0
   BOLKGB1       1      01:31       74     0.0      0.0      0.0          0.0      0.0         0.0        213.2     328.5      74.0
   BOLRSD5       1      02:54       13     1.0     73.0     73.0        912.0      5.6         2.5       9867.3    9867.3      13.0
   BOLSMR2       2      00:33       66     1.0    494.5    494.5        264.5      1.9         0.1       6979.4    6981.9      33.0
   BOLSMR3       2      00:31       94     0.0      0.0      0.0        308.0      1.6         0.3       5845.0    6259.1      47.0
   BSDMXM1       1      00:01        3     0.0      0.0      0.0          0.0      0.0         0.0       9347.5    9347.5       3.0
   BSDMXM2       1      00:00        3     0.0      0.0      0.0          0.0      0.0         0.0       4736.7    4736.7       3.0
   PPE2          1      04:10      141     0.0      0.0      0.0        986.0      0.3         1.1       8957.6    9045.9     141.0
   ROCPXA1       1      01:48       66     1.0    149.0    149.0       1891.0      7.4         3.7       9233.0    9233.0      66.0
     TOTAL      25               2,931
   AVERAGE/SESSION                         0.7    407.1    407.1        707.5      0.7         0.5       5529.9    5547.8     117.2
```

The number of users reported on is controlled by the LIMIT parameter in the Analyzer TSOUSER control statement. If the number of users is limited, the data displayed is for those users with the most SRM service consumption.

TSO User Summary Report field descriptions

Table 112 describes each field in the TSO User Summary Report.

Table 112 Field descriptions for the TSO User Summary Report

Field	Description
TSO USER	logon ID of any TSO user who logged off during the measurement interval, displayed in ascending alphanumeric order
TOTAL SESSIONS	total number of TSO sessions for user during the measurement interval
AVERAGE SESSIONS	average length of the user TSO session displayed in hours:minutes This value is 00:00 if user is logged on for under one minute.
TOTAL TRANS	total number of TSO transactions for all this user’s TSO sessions
AVERAGE SWAPS	user’s average number of address space swap sequences per session A swap sequence is one address space swap-out and swap-in.
PAGES/SWAP-IN	user’s average number of pages per swap-in
PAGES/SWAP-OUT	user’s average number of pages per swap-out
PAGES STOLEN	user’s average number of pages stolen per session
PAGE-INS/TRANS	user’s average number of page in operations per transaction
PAGE-OUTS/TRANS	user’s average number of page out operations per transaction
SERVICE RATE	rate at which user consumed service while transactions were active but not necessarily in storage
ABSORPTION RATE	rate at which user consumed service while transactions were resident in central storage
AVERAGE TRANS	user’s average number of transactions per TSO session

Virtual Storage Activity Report

The Virtual Storage Activity Report is produced by using the VIRTSTOR Analyzer control statement (see “VIRTSTOR” on page 318). The data is obtained by using the VSMDATA Extractor control statement (see “VSMDATA” on page 200).

The Virtual Storage Activity Report is divided into five sections, which can be produced independently or concurrently. The report sections are as follows:

- **Virtual Storage Map**

  This section describes the static boundaries of system areas and the private region both above and below 16 MB. The map also includes values regarding space usage in the PLPA. (See Figure 94 on page 540 for an example of this report.) This map can be suppressed by including the MAP=NO parameter.

  Because some values are reset at IPL time, the virtual storage map is printed once for each IPL that occurred within the DATETIME range.

- **Common Area storage summary**

  This section reports on storage usage by CSA and SQA, both above and below 16 megabytes. The Summary Report also includes information about allocated storage, free storage, and allocated area size. CSA allocated storage information is detailed by storage key. (See Figure 95 on page 542 for an example of this report.)

- **Common Area storage detail**

  This section details CSA and SQA allocated storage information by subpool and, for CSA, by storage key. The Detail Report includes information only about storage below 16 MB. (See Figure 96 on page 544 for an example of this report.) This report is generated by the DETAIL=YES parameter.

- **Private Area storage summary**

  This section reports on storage usage within the private area, both above and below 16 megabytes. Includes a storage map of the private area, indicating allocated area boundaries and GETMAIN limits. Storage usage data is reported for the authorized region (LSQA, SWA, and subpools 229 and 230) and the user region (subpools 0 through 128, 251 and 252). Information about free storage, allocated storage, and allocated area size is reported for both regions in minimum, maximum, and average values. (See Figure 97 on page 546 for an example of this report.)

- **Private Area storage detail**

  This section reports private area storage usage by subpool. Minimum, maximum, and average values are reported for both user and authorized subpools. (See Figure 98 on page 548 for an example of this report.)

Storage addresses are reported as hexadecimal byte addresses. Storage area lengths are given in decimal K bytes (for values up to 9999 K) and in decimal megabytes (for values greater than 9999 K).

Date / time stamps, included for all minimum and maximum values, are in this format:

- in summary reports—(*ddmmm/hh.mm)*
- in detail reports—(*dd mmm yy / hh.mm.ss*)

Because virtual storage data is obtained by sampling, the minimum and maximum values reported might not represent the actual usage limits experienced during the reporting interval.

Virtual Storage Map section

An example of the Virtual Storage Map section is shown in Figure 94.

Figure 94 Virtual Storage Map section

```
 PRODUCED BY CMF ANALYZER (v.r.mm)                    VIRTUAL STORAGE ACTIVITY REPORT                    RPTSEQ     3 PAGE      5
 BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 16.53
 ACTL 10 JUN YY 16.33.43  10 JUN YY 16.49.53              WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  78-2/4/119/0.27
 ------------------------------------------------------  VIRTUAL STORAGE MAP  -----------------------------------------------------
              -------- BELOW 16M --------                                     -------- ABOVE 16M --------
     DESCRIPTION          START ADDRESS        LENGTH               DESCRIPTION          START ADDRESS           LENGTH
   ---------------        ---------------     --------            ---------------        ---------------        --------
     NUCLEUS                   00FC8000          221K          EXTENDED NUCLEUS                 01000000          7,377K
     SQA                       00E08000        1,792K          EXTENDED SQA                     01735000           17.6M
     PLPA                      00BBB000        2,356K          EXTENDED PLPA                    028C8000           51.1M
     FLPA                        N/A              0K           EXTENDED FLPA                    05BE0000             12K
     MLPA                      00BA9000          72K           EXTENDED MLPA                    05BE3000            636K
     CSA                       00800000        3,748K          EXTENDED CSA                     05C82000            200M
     PRIVATE AREA              00001000        8,188K          EXTENDED PRIVATE AREA            12500000           1755M
     PSA                       00000000           4K
                                                   BELOW 16M        ABOVE 16M           TOTAL
                                                  -----------      -----------         -------
                  PLPA INTERMODULE SPACE    -             7K              88K             95K
     PLPA SPACE REDUNDANT WITH MLPA/FLPA    -             3K              10K             12K
            MAXIMUM POSSIBLE USER REGION    -         7,992K            1745M          1,753M
```

Virtual Storage Map section field descriptions

Table 113 describes each field in the Virtual Storage Map section of the Virtual Storage Activity Report.

Table 113 Field descriptions for the Virtual Storage Map section

Field	Description
BELOW 16M – DESCRIPTION	named area
BELOW 16M – START ADDRESS	address of beginning of the named area; reported in hexadecimal format If an N/A is displayed in this column, the named area was not allocated space. This is space that was not occupied by either the load modules or the link pack directory.
BELOW 16M – LENGTH	length of the named area, in decimal K bytes
ABOVE 16M – DESCRIPTION	named area
ABOVE 16M – START ADDRESS	address of the beginning of the named area; reported in hexadecimal format If an N/A is displayed in this column, the named area was not allocated space. This space was not occupied by either the load modules or the link pack directory.
ABOVE 16M – LENGTH	length of the named area, in decimal K bytes
PLPA INTERMODULE SPACE	PLPA space not occupied by modules or the LPA directory The value for this field is determined by subtracting the link pack directory size and the sum of sizes of all PLPA modules from the total PLPA size.
PLPA SPACE REDUNDANT WITH MLPA/FLPA	space occupied by modules in PLPA that are never accessed because a module with the same name exists in MLPA or FLPA
MAXIMUM POSSIBLE USER REGION	maximum amount of space available in the user region This amount is below the size of the user region and is obtained at start time for the CMF MONITOR Extractor by subtracting the address of the lowest available block in the user region from the end address of the region.

Common Area storage summary section

An example of the Common Area storage summary section is shown in Figure 95.

Figure 95 Common Area storage summary

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    VIRTUAL STORAGE ACTIVITY REPORT                    RPTSEQ     3 PAGE      6
 BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 16.53
 ACTL 10 JUN YY 16.33.43  10 JUN YY 16.49.53             WORLDWIDE HEADQUARTERS                          SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  78-2/4/119/0.27
 --------------------------------------------------  COMMON AREA REPORT - SUMMARY  -------------------------------------------------
                                                        ALLOCATED CSA / SQA
                                                       ---------------------
                     -------- BELOW 16M --------                       |          -------- ABOVE 16M --------
   AREA          MINIMUM               MAXIMUM              AVG        |       MINIMUM               MAXIMUM              AVG
  ------       ---------             ---------             -----       |      ---------             ---------            -----
   CSA       1,008K (10 JUN/16.34)  1,056K (10 JUN/16.47)  1,020K        |       58.7M (10 JUN/16.37)   62.5M (10 JUN/16.46)   59.9M
   SQA       1,128K (10 JUN/16.34)  1,132K (10 JUN/16.43)  1,129K        |       11.4M (10 JUN/16.43)   11.9M (10 JUN/16.45)   11.6M
  ALLOCATED CSA BY STORAGE KEY                                         |
                     -------- BELOW 16M --------                       |          -------- ABOVE 16M --------
   KEY           MINIMUM               MAXIMUM              AVG        |       MINIMUM               MAXIMUM              AVG
  -----        ---------             ---------             -----       |      ---------             ---------            -----
    0          584K (10 JUN/16.34)    620K (10 JUN/16.46)    593K        |       19.0M (10 JUN/16.35)   20.4M (10 JUN/16.47)   19.3M
    1           52K (10 JUN/16.34)     56K (10 JUN/16.47)     52K        |        924K (10 JUN/16.34)    928K (10 JUN/16.46)    925K
    2           36K (10 JUN/16.34)     36K (10 JUN/16.34)     36K        |      1,584K (10 JUN/16.34)  1,584K (10 JUN/16.34)  1,584K
    3            4K (10 JUN/16.34)      4K (10 JUN/16.34)      4K        |         40K (10 JUN/16.34)     40K (10 JUN/16.34)     40K
    4           52K (10 JUN/16.34)     52K (10 JUN/16.34)     52K        |       20.8M (10 JUN/16.35)   21.0M (10 JUN/16.44)   20.9M
    5            4K (10 JUN/16.34)      4K (10 JUN/16.34)      4K        |      2,736K (10 JUN/16.34)  2,736K (10 JUN/16.34)  2,736K
    6           76K (10 JUN/16.34)     76K (10 JUN/16.34)     76K        |      7,320K (10 JUN/16.38)  7,352K (10 JUN/16.46)  7,333K
    7           24K (10 JUN/16.34)     32K (10 JUN/16.46)     26K        |      5,780K (10 JUN/16.37)  8,172K (10 JUN/16.46)  6,549K
   8-F         176K (10 JUN/16.34)    176K (10 JUN/16.34)    176K        |        952K (10 JUN/16.34)    956K (10 JUN/16.46)    952K
  SQA EXPANSION INTO CSA                                               |
                   0K                    0K                    0K      |        0K                    0K                    0K
                                                       UNALLOCATED CSA / SQA
                                                      ----------------------
                                 -------- BELOW 16M --------                        -------- ABOVE 16M --------
       DESCRIPTION           MINIMUM               MAXIMUM              AVG        MINIMUM               MAXIMUM              AVG
      -------------        ---------             ---------             -----      ---------             ---------            -----
  CSA
    FREE PAGES (BYTES)    2,692K (10 JUN/16.47)  2,740K (10 JUN/16.34)  2,728K       138M (10 JUN/16.46)    142M (10 JUN/16.37)    141M
    LARGEST FREE BLOCK    2,692K (10 JUN/16.47)  2,740K (10 JUN/16.34)  2,728K       138M (10 JUN/16.46)    141M (10 JUN/16.34)    140M
    ALLOCATED AREA SIZE   1,008K (10 JUN/16.34)  1,056K (10 JUN/16.47)  1,020K      59.1M (10 JUN/16.34)   62.5M (10 JUN/16.46)   60.0M
  SQA
    FREE PAGES (BYTES)      660K (10 JUN/16.43)    664K (10 JUN/16.34)    663K     5,776K (10 JUN/16.45)  6,308K (10 JUN/16.43)  6,119K
    LARGEST FREE BLOCK      512K (10 JUN/16.34)    512K (10 JUN/16.34)    512K     5,636K (10 JUN/16.46)  5,872K (10 JUN/16.34)  5,813K
    ALLOCATED AREA SIZE   1,280K (10 JUN/16.34)  1,280K (10 JUN/16.34)  1,280K      11.8M (10 JUN/16.34)   12.1M (10 JUN/16.46)   11.9M
```

Common Area storage summary section field descriptions

Table 114 describes each field in the Common Area storage summary section of the Virtual Storage Activity Report.

Table 114 Field descriptions for the Common Area storage summary section

Field	Description
ALLOCATED CSA/SQA	amount of storage acquired through GETMAIN service from CSA and SQA A 4-K block is considered allocated if any portion of it is allocated.
MINIMUM	smallest amount of CSA / SQA allocated during the report interval
MAXIMUM	largest amount of CSA / SQA allocated during the report interval
AVG	average amount of CSA / SQA allocated during the report interval
KEY	storage key
SQA EXPANSION INTO CSA	amount of storage allocated within CSA for use by SQA
FREE PAGES	amount of storage not allocated Only blocks equal to or greater than 4 K in length are included in this total. Consequently, free storage blocks less than a page in length are not represented here.
LARGEST FREE BLOCK	largest contiguous block of free storage available
CSA ALLOCATED AREA SIZE	size of the area bounded by all the allocated blocks within CSA The difference between this figure and the amount of storage used within CSA gives an indication of fragmentation within CSA. This figure is computed as (high address of CSA) (address of lowest allocated block in CSA)
SQA ALLOCATED AREA SIZE	size of the area bounded by all the allocated blocks within SQA The difference between this figure and the amount of storage used within SQA gives an indication of fragmentation within SQA. This figure is computed as (size of SQA) (size of lowest free block within SQA)

Common Area storage detail section

An example of the Common Area storage detail section is shown in Figure 96

Figure 96 Common Area storage detail section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                  VIRTUAL STORAGE ACTIVITY REPORT                    RPTSEQ     3 PAGE      7
 BMC SOFTWARE, INC.                                           XYZ COMPANY                              REPORT DATE: DD MMM YY 16.53
 ACTL 10 JUN YY 16.33.43  10 JUN YY 16.49.53             WORLDWIDE HEADQUARTERS                        SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  78-2/4/119/0.27
 --------------------------------------------------  COMMON AREA REPORT - DETAIL  --------------------------------------------------
                             ALLOCATED CSA (BELOW 16M) BY SUBPOOL AND STORAGE KEY
                            ------------------------------------------------------
   CSA - SUBPOOL 227
    KEY          MINIMUM                          MAXIMUM                          AVG
   -----        ---------                        ---------                        -----
     0           28K (10 JUN YY / 16.34.14)       28K (10 JUN YY / 16.34.14)       28K
     1            0K (10 JUN YY / 16.34.14)        4K (10 JUN YY / 16.47.56)        0K
     2            0K                               0K                               0K
     3            0K                               0K                               0K
     4            0K                               0K                               0K
     5            0K                               0K                               0K
     6           12K (10 JUN YY / 16.34.14)       12K (10 JUN YY / 16.34.14)       12K
     7            0K                               0K                               0K
    8-F           0K                               0K                               0K
    ALL          40K (10 JUN YY / 16.34.14)       44K (10 JUN YY / 16.47.56)       40K
   CSA - SUBPOOL 228
    KEY          MINIMUM                          MAXIMUM                          AVG
   -----        ---------                        ---------                        -----
     0           32K (10 JUN YY / 16.34.14)       40K (10 JUN YY / 16.45.31)       32K
     1            4K (10 JUN YY / 16.34.14)        4K (10 JUN YY / 16.34.14)        4K
     2            0K                               0K                               0K
     3            0K                               0K                               0K
     4           52K (10 JUN YY / 16.34.14)       52K (10 JUN YY / 16.34.14)       52K
     5            0K                               0K                               0K
     6           12K (10 JUN YY / 16.34.14)       12K (10 JUN YY / 16.34.14)       12K
     7           12K (10 JUN YY / 16.34.14)       16K (10 JUN YY / 16.46.19)       13K
    8-F          32K (10 JUN YY / 16.34.14)       32K (10 JUN YY / 16.34.14)       32K
    ALL         144K (10 JUN YY / 16.34.14)      152K (10 JUN YY / 16.45.31)      145K

                             ALLOCATED SQA (BELOW 16M) BY SUBPOOL
                            --------------------------------------
    SUBPOOL      MINIMUM                          MAXIMUM                          AVG
   ---------    ---------                        ---------                        -----
    226          56K (10 JUN YY / 16.34.14)       56K (10 JUN YY / 16.34.14)       56K
    239         100K (10 JUN YY / 16.34.14)      100K (10 JUN YY / 16.34.14)      100K
    245         972K (10 JUN YY / 16.34.14)      976K (10 JUN YY / 16.43.59)      973K
```

Common Area storage detail section field descriptions

Table 115 describes each field in the Common Area storage detail section of the Virtual Storage Activity Report.

Table 115 Field descriptions for the Common Area storage detail section

Field	Description
ALLOCATED CSA (BELOW 16M)	amount of storage acquired through GETMAIN service from CSA, viewed by subpool and storage key This data is available only for the lower CSA (below 16 MB).
KEY	storage key
MINIMUM	smallest amount of CSA allocated during the report interval
MAXIMUM	largest amount of CSA allocated during the report interval
AVG	average amount of CSA allocated during the report interval
ALLOCATE SQA (BELOW 16M)	amount of storage acquired through GETMAIN service from SQA, viewed by subpool All SQA is acquired in storage key 0. This data is available only for the lower SQA (below 16 MB).

Private Area storage summary section

A sample of the Private Area storage summary section is shown in Figure 97.

Figure 97 Private Area storage summary section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                   VIRTUAL STORAGE ACTIVITY REPORT                    RPTSEQ     3 PAGE      9
BMC SOFTWARE, INC.                                            XYZ COMPANY                              REPORT DATE: DD MMM YY 16.53
ACTL 10 JUN YY 16.33.43  10 JUN YY 16.49.53             WORLDWIDE HEADQUARTERS                         SYSTEM ID: SJSE     Z v.rr.n


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  78-2/4/119/0.27
--------------------------------------------------  PRIVATE AREA REPORT - SUMMARY  ------------------------------------------------
                     JOB NAME     - XTSTKPAS                                 REGION REQUESTED (IN JCL)   -     32.0M
                     STEP NAME    - XTSTKPAS                                 REGION ASSIGNED (BELOW 16M) -    8,104K
                     PROGRAM NAME - BBM9DA00                                 REGION ASSIGNED (ABOVE 16M) -     32.0M
                     SAMPLE COUNT -    119
                                                      PRIVATE AREA MAP
             -------- BELOW 16M --------                                -------- ABOVE 16M --------
    DESCRIPTION          START ADDRESS        LENGTH               DESCRIPTION           START ADDRESS           LENGTH
  ---------------       ---------------     --------             ---------------       ---------------         --------
    AUTH REGION ALLOC AREA    007AE000          328K          EXT AUTH REGION ALLOC AREA       7E28A000            29.5M
    AUTH REGION UNUSED AREA   00800000            0K          EXT AUTH REGION UNUSED AREA      14500000            1694M
    AUTHORIZED REGION         00800000            0K          EXT AUTHORIZED REGION            14500000            1723M
  ----------  GETMAIN LIMIT : 00800000  -------------        ---------------  GETMAIN LIMIT :  14500000  ---------------
    USER REGION UNUSED AREA   001D7000        6,308K          EXT USER REGION UNUSED AREA      13008000            21.0M
    USER REGION ALLOC AREA    00006000        1,532K          EXT USER REGION ALLOC AREA       12500000            11.0M
    USER REGION               00006000        8,168K          EXT USER REGION                  12500000            32.0M
    SYSTEM REGION             00001000          20K
                                                PRIVATE AREA STORAGE SUMMARY
                                   -------- BELOW 16M --------                      -------- ABOVE 16M --------
        DESCRIPTION         MINIMUM              MAXIMUM            AVG        MINIMUM              MAXIMUM              AVG
       -------------       ---------            ---------          -----      ---------            ---------            -----
 AUTHORIZED REGION (LSQA/SWA/229/230)
    FREE PAGES (BYTES)       0K                   0K                 0K      1694M (10 JUN/16.47)   1694M (10 JUN/16.44)  1,694M
    LARGEST FREE BLOCK       0K                   0K                 0K      1694M (10 JUN/16.45)   1694M (10 JUN/16.34)  1,694M
    PAGES ALLOCATED
       (IN BYTES)          304K (10 JUN/16.34)    324K (10 JUN/16.45)    304K      29.2M (10 JUN/16.44)   29.4M (10 JUN/16.47)   29.3M
    ALLOCATED AREA SIZE    308K (10 JUN/16.34)    328K (10 JUN/16.45)    308K      29.3M (10 JUN/16.34)   29.5M (10 JUN/16.45)   29.4M
 USER REGION
    FREE PAGES (BYTES)   7,108K (10 JUN/16.45)  7,136K (10 JUN/16.34)  7,136K      21.0M (10 JUN/16.49)   21.7M (10 JUN/16.34)   21.5M
    LARGEST FREE BLOCK   6,308K (10 JUN/16.45)  6,328K (10 JUN/16.34)  6,328K      21.0M (10 JUN/16.47)   21.6M (10 JUN/16.34)   21.4M
    PAGES ALLOCATED
       (IN BYTES)          744K (10 JUN/16.34)    752K (10 JUN/16.45)    744K      10.3M (10 JUN/16.34)   11.0M (10 JUN/16.49)   10.5M
    ALLOCATED AREA SIZE  1,532K (10 JUN/16.34)  1,532K (10 JUN/16.34)  1,532K      10.4M (10 JUN/16.34)   11.0M (10 JUN/16.47)   10.6M

                                           -------------------- OVER 2G ---------------------
  HIGH VIRTUAL MEMORY               MINIMUM              MAXIMUM                AVG           PEAK
                                   ---------            ---------              ------         ------
       TOTAL                       1,024K (10JUN/16.34) 1,024K (10JUN/16.34)   1,024K         2,048K
       SHARED                          0                    0                      0              0
```

