Data Set Optimizer (DSO)

User Guide and Reference

Version 2.0

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

Figures

Tables

About This Book

This book contains detailed information about Data Set Optimizer (DSO) and is intended for storage administrators or other site personnel responsible for maintaining storage subsystem performance in an MVS environment. DSO produces a series of performance reports that show the optimal placement of data sets on the disk to reduce seek time.

To use this book, you should be familiar with Multiple Virtual Storage (MVS) systems, job control language (JCL), Interactive System Productivity Facility (ISPF) and the principles of storing data on DASD

How This Book Is Organized

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

Chapter/Appendix	Description
Chapter 1, “Overview of DSO,”	introduces DSO by giving you an overview of the product.
Chapter 2, “Collecting DSO Data,”	describes how to run the CMF Extractor, produce reports, and reorganize your volumes
Chapter 3, “Producing DSO Reports,”	
Chapter 4, “Evaluating DSO Reports,”	
Chapter 5, “CMF Extractor,”	describes CMF Extractor control statements
Chapter 6, “DSO Analyzer,”	describes DSO Analyzer control statements
Chapter 7, “DSO Reports,”	show examples of DSO reports
Chapter 8, “DSO Output JCL,”	show examples of output control statements
Appendix A, “Report Calculations,”	shows the variables and equations that are used to calculate the performance values that appear in DSO reports.

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
installation documents	*OS/390 and z/OS Installer Guide*	gives procedures to install and maintain MAINVIEW products
core document	*CMF MONITOR Batch User Guide*	describes how to use the CMF Extractor that collects performance data for DSO and other BMC Software products

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. The reader is provided at no cost, as explained in "To Access Online Books." You can also obtain additional printed books from BMC Software, as explained in "To Request Additional Printed Books."

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 Data Set Optimizer (DSO) product includes online Help. In the Data Set Optimizer (DSO) 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**.

Conventions

This section provides examples of the conventions used in this book, presents the control statements used by DSO, and explains how to read ISPF panel-flow diagrams and syntax statements.

General Conventions

This book uses the following general conventions:

BMC Software command processors gather and report information based upon user-defined control statements. A control statement has associated parameters that provide more detail about the operational characteristics of the command processor regulated by that control statement. This is a sample control statement:

[name] command parameter-list [- comment]

The following list explains general format conventions of all control statements:

- Up to four fields can appear in each control statement.
- The command field is required.
- The name and parameter-list fields can be optional, depending upon the particular control statement.
- The comment field is optional and does not affect the logical content of the control statement.

The following table explains the fields of a control statement

Field	Description
name	The first character of the name must start in column 1 and can be an alphabetic letter, a pound sign (#), an ampersand (@), or a dollar sign ($). The remaining characters can be alphabetic, numeric, #, @, or $. The name is one to eight characters long and is terminated with one or more blanks. Unless otherwise noted, the name field is treated as a comment.
command	The command field contains the name of a control statement. It is preceded and followed by one or more blanks. If the name field is omitted, the command field can begin anywhere except in column 1.
parameter(s)	A command may have no parameters, one parameter, or multiple parameters. The first parameter is separated from the command by one or more blanks. When a command has multiple parameters, they are listed consecutively and are separated by commas. There are two types of parameters:
keyword parameter	Consists of a value or parenthesized list of values preceded by a keyword= clause. Keyword parameters can appear anywhere in the list. The name of the keyword is defined within the command's parameter-list definition in either of two formats: keyword=value keyword=(value1,...,value*n*)
positional parameter	A value whose interpretation depends on its position in the parameter list with respect to other positional parameters. A positional parameter can be preceded by any number of keyword parameters, but it must retain its order in regard to any preceding positional parameters. If the default of a positional parameter is acceptable, or if customization is not required, a comma must be placed to mark the location of a positional parameter to maintain the proper order for subsequent positional parameters.
comment	A space-dash-space (outside of a quotation string in the parameter field) indicates that the rest of the statement is comment. An asterisk (*) in column 1 indicates that the entire statement is a comment.

This book uses the following types of special text:

Note: Notes contain important information that you should consider.

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

Tip: Tips contain useful information that may improve product performance or that may make procedures easier to follow.

Punctuation Conventions

The following punctuation conventions apply to all control statements:

- Nested parentheses are illegal.
- Fields can be coded in columns 1 through 71; columns 72 through 80 are ignored. Column 72 can contain a continuation character to maintain compatibility with Macro Assembler syntax.
- One or more blank spaces must separate the four different fields of a control statement.
- Two successive commas indicate an unspecified positional parameter and that the default is to be used.
- Commas between parameters in the parameter list are mandatory; if the parameter field of a previous statement ends in a comma, the next statement is treated as a continuation. Quotation strings are continued by omitting the closing quote. On the next statement, a single quote at or to the right of column 3 starts the continuation. A second single quote on a continuation statement terminates the string. Except for quotation strings, parameter fields can be continued from statement to statement only if the parameter list is broken after a comma. The continued statement must begin in column 3 or beyond.
- Inside a quote string, two successive quotes indicate the appearance of a single quote within the string. Such a double quote cannot be split between two statements (that is, the first of two successive quotes cannot be in column 71; such a quote would terminate the string). For example, the text: `'THIS IS A QUOTE(') EXAMPLE'`
  would be coded: `'''THIS IS A QUOTE('') EXAMPLE'''`

- A space-dash-space outside a quotation string indicates that the rest of the statement are comments. It is not possible to comment a statement containing a quote string that is to be continued on the next line. In the text below, '`1 SECOND SAMPLING`' is a comment.:

  ASMDATA SAMPLE=1000 - 1 SECOND SAMPLING

  An asterisk in column 1 indicates that the entire statement is a comment. An example of a comment statement on a single line is:

  * A COMMENT STATEMENT...

- If only one parameter of a parenthesized list of values is specified, the parentheses can be omitted. For example, DATA SYSTEM=(SYSA) is equivalent to DATA SYSTEM=SYSA.

- Conflicting or duplicate keyword usage of parameters within one command is not diagnosed. The specifications from the last appearance of the parameter are used.

- With the exception of the space-dash-space requirement for comments, JCL punctuation conventions can be used.

Chapter 1 Overview of DSO

DSO is an integrated software tool that enables you to improve the performance of your DASD subsystem. DSO analyzes the head move activity of your DASD volumes and recommends the most efficient placement of data sets. After your volumes have been reorganized according to the recommendations made by DSO, I/O operations can be substantially quicker because less time is needed to move read/write heads between data sets. This book explains how to use DSO to reduce the excessive head move activity characteristic of poorly organized DASD volumes.

DSO is offered as part of the BMC Software product CMF MONITOR, so you receive DSO when you purchase CMF MONITOR. These two products provide complementary functions for tuning your DASD subsystem. CMF MONITOR produces reports that show the average time spent in each phase of an I/O cycle for every volume that you study. You can easily identify volumes experiencing excessive head move times with these reports.

After the problem volumes are identified, you can use DSO to produce a series of reports that analyzes head move activity by individual data sets. Recommendations are made in a reorganization report that shows the potential reduction in head move times after your volumes have been reorganized.

As part of volume analysis, DSO also produces JCL and control statements that are used as an input job by DFDSS or COMPAKTOR. The input job is a plan to reorganize the volume based upon the recommendations made in the reorganization report.

