CMF® MONITOR

Batch User Guide

Version 5.5

June 30, 2003

Restricted Rights Legend

Contacting BMC Software

You can access the BMC Software Web site at **http://www.bmc.com**. From this Web site, 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

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 Support page on the BMC Software Web site or by contacting Customer Support by telephone or e-mail. To expedite your inquiry, please see “Before Contacting BMC Software.”

Support Web Site

You can obtain technical support from BMC Software 24 hours a day, 7 days a week at **http://www.bmc.com/support.html**. From this Web site, you can

- read overviews about support services and programs that BMC Software offers
- find the most current information about BMC Software products
- search a database for problems similar to yours and possible solutions
- order or download product documentation
- report a problem or ask a question
- subscribe to receive e-mail notices when new product versions are released
- find worldwide BMC Software 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. Outside the United States and Canada, please contact your local support center for assistance. To find telephone and e-mail contact information for the BMC Software support center that services your location, refer to the Contact Customer Support section of the Support page on the BMC Software Web site at **www.bmc.com/support.html**.

Before Contacting BMC Software

Before you contact BMC Software, have the following information available so that Customer Support can begin working on your problem 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 problem
- 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

About This Book 1-xiii

Chapter 1 About CMF MONITOR

CMF MONITOR Product Components 1-1
Relationship among the CMF MONITOR Product Components . . . 1-2
CMF MONITOR Extractor 1-3
CMF MONITOR Analyzer 1-3
CMF MONITOR Online 1-4
CMFMON 1-5
DSO Analyzer 1-5
CMF MONITOR and MAINVIEW 1-5
Required Address Spaces 1-5
CMF MONITOR and Other BMC Software Products 1-7
MAINVIEW for OS/390 1-8
CMF MONITOR Compatibility with IBM's RMF 1-9
Similarities 1-9
Differences 1-10
CMF MONITOR Functions Unavailable in RMF 1-10

Chapter 2 Collecting Report Data

Understanding CMF MONITOR Extractor 2-1
How the Extractor Works 2-2
How CMF MONITOR Components Use Extractor Data 2-3
When to Use CMF MONITOR Extractor CPM and IPM Modes . . . 2-3
Sampling Data and Producing Records 2-5
Four Methods of Sampling Data 2-6
How Extractor Statements Control Sampler Operations 2-6
Record Types 2-7
Writing Records 2-10
Customizing the Recording Interval. 2-11
Writing to SMF 2-11
Writing to the Cross-System Data Server (XDS) Buffers 2-12
Writing to CMF or DSO Data Sets 2-12
Archiving Your Data 2-14

Running CMF and RMF on the Same System 2-14
Using the Extractor Trace Facilities 2-15
Extractor Control Statements Used by BMC Software Products 2-16
CMF MONITOR Control Statements. 2-16
DSO Control Statements. 2-17
MAINVIEW for OS/390 Control Statements. 2-17
Defining Extractor JCL 2-17
Defining Extractor Control Statements 2-18
Default CPM and IPM Control Statement Sets 2-18
Using the MODIFY Command to Change Extractor Operation. 2-22
MODIFY Command Examples 2-23

Chapter 3 Producing and Using Reports

How Reports Are Generated 3-1
Generating JCL to Produce Analyzer Reports 3-2
Using the Report Generation Panels. 3-3
Panels for Generating CMF Analyzer JCL. 3-4
Selecting Analyzer Reports 3-7
Generating the JCL 3-10
Listing Previously Specified Members. 3-11
Defining Analyzer JCL Manually 3-12
Using the Default Analyzer JCL Member 3-13
Descriptions of Analyzer JCL Statements 3-14
Producing Graphics Reports on a JES2 System 3-17
Defining Your Report Output to a Laser Printer. 3-18
Setting Values of Region, DMSS Reserve, and CTRLSIZE 3-18
Defining Analyzer Control Statements 3-21
Using the Default Analyzer Control Statement Member 3-21
Specifying the CMF User Record Type 3-22
Specifying Records from CPM or IPM Monitoring Modes 3-23
Using General Control Statements 3-23
Using Report Control Statements 3-24
Using and Interpreting Reports 3-24
Knowing What Reports You Need 3-25
Interpreting Report Field Information 3-26
Understanding Report Headings. 3-27
Descriptions of Report Headings Fields 3-27
Writing Your Own Programs to Process Extractor Data. 3-28

Chapter 4 Preprocessing Extractor Data Sets

How to Use the CX10CVBS Utility 4-2
Defining the PARM Parameter 4-3
CX10CVBS Return Codes 4-6
How Data Is Copied Using the CX10CVBS Utility. 4-6

Chapter 5 Using the CMF MONITOR APIs

How XDS Works. 5-2
Activating XDS 5-2

Accessing Data Gathered by XDS 5-2
Specifying the Source of Input Data 5-3
Layout for Mapping an Answer Area for API Output 5-3
CX10XDQY 5-3
CX10XDRC 5-8
Sample DSECT for CX10XDRC Results 5-9
CX10XDGS 5-10
CX10XDGS Data Reduction Exit 5-13
Sample DSECT for the Common Answer Area Header 5-15
Return Codes for XDS APIs 5-17
Implementing the CX10GVID API 5-21
Customizing the Extractor to Get Data 5-21
Calling CX10GVID 5-22
General-Purpose Registers 5-23
Return Codes for CX10GVID 5-24
Additional CMF API Considerations 5-24
CMF APIs and OS/390 5-24
CMF APIs and SDSF 5-26

Chapter 6 CMF Analyzer Spreadsheet Converter

Installing the Spreadsheet Converter on to Your PC 6-1
Capturing CMF Analyzer and CMFMON Reports 6-2
Capturing CMF Analyzer Reports 6-2
Capturing CMFMON Reports 6-3
Downloading Reports to the PC 6-3
Running the Spreadsheet Converter 6-3
Spreadsheet Converter Output 6-6
A Brief Tutorial 6-7
Troubleshooting 6-12
Maintenance and Support 6-13
Maintenance 6-13
Customer Support 6-13

Chapter 7 Mapping CMF Records Created by CMF

Assembler MACROs in BBSAMP for Record Types 70-79 7-2
Assembler MACROs in BBSAMP for User Records 7-2
C Structures in BBSAMP 7-3
SAS Code in BBSAMP 7-3
Using CMF User Records with MXG 7-4

Chapter 8 Laser Printer Support

Types of Emphasized Data in Reports 8-1
Required JCL for Laser Printer Support 8-2
Examples 8-3
Selecting Fonts 8-3
Guidelines for Font Selection 8-4

Appendix A **Statistical Considerations**
Standard Deviation, the Mean, and the Mode A-1
Calculating Standard Deviation. A-2
Statistical Accuracy A-3

Appendix B **Workload Measurement in Goal Mode**
Service Definition B-2
Service Policy B-2
Classification Rules B-2
Resource Group B-2
Workload. B-2
Service Class. B-3
Performance Index B-4

Index

Figures

Figure 1-1 Relationship among CMF MONITOR Product Components 1-2
Figure 1-2 Communication between Address Spaces 1-6
Figure 2-1 Extractor Data Collection and Recording Functions 2-2
Figure 2-2 Extractor Control Statements in Starter Set CMFCPM00 2-19
Figure 2-3 Extractor Control Statements in Starter Set CMFIPM00 2-21
Figure 3-1 MAINVIEW Selection Menu 3-2
Figure 3-2 OS/390, z/OS, and USS Solutions Panel 3-3
Figure 3-3 Produce CMF Analyzer Batch Reports Main Menu 3-3
Figure 3-4 Panel for Setting Up CMF Analyzer JCL 3-4
Figure 3-5 Panel for Specifying Source of CMF Analyzer Data 3-5
Figure 3-6 Panel for Specifying CMF Analyzer Output Destination 3-6
Figure 3-7 Panel for Selecting CMF Analyzer Reports 3-7
Figure 3-8 Panel for Specifying CMFSUM Parameters 3-9
Figure 3-9 Panel for Specifying Filters on Input Data 3-10
Figure 3-10 Panel for Submitting Analyzer JCL 3-11
Figure 3-11 JCL Member List (left half) 3-11
Figure 3-12 JCL Member List (right half) 3-12
Figure 3-13 Example of Sample JCL Member CMFJANL 3-13
Figure 3-14 Storage Area Configuration 3-19
Figure 3-15 ANLYSAMP Control Statement 3-22
Figure 3-16 Example of a Standard Report Heading 3-27
Figure 4-1 Sample Execution JCL for CX10CVBS 4-2
Figure 6-1 Select BMC Software toolbar from the Excel Menu 6-4
Figure 6-2 Spreadsheet Converter Instruction Screen 6-4
Figure 6-4 BMC Software toolbar 6-5
Figure 6-6 Excel Add-Ins Available List Box 6-8
Figure A-1 Relationship of Large and Small Standard Deviations to the Mean A-2
Figure A-2 Equation for Calculating Standard Deviation A-3
Figure A-3 Confidence Levels for P=50% A-4
Figure A-4 Correction Factors for Confidence Intervals A-4

Tables

Table 2-1	Record Types with Corresponding Extractor Statements and Samplers	2-8
Table 2-2	Primary and Alternate Data Set DD Names	2-13
Table 2-3	Explanation of CMFCPM00 Control Statements	2-19
Table 2-4	Explanation of CMFIPM00 Control Statements	2-21
Table 2-5	CMF Monitor MODIFY Commands	2-22
Table 2-6	MODIFY Command Examples for CMF MONITOR	2-23
Table 3-1	JCL Control Statements for the CMF MONITOR Analyzer	3-14
Table 3-2	Field Descriptions for a Report Heading	3-27
Table 4-1	CX10CVBS Parameters	4-3
Table 4-2	Examples of PARM Values Used to Copy Records Selectively	4-5
Table 4-3	CX10CVBS Return Codes	4-6
Table 5-1	Parameters for Calling CX10XDQY	5-4
Table 5-2	XDRQSMFH SMF Record Header Format	5-7
Table 5-3	XDRQRMFI RMF Record Header Format	5-8
Table 5-4	Parameters for Calling CX10XDRC	5-9
Table 5-5	Parameters for Calling CX10XDGS	5-11
Table 5-6	Parameters for Calling the CX10XDGS Exit	5-14
Table 5-7	Return Codes for XDS APIs	5-17
Table 5-8	Extractor Samplers Required for SMF 79 Record Subtypes	5-22
Table 5-9	Returned Fields (APF-Authorized or Running in Supervisor State)	5-22
Table 5-10	Subtype 2 Real Storage Utilization Fields Affected by BBXS	5-23
Table 5-11	General-Purpose Register Values	5-23
Table 5-12	Return Codes	5-24
Table 7-1	CMF User Record Types Supported by MXG	7-4

About This Book

This book contains detailed information about about CMF MONITOR batch reporting operations and describes the operation of the Extractor and Analyzer components of CMF MONITOR 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.

How This Book Is Organized

This book is organized as follows. In addition, an index appear at the end of the book.

Chapter/Appendix	Description
Chapter 1, “About CMF MONITOR”	Presents overview of CMF MONITOR’s components and integration with other BMC Software products
Chapter 2, “Collecting Report Data”	Describes Analyzer control statements and their parameters.
Chapter 3, “Producing and Using Reports”	Provides report examples, data descriptions and field calculations
Chapter 4, “Preprocessing Extractor Data Sets”	Describes reasons for and how to preprocess Extractor data sets for producing reports
Chapter 5, “Using the CMF MONITOR APIs”	Describes the four application program interfaces provided by CMF MONITOR
Chapter 6, “CMF Analyzer Spreadsheet Converter”	Descibes how to turn CMF Analyzer report data into Microsoft Excel spreadsheets

Chapter/Appendix	Description
Chapter 7, “Mapping CMF Records Created by CMF”	Describes function of mapping CMF members to external programs using established standards
Chapter 8, “Laser Printer Support”	Describes the function that controls the appearance of the fields printed out in a report
Appendix A, “Statistical Considerations”	Contains information about statistical calculations used in reports.
Appendix B, “Workload Measurement in Goal Mode”	Contains information about workload management in MVS 5.1 or later systems running in goal mode.

Related Documentation

BMC Software products are supported by several types of documentation:

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

Category	Document	Description
Installing CMF MONITOR	*OS/390 and z/OS Installer Guide* *CMF MONITOR Online Getting Started*	Explains how to download product tape components, access AutoCustomization, and perform manual customization steps that are specific to CMF MONITOR. Describes allocation of data sets for the CMF MONITOR Extractor.
Using CMF MONITOR’s batch reporting components	*CMF MONITOR Batch Reference Guide* *CMF MONITOR Batch User Guide*	Explains how to use the Extractor and Analyzer components, and how to interpret the report data.
Using CMF MONITOR Online	*CMF MONITOR Customization Guide* *MAINVIEW Installation Requirements Guide* *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 presented.
Using CMFMON	*CMF MONITOR CMFMON User Guide*	Explains how to use CMFMON’s online facility and write facility, as well as how to generate CMFMON batch reports.
Using the DSO component	*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 for OS/390 product family	MAINVIEW for OS/390 User Guide, *Using MAINVIEW*	Other BMC Software products use the CMF Extractor as a base component to gather data for their reports and displays. The following table lists these products and their related documentation

Online and Printed Books

The books that accompany BMC Software products are available in online format and printed format. If you are a Windows or Unix user, you can view online books with Acrobat Reader from Adobe Systems.

To Access Online Books

Online books are formatted as Portable Document Format (PDF) files. You can view them, print them, or copy them to your computer by using Acrobat Reader 3.0 or later.You can access online books from the documentation compact disc (CD) that accompanies your product or from the World Wide Web.

In some cases, installation of Acrobat Reader and downloading the online books is an optional part of the product-installation process. For information about downloading the free reader from the Web, go to the Adobe Systems site at **http://www.adobe.com**.

To view any online book that BMC Software offers, visit the support page of the BMC Software Web site at **http://www.bmc.com/support_home**. Select a product to access the related documentation.

To Request Additional Printed Books

BMC Software provides printed books with your product order. To request additional books, go to **http://www.bmc.com/support_home.**

Online Help

The CMF MONITOR product includes online Help. In the CMF MONITOR ISPF interface, you can access Help by pressing **F1** 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 Software product. Release notes provide current information such as

- updates to the installation instructions

- last-minute product information

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

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, 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 1-1.