Private Area storage summary section field descriptions

Table 116 describes each field in the Private Area storage summary section of the Virtual Storage Activity Report.

Table 116 Field descriptions for the Private Area storage summary section

Field	Description
AUTH REGION ALLOC AREA	size of area bounded by allocated blocks in the region
AUTH REGION UNUSED AREA	area in the region from which no storage has been obtained The size of the UNUSED AREA is usually equal to the LARGEST FREE BLOCK.
AUTHORIZED REGION	size of the entire region in decimal K bytes
GETMAIN LIMIT	highest address at which storage can be obtained by using the GETMAIN service; equivalent to the top of the USER REGION.
MINIMUM	minimum value for a field during the reporting interval; date and time of minimum value
MAXIMUM	maximum value for a field during the reporting interval; date and time of maximum value.
AVG	average amount of storage allocated during the report interval
PEAK	maximum amount of storage allocated for the life of the job
FREE PAGES	amount of storage (in bytes) available as free 4-K blocks
LARGEST FREE BLOCK	largest contiguous block of free storage available
PAGES ALLOCATED	amount of storage (in bytes) in partially or wholly allocated 4-K blocks
ALLOCATED AREA SIZE	size of area bounded by all allocated blocks in private area storage
HIGH VIRTUAL MEMORY	TOTAL—64-bit storage above the 2-G bar SHARED—64-bit storage above the 2-G bar that is shared with other address spaces; this measure is applicable at z/OS 1.5 or later

Private Area storage detail section

A sample of the Private Area storage detail section is shown in Figure 98.

Figure 98 Private Area storage detail section

```
PRODUCED BY CMF ANALYZER (v.r.mm)                    VIRTUAL STORAGE ACTIVITY REPORT                  RPTSEQ     3 PAGE     10
 BMC SOFTWARE, INC.                                             XYZ COMPANY                            REPORT DATE: DD MMM YY 16.53
 ACTL 10 JUN YY 16.33.43  10 JUN YY 16.49.53             WORLDWIDE HEADQUARTERS                        SYSTEM ID: SJSE     Z v.rr.n


 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  78-2/4/119/0.27
 --------------------------------------------------  PRIVATE AREA REPORT - DETAIL  -------------------------------------------------
     JOB NAME - XTSTKPAS
                              ALLOCATED PRIVATE AREA STORAGE (BELOW 16M) BY SUBPOOL
                             -------------------------------------------------------
                        SUBPOOL             MINIMUM                           MAXIMUM                           AVG
                       ---------           ---------                         ---------                         -----
  AUTHORIZED REGION
                        229                   4K (10 JUN YY / 16.34.14)         8K (10 JUN YY / 16.46.44)        4K
                        230                 136K (10 JUN YY / 16.34.14)       136K (10 JUN YY / 16.34.14)      136K
                        236 (SWA)            60K (10 JUN YY / 16.34.14)        60K (10 JUN YY / 16.34.14)       60K
                        237 (SWA)            16K (10 JUN YY / 16.34.14)        16K (10 JUN YY / 16.34.14)       16K
                        255 (LSQA)           88K (10 JUN YY / 16.34.14)       108K (10 JUN YY / 16.45.19)       88K
  USER REGION
                          0                 288K (10 JUN YY / 16.34.14)       292K (10 JUN YY / 16.38.26)      288K
                          8                   4K (10 JUN YY / 16.34.14)         4K (10 JUN YY / 16.34.14)        4K
                         10                  20K (10 JUN YY / 16.34.14)        20K (10 JUN YY / 16.34.14)       20K
                         13                   4K (10 JUN YY / 16.34.14)         4K (10 JUN YY / 16.34.14)        4K
                         21                 144K (10 JUN YY / 16.34.14)       144K (10 JUN YY / 16.34.14)      144K
                         78                   8K (10 JUN YY / 16.34.14)        12K (10 JUN YY / 16.45.19)        8K
                        251 (MODULES)       124K (10 JUN YY / 16.34.14)       124K (10 JUN YY / 16.34.14)      124K
                        252 (REENTRANT)     152K (10 JUN YY / 16.34.14)       152K (10 JUN YY / 16.34.14)      152K
```

Private Area storage detail section field descriptions

Table 117 describes each field in the Private Area Storage Detail section of the Virtual Storage Activity Report.

Table 117 Field descriptions for the Private Area storage detail section

Field	Description
ALLOCATED PRIVATE AREA	amount of storage acquired by the job through GETMAIN service for the subpool; this data is available only for below 16 megabytes
SUBPOOL	subpool ID; can be accompanied by a description of the area
MINIMUM	smallest amount of storage allocated during the report interval
MAXIMUM	largest amount of storage allocated during the report interval
AVG	average amount of storage allocated during the report interval

Workload Manager Goal Mode Report

The Workload Manager Goal Mode Report provides information about the workloads in your service definition. These subsections of the Workload Manager Goal Mode Report are available:

- **Detail Section**

  This section contains detailed information about all aspects of each workload. To request this section of the report, specify TYPE=DETAIL on the WLMGL Analyzer control statement.

  The information in this section also contains the service classes specified by the PERFORM INCLUDE and PERFORM SCTYPE Analyzer control statements. By default, the report will include all service classes.

- **Goal Mode Delay Map Section**

  This section shows the most common delay reason for a particular workload for each service class period. To request this section of the report, specify TYPE=DELAY on the WLMGL Analyzer control statement.

- **Goal Activity Map Section**

  This section provides an overview of how well each workload met its predefined goals for each service class period.

  To request this section of the report, specify TYPE=ACTIVITY on the WLMGL Analyzer control statement.

  The data for this report is obtained by using the WORKLOAD Extractor control statement (see “WORKLOAD” on page 203).

Figure 99 on page 550 shows the Workload Manager Goal Mode Report, emphasizing the Activity Map and Delay Map sections.

Figure 99 WLM Goal Mode Report - Detail Section

```
PRODUCED BY CMF ANALYZER (v.r.mm)            WORKLOAD MANAGER GOAL MODE REPORT (DETAIL)               RPTSEQ     3 PAGE     16
BMC SOFTWARE, INC.                                          BMC ENGINEERING                           REPORT DATE: DD MMM YY 13.55
ACTL 25 JAN YY 09.40.31  25 JAN YY 10.10.00                                                           SYSTEM ID: SJSD     Z 1.06.1


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  72-3/288/6,900/0.49
----------------------------------------------------------- DETAIL SECTION ----------------------------------------------------------
- POLICY: BBPLEX01 - BMC Software Service Policy -----ACTIVATED: DDMMMYYYY  07:30:05 - INSTALLED: DDMMMYYYY  08:11:27 ID: CSTTXC   -
----- WORKLOAD: CICS ------- CICS Workload ---------------------- SERVICE CLASS: CICSNRM ---- CICS Normal Transactions ------------
PERIOD: 1   IMPORTANCE: 3   RESOURCE GROUP: ---
-TRANSACTIONS-  TRANS. TIME  HHH.MM.SS.TTT
AVG       0.00  ACTUAL               0.025
MPL       0.00  QUEUED               0.000
ENDED      611  EXECUTION            0.000
END/SEC   0.35  STD DEV              0.092
#SWAPS       0
EXECUTD      0

      -- RESPONSE TIME --   PERF   SUB    P    RESP   ------------------------- RESOURCE MANAGER STATES % ------------------------
      HH.MM.SS.TTT         INDX   TYPE        TIME %  --ACTIVE-- READY IDLE ---------------- DELAYS ----------------  -CONTINUATION-
GOAL  00.00.00.500    AVG         CICS  BTE           SUB  APPL             TOTAL UNKN                                LOC  REM SPLX
*ALL  00.00.00.025           0.1               558081  0.0   0.0   0.0  0.0 100.0  100                                0.0  0.0  0.0
 SJSD 00.00.00.025           0.1               558081  0.0   0.0   0.0  0.0 100.0  100                                0.0  0.0  0.0

TRANSACTION RESPONSE TIME DISTRIBUTION   (MSEC = MILLISECONDS   SECS = SECONDS   MINS = MINUTES   HRS = HOURS)
  PERCENTAGE OF GOAL------50------60------70------80------90-----100-----110-----120-----130-----140-----150-----200-----400---->400
  RESP TIME (MSEC)       250     300     350     400     450     500     550     600     650     700     750    1000    2000    >2000
  % IN BUCKET           98.9     0.3     0.0     0.2     0.0     0.0     0.0     0.0     0.0     0.0     0.3     0.0     0.3     0.0
  % CUMULATIVE          98.9    99.2    99.2    99.3    99.3    99.3    99.3    99.3    99.3    99.3    99.7    99.7     100     100
  # IN BUCKET            604       2       0       1       0       0       0       0       0       0       2       0       2       0
  # CUMULATIVE           604     606     606     607     607     607     607     607     607     607     609     609     611     611

PRODUCED BY CMF ANALYZER (v.r.mm)            WORKLOAD MANAGER GOAL MODE REPORT (DETAIL)               RPTSEQ     3 PAGE     17
BMC SOFTWARE, INC.                                          BMC ENGINEERING                           REPORT DATE: DD MMM YY 13.55
ACTL 25 JAN YY 09.40.31  25 JAN YY 10.10.00                                                           SYSTEM ID: SJSD     Z 1.06.1


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  72-3/288/6,900/0.49
----------------------------------------------------------- DETAIL SECTION ----------------------------------------------------------
- POLICY: BBPLEX01 - BMC Software Service Policy -----ACTIVATED: DDMMMYYYY  07:30:05 - INSTALLED: DDMMMYYYY  08:11:27 ID: CSTTXC   -
----- WORKLOAD: OMVS ------- Open MVS Workload ------------------ SERVICE CLASS: OMVSNRM ---- Open MVS Normal Transactions ---------
PERIOD: 1   IMPORTANCE: 3   RESOURCE GROUP: ---
-TRANSACTIONS-  TRANS. TIME  HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE UNITS---  -SERVICE SECONDS-  -----APPL%-----  -----STORAGE-----
AVG       1.01  ACTUAL               0.260  RATE     0.3  CPU         28,185  CPU           1.0  CP%         0.1  AVG          119.4
MPL       1.01  EXECUTION            0.258  RESP     0.6  SRB            120  SRB           0.0                   TOTAL        120.8
ENDED     0.00  QUEUED               0.002  CONN     0.4  I/O          8,990  RCT           0.0                   CENTRAL      120.8
END/SEC   0.05  R/S AFFINITY         0.000  DISC     0.0  MEM              0  I/O INT       0.0                   SHARED         0.0
#SWAPS     200  INELIGIBLE           0.000  PEND     0.2  TOTAL       37,295  HS SERV       0.0                   --PAGE-IN RATES--
EXECUTD      0  CONVERSION           0.000  IOSQ     0.0  SU/SEC          21  ZAAPonCP      0.0  ZAAPonCP%   0.0  SINGLE         0.0
AVG ENC      0  STD DEV              0.652                ABSRPTN       20.8  ZAAP          0.0  ZAAP%       0.0  BLOCK          0.0
REM ENC      0                                            TRX SRV       20.8  ZIIPonCP      0.0  ZIIPonCP%   0.0  SHARED         0.0
MS ENC       0                                                                ZIIP          0.0  ZIIP%       0.0  HSP            0.0
                                                                                                                  HSP MISS       0.0

EXECUTION VELOCITY MIGRATION:   I/O MGMT   8.9   INIT MGMT   8.9

      -- RESPONSE TIME --   EX   PERF   AVG   ----- USING % -----  ----- EXECUTION DELAYS % ----- %DLY %    %    -CRYPTO%- -RESCNT%-
      HH.MM.SS.TTT         VEL  INDX   # AS   CPU ZAAP ZIIP  I/O  TOTAL  CPU                      UNKN IDLE QUIE  USG  DLY  USG  DLY
GOAL                      60.0
*ALL                       8.9   6.8    3.0   0.0  0.0  0.0  0.0    0.2  0.2                       0.3 99.5  0.0  0.0  0.0  0.0  0.0
 SJSD                      8.9   6.8    3.0   0.0  0.0  0.0  0.0    0.2  0.2                       0.3 99.5  0.0  0.0  0.0  0.0  0.0

PRODUCED BY CMF ANALYZER (v.r.mm)            WORKLOAD MANAGER GOAL MODE REPORT (DETAIL)               RPTSEQ     3 PAGE     32
BMC SOFTWARE, INC.                                          BMC ENGINEERING                           REPORT DATE: DD MMM YY 13.55
ACTL 25 JAN YY 09.40.31  25 JAN YY 10.10.00                                                           SYSTEM ID: SJSD     Z 1.06.1


BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  72-3/288/6,900/0.49
----------------------------------------------------------- DETAIL SECTION ----------------------------------------------------------
- POLICY: BBPLEX01 - BMC Software Service Policy -----ACTIVATED: DDMMMYYYY  07:30:05 - INSTALLED: DDMMMYYYY  08:11:27 ID: CSTTXC   -
-------------------------------------------------------------------------------------------------------------------------------------
*TOTAL  SERVICE UNIT COEFFICIENTS: IOC  0.5  TCB  1.0  SRB  1.0  MSO  0.0000   I/O MGMT: NO   NORM.FACTORS: ZAAP 1.0000  ZIIP 1.0000
-TRANSACTIONS-  TRANS. TIME  HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE UNITS---  -SERVICE SECONDS-  -----APPL%-----  -----STORAGE-----
AVG     180.18  ACTUAL               1.698  RATE   167.5  CPU     14,476,860  CPU         490.4  CP%        20.4  AVG        4,876.5
MPL     180.17  EXECUTION            1.490  RESP     0.9  SRB      1,041,139  SRB          35.3                   TOTAL    878,602.3
ENDED    1,356  QUEUED               0.190  CONN     0.5  I/O      2,059,506  RCT           0.8                   CENTRAL  878,602.3
END/SEC   0.77  R/S AFFINITY         0.000  DISC     0.1  MEM              0  I/O INT       3.1                   SHARED     1,528.6
#SWAPS   2,937  INELIGIBLE           0.005  PEND     0.3  TOTAL   17,577,505  HS SERV       0.0                   --PAGE-IN RATES--
EXECUTD      0  CONVERSION           0.000  IOSQ     0.0  SU/SEC       9,942  ZAAPonCP     19.0  ZAAPonCP%   1.1  SINGLE         0.0
AVG ENC   1.02  STD DEV             49.085                ABSRPTN       55.1  ZAAP        169.7  ZAAP%       9.6  BLOCK          0.0
REM ENC      0                                            TRX SRV       55.1  ZIIPonCP      0.0  ZIIPonCP%   0.0  SHARED         0.0
MS ENC       0                                                                ZIIP          0.0  ZIIP%       0.0  HSP            0.0
                                                                                                                  HSP MISS       0.0

EXECUTION VELOCITY MIGRATION:   I/O MGMT   6.1   INIT MGMT   4.2

      -- RESPONSE TIME --   EX   PERF   AVG   ----- USING % -----  ----- EXECUTION DELAYS % ----- %DLY %    %    -CRYPTO%- -RESCNT%-
      HH.MM.SS.TTT         VEL  INDX   # AS   CPU ZAAP ZIIP  I/O  TOTAL  CPU                      UNKN IDLE QUIE  USG  DLY  USG  DLY
GOAL
*ALL                       4.2   ---  262.1   0.1  0.0  0.0  0.1    3.0  2.9                      10.8 86.0  0.0  0.0  0.0  0.0  0.0

SYSTEMS
  SYSID  -- START TIME --  --- END TIME ---  - DURATION -  - MVS --  OPT  -SU/SEC-
  SJSD   DDMMMYY 09.40.31  DDMMMYY 10.10.00      00.29.28  ZV010601  D1   29520.30
```

Figure 100 shows the Workload Manager Goal Mode Report, emphasizing the Activity Map and Delay Map sections.