Seek Time and the DASD I/O Cycle

Head moves are part of every DASD I/O that require the actuator to move to a different location on the disk to satisfy a request for a read or write. An actuator is a mechanical device that moves the read/write heads attached to it over the surface of a rotating disk. The actuator positions the heads over the precise cylinder location on the disk before a read or write operation begins. The time it takes to move heads during an I/O operation is called *seek* time.

Figure 1-1 Phases of a DASD I/O Cycle

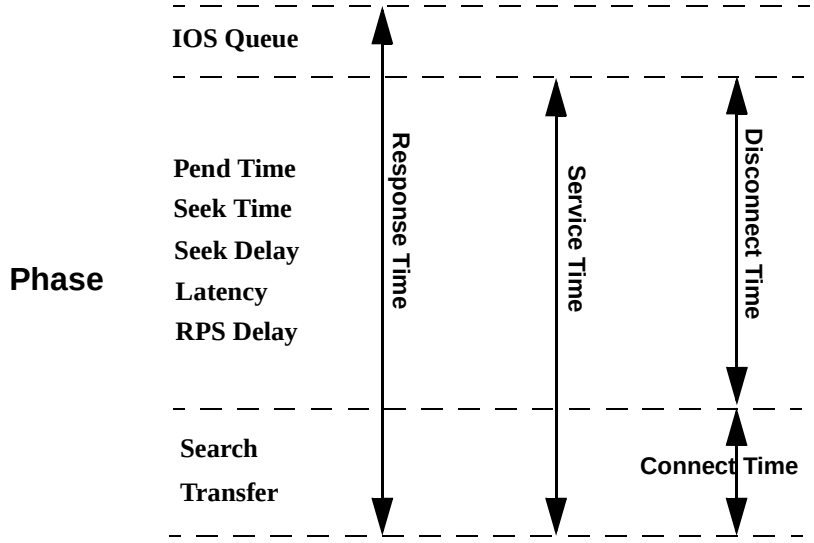

In recent years, improvements in MVS channel programs and hardware architecture have shortened each phase in the cycle. Storage resources are used much more efficiently and the technology is approaching an ideal environment that allows each I/O request to complete independently from another request. The characteristic delays caused by contention for storage resources are much less severe than in the past. RPS and seek delay are much smaller problems with the current generation of DASD than they were with older devices.

Seek time nevertheless remains a significant part of the I/O cycle. Despite faster actuators and increasing cylinder density, excessive seek time can become a serious problem, particularly if it reintroduces RPS and seek delay with every I/O request. The goal of your DASD tuning efforts is to keep the average seek time as small as possible to prevent it from causing other types of delay.

Seek time begins when the channel issues a SEEK command and ends when DASD responds with a SEEK COMPLETE signal to the controller. During a head seek, the actuator briefly accelerates and then rapidly decelerates as it approaches the destination cylinder. The seek completes by slowly moving the actuator to let the head settle over the target cylinder. The longest time spent in a typical seek is the period when the actuator creeps slowly to the correct cylinder.

Seek time is a variable length period in the I/O cycle. It is a function of the radial distance that the actuator travels across the disk surface. Modern DASD seek times are usually in the range of 8–20 milliseconds for an average length head move. You should examine the seek times of your volumes with CMF MONITOR to see if you have any that are out of range.

A lengthy seek time stems from two fundamental DASD data organization problems. The most common problem occurs when the volume frequently accesses data that is separated over large distances on the disk surface. As shown in Figure 1-2 on page 1-4, the actuator must continuously move back and forth between different cylinders for successive read/write operations. The second problem stems from the sheer number of required seeks to satisfy read/write requests. Almost all I/O requests require a head move instead of keeping related data on contiguous tracks of the same cylinder.

In either case, the essential problem is that the volume is poorly organized. Data must be relocated to keep the most active data sets as close together as possible to reduce the average seek time and the number of required seeks.

When to Use DSO

Figure 1-3 on page 1-5 shows where DSO can improve DASD performance. Figure 1-2 illustrates a poorly organized volume, where data sets A and B are active data sets and are responsible for most of the I/O activity on the volume. They are separated from each other by a relatively large cylinder distance. The actuator must travel back and forth between the two cylinders that hold the data sets to position the heads to read or write data for successive I/O operations. A large portion of each I/O operation must be devoted to seek time to move the actuator between the two cylinders. A volume is said to be *thrashing* when the actuator is in almost constant motion, seeking back and forth between cylinders.

Figure 1-2 Head Seeks between Data Sets A and B Before and After Reorganization

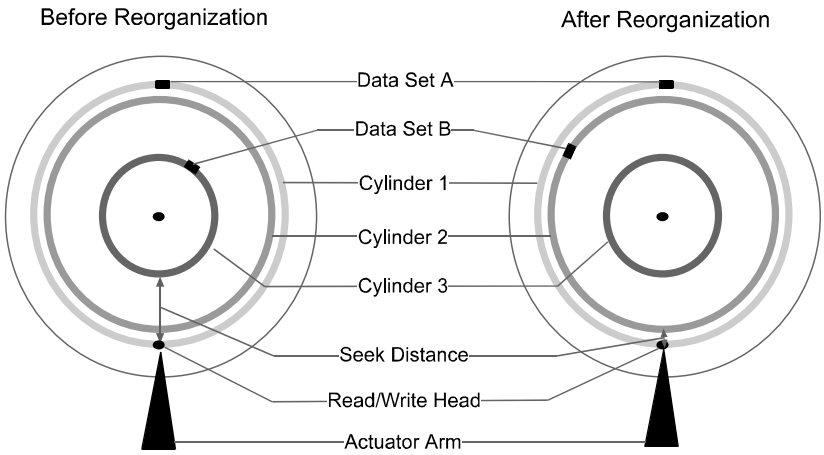

The drawing in Figure 1-2 shows the same volume after it has been reorganized according to the recommendations made by DSO. Data set B has been moved to a cylinder that is much closer to the cylinder holding data set A. The seek times for I/O operations are considerably shorter because the actuator does not have to travel as far to position the heads over the two data sets. In the most extreme case where the two data sets are separated by the maximum cylinder distance of the volume, over 20 milliseconds worth of seek time can be eliminated from each I/O by placing the data sets next to each other on contiguous cylinders.

An alternative strategy to improve the performance of a volume suffering from excessive seek times is to move one or more of the active data sets to another volume. The characteristic actuator thrashing between data sets can be eliminated or minimized. DSO produces two reports that show the number of head moves between data set pairs recorded during the sampling interval. These reports identify the data sets that are the best candidates for migrating to another volume if you do not want to go through the effort of completely reorganizing a volume.

How DSO Fits Together

Figure 1-3 shows the various components of DSO and how they fit together. The principal parts of DSO are the CMF Extractor and the DSO Analyzer. Most of your work with DSO involves setting up the Extractor to collect head move data and preparing the Analyzer to generate the information you need to reorganize your volumes.