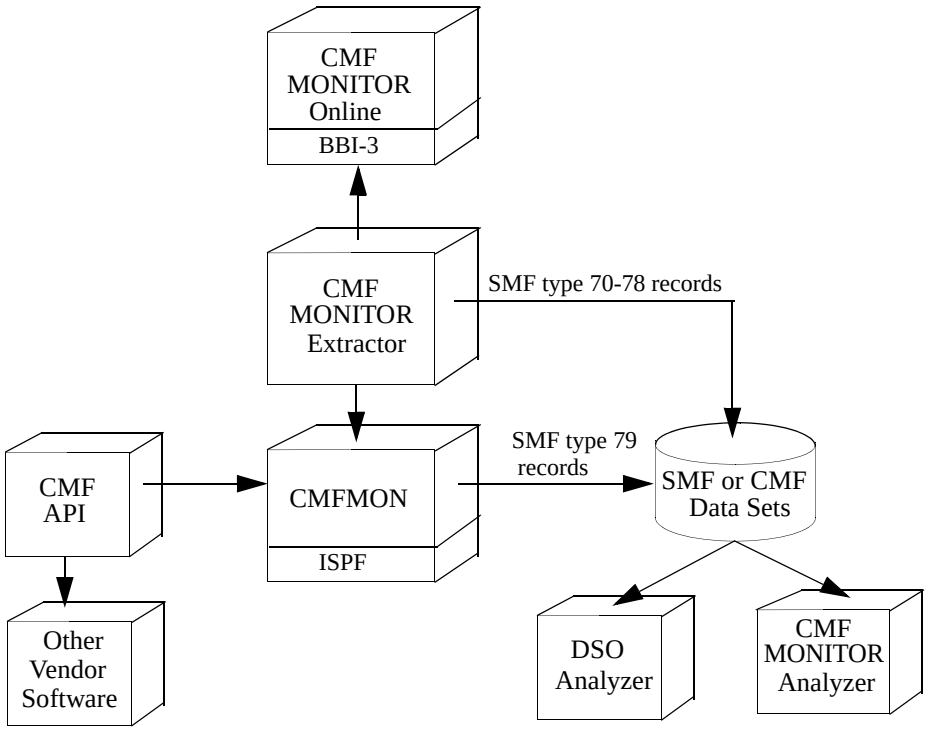

Figure 1-1 Relationship among CMF MONITOR Product Components

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 by the Analyzer to produce batch reports, by CMF MONITOR Online (and some CMFMON screens) to provide realtime bottleneck detection and analysis, and 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 OS/390 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 the section titled, "Extractor Control Statements" in the *CMF MONITOR Batch Reference Guide*.

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 the section titled, “Analyzer Control Statements” in the *CMF MONITOR Batch Reference Guide*. 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 3-2.

The reports produced by each Analyzer control statement or combination of statements are documented in the section titled, “CMF MONITOR Reports” in the *CMF MONITOR Batch Reference Guide*.

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 new 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 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 may only need to use the DSO Analyzer 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*, which is included in the CMF MONITOR package.

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 1-2 on page 1-6 illustrates the communication between address spaces under the MAINVIEW architecture.

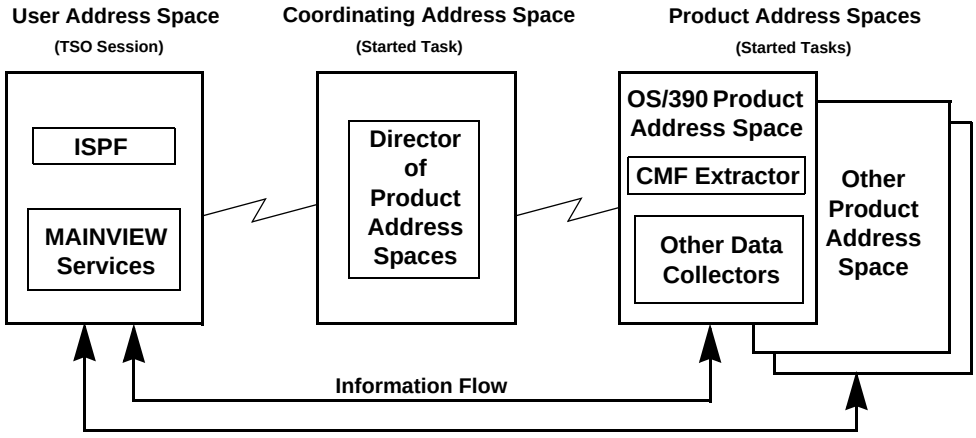

Figure 1-2 Communication between Address Spaces

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.

MVS Product Address Space (PAS)

The MVS PAS runs as a started task and contains both the CMF Extractor and the online component of CMF MONITOR—called CMF MONITOR Online. The MVS 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 MVS PAS on your system.

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

Note: Initialization of the MVS PAS data collectors is controlled by the DC={START|STOP|CPM|IPM} parameter. This parameter is located in the MVS PAS JCL and can be specified when starting the MVS PAS. See "Using the MODIFY Command to Change Extractor Operation" on page 2-22 for more information about starting and stopping the MVS PAS and using this parameter.

When you initialize the MVS PAS with

- 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 OS/390

The relationship among CMF MONITOR and other BMC Software products is shown in Figure 1-3.

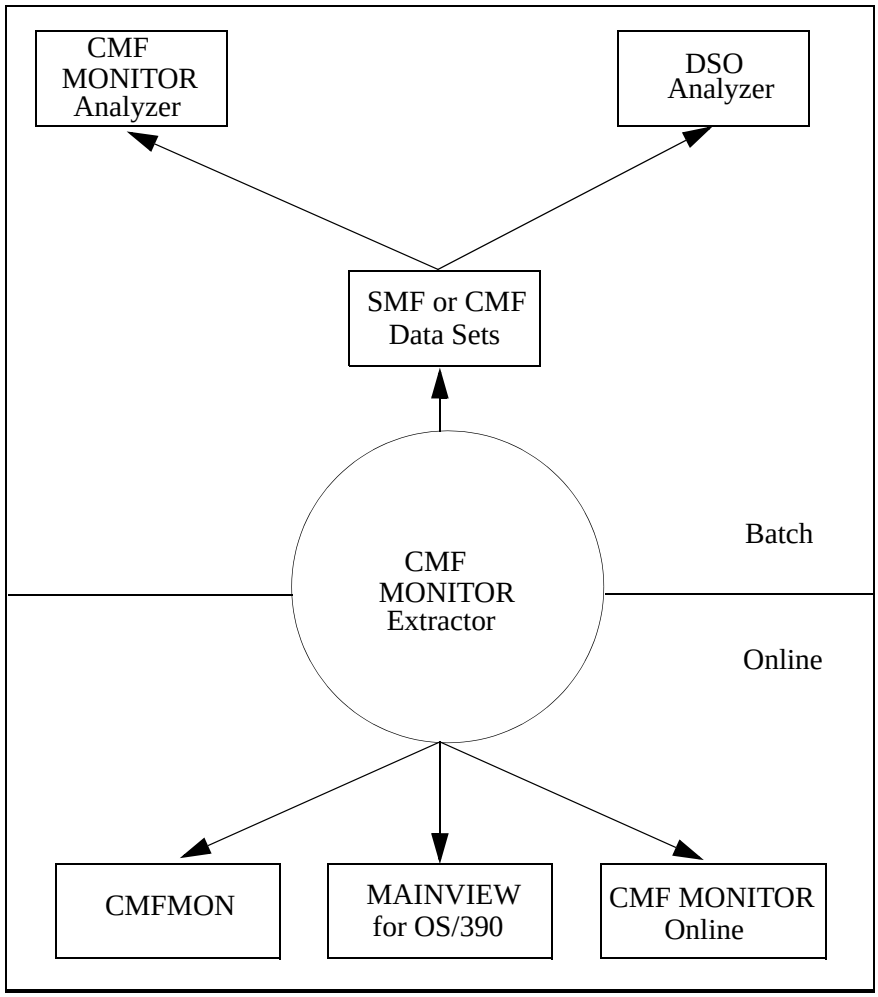

Figure 1-3 Interaction among CMF MONITOR and Other BMC Software Products

This section discusses products interaction with the CMF MONITOR Extractor, but does not discuss CMFMON or the CMF and DSO Analyzers (see “Relationship among the CMF MONITOR Product Components” on page 1-2 for more information about these CMF components).

MAINVIEW for OS/390

MAINVIEW for OS/390 is an MVS performance management monitoring product 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 OS/390 views. The views present this information in both realtime and historical timeframes, which allows for on-screen analysis of current and past resource performance.

The services offered through MAINVIEW for OS/390 provide you with the ability to instantly respond to any performance circumstances your system encounters. The set of services offered includes both System Programmer Services and Exception Monitoring Services.

MAINVIEW for OS/390 uses CMF MONITOR Extractor to collect some of the data for its online views. To collect the correct data for MAINVIEW for MVS, 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 2-16 for more information).

MAINVIEW for OS/390 operation and view data element descriptions and commands are documented in several books. See "Related Documentation" on page -xiv for more information.

CMF MONITOR Compatibility with IBM's 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 OS/390 1.2.0 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 OS/390 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 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 the Extractor DEVICE section in the *CMF MONITOR Batch Reference Guide*).

CMF MONITOR Functions Unavailable in RMF

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

- 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 2-3). RMF can run only one Extractor.

Chapter 2 Collecting Report Data

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 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 2-1 shows how the Extractor collects and records data.

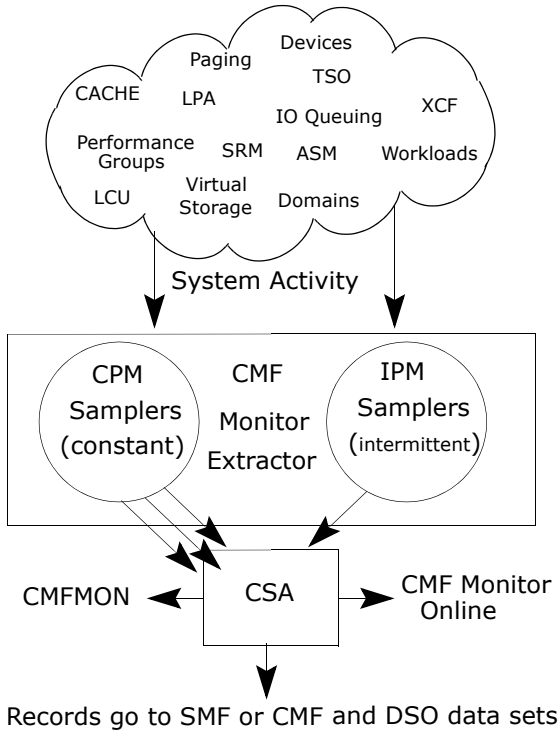

Figure 2-1 Extractor Data Collection and Recording Functions

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 2-5 for more information.

Controlling data collection

The function of each sampler running in your system is controlled by an Extractor control statement (see Table 2-1 on page 2-8 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 comprised 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 2-3 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 2-10 for more information.

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 3, “Producing and Using 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 2-22.

The following describes 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 2-18.

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 generally is used to sample head movement, individual devices, and the like.

When you detect a trend that may 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 either can use the MODIFY command (see "Using the MODIFY Command to Change Extractor Operation" on page 2-22 for more information), or change the default setting on the MVS PAS PROC to DC=IPM (see "Defining Extractor Control Statements" on page 2-18 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 2-20.

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 DIE (disabled interrupt exit), SRB (service request block), SRM (System Resources Manager), or TCB (task control block) methods of measuring the performance of a specific area of your system. 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 2-18 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 2-6 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 (see "When to Use CMF MONITOR Extractor CPM and IPM Modes" on page 2-3 for more information about monitoring modes), another function of the Extractor writes the data collected in CSA as records to SMF, or CMF, or DSO data sets (see "Writing Records" on page 2-10 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 this control through statement parameters is provided so you can tune the Extractor to collect only the data your site requires. This 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 the REPORT (GBLS sampler), HEADMOVE (HMOV sampler), and MACSCHAR (IOWS sampler) Extractor control statements. The GBLS, HMOV, and IOWS samplers write out records more frequently due to the volume of data they are recording.

Record Types

The Extractor samplers produce records that can be processed later by the CMF and DSO Analyzers or 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 2-1. User programs written to process SMF record types, such as MICS, MXG, or SLR, can process CMF MONITOR records.

Note: Refer to the Extractor control statement in the section titled, "DEVICE" in the *CMF MONITOR Batch Reference Guide* for more information on 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 the REPORT section in the *CMF MONITOR Batch Reference Guide* for more information).

Producing Your Own Reports Using Extractor Records

Records 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 in

- Assembler format is CMFREC*xx*
- C format is CMFC*xx*
- SAS format is 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 7, “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 2-1 shows record types used by CMF MONITOR in ascending numeric order.

Table 2-1 Record Types with Corresponding Extractor Statements and Samplers (Part 1 of 3)

Record Type	Description	Sampler	Control Statement
SMF 70	CPU activity	CPUS	CPU
SMF 71	Paging activity	PAGS	PAGING
SMF 72-1 (MVS 5.1 or later)	Workload activity This record is created in compatibility mode only.	WLMS	WORKLOAD
SMF 72-2	Resource usage and delay data by performance group/domain This sampler does not function unless the MVS PAS data collectors are active.	PGDS	PGDDLAY
SMF 72-3	Workload activity by service class This sampler is used in goal mode only.	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. In goal mode, service class is used.	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	CA03 CA13 CA23 CA3C	CACHE
SMF 75	Page/swap 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 78-1	I/O queuing data for 4381 and 3080 processors	IOQS	IOQ
SMF 78-2	Virtual storage data	VSMS	VSMDATA

Table 2-1 Record Types with Corresponding Extractor Statements and Samplers (Part 2 of 3)

Record Type	Description	Sampler	Control Statement
SMF 78-3	I/O queuing data for 3090, ES/9000 series, or later processors	IOQS	IOQ
SMF 79-1	Address space state data CMFMON	Not applicable	ASD
SMF 79-2	Address space resource data CMFMON	Not applicable	ARD
SMF 79-3	Central storage/processor/SRM activity CMFMON	Not applicable	SRCS
SMF 79-4	System paging activity data CMFMON	Not applicable	SPAG
SMF 79-5	Address space SRM data CMFMON	Not applicable	ASRM
SMF 79-6	Enqueue reserve data CMFMON	Not applicable	SENQR
SMF 79-7	Enqueue contention data CMFMON	Not applicable	SENQ
SMF 79-8	Transaction activity data CMFMON	Not applicable	TRX
SMF 79-9	Device activity data CMFMON	Not applicable	DEV
SMF 79-10	Domain activity data CMFMON	Not applicable	DDMN
SMF 79-11	Page and swap data set activity CMFMON	Not applicable	PGSPP PGSPS
SMF 79-12	Channel path activity CMFMON	Not applicable	CHANNEL
SMF 79-13	I/O queuing activity by logical control unit for the 308x and 4381 processors CMFMON	Not applicable	IOQ
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 240-04	Workload data This record is created in compatibility mode only.	WLMS	WORKLOAD

Table 2-1 Record Types with Corresponding Extractor Statements and Samplers (Part 3 of 3)

Record Type	Description	Sampler	Control Statement
CMF 240-05	Device data	DEVS	DEVICE
CMF 240-06	Extractor summary data This sampler does not run in IPM mode.	EXTS	EXTSUM
CMF 240-07	Performance group mapping 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 240-19	I/O workload record data	IOWS	MACSCHAR
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 240-27	Cache data records	CA03 CA13 CA23 CA3C	CACHE
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 2-14 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 4, "Preprocessing Extractor Data Sets", to recover and copy a damaged data set.