Figure 100 WLM Goal Mode Report - Activity and Delay Map sections

```
PRODUCED BY CMF ANALYZER (v.r.mm)                     WORKLOAD MANAGER MAP REPORT                      RPTSEQ     7 PAGE    180
 BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 13.26
 ACTL 10 JUN YY 17.00.00  11 JUN YY 17.00.00            WORLDWIDE HEADQUARTERS                           SYSTEM ID: **ALL**  COMB-MVS
 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  72-3/7,896/0/46.64
 ------------------------------------------------  GOAL ACTIVITY MAP SECTION (DAILY)  ------------------------------------------------
        POLICY INSTALLATION DATE-TIME / NAME / DESCRIPTION:  10JUNYYYY 10:50:03 / BBPOL002 / XYZ COMPANY Service Policy
 '+' = EXCEEDED GOAL, '.' = MET GOAL, '123...' = BELOW GOAL BY 10%, 20%, 30%...  ' ' = ZERO OR NO DATA
                   0   0   0   0   0   0   0   0   0   0   1   1   1   1   1   1   1   1   1   1   2   2   2   2
             TIME: 0   1   2   3   4   5   6   7   8   9   0   1   2   3   4   5   6   7   8   9   0   1   2   3
 WORKLOAD ----- BATCH ------- Batch Workload ---------------------------------------------------------------------
   BATNRM   P1                 +    + + +++ +++++++++++++++++++++++++++++++++++++++++++++++++++++  +     +     +
            P2                 +          + +++++++++++++++++++++++++++++++++++++++++++++++++++++
   BATPROD  P1           +                                   ++  ++ +                        +
 WORKLOAD ----- CICS -------- CICS Workload ----------------------------------------------------------------------
   CICSHOT  P1                      007++072027424717.+0 058+29467027330 907+ ++4  ++  1+. +6+ 5             1 +
   CICSNRM  P1                      ++3+++0++76+95++++++1+6++++++++++++3++++++++4++++   +09032
 WORKLOAD ----- IMS --------- IMS Workload -----------------------------------------------------------------------
   IMSNRM   P1                                     5 9     0        +   ++++++2++.2+995 +
 WORKLOAD ----- OMVS -------- Open MVS Workload ------------------------------------------------------------------
   OMVSNRM  P1                           90 7 88 9 94 726    70 1 535591577585 02 3  ++ 456977
 WORKLOAD ----- STC --------- STC Workload -----------------------------------------------------------------------
   GRS      P1      99 299  93194 7 99124942154 8719805086360579921632746931472950055580 928150365114969597890 2890
   STCNRM   P1      ++++++++++++++++++++++++++++++++++++++++++4++++++++++++1++++++++++++++++++++++++++++++++++++++++
   STCPAS   P1                                +++++++++1+++...+++++++++++++++++++++++++...5300+2142.26834.++++++++
   STCPROD  P1                      +   +     ++    6 + +++  2 +++  ++++   +++    +     + +         +
   STCSYS   P1      7 9              4       9                      4            + 9  548363545333545352333224333
 WORKLOAD ----- TSO --------- TSO Workload -----------------------------------------------------------------------
   TSONRM   P1           9 9++++++1.1231..31..2213133123112.1133212222121.1112.1..11+.+.1131.+.22231+.+++++++++++++
            P2            +333231.12131..1..112113111..1..+..12...11111.21....+111+++..+++.+++++.11++++     6 8+834
            P3            7 ++++.+3++.1++.+++++++++++++++++.++.+++++++++.+++++++++++++++++++++++++++++ +++ +  .+++
            P4           +  ++++ ++ ++++++++++++++++++++++++++++++++++51++2+2+++++++++++++++++++++++++++++++++++++++

 PRODUCED BY CMF ANALYZER (v.r.mm)                     WORKLOAD MANAGER MAP REPORT                      RPTSEQ     7 PAGE    181
 BMC SOFTWARE, INC.                                             XYZ COMPANY                              REPORT DATE: DD MMM YY 13.26
 ACTL 10 JUN YY 17.00.00  11 JUN YY 17.00.00            WORLDWIDE HEADQUARTERS                           SYSTEM ID: **ALL**  COMB-MVS

 BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:  72-3/7,896/0/46.64
 -----------------------------------------------  GOAL MODE DELAY MAP SECTION (DAILY)  ------------------------------------------------
        POLICY INSTALLATION DATE-TIME / NAME / DESCRIPTION:  10JUNYYYY 10:50:03 / BBPOL002 / XYZ COMPANY Service Policy
   C = CPU, A = CAPP, S = STOR, M = MPL, I = I/O, P = PAGE, O = OTHER, V = SERV, _ = IDLE, R = RSCE GRP, Z = QUIESCE,   = NO DATA
                   0   0   0   0   0   0   0   0   0   0   1   1   1   1   1   1   1   1   1   1   2   2   2   2
             TIME: 0   1   2   3   4   5   6   7   8   9   0   1   2   3   4   5   6   7   8   9   0   1   2   3
 WORKLOAD ----- BATCH ------- Batch Workload ---------------------------------------------------------------------
   BATNRM   P1                 OO      O OO_OOO____OOOOO__ OOOOOOOOOOOOOOOOOOOOOOOOOO_OOOOOOOOOOOO  O     O     O
            P2                 OO      O OO_OOO____OOOOO__ OOOOOOOOOOOOOOOOOOOOOOOOOO_OOOOOOOOOOOO  O     O     O
   BATPROD  P1           O                                   OO  OO O      O                 O
 WORKLOAD ----- OMVS -------- Open MVS Workload ------------------------------------------------------------------
   OMVSNRM  P1      ________________________________________________________________________________________________
 WORKLOAD ----- SLOW -------- Low Service ------------------------------------------------------------------------
   SLOW     P1                                                         CC__________
 WORKLOAD ----- STC --------- STC Workload -----------------------------------------------------------------------
   GRS      P1      OOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOO
   STCLOW   P1      OOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOOO
   STCNRM   P1      ________________________________________________________________________________________________
   STCPAS   P1                                _____________________OO_______________________________________________
   STCPROD  P1      ________________________________________________________________________________________________
   STCSYS   P1      ________________________________________________________________________________________________
 WORKLOAD ----- SYSTEM ------ THE SYSTEM WORKLOAD ----------------------------------------------------------------
   SYSSTC   P1      ________________________________________________________________________________________________
   SYSTEM   P1      ________________________________________________________________________________________________
 WORKLOAD ----- TSO --------- TSO Workload -----------------------------------------------------------------------
   TSONRM   P1           _OO________________________________________________________________________________________
            P2           _OO________________________________________________________________________________________
            P3           _OO________________________________________________________________________________________
            P4           _OO________________________________________________________________________________________
   TSOSUPP  P1                                         O
            P2                                         O
            P3                                         O
            P4                                         O
            P5                                         O
```

Workload Manager Goal Mode Report field descriptions

Table 118 describes each field in the Detail section of the Workload Manager Goal Mode Report.

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 1 of 9)

Field	Description
POLICY	name of policy used by the service definition For more information about service definition policies, see Appendix B, “Workload measurement.”
ACTIVATED	date and time service policy was activated
INSTALLED	date and time this policy was installed
ID	user ID of person who installed the policy
WORKLOAD	name and description of the workload in the service class definition in the workload manager For more information about goal mode workloads, see Appendix B, “Workload measurement.”
SERVICE CLASS, REPORT CLASS	name and description of the service class or report class in the service definition For more information about service classes, see Appendix B, “Workload measurement.”
PERIOD	period number of the service or report class For more information about service or report class periods, see Appendix B, “Workload measurement.”
IMPORTANCE	importance level of the performance goal, ranging from 1 (highest) to 5 (lowest) This field is not present for report classes and not applicable to service classes having system or discretionary goal.
RESOURCE GROUP	name of the resource group to which the service class is assigned Resource groups are used to set a minimum or limit the amount of CPU capacity available to one or more service classes. This field is not present for report classes.
HOMOGENEOUS REPORT PERIOD	a report class period is called homogeneous when transactions from a single service class contribute data to it; the report also shows the service class
HETEROGENEOUS REPORT PERIOD	a report class period is called heterogeneous when transactions from more than one service class contribute data to it
SERVICE UNIT COEFFICIENTS	numbers by which I/O, TCB, SRB, and storage service units are multiplied
I/O MGMT	whether WLM dynamically adjusts the priority of non paging DASD I/O requests to meet goals

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 2 of 9)

Field	Description
NORM.FACTORS	zAAP and zIIP Normalization Factors: used to convert zAAP and zIIP processor usage to equivalent CP usage ■ CPusage = zAAP usage * zAAP Normalization Factor ■ CPusage = zIIP usage * zIIP Normalization Factor These formulas allow for the possible differences in processor speed between CPs and zAAPs or zIIPs.
TRANSACTIONS	column heading of transaction counts
AVG	average number of active transactions (not necessarily in central storage)
MPL	average number of transactions resident in storage
ENDED	total number of transactions completed in the report interval
END/SEC	average number of transactions completed per second
# SWAPS	number of swap sequences during the interval
EXECUTD	number of execution phases completed as reported by subsystem work managers
AVG ENC	average number of independent enclaves
REM ENC	average number of foreign enclaves imported from a remote system
MS ENC	average number of multisystem enclaves exported to and executing on a remote system
TRANS. TIME HHH.MM.SS.TTT	column heading of average transaction times
ACTUAL	average transaction response time
EXECUTION	average transaction execution time
QUEUED	average queue delay time ■ For batch jobs, it is the time jobs spent waiting for an initiator. ■ For TSO users, it is a portion of the LOGON process. ■ For APPC, it is the time an APPC request spent on an APPC queue.
R/S AFFINITY	average time that batch jobs were ineligible to run because the resources the job had affinity to was unavailable
INELIGIBLE	average time that batch jobs spent on job queues (after JCL conversion) while ineligible to run for reasons other than resource affinity, including operator hold, delays due to duplicate job names, delays due to job class limits
CONVERSION	average time that batch jobs spent on JCL conversion
STD DEV	measure of how widely transaction response times varied from the average value
DASD I/O	column heading of non-paging DASD I/O requests
RATE	average number of non-paging DASD I/O requests completed per second
RESP	average response time of non-paging DASD I/O requests in milliseconds
CONN	average connect time in milliseconds, including data transfer and search time
DISC	average disconnect time in milliseconds
PEND	average pending time in milliseconds, including channel & control unit contention, device busy on another system

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 3 of 9)

Field	Description
IOSQ	average IOS queue time in milliseconds, that is, time waiting for a device while it is busy performing I/O
SERVICE UNITS	column heading of the number of service units of each type consumed in the interval
I/O	number of I/O service units used
CPU	number of CPU service units that were used by transactions while executing in TCB mode or preemptible SRB mode
SRB	number of CPU service units that were used by transactions while executing in non-preemptible SRBs
MEM	number of memory service units used
TOTAL	total number of I/O, TCB, MEM, and SRB service units used
SU/SEC	rate at which service units were used per second
ABSRPTN	service units used per second while transactions were resident in central storage
TRX SRV	service units used per second while transactions were active (not necessarily in central storage)
SERVICE SECONDS	column heading for CPU seconds that services consumed during the reporting interval; values shown are in seconds
CPU	number of CPU seconds that were used by transactions while executing in TCB mode or preemptible SRB mode This value reflects CPU time spent on all processor types.
SRB	number of CPU seconds that were used by transactions while executing in non-preemptible SRBs
RCT	region control task time in seconds
I/O INT	I/O interrupt processing time in seconds
HS SERV	hiperspace service time in seconds
ZAAPonCP	service time for work that was executed on standard CPs eligible to run on zAAPs; value is in seconds and is normalized to the equivalent time on a standard CP **Note:** zAAPs do not necessarily run at the same speed as standard CPs. To provide a proper basis of comparison, the zAAP time values are normalized to represent the amount of time the task would have taken if executed on a standard CP. **Note:** See section "zAAP and zIIP workload projection" on page 560.
ZAAP	zAAP time in seconds **Note:** Without zAAPs configured, N/A is displayed in this field.
ZIIPonCP	service time for work that was executed on standard CPs eligible to run on zIIPs; value is in seconds **Note:** See section "zAAP and zIIP workload projection" on page 560.
ZIIP	zIIP time in seconds **Note**: Without zIIPs online, N/A is displayed in this field.

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 4 of 9)

Field	Description
APPL%	column heading for percentage of the reporting interval that transactions used standard CPs, zAAPs, or zIIPs **Note:** It is possible for values shown in this column to exceed 100% when an LPAR has more than one logical processor or input contains data from more than one system.
CP%	percentage of the report interval that transactions used standard CPs; the sum of CPU, SRB, RCT, I/O INT, and HS SERV service seconds divided by the report interval
ZAAPonCP%	percentage of the report interval that service time was executed on standard CPs but was eligible to run on zAAPs **Note:** See section “zAAP and zIIP workload projection” on page 560.
ZAAP%	percentage of the report interval that zAAP time was used **Note:** Without zAAPs configured, N/A is displayed in this field.
ZIIPonCP%	percentage of the report interval that service time was executed on standard CPs but was eligible to run on zIIPs **Note:** See section “zAAP and zIIP workload projection” on page 560.
ZIIP%	percentage of the report interval that zIIP time was used **Note:** Without zIIPs online, N/A is displayed in this field.
STORAGE	column heading for storage frames allocated to address spaces
AVG	average central and expanded storage, in frames, allocated to an address space while resident in storage; it is page residency time divided by address space residency time
TOTAL	average central and expanded storage, in frames, allocated to all address spaces during the report interval; it is page residency time divided by report interval
CENTRAL	average central storage, in frames, allocated to all address spaces during the report interval
SHARED	average shared storage, in frames, allocated to all address spaces during the report interval
PAGE-IN RATES	column heading for paging, in pages per second
SINGLE	rate at which individual pages were transferred due to demand paging
BLOCK	rate of pages transferred (paged in) from auxiliary storage to central storage in blocks, as part of swap-in or directed page ins
SHARED	rate of shared pages transferred (paged in) from auxiliary storage to central storage
HSP	rate per second of hiperspace pages transferred (paged in) from auxiliary storage to central storage while transactions were active
HSP MISS	rate of expanded storage only (ESO) hiperspace read misses A miss occurs when a read to expanded storage is issued and the requested page resides in auxiliary storage.

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 5 of 9)

Field	Description
EXECUTION VELOCITY MIGRATION	helps you plan your choice of velocity goals The next two fields relate to Execution Velocity Migration.
I/O MGMT ***nn.n***	I/O Priority Management; a WLM facility that gives you the flexibility to control nonpaging DASD I/O requests that are queued because the device is busy This facility can be set to On or Off in the WLM ISPF application by using the Service Coefficient/Service Definition Options panel; the default is Off. ■ If set to On: WLM dynamically adjusts the I/O priority based on goals and I/O activity and includes the I/O information in calculating execution velocity. (For more information, see the IBM *z/OS MVS Planning Workload Management* manual.) ■ If set to Off: The number of samples of work using nonpaging DASD I/O resources and the number of samples that are delayed for nonpaging DASD I/O are excluded from the execution velocity calculation. The value in **I/O MGMT** is the achievable execution velocity, including nonpaging DASD I/O resources and delay samples (that is, the execution velocity if I/O Priority Management is set to On). After this line, you see the actual overall execution velocity of the same service class. If I/O Priority Management is enabled in the service definition, the value in **I/O MGMT** should match the execution velocity in the overall value (field **EX VEL** in the report). If your organization has not set I/O Priority Management to On, you can use the **I/O MGMT** value as a guide in your execution velocity planning exercise.
INIT MGMT ***xx.x***	Initiator Management; the value of achievable execution velocity if batch initiator delay samples were included in the velocity calculation Similar to I/O MGMT, this number helps you if you want to know what the execution velocity will be if batch initiator management is enabled. If you have Initiator Management enabled, the value in **INIT MGMT** will match the value in **EX VEL** in the report.
RESPONSE TIME	for average response time goal, the expected response time is formatted on the GOAL line and the actual response time is on the subsequent lines for sysplex and individual systems For response time with percentile goal, the expected percentage of transactions to be completed within the specified response time is formatted on the GOAL line and the actual percentage is on the subsequent lines for sysplex and individual systems. For system or discretionary goal, SYSTEM or DISCRETIONARY is printed. This column is not applicable to report classes.

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 6 of 9)

Field	Description	
EX VEL	for velocity goal, the expected execution velocity is formatted on the GOAL line The actual execution velocity is formatted on the subsequent lines for sysplex and individual systems, regardless of goal type. For more information about execution velocity, see Appendix B, “Workload measurement.”	
PERF INDX	performance index For more information about the interpretation of a performance index value, see Appendix B, “Workload measurement.” **Note**: This column is not applicable to report classes and service classes of system or discretionary goal.	
AVG # AS	average number of address spaces and enclaves that contributed delay and using samples	
USING %	**CPU**	ratio of samples that are using standard CPs compared with all using and delay samples
	ZAAP	ratio of samples using zAAPs compared with all using and delay samples
	ZIIP	ratio of samples using zIIPs compared with all using and delay samples
	I/O	ratio of nonpaging DASD I/O using samples compared with all using and delay samples
EXECUTION DELAYS %	**TOTAL**	percentage of all delay samples used in the calculation of execution velocity This field does not necessarily equal the sum of delay reason percentages listed (for example, if I/O priority management option is inactive, I/O delay reason is not included.)
	At most, five delay reasons are listed in descending order at the sysplex level:	
	CPU	CPU delay—work has been delayed while either waiting to be dispatched on a standard CP or waiting for the local lock
	CPU CAPP	CPU capping delay—work is nondispatchable because a resource group maximum is being enforced
	DASD	nonpaging DASD I/O This delay is not included in TOTAL if WLM does not manage I/O priority.
	MPL	MPL delay—work is delayed for storage due to multiprogramming level constraint
	PGIN COMM	page in for common storage
	PGIN EHSP	page in for ESO hiperspaces

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 7 of 9)

Field	Description
EXECUTION DELAYS % (*continued*)	**PGIN HSP** page in for standard hiperspaces **PGIN PRIV** page in for private storage **PGIN VIO** page in for VIO storage **PGIN XMEM** cross memory page in **SERV HSP** server hiperspace page in **SERV MPL** server MPL delay **SERV PRIV** server private area page in **SERV QUE** server queue delay—work is waiting for a server **SERV SWIN** server swap-in **SERV VIO** server VIO page in **SWAP IN** swap-in delay **ZAAP** zAAP delay—work that has been delayed while waiting to be dispatched on a zAAP **ZIIP** zIIP delay—work that has been delayed while waiting to be dispatched on a zIIP
%DLY UNKN	work is delayed, but none of the reasons listed under EXECUTION DELAYS % above apply
% IDLE	work is in STIMER wait, TSO terminal wait, APPC wait, or is an initiator waiting for work
% QUIE	work in this service class has been reset by the RESET jobname, QUIESCE command, which swaps out a swappable address space or gives the lowest possible performance characteristics to a nonswappable address space
CRYPTO% USG	crypto using state -- a TCB or SRB was found to be using a cryptographic asynchronous message processor (CAP) or an adjunct processor (AP)
CRYPTO% DLY	crypto delay state -- a TCB or SRB was found to be waiting for a cryptographic asynchronous message processor (CAP), an adjunct processor (AP) or a processor feature queue
RESCNT% USG	percentage of samples where work is holding a resource in contention, as reported by a resource manager by way of the IWMCNTN service
RESCNT% DLY	percentage of samples where work is waiting for a resource in contention, as reported by a resource manager by way of the IWMCNTN service
SUB TYPE	subsystem type, as specified in the classification rules
P	states sampled in the begin-to-end phase of a transaction (BTE) or in the execution phase (EXE)
RESP TIME %	percentage of the response time that a transaction spends in the begin-to-end or execution phase **Note**: Long-running and neverending transactions that were not completed during the report interval contributed state samples but not response time, causing this transaction to appear inflated, perhaps even registering above 100%.
RESOURCE MANAGER STATES %	ratio of the specified state samples compared with total state samples

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 8 of 9)

Field	Description
ACTIVE	a program was executing on behalf of the work request, from the perspective of the work manager; does not mean that the program is active from the MVS perspective **SUB** program was the subsystem itself **APPL** program was an application invoked by the subsystem
READY	a program was ready to execute on behalf of the work request, but the work manager has given priority to another work request
IDLE	work manager saw a transaction as idle
DELAYS	**TOTAL** percentage of all delay state samples At most, six delay reasons are listed in descending order at the sysplex level: **BPMI** buffer pool miss **CONV** waiting for conversation **DIST** waiting for distributed request; some function or data had to be routed prior to resumption of the work request **I/O** waiting for I/O **LTCH** waiting for new latch **LOCK** waiting for lock **LSES** waiting for a session to be established locally on the current MVS image **NSES** waiting for a session to be established somewhere in the network **PROD** waiting for another product **SSES** waiting for a session to be established somewhere in the sysplex **TIMR** waiting for a timer **UNKN** waiting for an unidentified resource **SSLT** waiting for an SSL thread **REGT** waiting for a regular thread **WORK** waiting for a registration worktable

Table 118 Field descriptions for the Workload Manager Goal Mode Report (part 9 of 9)

Field	Description
CONTINUATION	percentage of logical continuation state samples: **LOC** logical continuation on the local system; subsystem work managers function shipped a transaction to another component within the same system **REM** logical continuation on a remote system in the sysplex; subsystem work managers function shipped a transaction to another component on another system **SPLX** logical continuation somewhere within the network; subsystem work managers function shipped a transaction to another component within the network
TRANSACTION RESPONSE TIME DISTRIBUTION	for response time with percentile goal, this table shows the distribution of completed transactions
PERCENTAGE OF GOAL	percentage of goal reached for a specific column
RESPONSE TIME (*unit*)	response times of response time buckets; unit of measure is determined dynamically The unit can be milliseconds (MSEC), seconds (SECS), minutes (MINS), or hours (HRS).
% IN BUCKET	percentage of transactions completed in the reporting period represented by a response time bucket
% CUMULATIVE	percentage of transactions completed with a response time less than or equal to the value displayed in the RESP TIME line
# IN BUCKET	number of transactions completed in the reporting period represented by a response time bucket
# CUMULATIVE	number of transactions completed with a response time less than or equal to the value displayed in the RESP TIME line

zAAP and zIIP workload projection

If zAAPs and zIIPs are not online, N/A is reported in fields **zAAPonCP**, **zAAPonCP%**, **zIIPonCP**, and **zIIPonCP%**. However, these fields might be reported with values greater than zero even without this hardware, as a result of zAAP and zIIP *workload projection*.

zAAP and zIIP workload projection enables the system to project how much work could be offloaded from regular CPs to zAAPs and zIIPs. So even if you do not have zAAPs and zIIPs (and might therefore expect to see N/A in these fields), you could see a value in one or more of these fields if a zAAP- or zIIP-eligible workload exists. In this case, the values reported would be *projected* workload rather than real workload, and they would help you evaluate if the addition of zAAPs or zIIPs is justified.