Figure 1-3 Overview of DSO

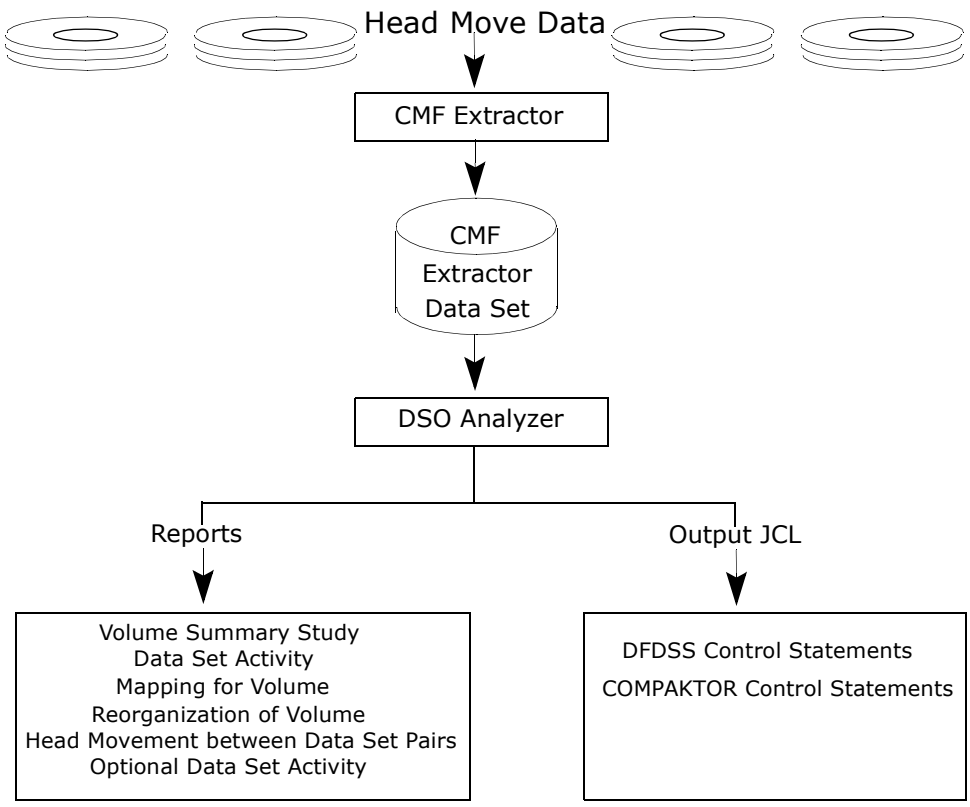

CMF Extractor

The CMF Extractor collects head move data and the VTOC map of each selected volume in your system over a defined interval. Additional information about the VSAM catalogs stored on your volumes can be collected by the Extractor. The Extractor writes this data to an SMF or CMF Extractor output data set.

DSO Analyzer

The DSO Analyzer processes the VTOC maps and head move data stored in the CMF Extractor data set to produce a series of reports. Nine different reports are available. Some reports are produced automatically; others are optional. You select optional reports by specifying parameters with the control statements that are included with the JCL to run the Analyzer.

DSO reports show head move activity by the following categories:

- Volume summary statistics

- Data set mapping
- Head moves between data set pairs
- Proposed volume reorganization
- Data set activity

At least one report for each category is produced automatically. Optional reports are available for the data set activity and head moves between data set pairs categories. The optional reports present similar information for each category, but list the data sets by a different value shown on the report. For example, four different data set activity reports list the data sets by name, busy time, error recovery, and the percentage of total head moves.

In addition to the reports, DSO can generate JCL that can be used with FDR COMPAKTOR or DFDSS to reorganize your volumes. With either product, the JCL is based upon the recommendations made in the Reorganization of Volume report. You select the type of volume reorganization you want by specifying parameters with your Analyzer control statements.

Getting Started with DSO

Performing a DSO study is a simple four-step job.

1. Collect head move data with the CMF Extractor

Run the Extractor to collect head move data on volumes that you suspect are suffering from excessive seek times.

2. Run the DSO Analyzer to generate reports and the volume reorganization job

Submit a batch job to produce DSO reports and the JCL to reorganize your volumes.

3. Analyze the DSO reports

Examine your DSO reports to see how you should reorganize your volumes to reduce DASD seek time.

4. Prepare the volume reorganization job

Edit the DSO JCL that reorganizes your volumes. Make any changes to suit the needs of your site.

Chapter 2 Collecting DSO Data

This chapter discusses collecting head move data, which is the first major task in reorganizing the volumes in your system. It shows you how to set up the CMF Extractor to collect all of the performance data needed by the DSO Analyzer. “Setting Up the Extractor JCL” on page 2-4, and “Running the Extractor” on page 2-10, describes the JCL needed to run the Extractor and some considerations you should keep in mind when collecting DSO data. “Altering Extractor Performance with the MODIFY Command” on page 2-13, and “Viewing Extractor Records” on page 2-15, shows you how to change an Extractor job with the MVS MODIFY command and look at the records in the Extractor output data set using PERUSE.

Figure 2-1 on page 2-2 gives you an overview of the entire process of collecting DSO data. Extractor samplers monitor the activity of your DASD volumes and collect data about head moves and the physical organization of information stored on the disks. The Extractor converts the head move counts and volume maps. DSO analyzes these records and creates reports that are a statistical representation of volume seek activity during the measurement interval.

Figure 2-1 Extractor Records Collected for DSO

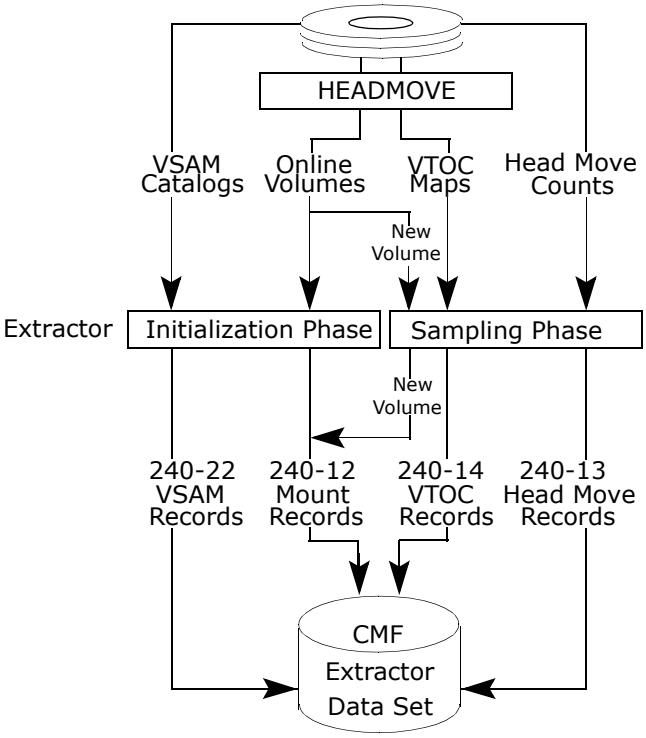

The Extractor samplers that collect seek data work under the control of the Extractor HEADMOVE control statement. You include parameters with the HEADMOVE control statement to select the information you want collected by the samplers. In “Setting Up the Extractor Control Statements” on page 2-8, the HEADMOVE control statement is shown with the recommended parameters to collect DSO data. “HEADMOVE” on page 5-6, gives you a complete description of all of the HEADMOVE parameters.

Figure 2-1 on page 2-2 shows the four HEADMOVE records that the Extractor creates during the Extractor initialization and sampling phases. In the initialization phase, the Extractor is collecting information that describes the architecture of all the components in the system. The Extractor creates VSAM maps and mount records for every selected volume that is online. In addition, the Extractor builds a table that lists all of the selected online volumes identified during the initialization phase.

The sampling phase monitors ongoing system activity of the selected volumes. The Extractor generates head move records and VTOC maps for the listed volumes.