Customizing the Recording Interval

You can customize the recording interval to be any amount of time up to a maximum of one hour, or, in SP4.3 or later systems, 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 the REPORT section in the *CMF MONITOR Batch Reference Guide* 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 CPM mode to write to SMF and 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 2-3 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 the REPORT section in the *CMF MONITOR Batch Reference Guide* 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

- 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 need to first 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 on how to empty CMF and DSO output data sets, see “Archiving Your Data” on page 2-14.

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.

Note: 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 2-2.

  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 are used.

Table 2-2 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 are used.

See the REPORT section in the *CMF MONITOR Batch Reference Guide* 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 CMF's COPYVBS utility is found in *hilevel*.UBBSAMP member CMFJCVBS. See Chapter 4, "Preprocessing Extractor Data Sets" for information about using this utility.

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

Once 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. Make sure you have archived your data if you wish to save 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's and RMF's 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 4, “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 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 the DEVICE section in the *CMF MONITOR Batch Reference Guide*).

The similarities and differences between CMF MONITOR and RMF are discussed in “CMF MONITOR Compatibility with IBM’s RMF” on page 1-9.

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 4KB. 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 the TRACE sections in the *CMF MONITOR Batch Reference Guide*.

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. The Extractor Control Statement section in the *CMF MONITOR Batch Reference Guide* provides detailed information about all Extractor control statements used by BMC Software products.

CMF MONITOR Control Statements

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

Control Statement	Sampler
ASMDATA	ASMS
CACHE	CA03
CHANNEL	CHNS
CPU	CPUS
DEVICE	DEVS
ENQUEUE	EQES
EXTSUM	EXTS
HEADMOVE	HMOV
IOQ	IOQS
LINKMAP	LPAM
MACSCHAR	IOWS
PAGING	PAGS
PGDDLAY	PGDS
REPORT	GBLS, RECD
TSODATA	TSOS
USER	USER
VSMDATA	VSMS
WORKLOAD	WLMS
XCFDATA	XCFS
CFDATA	CFTS
OMVS	OMVS
HFS	HFSS

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 the HEADMOVE and REPORT sections in the *CMF MONITOR Batch Reference Guide* 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 OS/390 Control Statements

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

Control Statement	Sampler
ASMDATA	ASMS
CACHE	CA03
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 OS/390, if your site uses this product (see information about the DC parameter in "Using the MODIFY Command to Change Extractor Operation" on page 2-22).

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 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. However, 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 you create for CPM mode must be named CMFCPM*xx* and any members you create for IPM mode 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 2-3 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, since SMF=YES is specified on the REPORT control statement.

Sample control statements in CMFCPM00 are shown in Figure 2-2. The statements are explained in Table 2-1.

Figure 2-2 Extractor Control Statements in Starter Set CMFCPM00

```
************************************************************************
*  CPM DEFAULT CONTROL STATEMENT PACKET FOR CMF MONITOR EXTRACTOR     *
*
*  THE EXTSUM STATEMENT SHOULD BE MODIFIED TO ADD INSTALLATION JOB    *
*  CLASSES AND PERFORMANCE GROUPS.  IF YOU HAVE CACHED CONTROLLERS,   *
*   SEE THE CACHE STATEMENT AT THE END.

************************************************************************
REPORT   CPM,INTERVAL=15,SYNCH=00,CSA=300,SMFRECID=240,
          SMF=YES,RUNTIME=1440
 ASMDATA  SAMPLE=2000
 CHANNEL
 CPU      SAMPLE=2000
 DEVICE   SAMPLE=2000,CLASS=DASD
 DEVICE   SAMPLE=2000,CLASS=TAPE,OFFLINE=YES
 ENQUEUE  MAJOR=SYSDSN
 EXTSUM   SAMPLE=2000,SPINOFF=NO,
          JOBCLASS=(JC=A,JD=CLASSA,
                    JC=B,JD=CLASSB),
          PERFORM=(PG=01,PP=1,PD=PGP1,
                   PG=01,PP=12,PD=PGP12)
 IOQ
 LINKMAP
************************************************************************
* THE MACSCHAR SAMPLER CAN HAVE HIGH CPU OVERHEAD, DEPENDING ON      *
* THE SIZE OF YOUR SYSTEM'S TRACE TABLE.
************************************************************************
*MACSCHAR IOWSAMP=40000
 PAGING   SAMPLE=6000
 TSODATA  SAMPLE=2000,LIMIT=50,USER=YES
 VSMDATA  SAMPLE=6000
 WORKLOAD
************************************************************************
* THE CACHE CONTROL STATEMENT NEEDS TO BE INVOKED FROM ONLY
* ONE SYSTEM IF THE CONTROL UNIT IS SHARED.                          *
************************************************************************
```

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

Table 2-3 Explanation of CMFCPM00 Control Statements (Part 1 of 2)

Statement	Explanation
REPORT	Required. CPM mode and a record time of 15 minutes are specified, synchronized on the hour. The ID of the SMF records is 240 (X'F0'). The amount of Common Storage Area (CSA) requested is 300K. Records will be written to SMF.
ASMDATA	Auxiliary storage manager data will be sampled every 2 seconds.
CHANNEL	Channel path activity will be sampled.
CPU	CPU activity will be sampled once every second.
DEVICE	All online DASDs and all tape devices (both online and offline) will be sampled every 2 seconds.
ENQUEUE	All enqueue activity for the major name SYSDSN will be sampled.
EXTSUM	Extractor Summary records will be written, but no spinoff reports will be generated.

Table 2-3 Explanation of CMFCPM00 Control Statements (Part 2 of 2)

Statement	Explanation
IOQ	I/O queuing activity will be sampled.
LINKMAP	Mapping data will be collected for the link pack area.
MACSCHAR	I/O and page fault data will be sampled every 4 seconds. You must remove the comment (*) from the MACSCHAR statement if you want to collect this data.
PAGING	Page and swap data will be sampled once every 6 seconds.
TSODATA	Activity for all TSO commands (up to 50) will be monitored, and TSO user statistics will be gathered as logoffs are encountered.
VSMDATA	SQA data by subpool and CSA data by subpool and key will be sampled once every 6 seconds, but private area data will not be collected.
WORKLOAD	WORKLOAD activity and storage data will be monitored.
CACHE	Cache subsystem data will be sampled for all devices.

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 2-3 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.

Note:

- 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.

Figure 2-3 Extractor Control Statements in Starter Set CMFIPM00

```
*************************************************************************
*                                                                       *
*    IPM DEFAULT CONTROL STATEMENT PACKET FOR CMF MONITOR EXTRACTOR     *
*                                                                       *
*    IF YOU HAVE CACHED CONTROLLERS, SEE THE CACHE STATEMENT AT THE     *
*    END.  REFER TO THE CMF MONITOR BATCH USER GUIDE AND REFERENCE,     *
*    TD-50J OR THE CMF MONITOR BATCH REFERENCE SUMMARY, TD-50Q FOR      *
*    INFORMATION ON ALL CMF EXTRACTOR CONTROL STATEMENTS AND THEIR      *
*    PARAMETERS.                                                        *
*                                                                       *
*************************************************************************
 REPORT IPM,INTERVAL=15,CSA=150,SMFRECID=240,RUNTIME=60,SMF=NO
 CHANNEL
 CPU      SAMPLE=500
 DEVICE   SAMPLE=500,CLASS=DASD
 DEVICE   SAMPLE=500,CLASS=TAPE,OFFLINE=YES
*************************************************************************
*    THE HEADMOVE SAMPLER CAN HAVE HIGH CPU OVERHEAD, DEPENDING ON      *
*    THE SAMPLE RATE AND THE NUMBER OF DEVICES MONITORED.               *
*************************************************************************
 HEADMOVE ALL,SAMPLE=33,VTOC=YES
 IOQ
*************************************************************************
*     THE CACHE CONTROL STATEMENT NEEDS TO BE INVOKED FROM ONLY         *
*     ONE CPU IF THE CONTROL UNIT IS SHARED.                            *
*************************************************************************
CACHE
```

Figure 2-3 shows a sample set. It does not include control statements for all possible options. The statements are explained in Table 2-4.

Table 2-4 Explanation of CMFIPM00 Control Statements

Statement	Explanation
REPORT	Required. IPM mode and a record time of 15 minutes are specified; the IPM mode will run for 60 minutes; 150K of CSA is requested; and the ID of the SMF records is 240 (X'F0'). Data will be written to the data sets specified in the //CMFIPM1 DD and //CMFIPM2 DD statements.
CHANNEL	All channel path activity will be sampled.
CPU	CPU activity will be sampled twice a second.
DEVICE	All tape and online DASDs will be sampled twice a second.
HEADMOVE	DASD head movement data will be sampled 30 times a second.
IOQ	I/O queuing activity will be sampled.
CACHE	Cache subsystem data will be sampled.

Using the MODIFY Command to Change Extractor Operation

At some point while running the Extractor in your system, you may 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]
```

Table 2-5 explains the MODIFY commands. Table 2-6 on page 2-23 provides examples of how to issue these commands.

Table 2-5 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*.
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*.

Table 2-5 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. This will result in a new set of CMF 240-14 records written to the extractor output dataset.
XDS	Specifies that the CMF Cross System Data Server (XDS) is to be started, stopped, or executed under a different control statement set.
*see the *CMF MONITOR Online User Guide* for more information	

MODIFY Command Examples

Examples of how to issue the attributes for the MODIFY command are provided in Table 2-6.

Table 2-6 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 OS/390 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,HMOV RESCAN=BOTH	Causes the MVS PAS HEADMOVE samplers for both IPM and CPM modes to initiate a VTOK scan and write CMF 240-14 records.

Chapter 3 Producing and Using 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 also can be used to affect specific reports (see “Using General Control Statements” on page 3-23 for more information).

Report Used to define specific reports or graphs to be produced (see “Using Report Control Statements” on page 3-24 for more information).

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

How to 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 the section titled, “Analyzer Control Statements” in the *CMF MONITOR Batch Reference Guide*.

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, and 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 3-1.

Figure 3-1 MAINVIEW Selection Menu

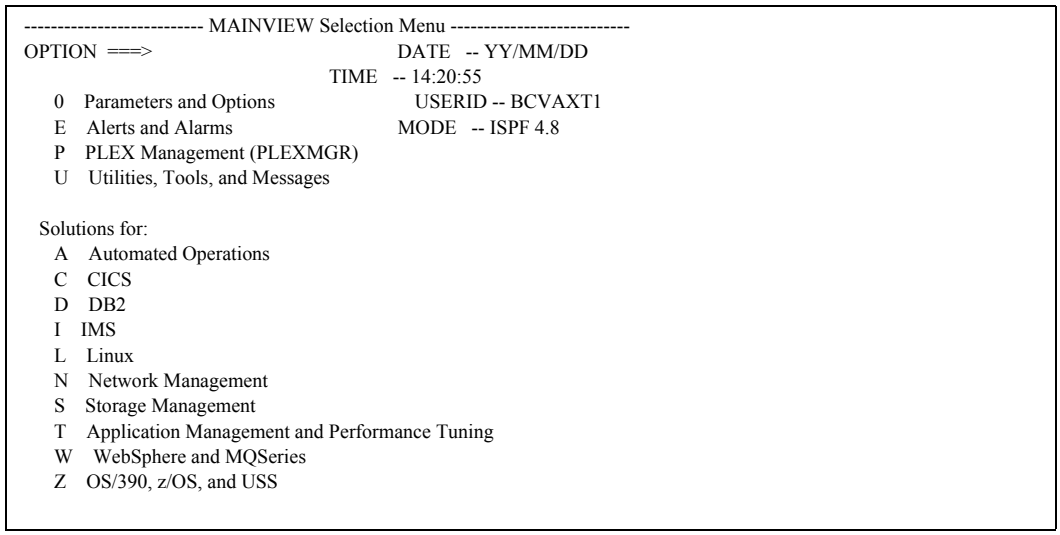

This is the panel from which all BMC Software MAINVIEW products are accessed. To display a menu that includes choices for the CMF Analyzer, select option Z (OS/390, z/OS, and USS). This displays the OS/390, z/OS, and USS Solutions panel, as shown in Figure 3-2.

Figure 3-2 OS/390, z/OS, and USS Solutions Panel

```
------------------- OS/390, z/OS, and USS Solutions -------------------------
OPTION  ===>                                  DATE   -- YY/MM/DD
                                            TIME   -- 14:22
  Performance                                USERID -- BCVAXT1
    1  MV390       MAINVIEW for OS/390              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 3-2 above, 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 3-3.

Figure 3-3 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 may select any of the options on the main menu at any time.

The first time 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”.

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 3-4 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 may change this to conform to your site’s naming conventions.

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

Specifying Analyzer Input and Output

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

Figure 3-5 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 5, “Using the CMF MONITOR APIs”.
- When DATASET is specified, an EXTDATA DD statement is generated for each data set 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 3-6.

Figure 3-6 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 is useful for generating reports to be converted to spreadsheets, as described in Chapter 6, “CMF Analyzer Spreadsheet Converter”.

Note: The Analyzer has been enhanced so that you can send data to both SYSOUT and to 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 3-7 is displayed when you select option 2 (Reports) from the main menu for the interface.

Figure 3-7 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
   CACHE       YES     Cache Subsystem Detail and Summary Reports        NO
   CFACT       NO      Coupling Facility Activity Report (MVS SP5)       NONE
   CMFSTAT     YES     CMF Records Statistics Report                     NONE
   CMFSUM      YES     CMF Summary Report                                YES
   COMBCONF    NO      Combinations of Configuration Activity Report     NO
   COMMSTOR    YES     Common Storage Usage Summary Report               NO
   CPU         YES     CPU Utilization Report                            NONE
   CPUCHAN     YES     CPU and Channel Path Activity Report              NO
   CPUCON      YES     Processor Concurrency Report                      NONE
   DASD        YES     Direct Access Report                              NONE
   DEVACT      YES     Device Activity Report                            NO
   DOMINO      NO      LOTUS DOMINO Server Report                        NONE
   ENQUEUE     YES     Enqueue Conflict Report                           NO
   EXCEPTS     NO      Exceptions Subreport                              NO
   EXCEPTS     NO      Exceptions Trace Detail Report                    NO
   GRAPH       NO      Various Graphical Reports                         NO
   GRAPH       NO      Graphics Trace Detail Report                      NO
   HFS         YES     HFS Statistics Report                             NONE
   HTTP        YES     HTTP Server Report                                NONE
   LCU         YES     Logical Control Unit Report                       NONE
   LINKPACK    NO      Link Pack Area Report                             NO
   MDF         NO      MDF Report or Logical Partition Report            NO
   OMVS        NO      OMVS Kernel Activity Report                       NONE
   PERFSUM     YES     Performance Summary Report                        NO
   PGDDLAY     NO      Domain Delay Analysis Report                      NO
   PROTKEY     YES     CPU Utilization by Protect Key Report             NO
   PRSM        YES     Logical Partition Report or MDF 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 (MVS SP5)       NO
   WORKLOAD    YES     System Workload Activity 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. Once 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 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 3-8 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 3-9 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 following panel 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 3-10 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 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 3-11 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 3-12.