Beginning with z/OS release 1.8, zAAP and zIIP workload projection is activated by the PROJECTCPU=YES option in parmlib member IEAOPT*xx*. For z/OS releases earlier than 1.8, this option is not available; however, the JVM runtime option **-xifa:force** will cause zAAP-eligible work to be projected and reported in the **zAAPonCP** and **zAAPonCP%** fields. No comparable Java runtime option for zIIP workload projection exists in z/OS releases earlier than 1.8.

Workload Manager Goal Mode Report - Activity Map and Delay Map sections

These sections provide an overview of workload performance over the course of a day, week, or month. By using the Activity map section, you can see the times during which workloads were running smoothly as well as the times they were subject to delays. By using the Delay Map section, you can see the main cause of delays during a particular time frame for each workload.

Either map can be specified as providing information on a daily, weekly, or monthly basis. If you request PERIOD=DAY on the WLMGL report control statement, each column of the map represents 15 minutes. If you request PERIOD=WEEK in the WLMGL report control statement, each column represents two hours. If you specify PERIOD=MONTH in the WLMGL report control statement, each column represents eight hours.

After the report header, the first line of a map section indicates the type of map being shown (Activity or Delay) and the period it covers (Daily, Weekly, or Monthly).

The next line contains information about the service policy.

Following the policy information is the report legend.

- For the Activity Map, the values for the legend are as follows:

+	indicates the goal was exceeded for the interval
.	indicates the goal was met for the interval
(blank)	indicates no data was found for the interval
1 2 3...9	indicates the goal was not met; the number shows the approximate percentage by which the workload did not achieve its goal A value of 1 indicates that the workload missed its goal by 10%, a value of 2 indicates that the workload missed its goal by 20%, and so on.

- For the Delay Map, the values for the legend are the following delay reasons:

A	capping
C	CPU
I	I/O
M	MPL
O	other delay reason
P	paging
Q	queuing
R	resource group cap
S	swapping
V	server
_	idle
z	quiesce
(blank)	zero or no data available

The next two lines are the column headers for DAY, WEEK, or MONTH reporting. The map is laid out as follows for each of the time frames:

- **DAY:** Twenty-four hours are displayed, from midnight to midnight. Each hour is broken into four 15-minute intervals starting at midnight, for a total of 96 intervals.

- **WEEK:** Seven days are displayed, from Sunday through Saturday. Each day is broken into twelve two-hour intervals starting 00:00 hours Sunday morning, for a total of 84 intervals.

- **MONTH:** Up to 31 days are displayed. Each day is divided into three eight-hour intervals starting at 00:00 hours on the first day of the month for a maximum of 93 intervals depending on the number of days in the month.

NOTE

Exercise caution in choosing data for this report to avoid the unexpected combining of data. For example, if PERIOD=DAY is used with data from more than one day, the report combines the data by each 15-minute interval, which generates a single report that contains the combined data from multiple days, rather than separate reports for each day.

To limit the data that is used for this report, use the Analyzer DATETIME parameter, as described in “DATETIME” on page 238.

A new report header is generated whenever a Workload Policy change is detected.

Part 3

Advanced topics

This part presents the following topics:

Chapter

9

Using the CMF MONITOR APIs

CMF MONITOR provides four application program interfaces (APIs) to SMF records: CX10GVID, CX10XDQY, CX10XDRC, and CX10XDGS. The CX10GVID API accesses only records collected on a single system. More information on the CX10GVID API is provided in “Implementing the CX10GVID API” on page 588.

The other three APIs provide access to SMF records across a sysplex.

- CX10XDQY queries buffered SMF data throughout the sysplex for any SMF record types.
- CX10XDRC requests buffered SMF data from the sysplex based on the results from the CX10XDQY query.
- CX10XDGS requests snapshot data throughout the sysplex using SMF type 79 records by invoking the CX10GVID API on the requested system or systems.

These three APIs constitute CMF MONITOR’s Cross-System Data Server (XDS). All records obtained with XDS are available for use by the CMF Analyzer, CMFMON (for batch reports), and other performance tools. In addition, XDS provides access to SMF data for application programs, and provides values to the SDSF DA screen if you are using SDSF version 1.5 or later in SYSNAME mode.

NOTE

If you are using a version of SDSF earlier than 1.5 or if you are using non-sysname mode, SMF data is accessed by means of the CX10GVID API.

How XDS works

Like the CMF Extractor, XDS collects and stores SMF records. Unlike the Extractor, these records are stored in a data space buffer, which is accessible to calls from other systems in the sysplex.

While the CMF Extractor writes only type 70-78 and type 240 records to SMF or CMF data sets, XDS can buffer and return record images for all SMF record types or for any subset that you specify.

Application programs that are written to use the RMF APIs (ERBDSQRY, ERBDSREC, ERB2XDGS) can also use the CMF APIs without modification.

Activating XDS

XDS is activated with the XDS parameter in the MVS PAS PROC. The default value of this parameter as distributed in XDS=00. To disable XDS, specify XDS=STOP or remove the XDS parameter from the PAS. XDS can be activated by specifying XDS=*xx*, where *xx* matches the suffix of a *hilevel*.BBPARM member CMFXDS*xx*. For example, specifying XDS=01 accesses member CMFXDS01, which is one of the sample members provided in *hilevel*.BBPARM.

For information on activating XDS in the MVS PAS, see the CMF MONITOR *Customization Guide*.

Accessing data gathered by XDS

You can access data gathered by XDS for any systems within a sysplex, when all of the following conditions exist:

- An MVS PAS is running in CPM mode with DC=START and XDS=*nn* on the system that you are currently using.
- The XDS parameter member specifies the data that you want to be able to access.
- An MVS PAS with the XDS parameter set is also running on the systems from which you want to access XDS data.
- The CAS on each system is defined to join the same XCF group (the default).

This data can be used by the CMF Analyzer and by CMFMON to generate batch reports for data across multiple systems in the sysplex, as long as TYPE=NONE is not specified.

Specifying the source of input data

If you are using the ISPF interface to set up your Analyzer JCL, follow these steps:

1. Select Option 1 on the main menu for producing Analyzer reports.
2. Specify XDS in the **SOURCE OF INPUT DATA** field.

If you are setting up your own JCL, the Analyzer automatically uses data that is collected by XDS, if XDS is available and the //EXTDATA DD statement is not present.

Layout for mapping an answer area for API output

When the XDS APIs are implemented, an application program calls programs CX10XDQY, CX10XDRC, and CX10XDGS, using standard MVS linkage conventions and passing the parameter list. The invoking program must provide a buffer into which the APIs return the requested data. The following sections provide

- the parameters that issue CALLs for each API
- sample DSECTS for each API
- the interface to the CX10XDGS data reduction exit
- a sample DSECT for the common header area

CX10XDQY

Call CX10XDQY (or ERBDSQRY) to request a directory of SMF record data available in the XDS buffers on each system in the sysplex.

To write a CALL to CX10XDQY, the following parameters must be coded in the specified order. Table 119 shows the format of the values that you assign to the parameters.

```
LINK EP=CX10XDQY,
      PARAM=(answer_area_addr
           ,answer_area_alet
           ,answer_area_length
           ,request_type
           ,start_time
           ,end_time
            ,smf_record_type_info
           ,smf_record_type_list
           ,smf_system_id_info
           ,smf_system_id_list
           ,time_out
           ,return_code
      ,reason_code)
```

Table 119 Parameters for calling CX10XDQY (part 1 of 3)

Parameter	Description	Format and length
answer_area_addr	address of the area where CMF returns the requested information The area can be in the caller's primary address space or in an address or data space that is addressable through a public entry on the caller's Dispatchable Unit Access List (DU-AL).	AL4
answer_area_alet	ALET of the answer area provided on the answer_area_addr parameter If the area resides in the caller's primary address space, answer_area_alet must be 0.	FL4
answer_area_length	length of the answer area provided on the answer_area_addr parameter. **Note**: If you do not provide enough length, CMF sets a return code and reason code, and places the necessary length in the answer_area_length parameter.	FL4

Table 119 Parameters for calling CX10XDQY (part 2 of 3)

Parameter	Description	Format and length
request_type	CX10XDQY request type; specify one of the following values: SMF — request information about SMF records of any type and subtype Information is returned about all SMF records whose time information, specified in the SMF record header, is within the time interval specified in the start_time and end_time (SMF *xx*DTE and SMF*xx*TME) parameters. RMF — request information about SMF records of any RMF type and subtype, that is, types 70-79 Information is returned about all SMF records whose projected CMF measurement interval end time is within the time interval specified in the start_time and end_time (SMF *xx*GIE and SMF*xx*LGO) parameters. Additional product section data is returned.	CL3
start_time	beginning of time interval for which information is requested To default to the *oldest* SMF time found in any of the data buffers at the time the service is called, pass a value of 14 blanks.	CL14 (in the sorted format *yyyy* | *mm* | *dd* | *hh* | *mm* | *ss*
end_time	Date and time of the end of time interval for which information is requested. To default to the *newest* SMF time found in any of the data buffers at the time the service is called, pass a value of 14 blanks.	CL14 (in the same sorted format as start_time)
smf_record_type_info	type of the list of SMF record types provided on the smf_record_type_list parameter; specify one of the following values: INCLUDE — The list of SMF record types provided on the smf_record_type_list parameter is an inclusion list. Information is requested for the listed SMF record types. EXCLUDE — The list of SMF record types provided on the smf_record_type_list parameter is an exclusion list. Information is requested for all but the listed SMF record types. ALL — The list of SMF record types provided on the smf_record_type_list parameter is ignored. Information is requested for all SMF record types. If you specify ALL, add four blanks to the right of the string.	CL7

Table 119 Parameters for calling CX10XDQY (part 3 of 3)

Parameter	Description	Format and length
smf_record_type_list	list of SMF record types for which information is requested The first fullword specifies the number of array elements. This is followed by an array of pairs of unsigned integers of length 2, where the first number of each pair specified the record type, and the second number of each pair specifies the record subtype. For record types without subtypes, specify a subtype of 0. If you have specified RMF for request_type, record types outside the range 70 to 79 are ignored.	array: FL4 + (HL2 + HL2 + (HL2 + HL2) + ...
smf_system_id_info	type of the list of SMF system IDs provided on the smf_system_name_list parameter; specify one of the following values: INCLUDE — The list of SMF system IDs provided on the smf_system_id_list parameter is an inclusion list. Information is requested for systems with the listed SMF system IDs. EXCLUDE — The list of SMF system IDs provided on the smf_system_id_list parameter is an exclusion list. Information is requested for all systems in the sysplex, except systems with the listed SMF system IDs. ALL — The list of SMF system IDs provided on the smf_system_id_list parameter is ignored. Information is requested for all systems in the sysplex. If you specify ALL, add 4 blanks to the right of the string.	CL7
smf_system_id_list	list of SMF system IDs for which information is requested The first fullword specifies the number of array elements. This is followed by an array of four-character SMF SYSIDs.	array: FL4 + CL4 + CL4 + CL4...
time_out	time interval in seconds If the time interval expires during the processing of the service, CMF returns to the caller with a corresponding return code, reason code, and partial data. The default value of 60 overrides any value not given in the specified time.	FL4 the value is a non-negative number
return_code	return code when CX10XDQY is completed	FL4
reason_code	reason code when CX10XDQY is completed	FL4

Sample DSECT for CX10XDQY results

When CX10XDQY has completed successfully and returns control to your program, the answer area contains the common header plus one directory entry for each requested SMF record. For information about the common header, see "Sample DSECT for the common answer area header" on page 581. The directory entry contains a record token that was created by CX10XDQY, which can be used for a subsequent call to CX10XDRC to request the actual SMF record itself, and the SMF record header.

The following DSECT is used for the answer area directory entry:

`XDRQ DSECT ,`	CX10XDQY data section
`XDRQRTKN DS XL8`	record token
`XDRQSMFH DS XL24`	SMF record header
`XDRQSSIZ EQU *-XDRQ`	size of CX10XDQY type=SMF entry
`XDRQRMFI DS XL32`	info from product section
`XDRQRSIZ EQU *-XDRQ`	size of CX10XDQY type=RMF entry

XDRQRTKN the record token provided by CX10XDQY to be used on subsequent calls to CX10XDRC

XDRQSMFH the SMF record header (24 bytes) as described in System Management Facility

For SMF record types without subtypes, which have a header only 18 bytes long, bytes 19 to 24 contain hex zeros. Table 120 describes the format of the header.

Table 120 XDRQSMFH SMF record header format

Name	Length	Format	Description
SMF*xx*LEN	2	Integer	SMF record length
SMF*xx*SEG	2	Integer	SMF segment descriptor
SMF*xx*FLG	1	Binary	SMF system indicator
SMF*xx*RTY	1	Integer	SMF record type
SMF*xx*TME	4	Integer	SMF record time (1 / 100 sec)
SMF*xx*DTE	4	0CYYDDDF	SMF record date
SMF*xx*SID	4	Char	SMF system ID
SMF*xx*SSI	4	Char	SMF subsystem ID
SMG*xx*STY	2	Integer	SMF record subtype

XDRQSMFH (*continued*)

For request_type = SMF, the directory entries are sorted by

SMF*xx*DTE: SMF record date
SMF*xx*TME: SMF record time
SMF*xx*RTY: SMF record type
SMG*xx*STY: SMF record subtype
SMF*xx*SID: SMF record system ID

XDRQRMFI

for request_type = RMF only, each directory entry contains additional information from the product section of the SMF record

Table 121 describes the format of the header.

Table 121 XDRQRMFI RMF record header format

Name	Length	Format	Description
SMF*xx*DAT	4	0CYYDDDF	actual interval start date
SMF*xx*IST	4	0HHMMSSF	actual interval start time
SMF*xx*INT	4	Integer	actual interval length
SMF*xx*OIL	2	Integer	synchronization length (seconds)
SMF*xx*SYN	2	Integer	synchronization value (seconds)
SMF*xx*LGO	8	(STCK)	offset GMT to local time
SMF*xx*GIE	8	(STCK)	projected interval end (GMT)

XDRQRMFI (*continued*)

For request_type = RMF, the directory entries are sorted by

SMF*xx*DAT: interval start date
SMF*xx*IST: interval start time
SMF*xx*RTY: SMF record type
SMG*xx*STY: SMF record subtype
SMF*xx*SID: SMF record system ID

CX10XDRC

Call CX10XDRC (or ERBDSREC) to request buffered SMF data from the sysplex based on the results from the CX10XDQY query.

To write a Call to CX10XDRC, the following parameters must be coded in the specified order:

```
LINK EP=CX10XDRC,
PARAM=(answer_area_addr
,answer_area_alet
,answer_area_length
,smf_record_token_list
,time_out
,return_code
,reason_code)
```

Table 122 shows the format of the values that you assign to the parameters.

Table 122   Parameters for calling CX10XDRC (part 1 of 2)

Parameter	Description	Format and length
answer_area_addr	address of the area where CMF returns the requested information The area can be in the caller's primary address space or in an address or data space addressable through a public entry on the caller's Dispatchable Unit Access List (DU-AL).	AL4
answer_area_alet	ALET of the answer area provided on the answer_area_addr parameter If the area resides in the caller's primary address space, answer_area_alet must be 0.	FL4
answer_area_length	length of the answer area provided on the answer_area_addr parameter If you do not provide enough length, CMF sets a return code and reason code and places the length you need in the answer_area_length parameter.	FL4
smf_record_token_list	list of record tokens for the requested SMF records The first fullword specifies the number of array elements. This is followed by an array of eight-character token values.	Array: FL4 + XL8 + XL8 + XL8 + ...

Table 122 Parameters for calling CX10XDRC (part 2 of 2)

Parameter	Description	Format and length
time_out	time interval in seconds If the time interval expires during the processing of the service, CMF returns to the caller with a corresponding return code, reason code, and partial data. The default value of 60 overrides any value not given in the specified format.	FL4 the value is a non-negative number
return_code	return code when CX10XDQY is completed	FL4
reason_code	reason code when CX10XDQY is completed	FL4

Sample DSECT for CX10XDRC results

When CX10XDRC returns control to your program after the service was completed successfully, the answer area contains the common header and one entry for each requested SMF record. For information about the common header, see “Sample DSECT for CX10XDRC results” on page 574. The entries appear in the order of the request, which is identical to the order of the tokens in the record token list. The entry for each record contains a data header, which is provided by CX10XDRC, and the SMF record itself.

The following DSECT can be used for the CX10XDRC data section:

```
XDRR      DSECT ,                       CX10XDRC Data Section
XDRRRL    DS    F                       Record length including header
XDRRRH    DC    A(XDRRSMFR-XDRR)         Length of this hdr
XDRRRC    DS    F                       Ret code for req of this rec
XDRRRC0   EQU   0                       . Data returned at XDRRSMFR
XDRRRC4   EQU   4                       . No data - request timed out
XDRRRC8   EQU   8                       . No data - bad record token
XDRRRSV1 DS     XL4                      Reserved
XDRRRTKN DS     XL8                     Record token from CX10XDQY
XDRRSMFR DS     0H                        SMF record start
```

XDRRRL length of this SMF record data entry, including the data header

XDRRRH length of this SMF record data header

XDRRRC return code for the request of this SMF record

XDRRRC0 return code 0; data returned

SMF record data follows this data header

XDRRRC4 return code 4; data not returned

Timeout occurred before the record was received from the remote system.

XDRRRC8	return code 8; data not returned The record token does not correspond to an existing SMF record in the sysplex.
XDRRRSV1	reserved
XDRRRTKN	record token for this SMF record (copied from input parameter)
XDRRSMFR	SMF record

CX10XDGS

Call CX10XDGS (or ERB2XDGS) to request type 79 data according to the specified SMF record type 79 subtype.