The Extractor creates all of the 240-14 VTOC records before head move sampling begins, except when you place a volume online during the sampling phase. In this case, the Extractor includes the volume and creates 240-12 mount and 240-14 VTOC records. The Extractor generates the VTOC map of the new volume and continues head move sampling for the remainder of the sampling phase.

240-12 Mount Record	A mount record contains the device address, device type, and volume serial number. The Extractor generates a mount record for every selected online volume during the initialization phase and when a new volume is encountered during a sampling period.
240-13 Head Move Record	The Extractor writes head move records to the Extractor data set when it observes a change in seek address or device status on a sampled volume. Each record contains the device address, seek address, device status, and sample count.
240-14 VTOC Record	The Extractor writes VTOC maps to the Extractor data set during the Extractor sampling phase. It creates the 240-14 record for every volume identified with a 240-12 mount record and whenever you bring a new volume online during the sampling phase. You must set the VTOC=YES parameter with the HEADMOVE control statement to collect VTOC records.
240-22 VSAM Record	VSAM records contain cluster information from VSAM master and user catalogs. The Extractor creates VSAM records during the initialization phase. VSAM records are optional with DSO. You must set the VSAM=YES parameter with the HEADMOVE control statement to collect these records.

DSO requires three of the four HEADMOVE records for every report. The 240-22 record is the optional fourth record that maps VSAM clusters within a VSAM data space. The Extractor builds this record during the initialization phase. Because of the complexity of extracting the VSAM catalog structure from a large number of volumes, the Extractor can take an excessive amount of time to create 240-22 records.

For example, it takes an excessive amount of time to create 240-22 records when you mount volumes that no longer exist but have entries in the catalog. In this case, the Extractor makes an attempt to mount non-existent volumes because the volumes have a catalog entry.

Problems can occur with the Extractor if it must scan every VSAM/ICF catalog in the system to build VSAM records. “Collecting VSAM Data” on page 2-11 describes an optional job that you can use to collect VSAM records separately from the other Extractor records.

In addition to the four HEADMOVE records, DSO uses two other records generated by the REPORT control statement. Both records describe the active environment when the Extractor is collecting data. The 240-00 record lists Extractor operating conditions and the 240-11 record describes the configuration and status of the system during the sampling interval.

240-00 SRM Record	The Extractor writes a system resources manager (SRM) record once every recording interval. The 240-00 record describes the active SRM constants, the installation performance specification (IPS) settings, Extractor control statements, and sampling rates.
240-11 Global Bit Map Record	The Extractor writes a global bit map record as needed and at least once every recording interval. The 240-11 record describes the configuration and status of the system CPU, the executing program, and a channel bit map.

Setting Up the Extractor JCL

Figure 2-2 on page 2-5 shows a complex job used to collect data when the Extractor runs with a coordinating address space (CAS) and product address space (PAS), which are necessary for CMF Online in a MAINVIEW environment.

The example shown in the figures is a generic job. The purpose of the generic job is to show the typical statements normally included to run the Extractor to collect DSO data. Your site may use different JCL, depending upon your needs and the release of the Extractor running at your site. Contact your system administrator about any local requirements before running the Extractor job.

A description of the JCL is in Table 2-1 on page 2-6. Carefully review and edit each statement to meet the requirements of your site:

- Change ?HILEVEL to the *high.level* qualifier you chose for the DSO libraries.
- Change ?BBASMFID to the SMF system ID of the target system.
- Delete the STEPLIB statement if BBLINK is in the link-list concatenation. Otherwise, change the ?BBLINK value to the name assigned to BBLINK.
- Change all ?BBSYSOUT to an acceptable SYSOUT class.
- Change all ?UBBPARM to the name you assigned to the user parameter library.
- Change ?SSID to the name of the CAS that is connected to this product address space.
- Delete the historical data set DD statements if they have not been allocated.

Figure 2-2 JCL to Collect DSO Data with CMF MONITOR Extractor

```
//****  SAMPLE JCL TO COLLECT DSO DATA WITH CMF MONITOR EXTRACTOR
//****
//****  NOTICE: SUPPLY INSTALLATION/USER VALUES FOR ALL PARAMETERS MARKED:
//****
//****                                  <===    OR  ?????
//****
//CMFPAS  PROC DBC=N,XDM=N,DC=START,
//             IPM=D2,
//             SSID=?SSID,
//             RGN=4M
//*
//CMFPAS  EXEC PGM=BBM9DA00,
//             PARM=('SSID=&SSID',
//             'DBC=&DBC,XDM=&XDM,DC=&DC',
//             'IPM=&IPM'),
               TIME=1440,
//             REGION=&RGN.M
//*
//STEPLIB  DD  DISP=SHR,DSN=?BBLINK
//*
//CMFIPM1  DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.IPMOUT1
//*CMFIPM2  DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.IPMOUT2
//*CMFCPM1  DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.CPMOUT1
//*CMFCPM2  DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.CPMOUT2
//CMFDUMP  DD  SYSOUT=?BBSYSOUT
//CMFLOG   DD  SYSOUT=?BBSYSOUT
//CMFMSG   DD  SYSOUT=?BBSYSOUT
//SYSUDUMP DD  SYSOUT=?BBSYSOUT
//SYSABEND DD  SYSOUT=?BBSYSOUT
//*
//* EXTRACTOR OPERATING PARAMETERS.
//*
//PARMLIB  DD  DISP=SHR,DSN=?UBBPARM
//         DD  DISP=SHR,DSN=?BBCHILV.BBPARM
//*
//* VIEW DEFINITIONS.
//*
//BBVDEF   DD  DISP=SHR,DSN=?HILEVEL.BBVDEF
//*
//* ACTION DEFINITION TABLES.
//*
//BBACTDEF DD  DISP=SHR,DSN=?HILEVEL.BBACTDEF
//*
//* HISTORICAL PROCESSING AND WORKLOAD CONFIGURATION FILES.
//*
//PARMFILE DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.PARMFILE
//WKLDFILE DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.WKLDFILE
//*
//* HISTORY DATA SETS.
//*
```

```
//* - THESE DATA SETS DEFINITIONS MAY BE DELETED IF YOU HAVE NOT
//*   DEFINED THEM.
//*
//HISTDS00 DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.HISTDS00
//HISTDS01 DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.HISTDS01
//HISTDS02 DD  DISP=SHR,DSN=?HILEVEL.?BBASMFID.HISTDS02
//*
```