Figure 3-12 JCL Member List (right half)

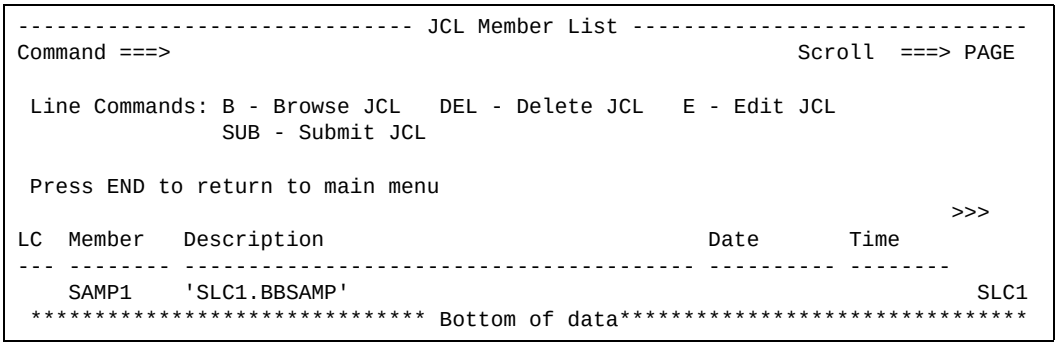

This panel allows you to select a report from a list of those previously set up. If you have a set of reports that you want to run multiple times, you can use this panel to submit your job directly, without having to use other panels in this interface.

Defining Analyzer JCL Manually

If you decide not to use the JCL generator described in “Generating JCL to Produce Analyzer Reports” on page 3-2, 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”. This sample member contains all but two of the JCL statements 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 3-17 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 3-18 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 Figure 3-10 on page 3-11.

The JCL statements are described in Table 3-1 on page 3-14.

Figure 3-13 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 ?BBASMFID TO THE SMF ID (SYSTEM ID) OF THE TARGET          00130000
//*   SYSTEM.                                                           00140000
//*                                                                     00150000
//*-------------------------------------------------------------------- 00160000
//*                                                                     00170000
//CMFRPTS  EXEC PGM=CMFANLYZ,REGION=6M                                  00180000
//*                                                                     00190000
//STEPLIB  DD   DISP=SHR,                 - ANALYZER LOAD LIBRARY       00200000
//         DSN=?BBCHILV.BBLINK                                          00210000
//DMSSMAIN DD   UNIT=SYSDA,               - DMSS WORK FILE              00220000
//         SPACE=(CYL,(10)),                                            00230000
//         DISP=NEW                                                     00240000
//EXTDATA  DD   DISP=SHR,                 - EXTRACTOR INPUT DATA        00250000
//         DSN=?BBCHILV.SYS?BBASMFID.CPMOUT1                            00260000
//*IPSLIB  DD   DISP=SHR,                 - INPUT FOR GRAPH TYPE=DOMAIN 00270000
//*        DSN=SYS1.PARMLIB                                             00280000
//SYSIN    DD   DISP=SHR,                 - ANALYZER CONTROL STATEMENTS 00290000
//         DSN=?BBCHILV.SYS?BBASMFID.UBBPARM(ANLYSAMP)                  00300000
//*PUNCH   DD   DISP=SHR,                 - OUTPUT FILE, IF USING       00310000
//*        DSN=?BBCHILV.MACSOUT           - MACSCHAR STATEMENT          00320000
//*CMXREC  DD   DISP=SHR,                 - COPY FILE, IF USING         00330000
//*        DSN=?BBCHILV.SAVEDATA          - CMFREC STATEMENT            00340000
//RPTCONTS DD   SYSOUT=*                  - REPORT TABLE OF CONTENTS    00350000
//CMFLOG   DD   SYSOUT=*                  - COLLECTION PHASE LOG        00360000
//SYSPRINT DD   SYSOUT=*                  - REPORTS                     00370000
//CMFPRINT DD   DUMMY                     - REPORTS                     00370000
//CMXTRACE DD   SYSOUT=*                  - TRACE DATA                  00380000
//SNAPS    DD   SYSOUT=*                  - ANALYZER SNAPS              00390000
//SNAPVBS  DD   SYSOUT=*                  - INVALID RECORD SNAPS        00400000
//SYSUDUMP DD   SYSOUT=*                  - ANALYZER ABEND              00410000
//                                                                      00420000
```

Descriptions of Analyzer JCL Statements

All of the JCL statements shown in Figure 3-10 on page 3-11 are described in Table 3-1.

Table 3-1 JCL Control Statements for the CMF MONITOR Analyzer (Part 1 of 4)

JCL Control Statement	Description
//CMFRPTS EXEC	• Required; specifies the program name (CMFANLYZ) for the Analyzer, the region size, and other processing parameters. • BMC Software recommends a region size of 6MB. • The PARM field defines either CTRLSIZE or NLOG, or both. • PARM='CTRLSIZE=*xxx*K defines the amount of dynamic work area 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 3-18. • PARM='NLOG' eliminates printing of Extractor characteristics, the IPS, and the SRM Constants Report of the Collection Phase Log. See the section titled, "Preliminary Reporting Information" in the *CMF MONITOR Batch Reference Guide* 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. If omitted, the Analyzer will use a hiperspace for its work file. 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 specified as DD DUMMY, a hiperspace is used as the work file.
//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: SORT FIELDS=(15,4,CH,A,11,4,BI,A,7,4,BI,A),EQUALS This will re-order the records by SYSID, Date, and Time. • This may 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 the SORT statement shown above fails to re-order the records, you must return to the original data if it is available.

Table 3-1 JCL Control Statements for the CMF MONITOR Analyzer (Part 2 of 4)

JCL Control Statement	Description
//IPSLIB DD	• Optional; defines a partitioned data set that contains an IPS member to be used for domain graphics (see the section titled, “Performance Objectives by Domain Graph” in the *CMF MONITOR Batch Reference Guide* for more information). The Analyzer dynamically allocates SYS1.PARMLIB if this DD is omitted, and if GRAPH TYPE=DOMAIN,IPS=*xx* is requested. • If you need this statement defined, you must remove the comment (*) character from the CMFJANL member.
//SYSIN DD	Required; 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 used, the //RPTCONTS DD statement 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 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*.
//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 sometimes are 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 3-1 JCL Control Statements for the CMF MONITOR Analyzer (Part 3 of 4)

JCL Control Statement	Description
//SYSPRINT DD	• Required; 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 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 may 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.
* see Chapter 8, “Laser Printer Support” for more information	

Table 3-1 JCL Control Statements for the CMF MONITOR Analyzer (Part 4 of 4)

JCL Control Statement	Description
//SYSUDUMP DD	Required; provides for a dump if a program fails.
//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 the section titled, "SHIFT" in the *CMF MONITOR Batch Reference Guide* for more information on 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 of defining these DD statements: • 1. Define BOTH DD statements with: 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, a user may 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 are IEC data set open/close DFP messages and may be dependent on (1) how the units are managed (for example, are the volumes SMS managed?) (2)are theymounted PRIVATE or STORAGE? (3) other vendor software hooks relating to Open/Close processing. (4) the DFP maintenance release level in place, etc. When this happens, the best way is to preallocate the staging data set with method (3) above, 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 (see Chapter 8, "Laser Printer Support" for more information). 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 16MB 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 3-14 illustrates how these storage areas are configured.

Figure 3-14 Storage Area Configuration

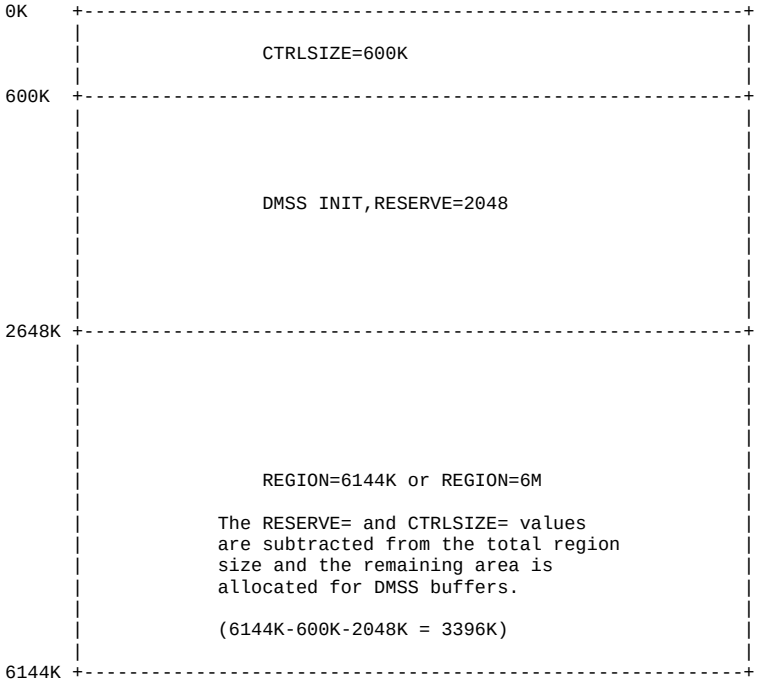

The default values for RESERVE and CTRLSIZE are dynamically computed by the Analyzer as follows:

- CTRLSIZE is set to about 20% of the OS Region size, or to what storage 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 usually will end up being about 33-38% of OS Region size.

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

BMC Software recommends a minimum region size of 1MB, with at least 3MB 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 may not be big enough. This can result in user abends U008, U100, or U0999, or system abends S80A, S878, or S106.

If you experience any of the noted user abends, use the following guidelines:

- First, set 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 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 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 1K 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, an additional 100K should be added.

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, so 2880 (96 * 20) reports are produced. This means you will need a CTRLSIZE value of at least 2.9MB.

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, that 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.

Figure 3-15 on page 3-22 is an example of the ANLYSAMP member.

The reports produced by and function of each statement defined in ANLYSAMP, shown in Figure 3-15 is discussed in *CMF MONITOR Batch Reference Guide*.

Figure 3-15 ANLYSAMP Control Statement

```
*********************************************************************
*    GENERAL CONTROL STATEMENTS                                     *
*********************************************************************
*
     RECTYPE 240             -  SET USER RECORD TYPE(DEFAULT)
 *
*********************************************************************
*    REPORT CONTROL STATEMENTS                                      *
*********************************************************************
*
*******> GENERAL
*
*
     CMFSTAT
     CMFSUM
     PERFSUM
*
*******> STORAGE
*
     AUXSTOR
     COMMSTOR REPORT=SUMMARY
     SRM
     STORAGE
     VIRTSTOR DETAIL=YES
*
*******> CPU
*
     CPU
     CPUCHAN
     CPUCON
     ENQUEUE THRESHLD=100
     PROTKEY
     PRSM
     TRACE
     XCF
*
*******> PERIPHERALS
*
     CACHE
     DASD
*     VOLSER VVVVVV,WWWWWW  <=== ENTER VOLUME(S) FOR DETAIL DASD REPORT
     DEVACT
     LCU
     SHARDEV TYPE=DASD
*
*******> 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,TYP70CPU),
        LIMIT=(10000,100000),CPU=ALL
*
*******> TSO
*
     TSOPERF
     TSOUSER
*
*******> WORKLOADS
*
     WKLDCHAR
     WORKLOAD TYPE=(DOMAIN,SUMMARY,PERFGRP)
```

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 the section titled, “REPORT” in the *CMF MONITOR Batch Reference Guide* 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 the section titled, "RECTYPE" in the *CMF MONITOR Batch Reference Guide* 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 may 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 the *CMF MONITOR Batch Reference Guide*.

The following list shows all Analyzer general control statements:

CMFREC	HEADERS
COMBINE	PERIOD
CYCLE	RECTYPE
DATA	SEVERITY
DATETIME	SHIFT
DMSS	SYSPLEX

Using Report Control Statements

Report control statements appear after the general control statement set. Most of the Analyzer's report control statements define specific reports to be produced and provide parameters for organizing or filtering report contents. The PERFORM statement, however, does not cause a report to be produced but modifies other report statements.The Report control statements are fully described in the General Control Statements section in the *CMF MONITOR Batch Reference Guide*.

The following list shows all report control statements:

AUXSTOR	CPUCON	LINKPACK	STORAGE
CACHE	DASD	MDF	TRACE
CFACT	DEVACT	OMVS	TSOPERF
CMFSTAT	DOMINO	PERFORM (modifier)	TSOUSER
CMFSUM	ENQUEUE	PERFUM	VIRTSTOR
COMBCONF	EXCEPTS	PGDDLAY	VOLSER
COMBINE (modifier)	GRAPH	PROTKEY	WKLDCHAR
COMMSTOR	HFS	PRSM	WLMGL
CPU	HTTP	SHARDEV	WORKLOAD
CPUCHAN	LCU	SRM	XCF

Using and Interpreting Reports

The information in CMF MONITOR's reports can be used to define thresholds for key system resources. They summarize performance data by job class, performance group, or domain; they report on DASD head movement activity, CPU usage, link pack area activity, and TSO usage data classified by command or by performance group.

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 the section regarding "Preliminary Reporting Information" in the *CMF MONITOR Batch Reference Guide* for more information.

Knowing What Reports You Need

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

The following sections show the five categories of reports and the Analyzer control statements that produce reports belonging to each category. See the section titled, “Report Control Statements” in the *CMF MONITOR Batch Reference Guide* for more information about each control statement and the reports it produces.

WORKLOAD Reports These reports include the following:

OMVS	TSOPERF	WKLDCHAR	WORKLOAD
PGDDLAY	TSOUSER	WLMGL	

CPU Reports These reports include the following:

CPU	CPUCHAN	CPUCON	MDF	PROTKEY	PRSM

SYSTEM RESOURCE Reports These reports include the following: S

AUXSTOR	ENQUEUE	SRM	TRACE	XCF
COMMSTOR	LINKPACK	STORAGE	VIRTSTOR	

DEVICE Reports These reports include the following:

AUXSTOR	CACHE	DEVACT	DASD	LCU	SHARDEV

Web-Related Reports These reports include the following:

DOMINO	HTTP

Miscellaneous Reports These reports include the following:

CMFSTAT	CMFSUM	EXCEPTS	GRAPH	HFS	TRACE

An example of each CMF MONITOR report is provided in the section titled, “CMF MONITOR Reports” in the *CMF MONITOR Batch Reference Guide*, as well as field descriptions and calculations. For instructions on capturing these reports, go to “Capturing CMF Analyzer and CMFMON Reports” on page 6-2.

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 this, together with the field descriptions and calculations for all report data, are provided in the section titled, “CMF MONITOR Reports” in the *CMF MONITOR Batch Reference Guide*.

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 E’s.

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 3-16.

Figure 3-16 Example of a Standard Report Heading