To write a CALL to CX10XDGS, the following parameters must be coded in the specified order. For parameters that CX10XDGS uses to obtain input values, assign values that are acceptable to CX10XDGS. Table 123 on page 576 shows the format of the values that you assign to the parameters.

```
LINK EP=CX10XDGS,
      PARAM=(answer_area_addr
           ,answer_area_alet
           ,answer_area_length
           ,system_id
           ,data_gatherer_parm
           ,data_gatherer_parm_length
           ,exit_name
           ,exit_parm
           ,exit_parm_length
           ,time_out
           ,return_code
           ,reason_code)
```

Table 123 Parameters for calling CX10XDGS (part 1 of 2)

Parameter	Description	Format and length
answer_area_addr	address of the area where CMF returns the requested information The area can be in the caller’s primary address space or in an address or data space addressable through a public entry on the caller’s Dispatchable Unit Access List (DU-AL).	AL4
answer_area_alet	ALET of the answer area provided on the answer_area_addr parameter If the area resides in the caller’s primary address space, answer_area_alet must be 0.	FL4
answer_area_length	length of the answer area provided on the answer_area_addr parameter If you do not provide enough length, CMF sets a return code and reason code and places the length you need in the answer_area_length parameter.	FL4
system_id	ID of the system for which you are requesting information This value is the four-character SMF system identification (SID). Specify *ALL to request information from all systems in the sysplex.	CL4
data_gatherer_parm	parameters for the type 79 data gatherer on each system The first variable is the two-byte record type; this is followed by the two-byte subtype, which is followed by options for the type 79 data gatherer for the specified SMF record type and subtype as a character variable with a maximum length of 32.	array: FL2 + FL2 + CL *n* where *n* has a maximum length of 32
data_gatherer_parm_length	length of the parameter string data_gatherer_parm	FL4

Table 123 Parameters for calling CX10XDGS (part 2 of 2)

Parameter	Description	Format and length
exit_name	name of a data reduction exit routine that is invoked by CMF on each system from which data is requested After the type 79 data has been retrieved by CMF, this exit may move selected areas from the data to the answer area provided by CMF. These data areas are then combined into the answer area provided by the caller on the requesting system. The data reduction exit routine (CX10XDGX, with alias ERB2XSMF) provided by CMF copies the entire record produced by CX10GVID (SMF record type 79) to the answer area. See “Implementing the CX10GVID API” on page 588 for more information about CX10GVID.	CL8
exit_parm	parameter area to be passed to the routine specified in exit_name Use this parameter to control the selection of type 79 data areas to be returned to the caller.	XL*n* where *n* is a value in the range of 0 to 32767
exit_parm_length	Llength of the parameter string exit_parm that is passed to the routine specified in exit_name	FL4 the value is between 0 and 32767
time_out	time interval in seconds If the time interval expires during the processing of the service, CMF returns to the caller with a corresponding return code and reason code and partial data. The default value of 60 overrides any value not given in the specified format.	FL4 the value is a non-negative number
return_code	return code when CX10XDGS completes	FL4
reason_code	reason code when CX10XDGS completes	FL4

Sample DSECT for CX10XDGS

When CX10XDGS returns control to your program after the service was completed successfully, the answer area contains the common header and one or more data sections. For information about the common header, see "Sample DSECT for the common answer area header" on page 581. Each data section contains a data header followed by the data itself.

A sample DSECT for the CX10XDGS data section is shown here:

```
XDRG      DSECT ,                       CX10XDGS Data Section
XDRGDEL   DS    F                       Data section len (with hdr)
XDRGHDL   DC    A (XDRGREC-XDRG)        Data section header len
XDRGRTN   DS    F                       Data retrieval return code
XDRGRSN   DS    F                       Data retrieval reason code
XDRGCPU   DS    F                       System CPU Utilization
XDRGPRT   DS    F                       System paging rate
XDRGDRC   DS    CL12                    Abend code : tcccrrrrrrrr
XDRGSRM   DS    F                       MVS/SRM Cpu Util
XDRGREC   DS    0X                      Type 79 record
```

XDRGDEL length of this data section

XDRGHDL length of this data header

XDRGRTN Data Retrieval return code

XDRGRSN Data Retrieval reason code

XDRGCPU System CPU Utilization

XDRGPRT System Paging Rate

XDRGDRC Data Reduction exit completion code, if the exit ended abnormally

The completion is in the format TCCCRRRRRRRR, where

- T is S or U for a system or user completion code, respectively.
- CCC is the hexadecimal completion code. The highest possible user completion code is x'FFF'.
- RRRRRRRR is the hexadecimal reason code associated with the completion code.

XDRGSRM MVS view of CPU utilization if CMF Extractor CPU gathering is active, otherwise the SRM view of the CPU utilization (CCVUTILP)

CX10XDGS data reduction exit

The exit routine specified in the exit_name parameter for CX10XDGS is invoked on each system receiving a call from CX10XDGS. The exit routine is assumed to have the following attributes:

Location	JPA
State	Problem
Key	Any
Amode	31
Rmode	Any
Dispatchable unit mode	Task
Address space control mode	AR
Cross Memory Mode	PASN=SASN=HASN
Serialization	Enabled, unlocked
Type	Reentrant, Refreshable

When CX10XDGS calls your data reduction exit, the following parameters are passed in the format shown in Table 124 on page 580.

```
CALL exit_name,
     (answer_area_addr
     ,answer_area_alet
     ,answer_area_length
     ,output_area_length
     ,input_data_address
     ,exit_parm
     ,exit_parm_length)
```

Table 124 Parameters for calling the CX10XDGS exit

Parameter	Description	Format and length
answer_area_addr	address of the area where the exit routine can return the selected information The area resides in a data space owned by the MVS PAS.	AL4
answer_area_alet	ALET of the answer area provided on the answer_area_addr parameter	FL4
answer_area_length	length of the answer area provided on the answer_area_addr parameter CMF MONITOR provides an answer area in the length of the answer area the caller provided to CX10XDGS, rounded to the next multiple of 4096. However, the data returned by the data reduction exit routine must fit into the answer area the caller provided to CX10XDGS, including the common header and data headers created by CMF MONITOR.	FL4
output_area_length	length of the data the exit routine provided If this value is larger than answer_area_length, a return code and reason code are set, indicating that the length of the answer area was not sufficient.	FL4
input_area_address	address of the SMF record type 79 image in storage	FL4
exit_parm	parameter provided for the exit routine by the caller of CX10XDGS	XL4
exit_parm_length	length of the parameter area exit_parm passed to the exit routine	FL4

Sample DSECT for the common answer area header

The following DSECT defines the common callable service answer area header. It precedes the DSECTs of the other API result areas.

```
XDRH     DSECT ,                         Common Answer Area Header
XDRHNAM  DC    CL4'DSQA'                 Acronym: DSQA, DSRA or XDGH
XDRHVER  DC    F'1'                       Version: 1
XDRHLEN  DC    F'0'                      Length of returned data
XDRHTLEN DC    F'0'                      Length needed for all data
XDRHPLX  DC    CL8'PLEXNAME'              Name of sysplex
XDRHSOF  DC     A(XDRHSYS1-XDRH)            Offset from hdr to 1st sys
entry
XDRHSLN  DC    F'16'                     Length of one sys entry
XDRHSNO  DC    A(0)                      Number of sys entries
XDRHDOF  DC    A(0)                     Offset from hdr to 1st data sec.
XDRHDLN  DC    F'0'                      Length of one data section
XDRHDNO  DC    A(0)                      Number of data sections
XDRS     DSECT ,                        System entry: 1 per sys in plex
XDRSSNM  DC    CL8'SYSNAME'               MVS system name
XDRSSID  DS    CL4                        SMF system ID
XDRSRMF  DS    X                          Status flags
XDRSCMAC EQU   X'80'                      CMF active on system
XDRSDBAC EQU   X'40'                      XDS Active on system
XDRSRSVF EQU   X'3F'                      Reserved bits
         DS    XL3                        Reserved
*
XDRHSYS1 EQU   XDRH+XDRHSIZE,XDRSSIZE  1st system entry
```

XDRHNAM four-character acronym of the common header as follows:

- DSQA for CX10XDQY
- DSRA for CX10DSRC
- XDGH for CX10XDGS

XDRHVER version of the common header (initially set to 1)

XDRHLEN total length of the returned data

XDRHTLEN total length of the answer area needed to contain all of the requested data

XDRHPLX name of the sysplex on which the calling application is running

XDRHSOF offset from the header to the first system list entry SNM

XDRHSLN length of one system list entry (SNM,SID,CMF)

XDRHSNO number of system list entries (SNM,SID,CMF)

XDRHDOF offset from the header to the first data section

For the detailed layout, refer to the individual data section explanations.

XDRHDLN	length of one data section For a variable length data section, this field is zero. In this case, the length is stored in the individual data section header.
XDRHDNO	number of returned data sections
XDRSSNM	eight-character system name
XDRSSID	four-character SMF system ID If CMF MONITOR Online is not active on this system, this field contains hex zeros.
XDRSCMF	32-bit CMF status indicator containing XDRSCMAC, XDRSDBAC, and XDRSRVF The values of XDRSCMAC and XDRSDBAC are currently identical (either 0 or 1).
XDRSCMAC	bit 0 (high-order bit) indicates the status of CMF MONITOR Online on this system ('1'B = active)
XDRSDBAC	bit 1 indicates the status of XDS on this system ('1'B = active)
XDRSRVF	bits 2 to 31 are reserved

NOTE

The XDRS DSECT repeats once for each system in the sysplex.

Return codes for XDS APIs

Table 125 contains all return codes and reason codes for each of the three XDS APIs.

Table 125 Return codes for XDS APIs (part 1 of 6)

Return code	Reason code	API	Meaning
0	0	CX10XDQY CX10XDRC CX10XDGS	operation successful The answer area contains the requested data.
8	30	CX10XDQY CX10XDRC	warning - Timeouts detected Due to timeout situations, CX10XDQY or CX10XDRC could not return all the requested information. Request a smaller amount of information on one call of the API service.
8	31	CX10XDRC	warning - No such record One or more requested SMF records were not available for CX10XDRC. Either the SMF record data was overwritten by the wraparound management of the data buffer, or it never existed. Make sure that the elapsed time is not too large between calls to CX10XDQY and CX10XDRC, and that a valid token list is passed to CX10XDRC.
8	35	CX10XDGS	warning - Defaults taken
8	70	CX10XDQY CX10XDRC	warning - Answer area too small The answer area provided by the calling program was too small for the service to return all the requested information. The variable answer_area_length contains the length of the answer area provided for this CX10XDQY or CX10XDRC request. Provide an answer area large enough to contain all of the requested information.
12	0	CX10XDQY CX10XDRC CX10XDGS	error - XDS is not active You must start the MVS PAS with the XDS parameter and DC=START or DC=CPM on the local system.

Table 125 Return codes for XDS APIs (part 2 of 6)

Return code	Reason code	API	Meaning
12	1	CX10XDQY CX10XDRC CX10XDGS	error - System(s) inactive None of the systems specified for the CX10XDQY or CX10XDGS were active in the sysplex. For CX10XDRC, none of the record tokens specified belong to SMF records collected on systems that are currently active in the sysplex. Check the system ID list (smf_system_id_list, for CX10XDQY), record token list (smf_record_token_list, for CX10XDRC), or the system ID (system_id, for CX10XDGS) parameter and rerun the program.
12	5	CX10XDGS	error - Extractor interval ended The Extractor interval ended during the data-gathering phase while processing the CX10XDGS request. Rerun the program.
12	6	CX10XDGS	error - No CMF data available No data is currently available that matches the specification in the data_gathering_parm parameter of the CX10XDGS service. Check the parameters of CX10XDGS and rerun the program.
12	7	CX10XDGS	error - No Extractor data The CMF MONITOR Extractor was not active or was not running CXEN=Y. However, for the data gathering of certain SMF record subtypes (record type 79, subtypes 9, 11, 13, and 14) specified for the CX10XDGS service, an active Extractor with CXEN=Y is required. Verify that the Extractor is active on the systems from which data is requested, and rerun the program.
12	25	CX10XDGS	error - SRM STCPS facility not available The system resource manager (SRM) Store Channel Path Status (STCPS) facility is not available.
12	27	CX10XDGS	error - Transaction data not available Therefore, the transaction activity data (record type 79, subtype 8) cannot be returned.
12	30	CX10XDGS	error - Timeout Due to a timeout situation, CX10XDGS could not return the requested information. Request a smaller amount of information on one call of the CX10XDGS service.

Table 125 Return codes for XDS APIs (part 3 of 6)

Return code	Reason code	API	Meaning
12	36	CX10XDQY	error - No data returned by CX10XDQY No SMF data was found in the sysplex matching the specification provided by the smf_start_time, smf_end_time, smf_record_type_info, smf_record_type_list, smf_system_id_info, and smf_system_id_list parameters of the CX10XDQY service. Check the parameter specifications.
12	37	CX10XDQY CX10XDRC	error - XDS is inactive on all systems specified on the smf_system_id_info and smf_system_id_list parameters of the CX10XDQY service For CX10XDRC, an attempt was made to request SMF records from a system on which XDS is inactive. Start XDS on one or more systems in the sysplex. Check the list of system IDs passed to the CX10XDQY service.
12	70	CX10XDGS	error - Answer area too small The answer area provided by the calling program was too small for the service to return all the requested information. The variable answer_area_length contains the length of the answer area provided for this CX10XDGS request. Provide an answer area large enough to contain all the requested information.
16	0	CX10XDQY CX10XDRC CX10XDGS	severe error - CMF encountered a severe error This situation is normally accompanied by error messages in the PAS address space, a dump, or both.
16	41	CX10XDQY	severe error - The calling program specified an invalid value for the request type (request_type) parameter for CX10XDQY Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	42	CX10XDQY	severe error - The calling program specified an invalid value for the interval/range start or end time (start_time or end_time) parameter (YYYYMMDDHHMMSS) on the CX10XDQY service; this includes wrong-formatted parameters and out-of-range or invalid dates, for example: '19930000...' or '19930229...' Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.

Table 125 Return codes for XDS APIs (part 4 of 6)

Return code	Reason code	API	Meaning
16	43	CX10XDQY	severe error **-** The calling program specified an invalid value for the SMF record type (smf_record_type_info) parameter (INCLUDE / EXCLUDE / ALL) of the CX10XDQY service Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	44	CX10XDQY	severe error **-** The calling program specified an invalid value for the SMF system ID (smf_system_id_info) parameter (INCLUDE / EXCLUDE / ALL) of the CX10XDQY service Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	46	CX10XDGS	severe error **-** A bad SMF record type or subtype (rty or sty) was specified for the CX10XDGS service Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	53	CX10XDQY	severe error **-** An invalid SMF record type or subtype was specified in the record type list (smf_record_type_list) for the CX10XDQY service Either the length of the list was negative or a record type was out of the range of 0 to 255. Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	54	CX10XDQY	severe error **-** An invalid SMF system ID was specified in the system ID list (smf_system_id_list) for the CX10XDQY service, or the length of the list was negative Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	55	CX10XDQY	severe error **-** An invalid data time interval (start_time or end_time) was specified for the CX10XDQY service, that is, the start time is greater than or equal to the end time Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	56	CX10XDQY	severe error **-** An empty SMF record type and subtype list (smf_record_type_list and smf_record_type_info = INCLUDE) was specified for the CX10XDQY service Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.

Table 125 Return codes for XDS APIs (part 5 of 6)

Return code	Reason code	API	Meaning
16	57	CX10XDQY	severe error - An empty SMF system ID list (smf_system_id_list and smf_system_id_info = INCLUDE) was specified for the CX10XDQY service Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	58	CX10XDRC	severe error - An empty record token list (smf_record_token_list) was specified for the CX10XDRC service Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	60	CX10XDQY CX10XDRC CX10XDGS	severe error - CMF could not access one or more of the parameters Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	61	CX10XDQY CX10XDRC CX10XDGS	severe error - CMF could not access the answer area through the specified ALET (answer_area_alet) Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	70	CX10XDQY CX10XDRC CX10XDGS	severe error - The answer area provided by the calling program (answer_area_addr and answer_area_length) was too small to contain even the header information Examine your program to locate the CALL that caused the error condition. Correct the wrong statements and rerun your program.
16	80	CX10XDQY CX10XDRC CX10XDGS	severe error - The user is not authorized to call XDS for SMF data (CX10XDQY and CX10XDRC) or type 79 data (CX10XDGS) Contact your local security administrator. See the CMF MONITOR *Customization Guide* for more information about security.
16	81	CX10XDQY CX10XDRC CX10XDGS	severe error - The calling program is not in task mode Rerun your program in the correct mode.
16	82	CX10XDQY CX10XDRC CX10XDGS	severe error - The calling program is not enabled Rerun your program in the correct mode.
16	83	CX10XDQY CX10XDRC CX10XDGS	severe error - The calling program is not unlocked Rerun your program in the correct mode.
16	90	CX10XDQY CX10XDRC CX10XDGS	severev error - CMF encountered a severe error when calling the service routine. This may be caused by a terminating PAS Restart the PAS and rerun your program.

Table 125 Return codes for XDS APIs (part 6 of 6)

Return code	Reason code	API	Meaning
16	91	CX10XDGS	severe error - CMF encountered a severe error when loading the service exit routine. The routine was not found Make sure the exit routine is properly installed on all systems to which the request is directed. Rerun your program.
16	92	CX10XDGS	severe error - CMF recognized a severe error when executing the service exit routine The exit completion code is provided in the answer area returned by the service. Correct the exit routine problems and rerun your program.
20	0	CX10XDQY CX10XDRC CX10XDGS	unrecoverable error - An unrecoverable CMF error was encountered during the processing of the requested service This situation is normally accompanied by error messages sent to the system console, a dump, or both.

Implementing the CX10GVID API

CX10GVID returns control to your program after the service was completed successfully. The answer area contains an SMF type 79 record.

When CMF's CX10GVID API is implemented, the values appearing on the SDSF DA screen are calculated by the API. The API prepares either a type 79-1 or a type 79-2 SMF record image, whichever type value is requested, and returns it to the caller—in this case, SDSF.

To support CX10XDGS, as well as other MVS performance tools, CX10GVID can also return record images for all other SMF type 79 subtypes. The CMFMON write facility (CX10GV79), not the CMF Extractor, writes the type 79 records to SMF or CMF data sets.

The CMF type 79 API is distributed in *hilevel*.BBLINK with the module name of CX10GVID and an alias name of ERBSMFI assigned to it. ERBSMFI is the name of the RMF type 79 API.

Application programs written to use the RMF ERBSMFI API can also use the CMF CX10GVID API. To do this, you must make the *hilevel*.BBLINK library available to the application by including it in a //STEPLIB or system linklist data set. Alternatively, you can copy or link edit CX10GVID to a library that is accessible to the desired application. If RMF is also present on your system, you must make sure that the intended API routine (either ERBSMFI distributed with RMF or ERBSMFI distributed with CMF) is available to the applications requesting it. If both modules are accessible, the MVS control program selects the first ERBSMFI found according to the MVS rules of load module search. For more information, see “CMF APIs” on page 592.

Customizing the Extractor to get data

Some data is available only when the MVS PAS is running with CXEN=Y and the appropriate sampler is active in CPM mode. For example, in a PR/SM environment, the CPU sampler must be active in CPM mode for the API to return the system CPU utilization. When the CPU sampler is not active (in CPM mode) in a PR/SM environment, the API returns a value of -1 (x’FFFFFFFF) instead of the CPU utilization value.

Table 126 lists the Extractor samplers that are required for SMF 79 record subtypes:

Table 126 Extractor samplers required for SMF 79 record subtypes

Subtype	Data description	Required CMF Extractor control statement
8	Transaction Activity	WORKLOAD
9	Device Activity	DEVICE
11	Paging Activity	ASMDATA
13	I/O Queuing Activity for 438*x* and 308*x* processors	IOQ
14	I/O Queuing Activity for ES/9000 and 3090 processors	IOQ
All samplers must be active in CPM mode.		

Calling CX10GVID

When CMF's CX10GVID API is implemented, an application program calls program CX10GVID using standard MVS linkage conventions and passing the parameter list. The invoking program must provide a buffer into which the API returns the requested data.