Table 2-1 Extractor JCL Statement Descriptions

JCL statement	Description
PROC	Identifies the procedure name to start the CMF Extractor product address space. CMFPAS is the default procedure name. XDMIndicates whether the CAS should execute in extended diagnostic mode to generate additional error messages. The default is XDM=N. DCSpecifies whether CMF MONITOR Online data collectors are started for this Extractor job. The default is DC=START. The DC=START JCL statement requires an active MAINVIEW CAS. IPMSpecifies a two-character suffix that identifies the member that holds the control statements used by the Extractor. The member that holds the default DSO control statements is CMFIPMD2. SSIDSpecifies the subsystem ID used to identify the MAINVIEW CAS that is connected to the Extractor PAS if DC=START. RGNDeclares the region size of the Extractor dynamic work area in kilobytes.
EXEC	Specifies the program name and assigns a region size of 1 megabyte. CMFEXT is the program to run the Extractor when a PAS and CAS are not necessary. BBM9DA00 is the PAS driver program to run the Extractor within PAS and with a CAS. The Extractor dynamically acquires and releases storage. The recommended 1024K region size should not adversely affect paging or real storage usage. If you specify 1024K with the REGION= parameter, it provides enough virtual storage to reduce the incidence of abends caused by insufficient virtual storage. Set the TIME= parameter to be large enough to prevent the possibility of a timeout abend. The recommended time is 1440, which indicates no time limit. You can hold Extractor control statements that collect data for DSO in members you specify with the PARM= field. Separate parameter keywords, CPM= and IPM=, identify the members that contain the appropriate control statements for the sampling mode. The control statement members must be named CMF*abcxx*, where abcSpecifies CPM or IPM parameter keywords xxSpecifies the two-character suffix referred to by the CPM or IPM keywords in the PARM= parameter. If you do not specify the PARM field, the Extractor uses the control statements from the default member CPM=00. The control statements held in a member of the PARMLIB are overridden by any instream control statements specified with the SYSINIPM or SYSINCPM DD statements.
STEPLIB DD	Defines a PDS that contains the CMF Extractor load modules. The PDS must be an authorized library. (BBLINK)
PARMLIB DD	Defines a PDS that holds the control statement members that regulate the operational characteristics of the Extractor. The PDS must be a fixed-block data set with a logical record length of 80. (UBBPARM)

Table 2-1 Extractor JCL Statement Descriptions (continued)

JCL statement	Description
SYSUDUMP DD	Specifies a dump listing of all of the virtual storage allocated to the user region of the Extractor private address space in the event of failure. As a rule, output is not generated for this data set because abends are trapped by ESTAEs and output is directed to the CMFDUMP data set.
CMFDUMP DD	Defines a print data set for snapshot dumps. When a program fails in DSO, the abend is intercepted by an ESTAE routine and a formatted dump of the failing component is printed.
CMFLOG DD	Defines an optional print data set containing initialization messages for the CMF Control Card Log. The CMFLOG data set allows messages to be inspected immediately, regardless of the length of the CMF Extractor job run. If this DD statement is not specified, the CMF Control Card Log is written to //SYSPRINT. In most cases, a DD statement for the CMFLOG data set should not be included in the JCL. The CMFLOG data set is dynamically allocated and deallocated as needed when no //CMFLOG DD statement is present. Dynamic allocation is to the MSGCLASS= specified on the JOB statement for batch jobs or CLASS A for started tasks.
CMFMSG DD	Defines a print data set containing noninitialization messages for the system. In most cases, a DD statement for this data set is not necessary. When no DD statement is present, the CMFMSG data set is dynamically allocated to SYSOUT class A. The data set is printed when you issue the MSGFREE command with the MVS MODIFY command. The data set is dynamically deallocated and printed, and another CMFMSG data set is immediately allocated.
CMFIPMxx DD or CMFCPMxx DD	Specifies a maximum of 101 DD statements that define output data sets for holding CMF Extractor data from either IPM or CPM sampling modes. The CMFIPM1 and CMFIPM2 (IPM mode) and CMFCPM1 and CMFCPM2 (CPM mode) DD names in the sample JCL illustrate the definition of four output data sets. You can define CMFIPM*xx* and CMFCPM*xx* DD statements to have the Extractor write simultaneously to the same data set, where *xx* contains any one or two legal characters (alpha, numeric, or a combination of both) you specify for 1–101 DD names. Data is written to data sets in the order in which you specify them. When the last data set is full, the Extractor attempts to locate enough space in the output data sets to write more data. The writing of output is suspended when no more space is found. You can clear a data set by executing the CMFOCLR procedure in batch mode or at the console. For more information, see the *CMF MONITOR Customization Guide.* There are other ways to specify IPM and CPM output data sets: Omit the CMF*x*PM*xx* DD names and specify the data set names with the REPORT control statement, included as DD * data in the JCL, or specify the output data sets in a member of the PARMLIB data set. The data sets are dynamically allocated. Code SMF=YES on the REPORT control statement, included as DD * data in the JCL, or specify the output data sets in a member of the PARMLIB data set. Do not specify data set names or DD names. All output is written to the SMF data set. The output data set has a physical sequential (PS) organization with a variable block spanned (VBS) record format. The record length is 32762. The block size is variable, depending upon the type of DASD storing the data set.
SYSINIPM DD or SYSINCPM DD	Defines optional data sets that contain control statements to run the CMF Extractor in either IPM or CPM sampling modes. If these DD statements are in the JCL, they override IPM or CPM specified with the PARM= option on the EXEC statement. These DD statements can point to sequential data sets, the control statements can be included as DD * data in the JCL, or you can specify the output data sets in a member of the PARMLIB data set. Once the monitor has started, any modifications to DSO from the console to restart CPM or IPM are obtained from the PARMLIB data set.

Setting Up the Extractor Control Statements

Figure 2-3 shows the contents of a member named CMFIPMD2 from your CMF MONITOR BBPARM library. The member contains a sample of the two control statements required to run the Extractor to collect DSO data. REPORT and HEADMOVE are the only control statements necessary to sample DASD volumes and create the records required for DSO. The recommended parameters for both control statements are also shown in Figure 2-3.

This section is intended to get you started quickly collecting the records you need for the DSO Analyzer. It briefly explains setting up the two control statements and their parameters. After you have learned the basics of collecting DSO data, refer to Chapter 5, “CMF Extractor,”. It contains a complete description of both control statements and the other parameters you can include with them.

The REPORT control statement must always precede the HEADMOVE control statement whether you keep them together in a separate member or place them as part of an instream JCL job to run the Extractor. Figure 2-3 shows the order in which both control statements are coded. Both the REPORT and HEADMOVE control parameters are described in Table 2-2 on page 2-9 and Table 2-3 on page 2-10.

The parameters included with both control statements are nominal values and are appropriate for most sites. The DSO data is written to the data sets specified in the Extractor JCL job.

Figure 2-3 Sample Extractor Control Statements to Collect DSO Data

```
*--------------------------------------------------------------------*
*  SAMPLE CMF Extractor CONTROL STATEMENTS FOR USE WITH DATA SET     *
*  OPTIMIZER.  THE DATA IS WRITTEN TO THE DATA SET SPECIFIED ON      *
*  THE CMFxPMxx OR CMFxDSxx DD STATEMENTS IN THE EXTRACTOR JCL JOB.  *
*                                                                    *
*  USE THE 'DSN=' OR 'DSNLIST=' PARAMETER TO WRITE DATA TO           *
*  DATA SETS OTHER THAN THE DATA SET SPECIFIED ON THE THESE          *
*  DD STATEMENTS.                                                    *
*                                                                    *
*  YOU ALSO CAN CHANGE RANGE NUMBERS (DEVICE ADDRESSES).             *
*  MEMBER CMFANLD1 HAS SAMPLE JCL FOR RUNNING THE DSO ANALYZER.      *
*--------------------------------------------------------------------*
 REPORT IPM,INTERVAL=QTR,CSA=180,SMFRECID=240,RUNTIME=60,DISP=NEW
 HEADMOVE ALL,SAMPLE=33,VTOC=YES
```

See “Examples” on page 5-16 for more REPORT control statement examples. You also can use the DSN= and DSNLIST= parameters to define output data sets.

REPORT Extractor Control Statement and Parameters

The REPORT Extractor control statement sets global conditions under which the Extractor operates when it collects data. The length of the sampling interval, the format of the Extractor records, and the sampling mode are some of the conditions set with the REPORT control statement. Refer to “REPORT” on page 5-12 for more details.