```
PRODUCED BY CMF MONITOR (5.3.0)              DOMAIN DELAY ANALYSIS REPORT
RPTSEQ     11  PAGE    31
BMC SOFTWARE, INC.                                  XYZ COMPANY
REPORT DATE: 29 MAR 99 13.53
REQD 29 MAR 99 09.00.00  29 MAR 97 16.59.59      WORLDWIDE HEADQUARTERS
SYSTEM ID:    SYSA  SP4.3
SHFT 09.00.00 17.00.00 TH
REPORT CYCLE: DAILY
ACTL 29 MAR 99 10.01.34  29 MAR 99 16.26.18
BASED ON REC TYPE/# RECS/# SAMPLES/REC HOURS:    72-2/140/15K/7.42
```

Descriptions of Report Headings Fields

A description of each field in a Report Heading is included in Table 3-2.

Table 3-2 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 3-16). • 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.
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)

Table 3-2 Field Descriptions for a Report Heading (Part 2 of 2)

Field Name	Description
(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 ID number. • If the Analyzer SYSPLEX control statement is not specified, the first system ID encountered in the input data set is the one that is reported. The system ID is the same ID value used by SMF. • 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* is printed in the system ID field of each report heading.
SP	Release number of the MVS system. • If the report contains data from a single MVS release, the release number appears in this field. • If the report contains data from multiple MVS releases, COMB-MVS appears in this field.
REPORT CYCLE	Report cycle based on CYCLE control statement.

Writing Your Own Programs to Process Extractor Data

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

There are three groups of members in BBSAMP that provide

- SAS
- C structures
- Assembler MACROs

For more information about CMF MONITOR's SMF record formats, see Chapter 7, "Mapping CMF Records Created by CMF".

Chapter 4 Preprocessing Extractor Data Sets

There are several reasons you may 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 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 may 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 may 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 4-6 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 4-1 on page 4-2, to execute the CX10CVBS utility.

Once 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 4-1 Sample Execution JCL for CX10CVBS

```
//JOBCARD JOB                                                           00010000
//*                                                                     00020000
//*--------------------------------------------------------------------- 00030000
//*                                                                     00040000
//* SAMPLE JCL FOR EXECUTING THE COPYVBS UTILITY.                       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 ?BBASMFID TO THE SMF ID (SYSTEM ID) OF THE TARGET          00130000
//*   SYSTEM.                                                           00140000
//*                                                                     00150000
//*   SPECIFY COPY OPTIONS IN THE PARM= FIELD OF THE EXEC STATEMENT.    00160000
//*                                                                     00170000
//*   CHANGE THE SYSUT1 DD STATEMENT TO POINT TO THE DATA SET YOU       00180000
//*   WISH TO COPY RECORDS FROM.                                        00190000
//*                                                                     00200000
//*   CHANGE THE SYSUT2 DD STATEMENT TO POINT TO THE DATA SET YOU       00210000
//*   WISH TO COPY RECORDS INTO.                                        00220000
//*                                                                     00230000
//*--------------------------------------------------------------------- 00240000
//*                                                                     00250000
//CMFCVBS  EXEC PGM=CX10CVBS,REGION=4096K,PARM='TYPE=CPM'               00260000
//*                                                                     00270000
//STEPLIB  DD   DISP=SHR,                 - BBLINK LOAD LIBRARY         00280000
//         DSN=?BBCHILV.BBLINK                                          00290000
//SYSUT1   DD   DISP=SHR,                 - "COPY FROM" DATA SET        00300000
//         DSN=SYS1.MAN1                                                00310000
//SYSUT2   DD   DISP=SHR,                 - "COPY INTO" DATA SET        00320000
//         DSN=?BBCHILV.SYS?BBASMFID.CPMOUT1                            00330000
//                                                                      00340000
```

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 4-1 on page 4-3 defines the CX10CVBS parameters. Examples of how to define PARM parameter values are shown in Table 4-2 on page 4-5.

Table 4-1 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 is 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. **Type** — SMF*rec*ID RMF — 70 through 79. If the *SMFrecID* value is omitted for TYPE=RMF, all RMF records (70 through 79) are copied. SMF — 1 through 255. If the *SMFrecID* value is omitted for TYPE=SMF, all SMF records (1 through 69 and 80 through 127) are copied. CPM — 70 through 79, and 128 through 255. If the *SMFrecID* value is omitted for TYPE=CPM, all CPM records are copied.

Table 4-1 CX10CVBS Parameters (Part 2 of 2)

Parameter	Definition
	IPM 70 through 79, and 128 through 255. If the *SMFrecID* value is omitted for TYPE=IPM, all IPM records are copied. CPR 70 through 79 CPM records, formatted to be like RMF records. IPR 70 through 79 IPM records, formatted to be like RMF records.
SUB=*nn*	This parameter is 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.
STARTAFT=*n*	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.
STOPAFT=*n*	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.

Because of the many different record types and ID combinations that can be specified, examples are shown in Table 4-2.

Table 4-2 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 4-3 describes the return codes issued by the I/O routine and control statement parsing routines of the CX10CVBS utility.

Table 4-3 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

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 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. This 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 5 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 is provided in “Implementing the CX10GVID API” on page 5-21.

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 below 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 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. This 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 system(s) 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 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 system(s) 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:

Step 1 Select Option 1 on the main menu for producing Analyzer reports.

Step 2 Specify XDS in the **SOURCE OF INPUT DATA** field.

If you are setting up your own JCL, the Analyzer automatically uses data 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 5-1 shows the format of the values 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 5-1 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 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
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.	CL3
	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.	

Table 5-1 Parameters for Calling CX10XDQY (Part 2 of 3)

Parameter	Description	Format and Length
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 4 blanks to the right of the string.	CL7
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:	CL7

Table 5-1 Parameters for Calling CX10XDQY (Part 3 of 3)

Parameter	Description		Format and Length
	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.	
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 4-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 5-15. The directory entry contains a record token created by CX10XDQY, which may 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 5-2 describes the format of the header.

Table 5-2 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

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 5-3 describes the format of the header.

Table 5-3 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)

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 5-4 shows the format of the values you assign to the parameters.

Table 5-4 Parameters for Calling CX10XDRC

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 8-character token values.	Array: FL4 + XL8 + XL8 + XL8 + ...
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 the Common Answer Area Header" on page 5-15. 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 may 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 5-5 shows the format of the values 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 5-5 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 is the 4-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 2-byte record type; this is followed by the 2-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 5-5 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 5-21 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	Length 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 5-15. Each data section contains a data header followed by the data itself.

A sample DSECT for the CX10XDGS data section is shown below:

```
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 5-6.

```
CALL exit_name,
   (answer_area_addr
   ,answer_area_alet
   ,answer_area_length
   ,output_area_length
   ,input_data_address
   ,exit_parm
   ,exit_parm_length)
```

Table 5-6 Parameters for Calling the CX10XDGS Exit

Parameter	Description	Format and Length
answer_area_addr	Address of the area where the exit routine may 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 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.	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 the requested data.

XDRHPLX Name of the sysplex the calling application is running on.

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 5-7 contains all return codes and reason codes for each of the three XDS APIs.

Table 5-7 Return Codes for XDS APIs (Part 1 of 4)

Return Code	Reason Code	API	Meaning
0	0	CX10XDQY CX10XDRC CX10XDGS	The operation was 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 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.
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.

Table 5-7 Return Codes for XDS APIs (Part 2 of 4)

Return Code	Reason Code	API	Meaning
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 8, 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	26	CX10XDGS	Error - System in WLM GOAL mode. The system is in MVS Workload Manager (WLM) GOAL mode. Therefore, the domain activity or transaction activity data (record type 79, subtypes 8 or 10) cannot be gathered.
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.
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.

Table 5-7 Return Codes for XDS APIs (Part 3 of 4)

Return Code	Reason Code	API	Meaning
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.
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.
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.

Table 5-7 Return Codes for XDS APIs (Part 4 of 4)

Return Code	Reason Code	API	Meaning
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	Severe 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.
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. CMFMON's 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 and OS/390" on page 5-24.

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 5-8 lists the Extractor samplers that are required for SMF 79 record subtypes:

Table 5-8 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 following fields are returned only when the caller is running in Supervisor state or is APF-authorized:

Table 5-9 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/swap data set device name Page/swap 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 real storage utilization fields of subtype 2 (listed below) contain null data.

- Bit R792RSM of byte R792FLG is on; this signifies that the fields are invalid.

Subtype 2 real storage utilization fields affected by BBXS are shown in Table 5-10.

Table 5-10 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 below:

Table 5-11 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 2-12 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 5-12 Return Codes

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
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 and OS/390
- CMF APIs and SDSF

CMF APIs and OS/390

As of OS/390 Release 2, IBM is shipping 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

At OS/390 2.6.0 or later, RMF APIs reside in SYS1.SERBLINK. Ensure that MVS rules of module search find the APIs in the *hilevel*.BBLINK library before SYS1.SERBLINK.

At OS/390 2.5 or earlier, RMF APIs are shipped in SYS1.LINKLIB. This situation, disabled RMF APIs found from SYS1.LINKLIB, is particularly noticeable when using SDSF (see “CMF APIs and SDSF” on page 5-26) since SDSF uses the RMF instead of the CMF API names. A common symptom of this problem is getting either the RMF NOT ENABLED or the RMF SYSPLEX NOT ACTIVE message. For more information, see “CMF APIs and SDSF” on page 5-26. IBM has agreed that this situation results from an OS/390 packaging problem that will be resolved; however, the solution is not available currently.

Perform one of the following actions to locate the CMF version (instead of the distributed disabled version) of these API routines.

- Use the new SYSLIB facility to logically replace LINKLIB with a data set other than SYS1.LINKLIB. Then, place the BBLINK data set ahead of SYS1.LINKLIB in the linklist. SYS1.LINKLIB must be explicitly included. An IPL of the system is required to implement this change.

- Copy the CX10GVID, CX10XDU0, and CX10XDGX and all of their aliases to a new data set. Then, use the SYSLIB facility to add that data set ahead of SYS1.LINKLIB in the linklist. This option, however, requires that any maintenance to CX10GVID, CX10XDU0, and CX10XDGX be performed in both BBLINK and the copied data set.

- Rename the ERB* alias names in SYS1.LINKLIB. If you choose this option, consider applying an SMP USERMOD to the members. This way, you will be informed if IBM maintenance is applied that relinks the API routines, and you can rename the aliases again after you remove the USERMOD and apply the IBM maintenance.

- Delete the ERB* alias names mentioned above from SYS1.LINKLIB. Consider applying SMP USERMOD as described in Option 3.

Until IBM resolves the problem, repeat the action taken for new OS/390 installations.

CMF APIs and SDSF

Various releases of IBM's 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 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 may 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 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 and OS/390" on page 5-24 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 and/or ERB2XDGS aliases to obtain data.

Step 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.

Step 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.

Step 3 Ensure that the BBLINK data set is available in the linklist (see “CMF APIs and OS/390” on page 5-24 for additional considerations), or ensure that the BBLINK data set is available in the TSO logon procedure STEPLIB.

Step 4 Specify CPU sampling in the CMFCPMxx PARMLIB member.

Step 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 CASes.
- An MVS PAS with XDS active. (By default, XDS=00 is distributed in the MVSPAS product.)
- The MVS PASes must all have DC=START in effect.
- The MVS PASes must have CXEN=Y in effect.

Chapter 6 CMF 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 may 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 97 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 on using Excel should be directed to Microsoft Corporation.

Installing the Spreadsheet Converter on to Your PC

The spreadsheet converter is distributed with CMF as member CX98SSCX of the BBSAMP and UBBSAMP libraries. To install the converter:

Step 1 Transfer the file CX98SSCX to your PC by 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.

Step 2 Rename the file CX98SSCX.XLA

Step 3 You may want to save the file in its own directory or folder for easy access.

Note: Information on 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.

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 may 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. This 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 *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 may 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 on the CX98SSCX.XLA file in the Explorer or File Manager.

Step 1 Once the spreadsheet converter is loaded, begin execution by selecting the BMC Software toolbar from the View menu item, as shown in Figure 6-1.

Figure 6-1 Select BMC Software toolbar from the Excel Menu

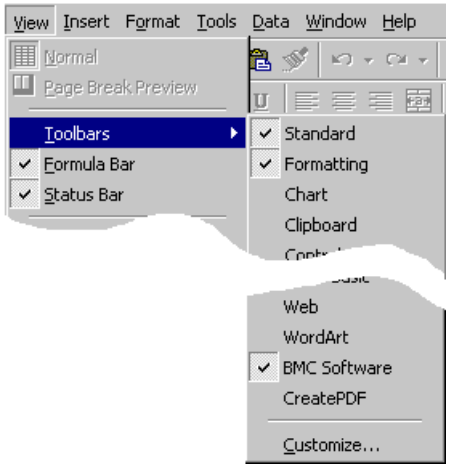

Step 2 When you select **Instructions**, the BMC Software Spreadsheet Converter Instruction screen is displayed, as shown in Figure 6-2.

Figure 6-2 Spreadsheet Converter Instruction Screen

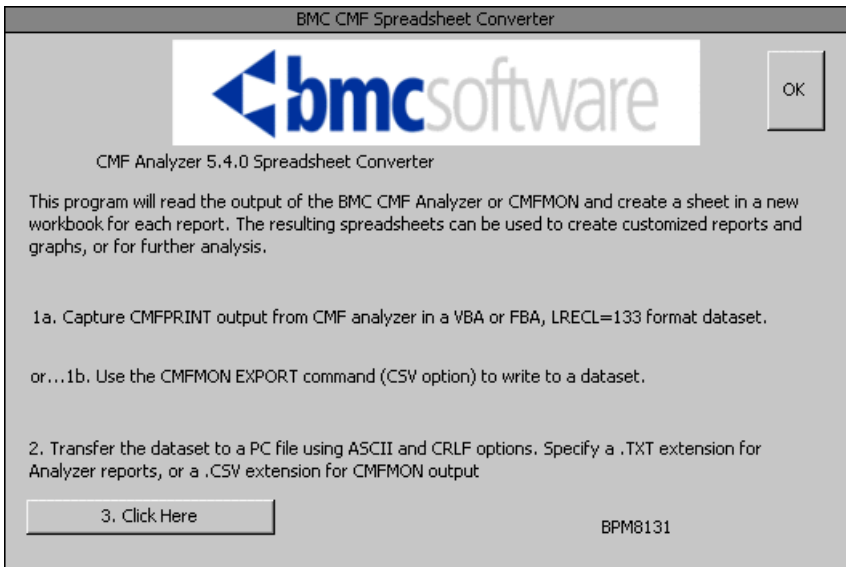

Step 3 When you select **Click Here**, the spreadsheet converter displays a list of the files you downloaded to your PC, as shown in Figure 6-3.