CX10GVID must be called in 31-bit addressing mode. It can be invoked by unauthorized programs, but the fields in Table 127 are returned only when the caller is running in Supervisor state or is APF-authorized:

Table 127 Returned fields (APF-authorized or running in Supervisor state)

Subtype	Field	Description
9	R799CUB R799DVB R799DPB	control unit busy delay time device busy delay time director port busy delay time
11	R79BDEVN R79BCU	page data set device name page data set control unit name

When a caller requests subtype 2 record images, the BBX subsystem (BBXS) must be active or the subtype 2 record images returned are incomplete. If BBXS is not active, the following occurs:

- The real storage utilization fields of subtype 2 (listed below) contain null data.
- Bit R792RSM of byte R792FLG is on, signifying that the fields are invalid.

Subtype 2 real storage utilization fields affected by BBXS are shown in Table 128.

Table 128 Subtype 2 real storage utilization fields affected by BBXS

Field	Description
R792PRFX	total fixed frames
R792FXBL	fixed frames below the 16-MB line
R792NLQF	non-LSQA fixed frames
R792LSQA	LSQA fixed frames
R792SLQR	LSQA real storage pages
R792LSQE	LSQA expanded storage pages

When a caller requests subtype 3 record images, BBXS must be active or the subtype 3 record cannot be returned.

NOTE

See the *MAINVIEW Administration Guide* for information about initializing BBXS.

General-Purpose registers

At entry to CX10GVID, the following general-purpose registers must contain the values described in Table 129:

Table 129 General-Purpose register values

Register	Value
1	address of the parameter list
13	address of a standard, 72-byte save area
14	return address
15	entry-point address of CX10GVID
Note: Registers 2 through 12 are preserved by CX10GVID.	

Return codes for CX10GVID

Return codes for CX10GVID are different from those of the XDS APIs. After return from CX10GVID, register 15 contains one of the following return codes:

Table 130 Return codes (part 1 of 2)

Return code	Description
0	all requested data in the buffer was returned
4	invalid syntax, no buffer address
8	operating system not supported
16	no data currently available
20	ESTAE macro failed
24	API abend, or GETMAIN failure
28	data would not all fit in buffer; partial data returned
32	data not available; CMF sampler not active
36	data not available; sampler is recording

Table 130 Return codes (part 2 of 2)

Return code	Description
40	channel measurement not active
44	requested subtype is not applicable in goal mode
100	invalid input record type or subtype
104	record buffer too short; no data returned
108	request type not known

Additional CMF API considerations

The following sections provide additional information about CMF API routines:

- CMF APIs
- CMF APIs and SDSF

CMF APIs

IBM ships the ERB* API routines even if you do not license RMF; however, these API routines are disabled by way of the IFAPRD dynamic product enablement facility.

API names/aliases at issue are as follows:

CMF API name	RMF API name	CMF module
CX10GVID	ERBSMFI	CX10GVID
CX10XDGS	ERB2XDGS	CX10XDU0
CX10DSQY	ERBDSQRY	CX10XDU0
CX10DSRC	ERBDSREC	CX10XDU0
CX10XDU0	ERB3XDRS	CX10XDU0
CX10XDGX	ERB2XSMF	CX10XDGX

RMF APIs reside in SYS1.SERBLINK. Ensure that MVS rules of module search find the APIs in the *hilevel*.BBLINK library before SYS1.SERBLINK.

CMF APIs and SDSF

Various releases of the IBM SDSF product use the type 79-record APIs to get system- and job-related performance information. The two type 79-record APIs are CX10GVID (ERBSMFI)—used to obtain local system data, and CX10XDGS (ERB2XDGS)—used to obtain data from systems anywhere in the sysplex.

Prior to SDSF Release 1.3.3, SDSF did not call any external API to get performance information. Instead, SDSF acquired all information displayed with the DA subcommand directly.

As of SDSF Release 1.3.3, SDSF began calling ERBSMFI, if available, to get performance information for the DA subcommand. CMF MONITOR release 4.3.1 shipped an ERBSMFI replacement called CX10GVID (with alias ERBSFI) that SDSF used, if available. If ERBSFI is not found—because of an installation error, for example—SDSF continues to function by obtaining the data itself. However, certain data might be inconsistent or missing.

As of SDSF Release 1.5.1, SDSF added a SYSNAME subcommand that indicates which system's data should be displayed by the DA command. SYSNAME results in calls to the ERB2XDGS API, which was implemented by CMF MONITOR release 5.2.1 and later. SDSF continues to use the ERBSMFI interface if the system being displayed is not in SYSNAME mode—that is, if the system is the local system. If ERB2XDGS is not found and SYSNAME is used to access another system, the message `RMF SYSPLEX NOT ACTIVE` results. See "CMF APIs" on page 592 for other situations in which this message appears. If ERBSMFI is not found, SDSF continues to function by obtaining the data itself, as in earlier releases.

Installation requirements for SDSF use of CMF APIs

Several installation tasks must be performed to have SDSF use the CMF ERBSMFI or ERB2XDGS aliases to obtain data.

1 Ensure that a release of SDSF/JES supporting the RMF and CMF APIs is running on your system:

- SDSF release 1.3.3 for local system use
- SDSF release 1.5.1 for remote system use through SYSNAME

2 Install the appropriate release of CMF to provide the APIs:

- CMF release 4.3.1 for ERBSMFI
- CMF release 5.2.1 or later for ERB2XDGS

3 Ensure that the BBLINK data set is available in the linklist (see "CMF APIs" on page 592 for additional considerations), or ensure that the BBLINK data set is available in the TSO logon procedure STEPLIB.

4 Specify CPU sampling in the CMFCPM*xx* PARMLIB member.

5 For SYSNAME (remote system) use, ensure that the following additional requirements have been met for the local (TSO) system and all target (SDSF DA SYSNAME) systems:

- A CAS is running with SPCF active and in the same XCF group as the other CASs
- An MVS PAS with XDS active. (By default, XDS=00 is distributed in the MVSPAS product.)
- The MVS PASs must all have DC=START in effect.
- The MVS PASs must have CXEN=Y in effect.

Chapter

10

Mapping CMF records created by CMF

The CMF MONITOR Extractor and the CMFMON Write Facility prepare SMF records. These records match the type 70-79 records defined in the *IBM System Management Facilities (SMF)* manual. They are supplemented by the CMF *user* records. All of these records are used by the CMF Analyzer, the MXG program product, and the MICS program product. You can also use these records to write your own reports.

This chapter discusses how to use members in the CMF MONITOR BBSAMP data set to map record formats created by CMF, as well as SMF record types 70-79. This chapter lists the CMF user record types that version 10.10 of Merrill Consultants' MXG supports.

CMF MONITOR provides enhanced *hilevel*.BBSAMP members that contain improved format information about CMF MONITOR's unique CMF user records, as well as SMF record types 70-79. The SMF type 70-79 record information in BBSAMP is provided in Assembler MACRO format.

CMF MONITOR provides a default SMF ID of 240 for its user records, but a different value could have been defined at the SMFRECID= parameter on the Extractor REPORT control statement (see "REPORT" on page 176).

The user record information in BBSAMP is provided in these three formats:

- Assembler MACROs
- C Structures
- SAS code

Assembler MACROs in BBSAMP for record types 70-79

IBM 370/ESA assembly language MACROs, which are included in BBSAMP, map SMF record types 70-79. These MACROS replace the equivalent RMF MACROs in your programs.

CMFSMF7x	replaces ERBSMF7*x*, where *x* is the final digit of the SMF record type

NOTE

Make sure that you specify a binary file transfer, since the the file is already in PC format in BBSAMP.

- These MACROs are used by CMF code. The labels generated will not be identical to the labels generated by the RMF MACROs. BMC Software recommends that you modify your existing programs to use the new MACRO names and the CMF-style labels.

  Alternately, you can add the parameter labels=rmf to the MACRO to generate RMF-compatible labels.
- These MACROs, by default, require HASM version 2 to compile. If you need to use Assembler XF or HASM version 1, you must add the parameter asm=xf to the MACRO.

Assembler MACROs in BBSAMP for user records

z/OS assembly language and the MACROs are included in BBSAMP map CMF user records. Two of the members provide the product section and record header information for all CMF user records.

CMFREC*xx*	where *xx* is the user record subtype
CMFRECCV	an internal MACRO for using the assembly language MACROs
CMFRECHD	the header format for all user records
CMFRECPR	the CMF MONITOR product section

C structures in BBSAMP

C structures are included in BBSAMP, which map CMF user records. These structures reduce the need to translate IBM 370/ESA assembly language definitions (DSECTs) of CMF user records into C.

CMFC*xx*	where *xx* is the user record subtype
CMFCSAMP	a sample program that reads subtype 4 records This C member is intended to be used as a guide to show how to use any of the CMFC*xx* members to write your own programs for processing these records.

SAS code in BBSAMP

SAS code is included in the BBSAMP data set for reading CMF user records. This code reads CMF records and creates corresponding SAS data sets for further processing. This code reduces the need to translate IBM 370/ESA assembly language definitions (DSECTs) of CMF user records into SAS input statements.

CMFS*xx*	where *xx* is the user record subtype This member is used with its associated CMFSK*xx* member. Each CMFSK*xx* member contains the SAS code for the CMF *xx* user record subtype. Descriptions of all fields read from the *xx* record are contained in each member.
CMFSHD	the record header format for CMF user records This member contains SAS code for reading the standard record header used for all CMF record subtypes. It is always used for reading CMF records.
CMFSK*xx*	where *xx* is the user record subtype These members contain a list of the variables kept in each of the SAS data sets created for each of the CMF user record subtypes. Each member is used with its associated CMFS*xx* member.
CMFSML	sample JCL and SAS code for reading and analyzing CMF user records Instructions for using this job are contained in this member.
CSMAPSAS	the COMMON STORAGE MONITOR records (subtype 29), containing a very detailed level of data This SAS member is intended to be used as a guide for writing customized in-house reports using these records. The format of CMF MONITOR subtype 29 record data is documented in the BBSAMP member CMFREC29.

Using CMF user records with MXG

Merrill Consultants' MXG version 10.10 product fully supports the CMF user record types shown in Table 131.

Table 131 CMF user record types supported by MXG

CMF user record type	Record description	CMF Extractor statement	MXG data set name
240-00	SRM constants, installation performance specifications, and Extractor control cards data	REPORT	CMFDEVIC CMFDOM CMFIPS CMFOBJ CMFPG CMFSRMC CMFEXTCC CMFEXTPG CMFEXTRT
240-01	CPU data	CPU	CMFCPUQ CMFCPUS
240-02	ASM data	ASMDATA	CMF02PSD
240-03	paging data	PAGING	CMF03PGS
240-05	device data	DEVICE	CMF05DDS CMF05TDS
240-06	Extractor summary data	EXTSUM	CMF06GDA CMF06JDS
240-09	ASM data	ASMDATA	CMFASMQ
240-20	TSO command summary record data	TSODATA	CMF20CCS CMF20CSS
240-21	TSO user summary record data	TSODATA	CMF21USS
Note: For details on each CMF Extractor statement, see the related section in Chapter 6, "Extractor control statements."			
240-29	COMMON STORAGE MONITOR records	CSMON	CMF29COS CMF29CJS CMF29CDS

Part 4

Appendixes

This part presents the following topics:

A

Statistical considerations

Several of the reports produced by CMF MONITOR give standard deviations for various measures, such as device busy time, TSO response time, or CPU utilization. This appendix discusses how the standard deviation affects the statistical accuracy of data in reports.

It is important to understand the impact of the standard deviation because it is a factor to consider when using CMF MONITOR report data to tune your system.

Standard deviation, the mean, and the mode

The average of a measurement in a CMF MONITOR report is the mean value for that measurement. The standard deviation of a measurement in a CMF MONITOR report is a value signifying the degree of variation that can occur from the mean for that measurement.

A small standard deviation, or small degree of variation, indicates that most of the extracted measurement values are close to the average or mean value. A large standard deviation, or large degree of variation, indicates that the measurement values are widespread in relationship to the mean.

Figure 101 on page 604 shows the relationship of standard deviations to the mean.

Figure 101 Relationship of large and small standard deviations to the mean

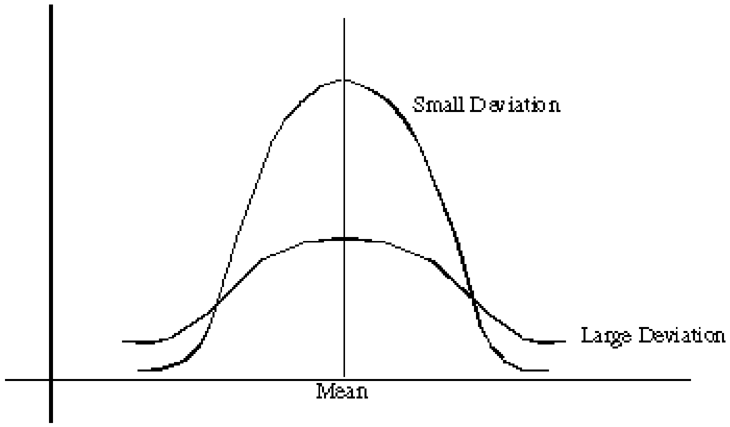

The standard deviation is particularly valuable when analyzing average TSO response time, where a high standard deviation can indicate irregular service to the end user.

A mode is generally used in reference to distribution graphs. Modes represent peaks in graphed values. A graph can have any number of modes. All that is required to graph a mode is for the preceding and following values to be less than the mode value.

Calculating standard deviation

The equation used to calculate the standard deviation is shown in Figure 102.

Figure 102 Equation for calculating standard deviation

$$\frac{\sqrt{N\sum_{i=1}^{N} X_i^2 - \left(\sum_{i=1}^{N} X_i\right)^2}}{N}$$

where

N	is the number of samples
$\mathbf{X_i}$	is the value of the variable for the *i*th sample
i	is the sample index

A record interval occurs when the CMF Extractor terminates data collection to write a record and start a new interval. This action is controlled by the INTERVAL parameter of the Extractor REPORT control statement. If there is only one sample, the standard deviation is zero.

Statistical accuracy

Due to the sampling technique used, accurate results are obtained when the number of samples is significant, such as 10,000 samples. Therefore, not only should the standard deviation of a measurement be considered when analyzing report data, but the number of samples counts should also be considered. The sample counts produced are shown at the top of the report. (See “Understanding report headings” on page 101.)

The measures reported by CMF MONITOR are a percentage (P) of the total number of samples taken (N) for which the measured conditions were true.

Statistical measures (with errors that are normally distributed) are usually expressed as a *percentage* (P) plus or minus a *confidence interval* (E) with a *confidence level* of (C).

- The confidence interval is an estimate of the maximum error from the true value of P.
- The confidence level is the probability that the difference between P and the true value is less than (E).

To calculate the statistical error, refer to Figure 103 and locate the following:

- number of samples taken by the Extractor (the N-axis)
- desired confidence level (one of the plotted diagonal lines)

The intersection point yields the uncorrected value for the confidence interval (E).

Figure 103 Confidence levels for P=50%

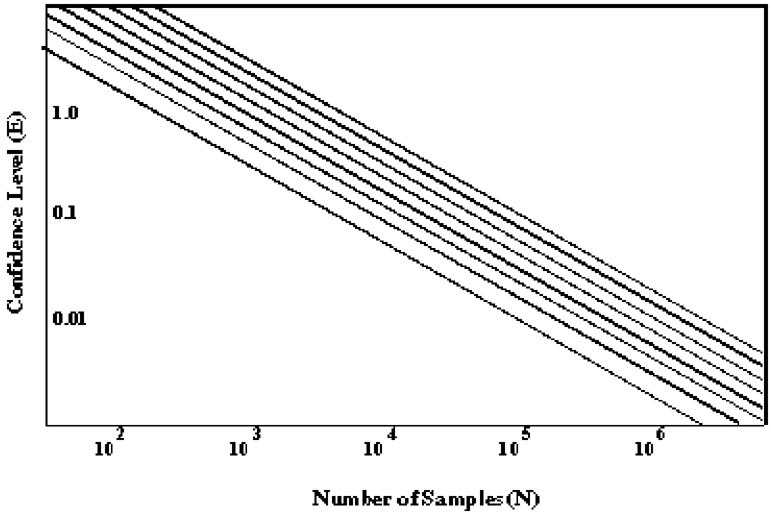

NOTE

Diagonal lines indicate confidence interval (E) with percentages (P) shown for each.

This confidence interval value is valid when the percentage (P) is 50%. A correction factor can be applied for other values of (P). To determine the correction factor to (E), refer to Figure 104 on page 607.

The true confidence interval is the product of the correction factor multiplied by the value of (E) determined above.

For example, if a measure reported by the Analyzer is 10%, the desired confidence level is 95%; if the Extractor took 5000 samples, the uncorrected confidence interval is plus or minus 1.5%. Since the correction factor for a 10% measure is 0.64, then the corrected confidence interval is 0.64 x 1.5% = 0.96%.

In other words, the analyst can expect only 1 chance in 20 (95% confidence level) that the actual value (reported as 10%) was less than 9.04% or greater than 10.96%.

Figure 104 Correction factors for confidence intervals

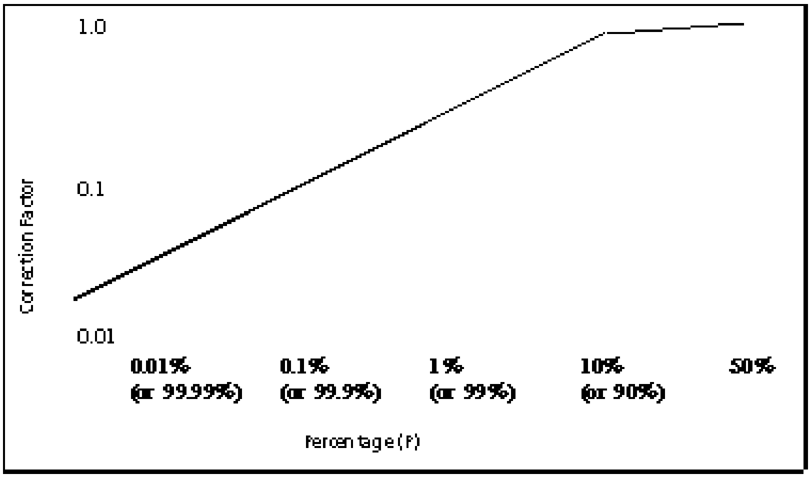

B

Workload measurement

The SRM uses the MVS Workload Manager (WLM) to perform resource management functions. With WLM, each sysplex is managed by a service definition. The service definition consists of one or more service policies, which are set up to define the workload goals for your system.

CMF MONITOR provides three reports that contain information gathered by WLM:

- Extractor Summary Report (Service Class Activity Section)
- Performance Summary Report (Service Class Activity Section)
- Workload Manager Goal Mode Report

These reports help you determine the extent to which your workload goals are being met. To interpret the values of these reports, you must be familiar with terms that describe goal mode performance. This appendix provides information about goal mode terminology, to assist you in interpreting your goal mode performance.

Service definition

The service definition you set up contains all the information required by WLM to manage your workloads. The service definition consists of

- Service policies
- Classification rules
- Resource groups
- Workloads
- Service classes

Service policy

A service policy is a named set of performance goals used by WLM. Different policies can be set up for different system requirements, but only one policy can be active for a particular sysplex at any one time. WLM uses the goals defined in the active service policy to schedule the appropriate resources for work in the sysplex.

Classification rules

Classification rules are the rules used to associate incoming work with a service class.

Resource group

A resource group is a service class or group of service classes that may be defined within a service policy. By defining a resource group, you can assign the amount of processing capacity across one or more MVS images for those service classes within the resource group. Defining a resource group within a service policy is not required.

Workload

A workload is a group of service classes that are tracked and managed as a unit, usually because they have something in common.

Service class

A group of work that has the same performance goals, resource requirements, or business importance. Up to 1000 service classes can be defined, and you can assign a performance goal for each service class.

Service class period

A service class period has a service goal and importance level assigned to a service class for a specific duration. Generally, multiple periods are created for workloads that have changing resource requirements, such as TSO. As a service class consumes more resources than are defined for its duration, it moves to a different period within the service class. Each service class can have up to 8 periods defined.