Table 2-2 provides a list and description of the REPORT Extractor control parameters.

Table 2-2 REPORT Extractor Control Parameters and Description

REPORT Extractor Parameters	Description
IPM	The positional keyword that declares that all Extractor control statements and their parameters are intended for the IPM sampling mode. The Extractor should be run in IPM sampling mode to collect DSO data. DASD head moves are rapid events that are accurately sampled during the short periods of intense measurement that the IPM mode provides.
INTERVAL=QTR	A large number of head move samples accumulates quickly in most DASD subsystems. The INTERVAL parameter sets the maximum length of time between recording intervals in minutes. The recording interval is the period when the Extractor gathers data collected by the samplers and writes it to a data set(s). A 15-minute data collection interval is the recommended parameter value for DSO.
CSA=180	A common storage area of 180 kilobytes (K) is large enough for most DSO Extractor jobs.
SMFRECID=240	The Extractor writes records to the data set in SMF format. The 240 record type is the default SMF record number used by the Extractor.
RUNTIME=60	A 60-minute period is the recommended length to collect data with the IPM sampling mode.
DISP=NEW	New Extractor sampling records are written into the data set from the beginning, rather than concatenating them to existing records.

HEADMOVE Extractor Control Statement and Parameters

The HEADMOVE Extractor control statement sets the conditions for collecting data from DASD volumes. The four Extractor records used by the Analyzer to create DSO reports are defined by this control statement. Refer to “HEADMOVE” on page 5-6 for more details.

Table 2-3 provides a list and description of the HEADMOVE Extractor control statements.

Table 2-3 HEADMOVE Extractor Control Parameters and Description

HEADMOVE Extractor Parameters	Description
ALL	ALL is a positional keyword parameter that specifies that every volume in your system is sampled for head move activity **Caution:** Collecting head move data from every online volume in your DASD subsystem may impose a severe performance burden on your system by the Extractor. Head move data from all of the volumes in your system should be collected only if your site has fewer than 100 volumes. Otherwise, restrict Extractor sampling to the volumes identified by CMF MONITOR as suffering from excessive seek times. Chapter 5, “CMF Extractor,”, shows you how to use several other HEADMOVE parameters to declare specific volumes for sampling.
SAMPLE=33	DASD head moves occur quickly and often. Head move data needs to be collected over relatively short intervals. In this case, the data is collected every 33 milliseconds during the sampling period. The default 33-millisecond interval is the recommended value for collecting DSO data.
VTOC=YES	VTOC maps are required from every sampled volume. VTOC=YES must always be included with the HEADMOVE control statement to create the 240-14 records.

Running the Extractor

One of the first decisions you must make when you run the Extractor is which volumes to study. CMF MONITOR produces reports that show the average seek time of DASD volumes. These reports are a good place to find the volumes that are experiencing excessive seek time. The volumes identified in the reports are usually the best candidates to collect Extractor records for DSO.

Characteristically, the best candidate volumes to collect data are those that

- Hold several active data sets that are separated from each other by large cylinder distances
- Contain unrelated data sets that are accessed frequently by concurrent applications

In either case, DASD performance is in jeopardy because of the potential performance penalties that are paid when poorly organized volumes cause excessive seek times. The volume cannot reconnect to complete an I/O operation because seek times become inordinately large, effectively disconnecting the volume for longer periods. Typically, as the number of I/Os increase with the workload, excessive seek time causes RPS and seek delay. Your best strategy is to focus your extractions on these volumes.

As an alternative, it is possible to make an extraction on every volume in your system by selecting the ALL parameter of the HEADMOVE control statement. This strategy is recommended if your DASD subsystem has less than 100 volumes and you are willing to pay the cost of processing a large number of Extractor records. The *shotgun* approach of collecting records from every volume with one Extractor run allows you to get a consistent time-slice of head move activity. It is a good way to assess I/O activity across your entire system to find where performance bottlenecks are occurring.

Note: Be careful when you collect records from a large number of volumes. Scanning the VTOC of every volume can severely impact the performance of your system.

DASD head moves are short, transient events. Many head moves can be missed if the sampling interval is much larger than the seek time. If the Extractor collects head move data over long intervals, the validity of the data may be compromised because a significant number of short head moves were never recorded.

Accurate head move sampling occurs when the sampling interval approximates the average seek time. For this reason, the Extractor should collect head move records at the recommended 33-millisecond interval. If the Extractor overhead becomes excessive because of the number of data collection cycles, reduce the number of sampled volumes.

In summary, consider the following recommendations when you collect Extractor data for DSO:

- Focus your extractions on the volumes that you know are suffering from excessive seek time if the cost of running the Extractor is a consideration at your site.
- Run the Extractor in the IPM sampling mode to collect DSO data.
- Run the Extractor with the shortest sampling interval.
- Reduce the number of sampled volumes rather than lengthen the sampling interval if Extractor performance becomes a consideration.

Collecting VSAM Data

When you code the VSAM=YES parameter with the HEADMOVE control statement, the Extractor obtains cluster allocation information for VSAM clusters that are defined within a VSAM data space. A secondary reason for coding the VSAM=YES parameter is to provide the cataloged name of the cluster for the component name.

Since the advent of DF/EF VSAM, data sets with multiple clusters no longer exist. Coding VSAM=YES provides little benefit to DF/EF VSAM users. You should carefully weigh the benefits of collecting VSAM records against the cost of several common problems that can occur with VSAM sampling. The following list describes some of the problems that can occur with the Extractor when VSAM records are collected:

- Collecting VSAM data can degrade system performance during the Extractor initialization phase when VSAM records are created. The initialization phase is prolonged because each VSAM or ICF catalog from every sampled volume is scanned for the cataloged VSAM clusters. The delay is a function of the number of volumes and catalogs of VSAM data sets. In large systems, the delay can be quite long.

- CSA shortages are another characteristic problem of collecting VSAM records. Because of the number of catalogs that needs to be scanned, the amount of CSA needed by the Extractor can exceed the amount specified by the REPORT control statement. The Extractor abends with a CSA shortage failure during the initialization phase.

- Another common problem appears when VSAM catalogs are not kept current. Data sets are cataloged on volumes that no longer exist, resulting in catalog management requests for a mount of a nonexistent volume. Operators must manually cancel the mount request in response to the IEF244I message.

Potential problems can be minimized by collecting VSAM records separately with a different Extractor job. The performance burden and delays caused by VSAM records can be lessened if they are not collected at the same time as other records.

With this strategy, two Extractor jobs collect all of the records required by DSO and store them in separate data sets. The first job collects all DSO records except VSAM records by removing the VSAM=YES parameter from the HEADMOVE control statement in the Extractor JCL. The second Extractor job collects only the VSAM records.

Figure 2-4 on page 2-13 shows a sample batch job held in the DSOVSAM member of your CMF MONITOR BBPARM library. The program CX10DDRV invoked by the job is designed to collect only VSAM records. This job should be run separately during offshift hours to extract VSAM data set mapping information for all of the data sets on the volumes specified by the VSAM SYSIN DD statement.

The VSAMOUT data set can be concatenated in front of the other Extractor output data set. Together, both data sets are used as input to the DSO Analyzer.