Figure 6-3 Select the file containing the reports you want to convert to Excel.

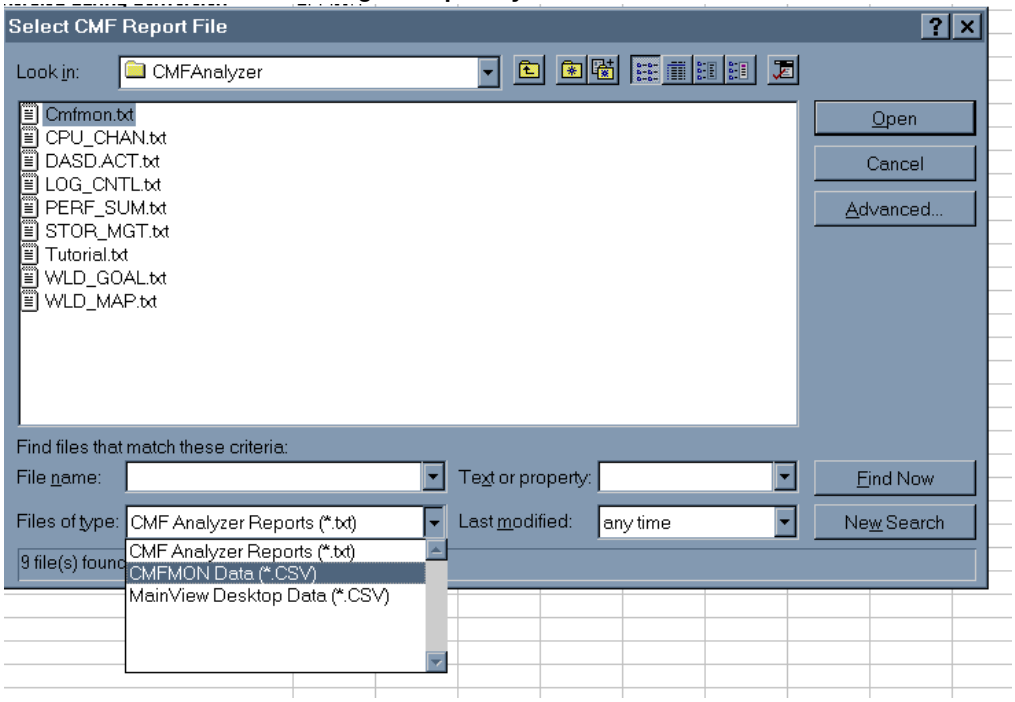

When you select a report, the spreadsheet converter

- Creates a new workbook in which to store the results of the conversion.
- Automatically converts the reports in the file you select.

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

Figure 6-4 BMC Software toolbar

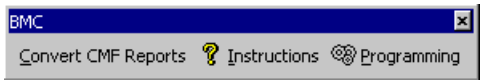

Step 4 When the file contents have been converted, the Conversion Log is displayed, as shown in Figure 6-5 on page 6-6.

Figure 6-5 Conversion Log for Converted Reports

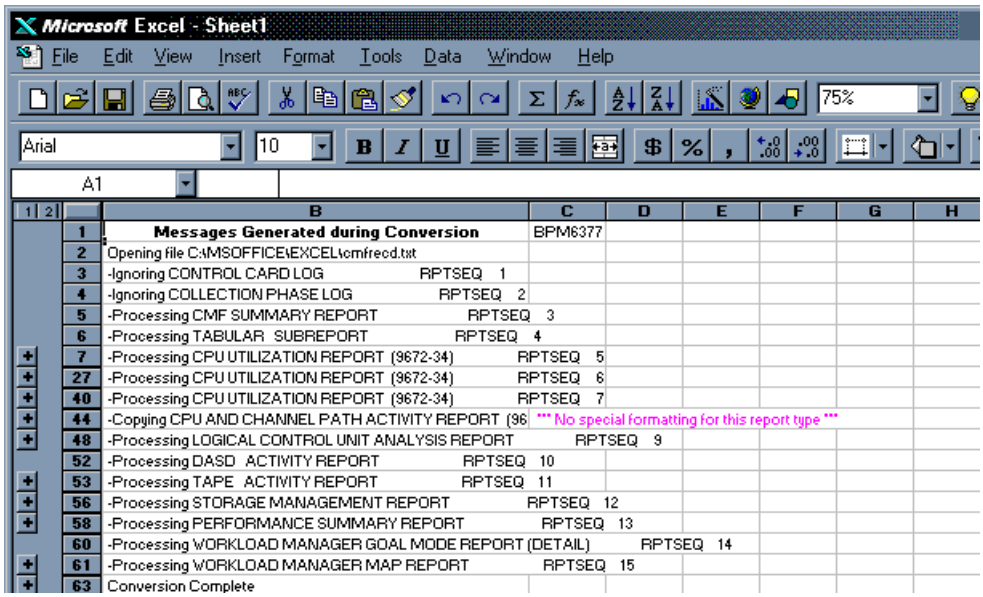

The Conversion Log is the last page in the workbook. The Conversion Log lists reports found in the input along with a description of how each one was processed. Any error messages produced during processing are also listed on this page. Additional information on any error messages you see can be displayed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

Step 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 further information on 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.

Step 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 to Your PC” on page 6-1 for additional information.

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

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

Step 4 Run the job.

This creates a file containing sample CMF Analyzer reports.

Step 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.

Step 6 Open the spreadsheet converter (CX98SSCX.XLA) in Excel by selecting **Add-ins** from the **Tools** menu.

Step 7 Select CMF Spreadsheet Converter from the **Add-Ins Available** list box shown in Figure 6-6.

Figure 6-6 Excel Add-Ins Available List Box

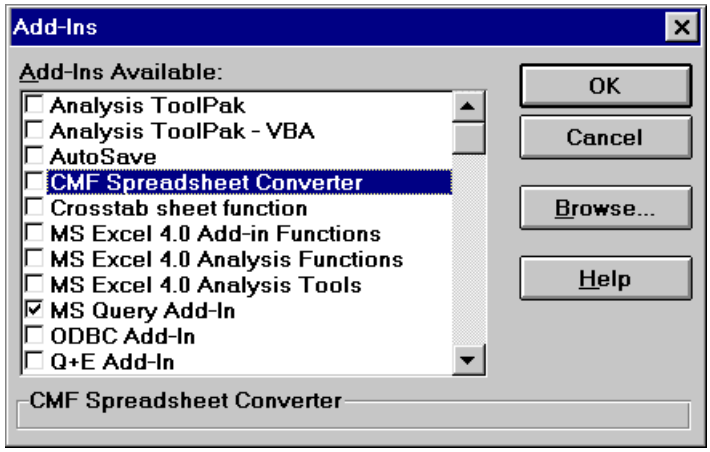

Step 8 Run the spreadsheet converter by selecting Convert CMF Reports. A dialog box opens to display a list of your report files.

Step 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.

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

Step 11 Use your mouse to select the CMF Summary Report. The report displayed should look similar to Figure 6-7. In this particular spreadsheet, a line graph was added to provide visual comparison between CPU, DASD, and TSO information.

Figure 6-7 Converted CMF Summary Report

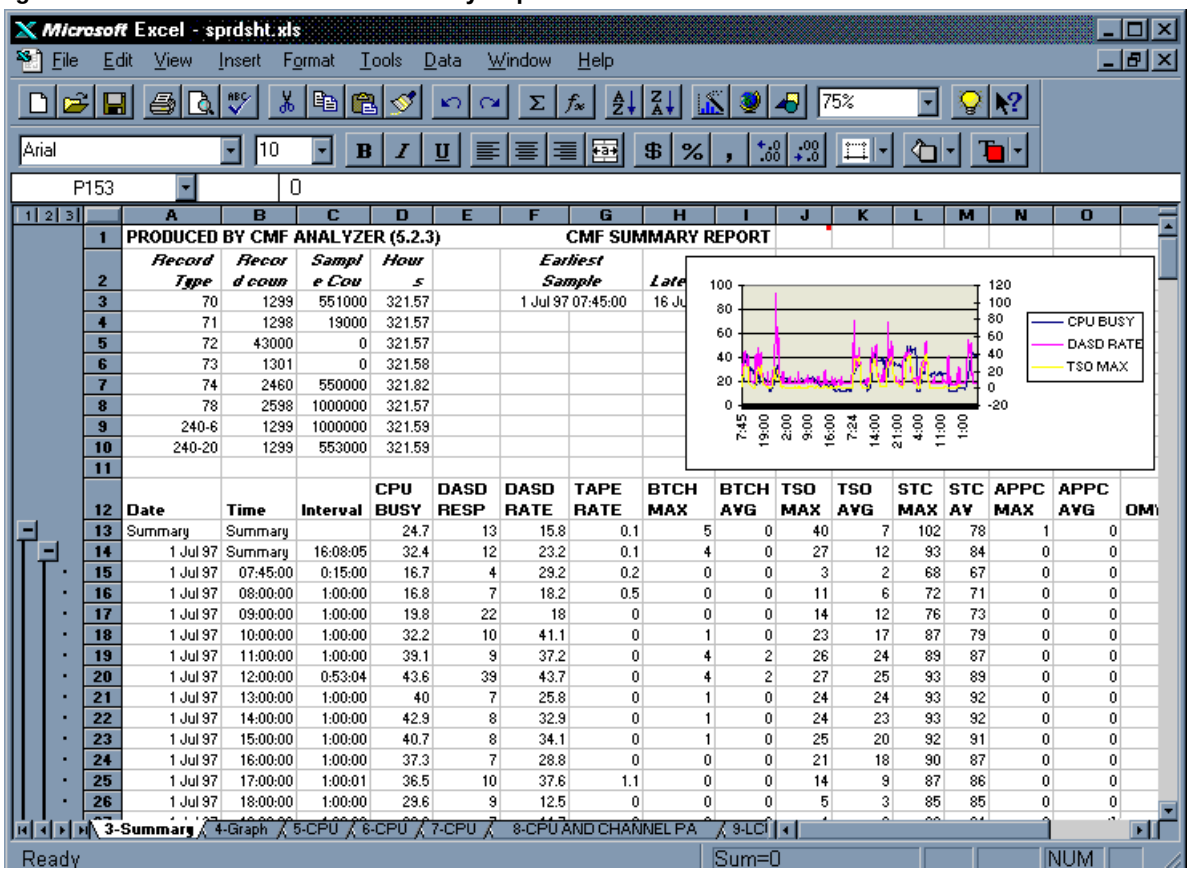

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1	PRODUCED BY CMF ANALYZER (5.2.3)						CMF SUMMARY REPORT								
2	*Record Type*	*Record coun*	*Sample Cou*	*Hours*		*Earliest Sample*		*Late*							
3	70	1299	551000	321.57		1 Jul 97 07:45:00		16 Ju							
4	71	1298	19000	321.57											
5	72	43000	0	321.57											
6	73	1301	0	321.58											
7	74	2460	550000	321.82											
8	78	2598	1000000	321.57											
9	240-6	1299	1000000	321.59											
10	240-20	1299	553000	321.59											
11															
12	Date	Time	Interval	CPU BUSY	DASD RESP	DASD RATE	TAPE RATE	BTCH MAX	BTCH AVG	TSO MAX	TSO AVG	STC MAX	STC AV	APPC MAX	APPC AVG
13	Summary	Summary		24.7	13	15.8	0.1	5	0	40	7	102	78	1	0
14	1 Jul 97	Summary	16:08:05	32.4	12	23.2	0.1	4	0	27	12	93	84	0	0
15	1 Jul 97	07:45:00	0:15:00	16.7	4	29.2	0.2	0	0	3	2	68	67	0	0
16	1 Jul 97	08:00:00	1:00:00	16.8	7	18.2	0.5	0	0	11	6	72	71	0	0
17	1 Jul 97	09:00:00	1:00:00	19.8	22	18	0	0	0	14	12	76	73	0	0
18	1 Jul 97	10:00:00	1:00:00	32.2	10	41.1	0	1	0	23	17	87	79	0	0
19	1 Jul 97	11:00:00	1:00:00	39.1	9	37.2	0	4	2	26	24	89	87	0	0
20	1 Jul 97	12:00:00	0:53:04	43.6	39	43.7	0	4	2	27	25	93	89	0	0
21	1 Jul 97	13:00:00	1:00:00	40	7	25.8	0	1	0	24	24	93	92	0	0
22	1 Jul 97	14:00:00	1:00:00	42.9	8	32.9	0	1	0	24	23	93	92	0	0
23	1 Jul 97	15:00:00	1:00:00	40.7	8	34.1	0	1	0	25	20	92	91	0	0
24	1 Jul 97	16:00:00	1:00:00	37.3	7	28.8	0	0	0	21	18	90	87	0	0
25	1 Jul 97	17:00:00	1:00:01	36.5	10	37.6	1.1	0	0	14	9	87	86	0	0
26	1 Jul 97	18:00:00	1:00:00	29.6	9	12.5	0	0	0	5	3	85	85	0	0

Step 12 Select the tab for the CPU Utilization Report. The report displayed should look similar to Figure 6-8 on page 6-10. In this particular spreadsheet, a bar chart was generated to compare performance between systems.