Goals

You can assign a goal to each service class. The assigned goals, along with the assigned importance value, help the WLM determine the distribution of resources. The four goal types are

- Average response time
- Percentage response time
- Execution velocity percentage
- Discretionary

A fifth type of goal, System, is assigned to workloads that need to receive the highest priority services.

Importance

Importance is the degree of importance of a service goal relative to other service class goals. The values for importance are

1	Highest
2	High
3	Medium
4	Low
5	Lowest

In some instances, a goal may not have an importance value.

- For a Discretionary goal, work is run using any system resources not required to meet the goals of other work.
- For a System goal, work is run for all address spaces requiring high priority service.

Performance index

Performance index is a relative calculation to determine how well your service classes are meeting their goals.

- A value of 1.0 shows it is exactly meeting its goal.
- A value of less than 1.0 shows the service class is exceeding its goal.
- A value greater than 1.0 shows the service class is not meeting its goal.

The way the performance index is calculated depends on the type of service class.

- **Discretionary Goal**

  Service classes whose importance is defined as *discretionary* do not have goals, so this type of service class does not have a performance index.

- **System Goal**

  Service classes whose importance is defined as *system* do not have goals, so this type of service class does not have a performance index.

- **Execution Velocity Percentage Goal**

  For a service class defined with an *execution velocity percentage* goal, the performance index is calculated by dividing the defined execution velocity percentage goal by the actual execution velocity percentage.

  For example, if the goal for a service class is defined as an execution velocity percentage of 90 and the actual velocity percentage is 50, the performance index for that service class is calculated by dividing 90 by 50, for a performance index of 1.8. This would indicate the service class is not meeting its goals, since the performance index is greater than 1.

- **Average Response Time Goal**

  The performance index for a service class defined with an *average response time* goal is calculated by dividing the actual response time by the defined average response time goal.

  For example, if the average actual response time is .25 seconds and the defined average response time goal is .5 seconds, the performance index for that service class is calculated by dividing .25 by .5, for a performance index of .5. This would indicate the service class is exceeding its goals, since the performance index is less than 1.

- **Percentage Average Response Time Goal**

  The performance index for a service class defined with a *percentage average response time* goal is calculated by dividing the actual response time by the response time goal. The actual response time must be calculated by determining which response time *bucket* (from the SMF 72-3 response time distribution data section) contains the transaction corresponding to the percentage goal. The number of transactions required to meet the goal is calculated as the total number of transactions during the interval multiplied by the percentage portion of the goal.

For example, suppose a service class has a percentage response time goal of 90% of transactions executing in less than .5 seconds. If the number of transactions for that service class is 50, the goal is for at least 45 transactions (90% of 50) to finish within .5 seconds.

The actual response time is calculated in the following way:

A response time distribution is set up, which contains 14 *buckets*. In this example, they have the following percentages and values:

1. count of transactions completed with response times <50% of the goal (less than or equal to .25 seconds): 10
2. count of transactions completed with response times >50% and less than 60% of the goal (between .25 and .30 seconds): 5
3. count of transactions completed with response times >60% and less than 70% of the goal (between .30 and .35 seconds): 4
4. count of transactions completed with response times >70% and less than 80% of the goal (between .35 and .40 seconds): 4
5. count of transactions completed with response times >80% and less than 90% of the goal (between .40 and .45 seconds): 2
6. count of transactions completed with response times >90% and less than 100% of the goal (between .45 and .50 seconds): 10
7. count of transactions completed with response times >100% and less than 110% of the goal (between .50 and .55 seconds): 0
8. count of transactions completed with response times >110% and less than 120% of the goal (between .55 and .60 seconds): 3
9. count of transactions completed with response times >120% and less than 130% of the goal (between .60 and .65 seconds): 4
10. count of transactions completed with response times >130% and less than 140% of the goal (between .65 and .70 seconds): 0
11. count of transactions completed with response times >140% and less than 150% of the goal (between .70 and .75 seconds): 0
12. count of transactions completed with response times >150% and less than 200% of the goal (between .75 and 1.0 seconds): 0
13. count of transactions completed with response times >200% and less than 400% of the goal (between 1.0 and 2.0 seconds): 1

14. count of transactions completed with response times >400% of the goal (greater than 2.0 seconds): 2

The actual response time is the response time belonging to whichever *bucket* the 45th fastest transaction fell into. If the 45th fastest transaction finished in .63 seconds, the actual average response time would be .60, because the value of .63 is in the *bucket* with that value.

In this example, the performance index is calculated as 1.2 (the actual value of .6 divided by the goal value of .5).

Although 80% of the transactions finished in less than .50 seconds (and the average response time for the period may have been well below the goal of .50 seconds), the performance period in this example did not meet its goal, since the goal required that 90% of the transactions complete in less than .50 seconds. The fact that the goal was not met is reflected in the performance index, which is greater than 1.

Execution Velocity—Execution velocity is defined as the total number of samples where an address space was using CPU divided by the number of times an address space was using CPU plus the number of general execution delays the address space experienced. This value is then multipled by 100 to get a percentage execution velocity.

The following is a list of general execution delays, as defined by WLM:

CPU delay	TCB or SRB waiting to be dispatched or a TCB waiting for a local lock
CPU capping delay	TCB or SRB marked nondispatchable because a resource group maximum is being enforced
Swap-in delay	swap-in started, but not completed
MPL delay	ready, but swap-in not started
Aux page from private	delay waiting for a private page to be brought in from auxiliary storage
Aux page from common	delay waiting for a common page to be brought in from auxiliary storage
Aux page from cross-mem	delay waiting for a cross-memory page to be brought in from auxiliary storage
Aux page from VIO	delay waiting for a VIO page to be brought in from auxiliary storage
Aux page from std hiper	delay waiting for a standard hiperspace page to be brought in from auxiliary storage
Aux page from ESO hiper	delay waiting for an ESO hiperspace page to be brought in from auxiliary storage
Shared paging	delay waiting for a shared storage page to be brought in

DASD I/O delay samples	number of samples of work delayed for paging DASD I/O
Queue delay samples	number of samples of work that is waiting for a server
Server private area paging delay samples	number of samples delayed for private area paging for a server
Server space VIO paging delay samples	number of samples delayed for VIO paging for a server
Server hiperspace paging delay samples	number of samples delayed for hiperspace paging for a server
Server MPL delay samples	number of samples delayed for MPL for a server
Server swap-in delay samples	number of samples delayed for swap-in for a server
zAAP	zAAP delay; work that has been delayed while waiting to be dispatched on a zAAP
zIIP	zIIP delay; work that has been delayed while waiting to be dispatched on a zIIP

Appendix

C

Measure and trace values

The information presented in Values for EXCEPTS and GRAPH statements defines the acceptable values for the MEASURE parameter of the Analyzer EXCEPTS (see “EXCEPTS” on page 252) and GRAPH (see “GRAPH” on page 259) control statements.

The information presented in “Traceable data fields” on page 630 defines the acceptable values for the MEASURE parameter of the Analyzer EXCEPTS, GRAPH control statements and the FIELD parameter of the Extractor TRACE76 (see “TRACE” on page 312) control statement.

Values for EXCEPTS and GRAPH statements

Table 132 lists the allowable values that can be defined to the MEASURE parameter of the Analyzer EXCEPTS and GRAPH control statements.

NOTE

VSDB and VSDA labels used in the calculations in Table 132 are DSECT fields in VSDATA. They define areas within major SMF78RCD triplet areas associated with CSA and SQA data.

Table 132 Values for EXCEPTS and GRAPH statements (part 1 of 12)

Measure	Description	Record type
APPC-MAX	maximum number of APPC address spaces	SMF70
APPC-MIN	minimum number of APPC address spaces	SMF70
APQ	average amount of available page frames in K-bytes	SMF70
AVTSK	average number of address spaces in storage and ready to execute	SMF70
APPC-AVG	average number of APPC address spaces	SMF70
BTCH-AVG	average number of batch jobs	SMF70
BTCH-CPU	percentage of time during the measurement interval that the CPU was executing on behalf of a batch job	CMF-01
BTCH-MAX	maximum number of batch jobs	SMF70
BTCH-MIN	minimum number of batch jobs	SMF70
CHA-*an*	average rate of SSCH per second for the specified LCU; *nn* is the two-character logical control unit number in hexadecimal	SMF78
CHC-*nnnn*	percentage of requests that were satisfied from cache for control unit *nnnn* (number of hits / number of requests)	CMF-27
CHD-*nnnn*	percentage of requests that were satisfied from cache for device *nnnn* (number of hits / number of requests)	CMF-27
CHN-*nn*	percentage of time during the measurement interval that the specified channel path was observed to be busy; *nn* is the two-character channel path in hexadecimal	SMF73
CHS-*nn*	service time in milliseconds for the specified LCU; *nn* is the two-character logical control unit number in hexadecimal	SMF74
CMC-*nnnn*	percentage of requests that were not satisfied from cache for control unit *nnnn* (number of misses / number of requests)	CMF-27
CMD-*nnnn*	percentage of requests that were not satisfied from cache for device *nnnn* (number of misses / number of requests)	CMF-27
CNN-*nnn*	average number of milliseconds that the device was processing an SSCH instruction and transferring data; *nnn* is the three- or four-character device address	SMF74
COEFCPU	IPS CPU service definition coefficient	SMF72
COEFIO	IPS I / O service definition coefficient	SMF72
COEFMEM	IPS memory service definition coefficient	SMF72
COEFSRB	IPS SRB service definition coefficient	SMF72
COMRECL	common system area page reclaim rate per minute	SMF71
CPK-*n*	percentage of time during the measurement interval that the processor was observed to be executing in the specified key; *n* is a one-character hexadecimal number from 0 to F	CMF-01

Table 132 Values for EXCEPTS and GRAPH statements (part 2 of 12)

Measure	Description	Record type
CPU	PR/SM environment: the LPAR CPU busy percentage during the specified interval This figure is the percentage of processor dispatch capacity used by the home partition (the partition on which the records were extracted). The processor dispatch capacity is equal to the interval length multiplied by the number of logical processors assigned to the partition. The formula is **Wait completion on:** (dispatch time - wait time) / interval **Wait completion off:** dispatch time / interval Native mode: average CPU busy percentage during the specified interval The formula is (interval - wait time) / interval. In this environment, this measure has the same meaning as CPUBZMVS.	SMF70
CPUBZMVS	**PR/SM environment**: the average percentage of time during the specified interval that the CPU either was not dispatched or was dispatched and busy This value is calculated as (interval length - wait time) / interval length. This measure is not available if the SMF type 70 records were extracted from a partition running a version of MVS prior to 4.3 and at least one nondedicated processor has wait completion off. **non-PR/SM environment**: the average percent of time during the specified interval that the CPU was busy	SMF70
CPUCAP	percentage of the partition used by the home partition This figure is equal to the CPU busy for a PR/SM environment multiplied by the portion of the environment allotted to the home partition.	SMF78
CRC-*nnnn*	percentage of read requests that were satisfied from cache for control unit *nnnn* (read hits/read requests)	CMF-27
CRD-*nnnn*	percentage of read requests that were satisfied from cache for device *nnnn* (read hits/read requests)	CMF-27
CSA-AVG	average CSA allocated area size as a percentage of total CSA available above and below the 16-megabyte line; this area contains both used and free CSA storage	SMF78
CSA-MAX	maximum CSA allocated area size as a percentage of total CSA available above and below the 16-megabyte line; this area contains both used and free CSA storage	SMF78
CSA-MIN	minimum CSA allocated area size as a percentage of total CSA available above and below the 16-megabyte line; this area contains used and free CSA storage	SMF78

Table 132 Values for EXCEPTS and GRAPH statements (part 3 of 12)

Measure	Description	Record type
CSAUSE-*x*	average amount of CSA used: ■ above 16 MB, when *x* = **A** ■ below 16 MB, when *x* = **B** The CSA in use includes fragmentation caused by allocation of 4-K storage blocks.	SMF78
CSC-*nnnn*	percentage of read requests that were not satisfied from cache for control unit *nnnn* (read misses / read requests)	CMF-27
CSD-*nnnn*	percentage of read requests that were not satisfied from cache for device *nnnn* (read misses / read requests)	CMF-27
CSF-AVG	average free CSA (within the size of CSA allocated area) as a percentage of total CSA available above and below the 16-megabyte line	SMF78
CSF-MAX	maximum free CSA (within the size of CSA allocated area) as a percentage of total CSA available above and below the 16-megabyte line	SMF78
CSF-MIN	minimum free CSA (within the size of CSA allocated area) as a percentage of total CSA available above and below the 16-megabyte line	SMF78
CSU-AVG	average CSA used (within the size of CSA allocated area) as a percentage of the total CSA available above and below the 16-megabyte line	SMF78
CSU-MAX	maximum CSA used (within the size of CSA allocated area) as a percentage of total CSA available above and below the 16-megabyte line	SMF78
CSU-MIN	minimum CSA used (within the size of CSA allocated area) as a percentage of total CSA available above and below the 16-megabyte line	SMF78
CUB-*nnn*	average control unit busy delay time: the average number of milliseconds of delay for an I/O request because the control unit was busy If the device is shared at the control unit level, the contention can be caused by the sharing system. If the device is not shared at the control unit level, the contention is because of other device activity using the same or an alternate path serviced by the control unit. *nnn* is the three- or four-character device address.	SMF74
CWC-*nnnn*	percentage of write requests that were satisfied from cache for control unit *nnnn* (write hits / write requests)	CMF-27
CWD-*nnnn*	percentage of write requests that were satisfied from cache for device *nnnn* (write hits / write requests)	CMF-27
CXC-*nnnn*	percentage of write requests that were not satisfied from cache for control unit *nnnn* (write misses / write requests)	CMF-27
CXD-*nnnn*	percentage of write requests that were not satisfied from cache for device *nnnn* (write misses / write requests)	CMF-27
DEV-*nnn*	percentage of time during the measurement interval that the specified device was busy; *nnn* is the three- or four-character device address	SMF74
DIS-*nnn*	average number of milliseconds during which the device was processing an SSCH instruction, but was not transferring data; *nnn* is the three- or four-character device address	SMF74

Table 132 Values for EXCEPTS and GRAPH statements (part 4 of 12)

Measure	Description	Record type
DSO-*nnn*	average number of data sets opened on the specified device This figure excludes measurement intervals during which the designated device was varied online or offline, as well as multiple exposure devices; *nnn* is the three- or four-character device address.	SMF74
DSL-GSRB	percentage of CPU busy time that the processor was executing as a global SRB in the disabled state	CMF-01
DSL-LSRB	percentage of CPU busy time that the processor was executing as a local SRB in the disabled state	CMF-01
DSL-TCB	percentage of CPU busy time that the processor was executing in TCB mode in the disabled state	CMF-01
DSL-TOT	percentage of CPU busy time that the processor was executing in the disabled state	CMF-01
DVA-*nnn*	average rate of SSCH per second for the specified device	SMF74
DVB-*nnn*	average device busy delay time This figure is the average number of milliseconds of delay for an I/O request because the device was busy; *nnn* is the three- or four-character device address.	SMF74
DVS-*nnn*	service time in milliseconds for the specified device; *nnn* is the three- or four-character device address	SMF74
ENL-GSRB	percentage of CPU busy time that the processor was executing as a global SRB in the enabled state	CMF-01
ENL-LSRB	percentage of CPU busy time that the processor was executing as a local SRB in the enabled state	CMF-01
ENL-TCB	percentage of CPU busy time that the processor was executing in TCB mode in the enabled state	CMF-01
ENL-TOT	percentage of CPU busy time that the processor was executing in the enabled state	CMF-01
ERR-*nnn*	percentage of time during the measurement interval that the specified device was in error recovery; *nnn* is the three- or four-character device address	CMF-05
EXP-AAV	average number of available expanded storage 4-K page frames	SMF71
IOIRATE	rate per second at which the processor handled I/O interruptions	SMF70
IOTPI	percentage of total I/O interrupts for the processor that were handled by the TPI instruction	SMF70
LCAP-##	percentage of the logical processor utilization for a logical partition relative to its weight factor, where the partition is specified by ## Valid values are hexadecimal digits between 0 and FF. The partition number is displayed under the column **No**. of the CMF MONITOR Online view LPARSTAZ.	SMF70

Table 132 Values for EXCEPTS and GRAPH statements (part 5 of 12)

Measure	Description	Record type
LPU-##	average percentage of logical processor utilization for a logical partition relative to the weight factor of the partition specified by ## Valid values are hexadecimal digits between 0 and FF. The partition number is displayed under the column **No**. of the CMF MONITOR Online view LPARSTAZ.	SMF70
MAX3420	maximum number of 3420 type tape drives that were concurrently allocated	CMF05TDA
MAX3480	maximum number of 3480 type tape drives that were concurrently allocated	CMF05TD8
MAX3490	maximum number of 3490 type tape drives that were concurrently allocated	CMF05TD9
MAX3590	maximum number of 3590 type tape drives that were concurrently allocated	CMF05MDV
NUCPAGES	average number of page frames in the nucleus, in K-bytes	SMF71
NTU-AVG	average number of TSO users	SMF70
NTU-MAX	maximum number of TSO users	SMF70
NTU-MIN	minimum number of TSO users	SMF70
OMVS-AVG	average number of OMVS address spaces	SMF70
OMVS-MAX	maximum number of OMVS address spaces	SMF70
OMVS-MIN	minimum number of OMVS address spaces	SMF70
OUS-*nnn*	average number of users on the SRM queue ready to execute for the specified domain Ready-user average is stored by the SRM in the DMDT; *nnn* is the domain number.	CMF-04
PAGEDPR	demand paging rate from DASD	SMF71
PAGEINS	total number of pages transferred from auxiliary storage to central storage, including VIO, swap, hiperspace, and block paging	SMF71
PAGEIO	total number of pages transferred between auxiliary storage and central storage, including VIO, swap, hiperspace, and block paging	SMF71
PAGEOUTS	total number of pages transferred from central storage to auxiliary storage, including VIO, swap, and hiperspace paging	SMF71
PAGERECL	total number of page reclaims, including VIO	SMF71
PAGESEC	total paging rate in pages per second between auxiliary storage and central storage, including VIO, swap, hiperspace, and block paging	SMF71
PAGESMIG	average rate of page migration from expanded storage to auxiliary storage	SMF71
PAGESTRN	average paging transfer rate to expanded storage	SMF71
PAGETIME	average number of milliseconds needed to complete an I/O request to a local page data set	SMF75

Table 132 Values for EXCEPTS and GRAPH statements (part 6 of 12)

Measure	Description	Record type
PCAP-##	percentage of the physical processor utilization for a logical partition relative to its weight factor, where the partition is specified by ## Valid values are hexadecimal digits between 0 and FF. The partition number is displayed under the column **No**. of the CMF MONITOR Online view LPARSTAZ.	SMF70
PEN-*nnn*	average number of milliseconds an I/O request must wait for hardware (such as a channel path or control unit), and time between the SSCH pending at the channel and the device active on the subchannel The time also includes delays caused by another processor reserving this device in a shared DASD environment; *nnn* is the three- or four-character device address.	SMF74
PGINPROC	average ASM queue length as seen by the SRM; average value does not consider samples when there was no queue	CMF-09
PPB	percentage of CPU busy time that the processor was executing in problem state	CMF-01
PPPAGES	average number of page frames in the private area, in K-bytes; includes fixed non-LSQA frames	SMF71
PPU-##	average percent of physical processor utilization for a logica Valid values are hexadecimal digits between 0 and FF. The partition number is displayed under the column **No.** of the CMF MONITOR Online view LPARSTAZ.	SMF70
PRIPAGI	private area page-in rate per second, excluding VIO and SWAP	SMF71
PRIPAGO	private area page-out rate per second, excluding VIO and SWAP	SMF71
PRIRECL	private area page reclaim rate per minute	SMF71
QDP-*nnn*	maximum queue depth of the device; *nnn* is the three- or four-character device address	CMF-05

Table 132 Values for EXCEPTS and GRAPH statements (part 7 of 12)