Figure 2-4 Sample JCL to Collect VSAM Records

```
//JOBCARD JOB
//*-------------------------------------------------------------------
//* SAMPLE JCL FOR EXTRACTING VSAM CATALOG INFORMATION FOR DSO.
//*
//*   REVIEW THE JCL FOR APPLICABILITY TO YOUR INSTALLATION
//*   STANDARDS AND CHANGE ?BBCHILV TO THE HIGH-LEVEL QUALIFIER
//*   YOU CHOSE FOR THE CMF MONITOR LIBRARIES.
//*
//*   THIS JOB EXTRACTS VSAM DATA SET MAPPING INFORMATION
//*   FOR ALL OF THE DATA SETS ON THE VOLUMES SPECIFIED IN THE
//*   VSAMIN SYSIN STREAM.
//*
//*   THE PARM FIELD SPECIFIES THE RECORD ID AND THE MONITOR
//*   ID TO BE FILLED IN THE RECORD HEADERS.  THE OUTPUT DATA SET
//*   (VSAMOUT) WILL BE CONCATENATED IN FRONT OF THE DATA SET
//*   THAT THE DSO SAMPLER, HEADMOVE, WROTE ITS RECORDS TO.  BOTH
//*   DATA SETS MAY THEN BE USED AS INPUT TO THE DSO ANALYZER.
//*-------------------------------------------------------------------
//DDRV     EXEC PGM=CX10DDRV,REGION=1024K,PARM='240,CPM'
//STEPLIB  DD  DISP=SHR,DSN=?BBCHILV.BBLINK
//VSAMOUT  DD  DISP=OLD,DSN=?BBCHILV.VSAMDATA
//SYSUDUMP DD  SYSOUT=*
//VSAMIN   DD  *
VOL001
VOL002
VOL003
```

Review and edit the JCL to make it suitable for your site standards. Change ?BBCHILV to the high-level qualifier chosen for your CMF MONITOR libraries.

Altering Extractor Performance with the MODIFY Command

You can change the operating characteristics of the Extractor at any time with the MVS MODIFY (F) command. The MODIFY command allows you to dynamically change some aspect of the Extractor as it is running. You do not need to stop the Extractor to make changes.

The following diagram shows the syntax of the MODIFY command and several parameters that alter some aspect of CMF Extractor performance. Each MODIFY parameter is described in more detail following the command syntax. Some examples of the MODIFY command are shown on page 2-15.

F jobname
[,MSGFREE]
[,CPM={xx|STOP}]
[,IPM={xx|STOP}]
[,STATUS]
[,FLIP={IPM|ALTCPM}]
[,PROFILE]

jobname — Name of the CMF Extractor job altered by the MODIFY command.

MSGFREE — Requests that the CMF Extractor CMFMSG message file be dynamically deallocated and released for printing. The message file is reallocated immediately after deallocation with no loss of data.

CPM=, IPM= — Specifies that CPM or IPM sampling is stopped or executed under a different set of Extractor control statements.

- xx — Runs the Extractor using the control statements held in the CMFCPM*xx* or CMFIPM*xx* members of the PARMLIB library. The current sampling job is terminated before the new samplers start collecting data with the control statements requested by the MODIFY command.
- STOP — Terminates CMF Extractor CPM or IPM sampling. The CMF Extractor job is stopped if both CPM=STOP and IPM=STOP are specified with the MODIFY command.

STATUS — Displays status information about the Extractor sampling mode on the console.

FLIP= — Writes Extractor data to the next available CPM or IPM data set. The MODIFY command is rejected if you have not defined an alternate data set or if an alternate is not available.

If the MODIFY command is rejected, you may see error message CMFOWR10 or CMFOWR34.

Note: When flipping to a VSAM or BSAM data set, the output data sets must be empty.

PROFILE — Displays the system profile panel on the console.

Examples of MODIFY Commands

The following examples show some typical changes you can make to an Extractor job using the MODIFY command:

```
F CMF,IPM=04
```

IPM sampling is started using the control statements held in the CMFIPM04 member of the PARMLIB library.

```
F CMF,MSGFREE
```

CMF MSG data set is freed for printing and reallocated.

```
F CMF,IPM=XY
```

Invoke IPM sampling by using the control statements held in the CMFIPMXY member of the PARMLIB library.

```
F CMF,STATUS
```

Displays the CMF Status Display panel on the console.

```
F CMF,FLIP=CPM
```

Extractor writes data collected under CPM sampling mode to the next available output data set.

```
F CMF,PROFILE
```

CMF System Profile panel is displayed on the console.

CMF Extractor also can be terminated by issuing the MVS STOP (P) command followed by the CMF Extractor jobname.

```
P jobname
```

Viewing Extractor Records

You can examine your SMF or Extractor data sets to view the DSO records collected by the samplers. The Extractor has a PERUSE option that searches, displays, and prints the sampler records held in your data set. Looking at the data set is particularly useful to confirm the presence of all of the required Extractor records before processing them with the DSO Analyzer to produce reports.

You can use two different methods to select the PERUSE option. The method you choose depends upon whether you have the Extractor running under the PERFORMANCE MANAGER or MAINVIEW.

To select PERUSE in the PERFORMANCE MANAGER environment:

1. From the MVS PERFORMANCE MANAGER panel, select Option C, **CONTROL**. The Start or Stop CS Monitor and CMF Extractor panel is displayed.

2. Select the Extractor by choosing Option 2. The CMF Extractor panel is displayed.

3. Select Option 6, **PERUSE.**

To select PERUSE in the MAINVIEW environment:

1. From the MAINVIEW Selection Menu panel, select Option Z, **OS/390, z/OS, and USS**. The OS/390, z/OS, and USS panel is displayed.

2. Select Option P, **PERUSE**. The Peruse and Print CMF Records panel is displayed and shown in Figure 2-5.

Fill out the panel according to the descriptions of the panel fields explained in Table 2-4 on page 2-17.

Figure 2-5 Peruse and Print CMF Records Panel

```
---------------------- PERUSE AND PRINT CMF RECORDS ----------------------------
COMMAND ===>

Peruse current CMF (CPM or IPM) or SMF data set:
 DATA SET TYPE  ===> I       ( C - CPM, I - IPM, S -SMF data set)
                             ( NO or blank to peruse data set specified below)
 EXTRACTOR SEQ  ===> 1       (Change only if multiple Extractors are active;
                               use STATUS(S) to determine the sequence number)
 DISPLAY FORMAT ===> New   (OLD or NEW)

Peruse other CMF or SMF data set:
 DATA SET NAME  ===> 'HIGH.LEVEL.EXTRACT'
 VOLUME SERIAL  ===>  (Required if data set not cataloged)
 DASD UNIT NAME ===>  (Required if not cataloged and not default unit)

Set parameters for searching and printing:
 SEARCH LIMIT   ===> 500 (Maximum records to search per request
 SYSOUT CLASS   ===> A      (For print option)

To search for, peruse, and optionally print a record from a CMF or SMF data set,
first enter the necessary information above. Another panel will then appear; on
it you can view the records and enter line commands to search for and print
records.
```