Figure 6-8 CPU Utilization Report

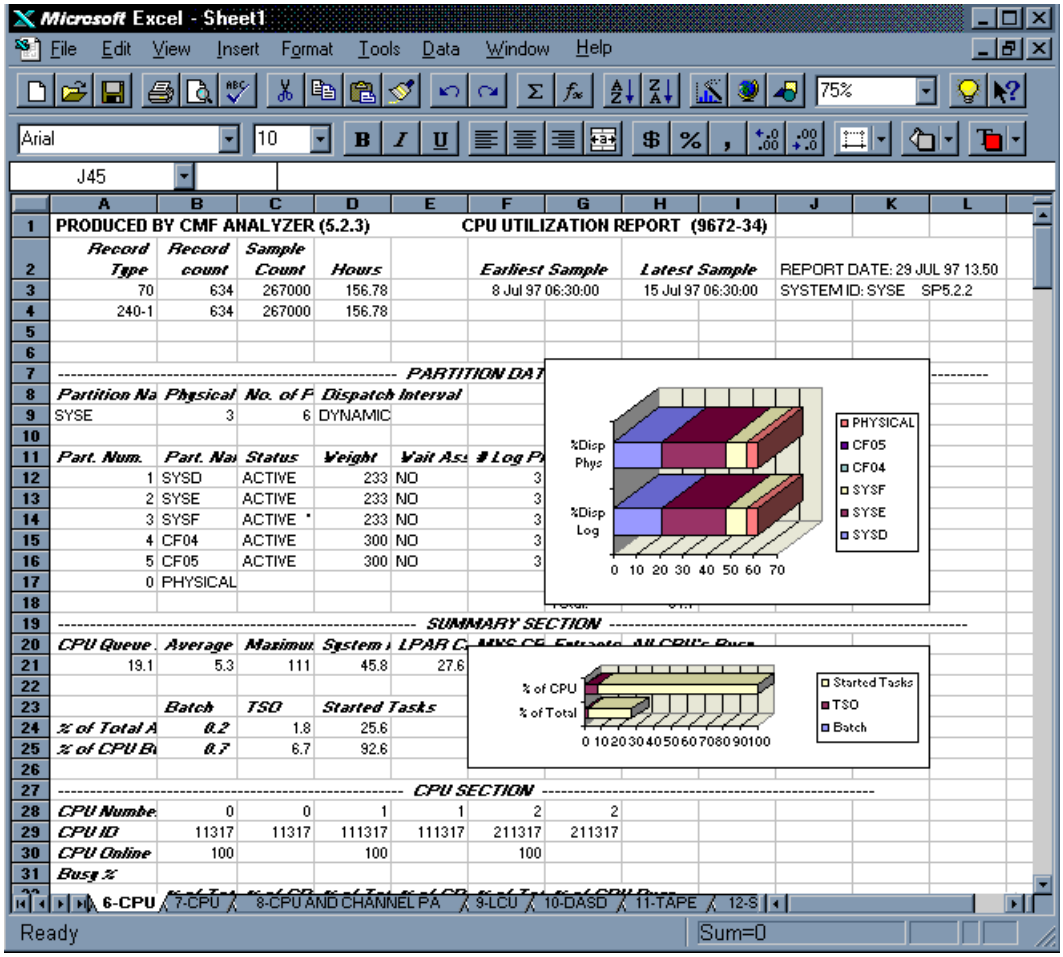

Step 13 Select the tab for the Workload Manager Map Report. The report displayed should look similar to Figure 6-9 on page 6-11. It is difficult to display time-sensitive data as columns of figures. In this example, the polar plot graphic quickly shows the times when goals were not met.

Figure 6-9 Workload Manager Map Report

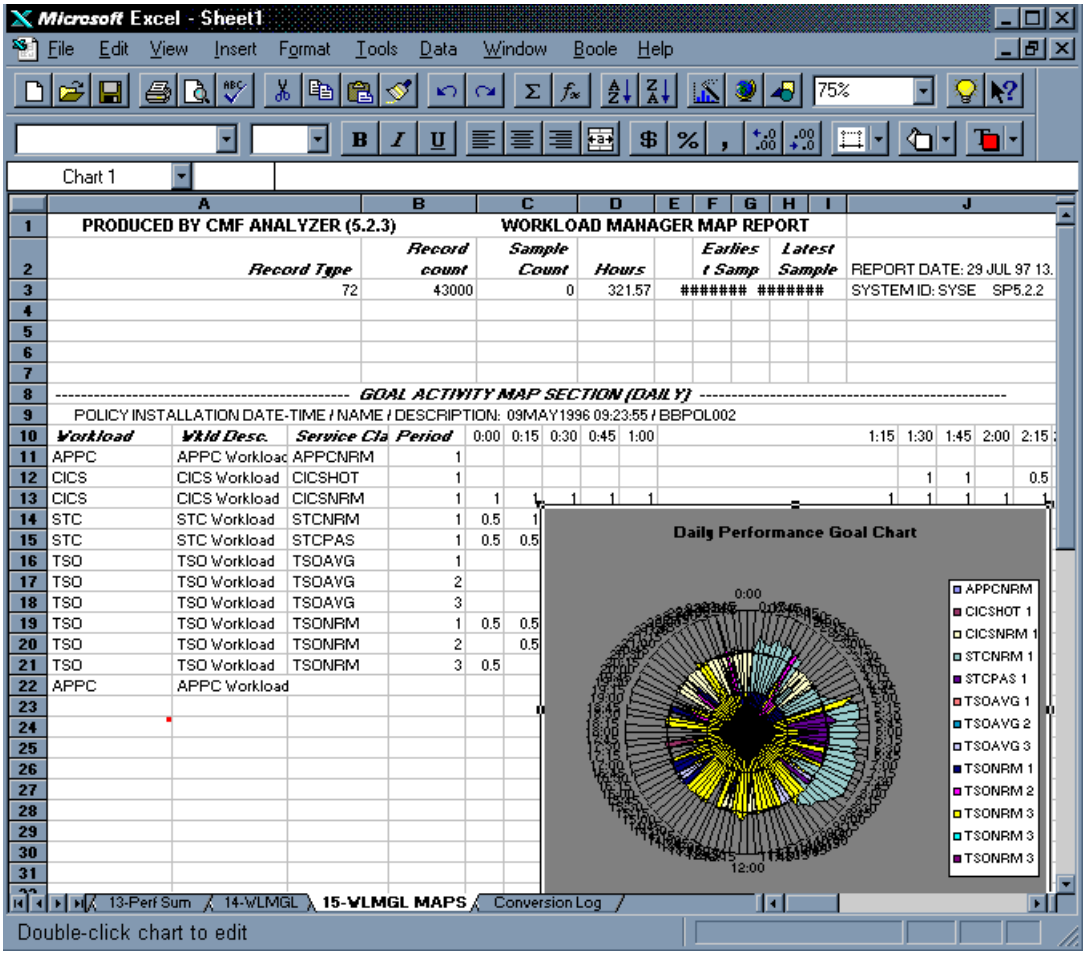

Step 14 Select other reports to see how the data has been enhanced by using the Excel tools.

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

Troubleshooting

While 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 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 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 that:
   - 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 to Your PC” on page 6-1.

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.

Chapter 7 Mapping CMF Records Created by CMF

The CMF MONITOR Extractor and the CMFMON Write Facility prepare SMF records. The records prepared match the type 70-79 records defined in the *IBM System Management Facilities (SMF)* manual. These 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 CMF MONITOR's 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 containing 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 the *CMF MONITOR Batch Reference Guide*, "REPORT" for more information).

The user record information in BBSAMP is provided in these three formats:

- Assembler MACROs
- C Structures
- SAS code

This chapter includes the following sections:

- "Assembler MACROs in BBSAMP for Record Types 70-79" on page 7-2
- "Assembler MACROs in BBSAMP for User Records" on page 7-2

- "C Structures in BBSAMP" on page 7-3
- "SAS Code in BBSAMP" on page 7-3
- "Using CMF User Records with MXG" on page 7-4

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:

- 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

OS/390 assembly language (and above) MACROs, which 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 that 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 showing how to use any of the CMFC*xx* members to write your own programs to process 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. Contains SAS code for reading the standard record header used for all CMF record subtypes. This member 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.

CSMAPSASThe COMMON STORAGE MONITOR records (subtype 29) contain 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's 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 following CMF user record types shown in Table 7-1.

Table 7-1 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
240-04	Workload data	WORKLOAD	CMF04WOR
240-05	Device data	DEVICE	CMF05DDS CMF05TDS
240-06	Extractor summary data	EXTSUM	CMF06GDA CMF06JDS
240-09	ASM data	ASMDATA	CMFASMQ
240-19	I/O workload record data	MACSCHAR	CMFTRACE
240-20	TSO command summary record data	TSODATA	CMF20CCS CMF20CSS
240-21	TSO user summary record data	TSODATA	CMF21USS
240-27	Cache data records	CACHE	CMF27C93 CMF27CAR CMF27CSD
Note: For details on each CMF Extractor statement, see the related section in the *CMF MONITOR Batch Reference Guide*.			
240-29	COMMON STORAGE MONITOR records	CSMON (see the *CMF MONITOR Batch Reference Guide*)	CMF29COS CMF29CJS CMF29CDS

Chapter 8 Laser Printer Support

You can use laser printer support for CMF MONITOR reports to enhance report output by controlling any fields in a report that are emphasized when printed.

The laser printer support currently applies to these CMF MONITOR reports:

- Logical Partition Report
- Multiple Domain Facility Report

Laser printers produce printouts based on available laser fonts and the print characteristics associated with any one report field. The Analyzer can produce report output for laser printers that accommodates font changes for bolding or other special emphasis.

With an impact printer, there is only one typeface or font available. Report fields are highlighted by repeatedly printing over the same field of information—this is called overstriking. The number of overstrikes to a special emphasis report field can be controlled for impact printers with this printer enhancement.

This section discusses the different types of emphasized data that can appear in reports and the changes required for the JCL to format output for a laser printer. A discussion of font selection for laser printing is also included.

Types of Emphasized Data in Reports

Report fields are emphasized when printed, to call attention to certain data and to separate report field names from report data. Reports can contain three types of printed fields:

- Regular

- **Bold type**
- ***Special emphasis***

Most of a report printout typically has report lines that are printed in a regular font. Some report field names or certain data may be in bold type when printed.

With the advent of laser printers, special emphasis, due to the availability of multiple fonts, began appearing on report printouts. Normally, special emphasis fields stand out from other report fields because the font used is aesthetically different from the regular or bold fonts.

Required JCL for Laser Printer Support

To provide the extra printer control information required for output sent to a laser printer, there are two operands to add to the Analyzer JCL:

Warning! These operands are required only for laser printer output and should not be defined when directing output to an impact printer.

OPTCD=J Required. This is a standard JCL operand that defines the output to a laser printer. When OPCODE=J is specified, every line of the output contains a font control character as the second character. The CMF Analyzer is sensitive to this laser printing operand on three DD statements:

SYSPRINT The primary DD statement to which reports are written. An OPTCD=J operand must appear on the //SYSPRINT DD statement for reports that support the new printing function to print on a laser printer.

RPTCONTS The DD statement where the Report Table of Contents is written. An OPTCD=J operand must appear on the //RPTCONTS DD statement for the table of contents to print on a laser printer.

CMFLOG The DD statement where the Collection Phase Log report output is written. An OPTCD=J operand must appear on the //CMFLOG DD statement for the Collection Phase Log to print on a laser printer.

CHARS= Optional. This is a standard JCL operand used to define the names of the laser printer fonts to be used when the output is printed.

You can define this operand in one of two ways:

— Add a //*name* OUTPUT statement that includes the CHARS= operand.

— Add the CHARS= operand to the //SYSPRINT DD, //RPTCONTS DD, or //CMFLOG DD statements.

Examples

Following is an example of how you can add the OUTPUT JCL statement with the CHARS= operand:

```
//jobname JOB . . .
//ANAL EXEC PGM=CMFANLYZ, . . .
//name01 OUTPUT CHARS=(font1,font2,font3,font4)
//name02 OUTPUT CHARS=(font1,font2,font3,font4)
//SYSPRINT DD SYSOUT=class,OUTPUT=*.name01,OPTCD=J
//RPTCONTS DD SYSOUT=class,OUTPUT=*.name02,OPTCD=J
```

Note: The CMF Analyzer cannot recognize the TRC=YES operand of the //*name* OUTPUT statement. BMC Software recommends that if OUTPUT statements are used, that one OUTPUT statement be created for each DD statement. OPTCD=J must be coded on the appropriate DD statement.

Following is an example of how you can specify the CHARS operand to either the //SYSPRINT DD or the //RPTCONTS DD statements:

```
//jobname JOB . . .
//ANAL EXEC PGM=CMFANLYZ, . . .
//SYSPRINT DD
SYSOUT=class,CHARS=(font1,font2,font3,font4),OPTCD=J
//RPTCONTS DD
SYSOUT=class,CHARS=(font1,font2,font3,font4),OPTCD=J
```

Selecting Fonts

You define different fonts for your report output by specifying the names of up to four different fonts at the CHARS operand, in place of the following:

- font1
- font2
- font3
- *font4*

You should define at least two fonts; three if the report contains special emphasis fields. Each font name must be separated by a **,** (comma), and all the font names must be enclosed within **( )** (parentheses).

Note: The number, type, and names of fonts available to your laser printer are customized at your installation. Contact your system operator to find out the names of the font selections available for your laser printer at your installation.

Note: The font control numbers are not used by impact printers.

Note: For a report printout to contain *font3* or *font4*, the report must provide special emphasis characteristics for those fields requiring special emphasis, and the LASER option must be defined to the report control statement.

The laser printer fonts can be experimented with to see which ones work best for your particular requirements. BMC Software recommends

font1 Used most commonly for printing.

font2 Bold version of the first font in the same pitch. This is selected for reports that contain bold type print lines

font3 and font4 Bold or unique type face that is easy to read for special emphasis; BMC Software recommends a pitch lower than font1 or font2.

Guidelines for Font Selection

Here are a few guidelines for font selection that may be useful:

- It is possible to select a font with a pitch that is too low. The pitch of a font is the number of characters per inch that print in that font. A low pitch means the characters are large; a high pitch means smaller print characters because more of them fit into one inch of a print line. If the pitch is too low and the report data runs off the page when printed, you can increase the size of the page, reduce the margins, or increase the pitch of the characters.
- It is important to select compatible fonts at the same pitch for *font1* and *font2*. Reports that contain bolding but do not support the special laser option overstrike *font1* with *font2*
- Special emphasis fields are produced using only *font3* and *font4* when the LASER option is selected for a report supporting the special laser option.
- A step up or down in pitch may be very effective for *font3* and *font4*.
- If special emphasis is used, at least three fonts should be selected.

Appendix 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.

The following topics are discussed in this appendix:

- Standard Deviation, the Mean, and the Mode
- Calculating Standard Deviation
- Statistical Accuracy

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 A-1shows this relationship for standard deviations to the mean.

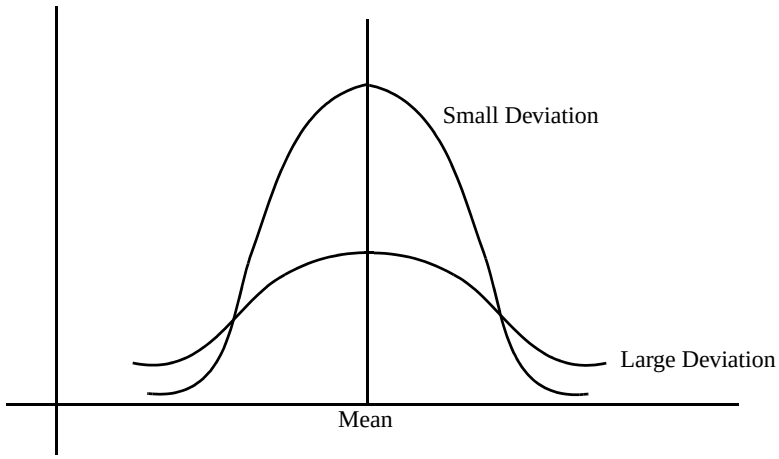

Figure A-1 Relationship of Large and Small Standard Deviations to the Mean

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 A-2.

$$\frac{\sqrt{N\sum_{i=1}^{N} X_i^2 - \left(\sum_{i=1}^{N} X_i\right)^2}}{N}$$

Figure A-2 Equation for Calculating Standard Deviation

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 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 *CMF MONITOR Batch User Guide*, Chapter 4. “Producing and Using Reports”, the section titled “Understanding Report Headings”).