Measure	Description	Record type
QL / *xx*-*nn*	queue length measure, where *xx* is one of the following valid address space types: BA batch users IR in and ready to execute IN in storage OR out of storage and ready to execute OW out of storage and waiting to execute LR logically out of storage and ready to execute LW logically out of storage and waiting to execute ST started users TS TSO users *nn* percentage of time or range that an address space type had a specified queue length; valid queue lengths are 0 through 99 You can also specify one of the following values: AV requests the average queue length MN requests the minimum queue length MX requests the maximum queue length **Note**: Because queue lengths are summarized in the SMF type 70 records, in ranges such as 16 through 20, a queue length request of QL / IN-16 yields the same results as a request for QL / IN-20. Refer to the IBM reference manual, *System Programming Library: System Management Facilities (SMF)*, for a full description of the queue length distributions in the SMF type 70 records.	SMF70
QTM-*nnn*	average number of milliseconds that an I / O request for device *nnn* must wait on an IOS queue before an SSCH instruction is issued	SMF74
RMCTADJC	time to process 1 / 16 of a service unit in microseconds	SMF72
RWC-*nnnn*	ratio of read requests to write requests from cache for control unit *nnnn* (read requests / write requests)	CMF-27
RWD-*nnnn*	ratio of read requests to write requests from cache for device *nnnn* (read requests / write requests)	CMF-27
SKA-*n*	average CSA usage by storage key ID (*n*) as a percentage of total CSA available above and below the 16-megabyte line Valid storage keys are (0,1,2,3,4,5,6,7,8), where 8 is 8-F storage keys.	SMF78

Table 132 Values for EXCEPTS and GRAPH statements (part 8 of 12)

Measure	Description	Record type
SKM-*n*	minimum CSA usage by storage key ID (*n*) as a percentage of total CSA available above and below the 16-megabyte line Valid storage keys are (0,1,2,3,4,5,6,7,8), where 8 is 8-F storage keys.	SMF78
SKX-*n*	maximum CSA usage by storage key ID (*n*) as a percentage of total CSA available above and below the 16-megabyte line Valid storage keys are (0,1,2,3,4,5,6,7,8), where 8 is 8-F storage keys.	SMF78
SPA-*nnn*	average CSA / SQA usage by subpool ID (*nnn*) as a percentage of total CSA / SQA available below the 16-megabyte line Valid CSA subpool IDs are 227, 228, 231, 241. Valid SQA subpool IDs are 226, 239, 245.	SMF78
SPM-*nnn*	percentage of total CSA / SQA that are available below the 16-megabyte line; *nnn* is the subpool ID Valid CSA subpool IDs are 227, 228, 231, 241. Valid SQA subpool IDs are 226, 239, 245.	SMF78
SPX-*nnn*	maximum CSA / SQA usage by subpool ID (*nnn*) as a percentage of total CSA / SQA available below the 16-megabyte line Valid CSA subpool IDs are 227, 228, 231, 241. Valid SQA subpool IDs are 226, 239, 245.	SMF78
SQA-AVG	average SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line; area contains used and free SQA storage	SMF78
SQA-MAX	maximum SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line; area contains used and free SQA storage	SMF78
SQA-MIN	minimum SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line; area contains used and free SQA storage	SMF78
SQF-AVG	average free SQA that is available within SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line	SMF78
SQF-MAX	maximum free SQA that is available within SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line	SMF78
SQF-MIN	minimum free SQA that is available within SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line	SMF78
SQU-AVG	average SQA that is used within SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line	SMF78
SQU-MAX	maximum SQA that is used within SQA allocated area as percentage of total SQA available above and below the 16-megabyte line	SMF78
SQU-MIN	minimum SQA that is used within SQA allocated area as a percentage of total SQA available above and below the 16-megabyte line	SMF78

Table 132 Values for EXCEPTS and GRAPH statements (part 9 of 12)

Measure	Description	Record type
SRV-ABS	rate per second at which service is consumed while transactions are resident in central storage	SMF72
SRV-CPU	average rate per second at which transactions that were active but not necessarily in storage used TCB and SRB service units	SMF72
SRV-I/O	average rate per second at which transactions that were active but not necessarily in storage used IOC service units	SMF72
SRV-MEM	average rate per second at which transactions that were active but not necessarily in storage used MSO service units	SMF72
SRV-SRB	average rate per second at which transactions that were active but not necessarily in storage used SRB service units	SMF72
SRV-TCB	average rate per second at which transactions that were active but not necessarily in storage used TCB service units	SMF72
SRV-TOT	average rate per second at which transactions that were active but not necessarily in storage used ALL service units	SMF72
SSC-*nnn*	start subchannel count, where *nnn* is the device address	SMF74
STC-AVG	average number of Started Tasks	SMF70
STC-CPU	percentage of time during the measurement interval that the CPU was executing on behalf of a Started Task	SMF70
STC-MAX	maximum number of started users	SMF70
STC-MIN	minimum number of started users	SMF70
SUP	percentage of CPU busy time that the processor was in the supervisor state	CMF-01
SUPPAGES	average number of CSA pageable frames, in K-bytes	SMF71
SUSEC	rate at which service units are processed (service units per second)	SMF72
SWAPAGI	page swap-in rate per second	SMF71
SWAPAGO	page swap-out rate per second	SMF71
SWAPMIN	swaps per minute	SMF71
SWP-APS	number of swap-outs due to APPC WAIT	SMF71
SWP-ASD	number of swap-outs to auxiliary storage directly	SMF71
SWP-AST	number of swap-outs to auxiliary storage total	SMF71
SWP-ATR	number of swap-outs to auxiliary storage by transition This figure includes logical swaps that became physical swaps to auxiliary storage and swaps to expanded storage that were migrated to auxiliary storage.	SMF71
SWP-AXS	number of swap-outs due to an auxiliary storage shortage Such a shortage occurs when a percentage of all local page data set slots is allocated. This percentage value is an SRM-modifiable constant in the MCT (field MCCASMT1); the default is 70.	SMF71

Table 132 Values for EXCEPTS and GRAPH statements (part 10 of 12)

Measure	Description	Record type
SWP-DW	number of swaps due to a detected wait SRM detects a wait when an address space in storage is not dispatched for a processor-dependent period of time.	SMF71
SWP-ENQ	number of swaps due to an enqueue exchange Such an exchange occurs when one address space tries to enqueue on a resource held by another address space. The address space holding the resource is treated as nonswappable for the SRM-defined interval in the **RMPT** (field name **RMPTERV**). This value is user-defined in member IEAOPT of SYS1.PARMLIB.	SMF71
SWP-ESD	number of swap-outs to expanded storage directly	SMF71
SWP-ESE	number of swap-outs to expanded storage effectively This figure includes only those swaps to expanded storage that were not migrated to auxiliary storage.	SMF71
SWP-EXC	number of exchange swaps based on the SRM recommendation value A user was swapped out so another user with a higher workload level in the same domain could be swapped in. This SRM action keeps the MPL within the target MPL for that domain.	SMF71
SWP-ICS	number of swap-outs due to the need to improve central storage usage	SMF71
SWP-IPS	number of swap-outs due to the need to improve the system paging rate	SMF71
SWP-ITR	number of candidates for logical swap due to a wait for input buffers at a terminal One of the following actions will occur: ■ logical swap-out followed by a logical swap-in ■ logical swap-out followed by a physical swap-out due to a detected long think time ■ physical swap-out due to input terminal wait	SMF71
SWP-LSE	number of effective logical swap-outs; includes only those logical swaps that did not result in a physical swap to any medium	SMF71
SWP-LSI	number of logical swap-ins When a logical swap-out does not result in a detected long think-time swap, it is logically swapped back in.	SMF71
SWP-LSO	number of logical swap-outs; the number of candidates for logical swap that are not physically swapped out	SMF71
SWP-LST	number of logical swap-outs total	SMF71
SWP-LTT	number of detected long think-time swaps A logical swap-out fails to be logically swapped in.	SMF71

Table 132 Values for EXCEPTS and GRAPH statements (part 11 of 12)

Measure	Description	Record type
SWP-LW	number of swaps due to long waits A long wait occurs when a program issues a WAIT,LONG=YES macro.	SMF71
SWP-MIG	number of swap-outs to expanded storage that were migrated to auxiliary storage	SMF71
SWP-MRS	number of swap-outs made to make room to swap-in a user who has been swapped out too long	SMF71
SWP-OTR	number of candidates for logical swap due to a wait for output buffers at a terminal One of the following actions will occur: ■ logical swap-out followed by a logical swap-in ■ logical swap-out followed by a physical swap-out due to a detected long think time ■ physical swap-out due to input terminal wait	SMF71
SWP-REQ	number of swaps caused by the REQSWAP SYSEVENT REQSWAP is issued when the CONFIG storage command is issued and the address space that occupies the storage to be taken offline must be swapped out. REQSWAP is also issued when the PPT marks an EOB as nonswappable so that the address space can be swapped out and its LSQA swapped back in to preferred storage.	SMF71
SWP-RSS	number of swaps due to shortage of pageable storage Either the SRM or the RSM can detect a shortage of available real page frames. Users with the most fixed frames are swapped out.	SMF71
SWP-TOT	total number of swap sequences A swap sequence is a swap-out and a swap-in of an address space.	SMF71
SWP-TWI	number of terminal I/O wait swap-ins	SMF71
SWP-TWO	number of terminal I/O wait swap-outs	SMF71
SWP-TWT	number of swap-outs due to input terminal wait; number of candidates for logical swap that resulted in a physical swap-out	SMF71
SWP-TXS	number of swaps due to transition to nonswappable A transition swap occurs when the TRANSWAP SYSEVENT is issued. An address space is swapped out so its LSQA can be swapped back in to preferred storage.	SMF71
SWP-UNI	number of unilateral swaps The SRM swaps out an address space without swapping in another address space. This occurs when the MPL of a domain exceeds the target MPL specified for the domain.	SMF71
SYSPAGI	system area page-in rate per second, excluding VIO	SMF71
SYSPAGO	system area page-out rate per second, excluding VIO	SMF71

Table 132 Values for EXCEPTS and GRAPH statements (part 12 of 12)

Measure	Description	Record type
SYST	percentage of CPU busy time that the processor was executing in key zero	CMF-01
TAPES	the highest of MAX3420, MAX3480, MAX3490, or MAX3590	CMF-05
TPL-*nnn*	average target multiprogramming level for the specified domain; the TMPL stored by the SRM in the domain table (DMDT); *nnn* is the domain number	CMF-04
TSO-ABS	rate at which TSO users who logged off during the measurement interval consumed service units per second while transactions were resident in central storage	CMF-21
TSO-COMM	total number of TSO commands	CMF-20
TSO-COMR	rate per minute at which TSO commands are issued when there are commands to measure in a record interval; TSO commands are monitored using SYSEVENT 0 (TSEVENT)	CMF-20
TSO-CPU	percentage of time during the measurement interval that the CPU was executing on behalf of a TSO user	CMF-01
TSO-PAGR	page rate per second for TSO users who logged off during the measurement interval	CMF-21
TSO-PAGS	average number of pages stolen per TSO session	CMF-21
TSO-RESP	average response time in seconds for all TSO commands timed during the measurement interval Response time is the difference between SYSEVENTs 0 (TSEVENT) and 34 (TPUT).	CMF-20
TSO-SERV	rate at which TSO users who logged off during the measurement interval consumed service units per second while transactions were active	CMF-21
TSO-SWAP	average number of swap sequences per TSO session	CMF-21
TSO-SWPR	swap rate per minute for TSO users who logged off during the measurement interval	CMF-21
TSO-TRNS	average number of transactions per TSO session	CMF-21
VIOPAGI	VIO page-in rate per second	SMF71
VIOPAGO	VIO page-out rate per second	SMF71
VIORECL	VIO page reclaim rate per minute	SMF71

Traceable data fields

Control block field names and descriptions for traceable data fields are outlined in the following tables. These field names are acceptable values for the FIELD parameter of the Extractor TRACE76 control statement and MEASURE parameter of both the Analyzer GRAPH and EXCEPTS control statements.

The following control blocks field names are discussed in this section:

- Auxiliary Storage Manager Vector Table (ASMVT)—see next section
- Operations Measurement Data Gatherer area (OMDG)—see page 631
- System Resource Management Data area (SRM)—see page 631
- RSM Address Space Block Extension (RAX)—see page 632
- RSM Control and Enumeration Area—see page 633
- System Management Facilities Data Area (SMCA)—see page 634

Auxiliary Storage Manager Vector Table

Table 133 lists the control-block field names used with the Analyzer EXCEPTS and GRAPH report control statements and with the Extractor TRACE76 statement.

Table 133 Auxiliary Storage Manager Vector Table (ASMVT)

Field	Description
ASMERRS	bad slots on local page data sets
ASMIORQC	count of I/O requests completed and returned to RSM
ASMIORQR	count of I/O requests received by I/O control
ASMNVSC	total local slots allocated for non-VIO private area
ASMSLOTS	total local slots (sum of slots in open local page data sets)
ASMVSC	total local slots allocated for VIO private area pages

Operations Measurement Data Gatherer area

Table 134 describes the fields that appear in the Operations Measurement Data Gatherer area (OMDG).

Table 134 Operations Measurement Data Gatherer area (OMDG)

Field	Description
OMDGAMRE	maximum queue size of the action message retention facility (AMRF)
OMDGCMDI	number of commands per second
OMDGOREB	maximum queue size of operator reply entries (OREs)
OMDGWQEB	maximum queue size of the message line (WTO) queue elements (WQEs)
OMDGWTLI	number of writes to the log (WTLs) per second
OMDGWTOI	number of message lines (WTOs) per second

System Resource Management data area

Table 135 describes the fields that appear in the System Resource Management data area (SRM).

Table 135 System Resource Management (SRM) data area (part 1 of 2)

Field	Description
CCVCPUCT	number of online CPUs
CCVENQCT	number of users nonswappable for enqueue reasons
CCVRBSTD	recent base time of day comp
CCVRBSWT	recent base system wait time
CCVUTILP	system CPU utilization
CCVWTDB	alternate wait management
LSCTCNT	current number of logically swapped users for terminal wait
LSCTCNTW	number of users logically swapped for a long wait or detected wait
LSCTMTE	maximum think time allowed for logical swap candidate
MCVFRCNT	number of pages needed to be stolen by force steal routine
MCVMGAGE	extended storage area migration
MCVSBLTF	long-term percentage of eligible storage that is actually fixed
MCVSIPR	common page-in rate
MCVSTCRI	highest system UIC
MCVTWSS	common target working set size
RCVAFQA	average available frame count
RCVAVQC	AVQ low count

Table 135 System Resource Management (SRM) data area (part 2 of 2)

Field	Description
RCVCPUA	CPU usage average * 16
RCVFXIOP	percentage of central storage that is fixed or allocated for paging
RCVMFXA	average number of fixed frames for the system
RCVPAGRT	total paging rate
RCVPTR	paging rate
RCVSWPTM	time, in milliseconds, used by ASM to process a request to transfer a group of pages to or from a data set
RCVUICA	UIC average
RMCAAWSC	APPC verb service request wait swap count
RMCADWSC	detected wait physical swap count
RMCAEXSC	exchange on recommendation value swap count
RMCAFHLD	number of swaps failed because of an outstanding HOLD SYSEVENT
RMCAICSC	improved central storage utilization swap count
RMCAIPSC	improved demand page-in rate swap count
RMCALWSC	long wait physical swap count
RMCAMRSC	make room for an out-too-long address space swap count
RMCANQSC	CPU enqueue exchange swap count
RMCAOISC	OMVS input wait
RMCAOOSC	OMVS output wait
RMCARQSC	requested swap count
RMCARSSC	central storage shortage swap count
RMCATISC	terminal input swap count
RMCATOSC	terminal output swap count
RMCATSSC	count of transition swaps
RMCAUSSC	unilateral swap-out count
RMCAXSSC	auxiliary storage shortage swap count
RMCTTRPC	number of pages used for transaction elements

RSM Address Space Block Extension (RAX)

Fields in the RSM Address Space Block Extension are described in Table 136.

Table 136 RSM Address Space Block Extension (RAX)

Field	Description
RAXESCT	number of pages on extended storage
RAXFMCT	number of frames currently in use by this address space

RSM Control and Enumeration area for MVS/XA

Fields in the RSM Control and Enumeration area (RCE) are described in Table 137.

Table 137 RSM Control and Enumeration area (RCE) (part 1 of 2)

Field	Description
RCEAEC	total number of expanded storage E-frames currently on the available ESTE queue
RCEAECLO	available ESTE queue low threshold
RCEAECOK	available ESTE queue satisfactory threshold
RCEAFCLO	available central storage frame queue low threshold
RCEAFC	total number of frames currently on all available frame queues
RCEAFCOK	available central storage frame queue satisfactory threshold
RCEBELFX	total number of currently fixed pages that are backed below 16 MB real: the sum of page-fixed LSQA, SQA (excluding reserved SQA), and V=R allocated pages
RCECOMAL	number of frames allocated to common
RCECOMPI	number of common area pages paged-in
RCECOMPO	number of common area pages paged-out
RCECOMRC	number of common area pages that have been reclaimed from an available frame queue
RCEDFRS	number of times a deferred frame allocation has been satisfied
RCEESINU	number of expanded storage E-frames in use
RCEESREA	number of pages (non-VIO) read from expanded storage
RCEESWRT	number of pages written out to expanded storage
RCEHSPEM	total number of hiperspace pages migrated from expanded storage to auxiliary storage
RCEHSPER	total number of hiperspace pages read from expanded storage into central storage
RCEHSPEW	total number of hiperspace pages written to expanded storage from central storage
RCEHSPPI	total number of hiperspace pages paged-in from auxiliary storage
RCEHSPPO	total number of hiperspace pages paged-out to auxiliary storage
RCELPAPI	number of PLPA and PLPA directory pages paged-in
RCELPARC	number of PLPA and PLPA directory pages that have been reclaimed from an available frame queue
RCEMVBEL	number of central storage pages moved from below the 16-megabyte line
RCENWSF	total number of changed nonworking set pages and secondary working set pages that have completed migration
RCEPAGMV	number of times a frame was moved from one frame to another
RCEPOOL	number of frames currently available to the system, including frames backing permanent storage (nucleus frames, hardware storage area frames, FLPA frames, or fixed BLDL frames), bad frames, and offline frames
RCESPFR	number of frames available by swap-out without requiring I/O
RCESWPPI	total number of pages requiring I/O to swap-in
RCESWPPO	total number of pages requiring I/O to swap-out

Table 137 RSM Control and Enumeration area (RCE) (part 2 of 2)

Field	Description
RCETOTFX	total number of pages currently fixed: the sum of page fixed LSQA, SQA (excluding reserved SQA), and V=R allocated pages
RCETOTPI	total number of pages paged-in excluding swap-in and VIO page-in
RCETOTPO	total number of pages paged-out, excluding swap-out, move-out of VIO pages, and page-out of VIO pages
RCETOTRC	total number of times a page was reclaimed from an available frame queue
RCEVIOME	number of VIO data set pages written out to expanded storage
RCEVIOMG	number of VIO data set pages migrated from expanded storage to auxiliary storage
RCEVIOPI	total number of VIO pages paged-in, excluding swap-in
RCEVIOPO	total number of VIO pages, excluding swap-out, moved out or paged out
RCEVIORE	number of VIO data set pages read from expanded storage
RCEVIORU	number of times a VIO data set page was reused
RCEWSDNE	total number of primary working set pages which have completed migration

System Management Facilities Control Area

Fields that appear in the System Management Facilities Control Area (SMCA) are described in Table 138

Table 138 System Management Facilities (SMF) Control Area (SMCA)

Field	Description
SMCABFLS	number of records lost because of a shortage of buffer
SMCABFWT	number of buffers written
SMCADSCT	number of records lost because of a full data set
SMCANMFL	current number of full buffers
SMCARCWT	number of records written

Index

A

B

C

D

E

F

G

H

I

J

K

L

P

Q

R

S

T

Notes