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 A-3 and locate the

- Number of samples taken by the Extractor (the N-axis)
- The desired confidence level (one of the plotted diagonal lines)

The intersection point yields the uncorrected value for the confidence interval (E).

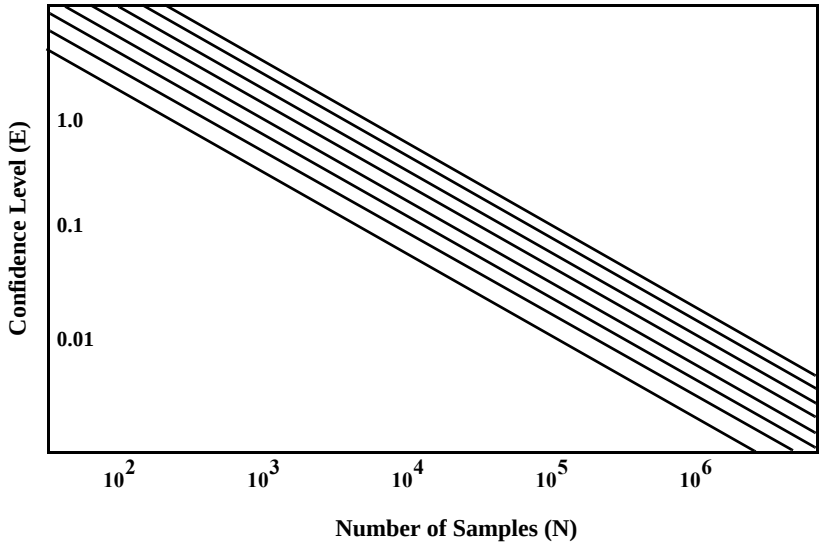

Figure A-3 Confidence Levels for P=50%

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 A-4.

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% and 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%.

This means that 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 A-4 Correction Factors for Confidence Intervals

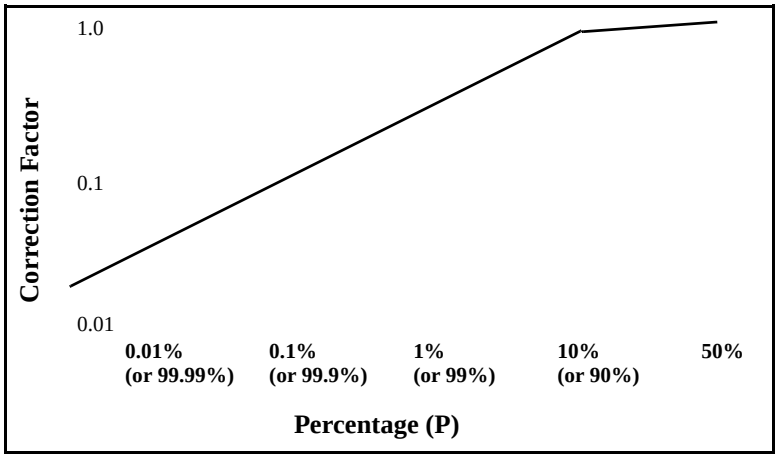

Appendix B Workload Measurement in Goal Mode

In MVS 5.1 and later, your system can run either in compatibility mode or in goal mode. In compatibility mode, the System Resource Manager (SRM) functions externally the same way it did on MVS 4.3 or earlier systems, using performance groups and domains to manage resources. If you are running MVS 5.1 or later in compatibility mode, CMF MONITOR provides additional information, such as execution velocity, for some compatibility mode reports. This information may assist you in determining appropriate goals when you begin planning a switch to goal mode.

In goal mode, 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.

Note: For OS/390 1.3, if you specify I/O priority management on the Service Coefficient/Service Definition Options panel, the following formula is used to calculate execution velocity:

$$\frac{\text{\# CPU using samples + \# I/O using samples}}{\text{\# CPU using samples + \# I/O using samples + \# general execution delays}}$$

The following is a list of general execution delays, as defined by WLM:

CPU delay:
TCB or SRB is waiting to be dispatched or a TCB is waiting for a the local lock.

CPU capping delay:
TCB or SRB is marked nondispatchable because a resource group maximum is being enforced.

Swap-in delay:
Swap-in has started, but not completed.

MPL delay:
Ready, but swap-in has 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.

For MVS OS/390 1.2 and later:

Shared paging:
Delay waiting for a shared storage page to be brought in.

For MVS OS/390 1.3 and later:

DASD I/O delay samples:
The number of samples of work delayed for paging DASD I/O.

Queue delay samples:
The number of samples of work that is waiting for a server.

Server private area paging delay samples:
The number of samples delayed for private area paging for a server.

Server space VIO paging delay samples:
The number of samples delayed for VIO paging for a server.

Server hiperspace paging delay samples:
The number of samples delayed for hiperspace paging for a server.

Server MPL delay samples:
The number of samples delayed for MPL for a server.

Server swap-in delay samples:
The number of samples delayed for swap-in for a server.

Index

A

abends
- S106 3-19
- S80A 3-19
- S878 3-19
- U008 3-19
- U0999 3-19
- U100 3-19

address space 1-5
alternate data sets 2-12
Analyzer
- storage 3-18

Analyzer storage 3-18
ANLYSAMP member
- control statements 3-24
- description 3-21 to 3-23

assembler MACROs 7-2

B

batch reports
- generation panels 3-3
- producing 3-3

BBSAMP
- assembler MACROs 7-2
- C structures 7-3
- SAS data set 7-3

BMC Software products 1-7
- MAINVIEW for OS/390 1-8

C

C structures 7-4
CMF MONITOR
- compatibility with IBM's RMF 1-9 to 1-10
  - differences 1-10
  - functions unavailable in RMF 1-10
  - running on the same system as RMF 2-14
  - similarities 1-9
- product components 1-1 to 1-5
  - Analyzer 1-5
  - DSO 1-5
  - Extractor 1-3
  - Online 1-4

CMF MONITOR Analyzer 1-3, 3-1 to 3-28
- defining JCL 3-12 to 3-24
  - ANLYSAMP member 3-21 to 3-22
  - general control statements 3-23
  - JCL control statements 3-12 to 3-20
  - report control statements 3-24
- JCL Generator 3-2
- JCL generator panels
  - filtering input data 3-10
  - listing previously specified members 3-11
  - selecting reports 3-7
  - setting up JCL 3-4
  - specifying CMFSUM parameters 3-9
  - specifying data source 3-5
  - specifying report output destination 3-5
  - submitting JCL 3-11
- reports

approaching data 3-24
CPU 3-25
device 3-25
field descriptions 3-27 to 3-28
generation 3-1
headings 3-27
interpreting 3-24 to 3-28
miscellaneous 3-25
system resource 3-25
using 3-24 to 3-28
web-related 3-25
workload 3-25
CMF MONITOR DSO 1-5
CMF MONITOR Extractor
alternate data sets 2-12
change operation 2-22
control statements 2-18 to 2-22
CMF MONITOR 2-16
CMFCPM00 2-18 to 2-20
CMFIPM00 2-20 to 2-21
DSO 2-17
MAINVIEW for OS/390 2-17
creating your own reports 3-27
defining JCL and control statements 2-17 to 2-23
CMFCPM00 control statement 2-18 to 2-20
CMFIPM00 control statement 2-20 to 2-21
MODIFY command 2-22
description 1-3
output 2-10
overview 2-1 to 2-16
CPM, mode 2-4
data set 2-12
how it works 2-2
IPM mode 2-4
monitoring modes 2-2
output 2-10
records 2-7
sampling methods 2-6
SRB trace facility 2-15
SRM trace facility 2-15
primary data sets 2-12
reports, creating your own 3-27
samplers 2-5 to 2-10
trace facilities 2-15
CMF type 79 API 5-21
CMFCPM00 control statement set 2-18 to 2-20
CMFIPM00 control statement set 2-20 to 2-21
CMFMON
overview 1-5
reports, exporting to spreadsheet converter 6-3
CMFPRINT 3-16
CMFSUM
parameters, specifying from JCL generator panels 3-9
collecting report data 2-1
command, HMOVRESCAN= 2-23
Control 2-17
control statements
Analyzer general 3-23
Analyzer JCL
CMFLOG DD 3-15
CMFPRINT DD 3-16
CMFRPTS EXEC 3-14
CMFSTAGE DD 3-17
CMFSTAGO DD 3-17
CMXREC DD 3-15
CMXTRACE DD 3-16
DMSSMAIN DD 3-14
EXTDATA DD 3-14
IPSLIB DD 3-15
OUTPUT 3-18
RPTCONTS DD 3-15
SNAPS DD 3-16
SNAPVBS DD 3-16
STEPLIB DD 3-14
SYSIN DD 3-15
SYSPRINT DD 3-16
SYSUDUMP DD 3-17
Analyzer reports 3-1 to 3-28
controlling data collection 2-2
CPM
mode 2-3
CPU
reports 3-25
cross-system API 5-1
CTRLSIZE 3-18
CX10CVBS
return codes 4-6
sample JCL 4-2
CX10GVID
API 5-1
general-purpose registers 5-23

CX10GVID API
calling 5-22
Extractor requirements 5-21
return codes 5-24
CX10XDGS
API 5-1
data reduction exit 5-13
sample DSECT 5-12
CX10XDQY
API 5-1, 5-3
sample DSECT 5-6
CX10XDRC 5-1, 5-8
CX98REPG file for spreadsheet converter tutorial 6-7
CX98SSCX 6-7
CX98SSCX.XLA
conversion log 6-7
workbook 6-7

D

data
archiving 2-14
sampling 2-6
data sets
alternate 2-12
CMF 2-12
DSO 2-12
primary 2-12
DEVACT
deviation, standard A-1
DEVICE
reports 3-25
DMSS Reserve 3-18
DSO, Extractor control statements 2-17
dynamic work area 3-14

E

ERB2XDGS 5-2, 5-10, 5-25
ERB2XSMF 5-25
ERB3XDRS 5-25
ERBDSQRY 5-3, 5-25
ERBDSREC 5-8, 5-25
ERBSMFI 5-21, 5-25
exponential notation 3-26
Extractor
See CMF MONITOR Extractor
Extractor data set, preprocessing 4-1

F

field descriptions
Analyzer 3-27 to 3-28
font selection 8-3 to 8-4
format, PARM field 4-3 to 4-5

G

generating Analyzer JCL 3-2

H

headings, report 3-27
HMOVRESCAN= modify command 2-23

I

IBM’s RMF
compatibility with CMF MONITOR 1-9 to 1-10
differences 1-10
functions unavailable in RMF 1-10
similarities 1-9
id=apixa.statistical
considerations A-1
IFASMFDP 4-7
IMP mode 2-3
IPM mode 2-4
ISPF interface, generating Analyzer JCL 3-2

J

JCL
Analyzer control statements
CMFLOG DD 3-15
CMFRPTS EXEC 3-14
CMFSTAGE DD 3-17
CMFSTAGO DD 3-17
CMXREC DD 3-15

CMXTRACE DD 3-16
DMSSMAIN DD 3-14
EXTDATA DD 3-14
IPSLIB DD 3-15
OUTPUT 3-18
RPTCONTS DD 3-15
SNAPS DD 3-16
SNAPVBS DD 3-16
STEPLIB DD 3-14
SYSIN DD 3-15
SYSPRINT DD 3-16
SYSUDUMP DD 3-17
CMF MONITOR Analyzer, generating 3-2
CX10CVBS sample 4-2
defining for Analyzer 3-12 to 3-20
defining for Extractor 2-17
laser printer support 8-1 to 8-3

L

laser printer support
emphasized data 8-1
font selection 8-3 to 8-4
JCL 8-2 to 8-3
reports 8-2

M

MACROs 7-2
MAINVIEW 1-2, 1-4
MAINVIEW for OS/390
JCL generator panel, Analyzer 3-3
panel, CMF MONITOR Analyzer 3-2
use of Extractor data 1-9
mean A-1
members 3-11
menu for producing CMF Analyzer batch reports 3-3
Microsoft Excel spreadsheets, Analyzer reports 6-1
mode A-1
modify command, HMOVRESCAN= 2-23
monitoring modes 2-2
MVS product address space (PAS) 1-4

O

output 2-10

P

PARM field format, copying records 4-3 to 4-5
preprocessing Extractor data set 4-1
primary data sets 2-12
printer support
See laser printer support
private area storage 3-18
producing reports 3-1 to 3-26
product components 1-1 to 1-5
Analyzer 1-3
CMFMON 1-5
DSO 1-5
Extractor 1-3
Online 1-4

R

recording interval, customizing 2-11
records
archiving 2-14
collecting data 2-1
copying 4-2
types 2-7
writing 2-3, 2-10
CMF 2-12
DSO 2-12
SMF 2-11
region size 3-14, 3-18
REPORT 2-14
reports
Analyzer control statements 3-21
Analyzer JCL 3-12
collecting data 2-1
create your own 3-28, 7-1 to 7-4
data 3-24
DEVICE 3-25
headings 3-27
interpreting 3-24 to 3-26
miscellaneous 3-25
printing output to a laser printer 8-1
producing 3-28

using 3-1 to 3-28
Web-related 3-25
WORKLOAD 3-25
return codes, CX10CVBS 4-6
RMF
CMF functions unavailable in RMF 1-10
compatibility with CMF monitor 1-9
similarities with CMF 1-9

S

sampler operations 2-6
sampling
data 2-2
methods 2-5
SAS data set 7-3
scientific notation 3-26
spreadsheet converter 6-1 to 6-12
Analyzer reports 6-2
capturing
analyzer reports 6-2
CMFMON reports 6-3
customer support 6-13
hints 6-12
installing on PC 6-1
maintenance 6-13
output 6-6
reports, downloading to the PC 6-3
running 6-3
tutorial, CX98SSCX.XLA 6-7
workbook 6-5
SRB 2-15
SRM
trace facility 2-15
standard deviation
calculating A-3
mean A-1
mode A-1
statistical
accuracy A-3 to A-4
considerations A-3 to A-4
storage area configuration 3-18
system abend
S106 3-19
S80A 3-19
S878 3-19
SYSTEM RESOURCE Reports 3-25

T

TCB 2-6
The 2-3
These 3-25
trace facilities 2-15

U

user abend
U008 3-19
U0999 3-19
U100 3-19
using reports 3-1 to 3-28

W

Web-related
reports 3-25
WORKLOAD
reports 3-25
writing records 2-10
CMF 2-12
DSO 2-12
SMF 2-11
XDS buffers 2-12

X

XDS
accessing data 5-2
activating 5-2
API output, mapping an answer area 5-3
buffers 2-12
common answer area
header 5-15
sample DSECT 5-15
CX10XDGS API 5-10
CX10XDGS data reduction exit 5-13
CX10XDQY API 5-3
CX10XDRC API 5-8
return codes 5-17

Notes

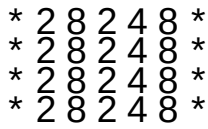

28248