Table 2-4 Peruse and Print CMF Records Panel Field Descriptions

Field	Description
DATA SET TYPE	Data set record type based upon the Extractor sampling mode. Leaving the field blank is the same as specifying NO. C Searches the active CPM data set, if CMF CPM mode is active. I Searches the active IPM data set, if CMF IPM mode is active. S Searches the active SMF data set, SYS1.MANx, where x is the designator for the currently active SMF data set. NO Searches the data set specified in the DATA SET NAME field.
DISPLAY FORMAT	Allows you to display records in either the old or the new format. The new format contains an additional command, F, which allows you to find a triplet, EBCDIC, or hexadecimal string in the record.
EXTRACTOR SEQ	When multiple Extractors are running in your system, you can define the sequence number assigned to the Extractor that gathered the records you want to view. If you do not know the sequence number of the Extractor, use the STATUS option to determine it. The sequence number of the Extractor coincides with the order in which its information is displayed on the STATUS panel.
DATA SET NAME	Name of the data set containing CMF, SMF, or RMF records. This field must be completed if NO or a blank is specified in the DATA SET TYPE field. The data set name is prefixed with your TSO user ID if you omit the enclosing apostrophes around the data set name.
VOLUME SERIAL	One- to six-alphanumeric identifier of the volume where the data set resides. This field is required if the data set is not catalogued.
DASD UNIT NAME	Generic four-digit DASD type that identifies the class of DASD device that stores the data set. The default DASD unit name from your TSO user ID profile is used if the data set is not catalogued and the VOLUME SERIAL and DASD UNIT NAME fields are left blank.
SEARCH LIMIT	Maximum number of data set records searched per request. The default value is 500. The maximum number of records searched per request is 5000.
SYSOUT CLASS	JES output class used to print a SNAP dump image of a data set record. The default SYSOUT class is A.

Press Enter after completing the fields of the Peruse and Print CMF Records panel. The first record in the data set is displayed when you invoke a PERUSE session.

Figure 2-6 on page 2-18 shows an example of a typical PERUSE output panel. A header area is displayed at the top of the PERUSE output panel for any CMF records and for RMF type 70 through 78 records. The record is displayed beneath a title bar that separates it from the header area. A 240-12 volume mount record is shown in Figure 2-6. The columns and record fields are described in Table 2-5 on page 2-18.

Figure 2-6 PERUSE Output Panel

```
PERUSE:  'TSG.SYSB.IPMDATA' ----------------------------  ROW 1 TO 15 OF 796
COMMAND ===>                                              SCROLL ===> PAGE
System  - SYSB   MVS     - ZV010100  Monitor  - IPM       Format   - ESA4PRSM
Record  - 240-12 Samples - 00002018  Cycle    - 0000250F  Duration - 0828905F
Release - 421F   Compat  - 5F        Date     - 01192F    Start    - 0135131F
Product - 0034   1st: 0068,005E,0001 2nd: 00C6,007C,0065  3rd: ,,

------ Index -- HeadMove volume mount data (DAV) -- -- *---- EBCDIC ----* ----
00000    11A80000 CFF00030 73430092 256FE2E8    *.Y...0.....K.?SY*
00010    E2C2C3D4 C6400001 00020000 0000002C    *SBCMF ..........*
00020    00340001 00000060 001C009E 430FC3D4    *.............CM*
00030    C661E7C1 40400084 912F0092 256F0000    *F.XA  .DJ..K.?..*
00040    000F0000 00000001 0C470000 40404040    *................*
00050    0000125F E2D7F34B F14BF340 03E00202    *....SP3.1.3 ....*
00060    200F3030 200EE2E8 E2D7F0F0 00780080    *......SYSP00....*
00070    05000052 0A6A2323 01000000 201F3030    *................*
00080    200EE2E8 E2D7F0F1 00780080 05000052    *..SYSP01........*
00090    0A6A2323 01000000 202F3030 200EE2E8    *..............SY*
000A0    E2D7F0F2 00780080 05000052 0A6A2323    *SP02............*
000B0    01000000 203F3030 200EE2E8 E2D7F0F3    *.........SYSP03*
000C0    00780080 05000052 0A6A2323 01000000    *................*
000D0    204F3030 200EE2E8 E2D7F0F4 00780080    *......SYSP04....*
000E0    05000052 0A6A2323 01000000 205F3030    *................*
```

Table 2-5 PERUSE Output Panel Field and Column Descriptions

Field	Description
System	System ID taken from the SMCA.
MVS	MVS operating system release level.
Monitor	CMF Extractor sampling mode used to collect the data shown in the displayed record.
Format	Description of the operating system that formatted the record.
Record	Sampler record type shown in the data area of the panel.
Samples	Count of samples taken in the recorded interval.
Cycle	Sampling rate in milliseconds.
Duration	Length of the sampling interval in minutes (mm), seconds (ss), and milliseconds (tttF).
Release	Recording product release level.
Compat	Data compatibility level.
Date	Julian start date of the recording interval.
Start	Start time of the recording interval in hours, minutes, and seconds (0hhmmssF).
Product	Hexadecimal offset to product section of the record.
1st	Hexadecimal offset to first data section, data length, and byte count.
2nd	Hexadecimal offset of second data section, data length, and byte count.
3rd	Hexadecimal offset of third data section, data length, and byte count.
Index	Hexadecimal offset of each line of data within the record.
Columns 2–5	Hexadecimal dump of the record; each line represents four full words.
Column 6	Eye-catcher area that shows the EBCDIC character representation of each line of data.

Displaying Specific Records in PERUSE

You can issue several commands from the PERUSE output panel to search and print Extractor records:

- GET (G)
- CONTINUE (C)
- PRINT (P)

GET Command

The GET command allows you to search and display specific records without scrolling forward or backward through the Extractor data set. You can narrow the focus of your search to display only the specific types of records that you request.

The GET command has several variations to search the data set by the type of record you need to find:

- Type **G** on the **COMMAND** line to display the next record of the type that is currently active for PERUSE.

- Type **G00** on the **COMMAND** line to display the next record, regardless of type, and nullify the current record type display restriction.

- Type **Gxxx** on the **COMMAND** line to display the next record of type *xxx* For example, G240 displays the next type 240 record in the data set.

- Include a hyphen and number to select a record subtype with the GET command. For example, G240-12 displays the next type 240-12 record in the data set.

- Type the **GET FIRST** command to display the first record in the data set of the type you requested. There are three forms of the GET FIRST command:

  — GF displays the first record in the data set of the current type.

  — GF00 displays the first record in the data set and specifies that the current record type display is null.

  — GF*xxx* displays the first record in the data set of type *xxx.* For example, GF240 displays the first 240-type record in the data set and specifies that the current record type display is for 240 records. If you type G again, the next 240 record in the data set is displayed.

CONTINUE Command

If the number of searched records exceeds the limit given from the SEARCH LIMIT field of the Peruse and Print CMF Records panel before the desired record type is found, the search is stopped and the message xxx RECORDS SEARCHED is displayed. Type the CONTINUE command, **C**, on the **COMMAND** line to continue the search. The search is continued for an additional number of records set from the SEARCH LIMIT field of the Peruse and Print CMF Records panel.

PRINT Command

Type the PRINT command, **P**, on the **COMMAND** line to generate a SNAP dump of the current record.

PERUSE Command Summary

Table 2-6 summarizes how the three PERUSE commands search and print records from the Extractor data set.

Table 2-6 PERUSE Option Commands

Command	Description
G	Gets the next record of the type currently displayed by PERUSE. The next record in the database that is the same as the current type of record is displayed.
G*xxx*	Gets the next record of type *xxx.* Type the record type, *xxx*, to change the type of displayed record. For example, G240 displays the next 240 record in the database. A 240 record becomes the active record type for the GET command.
GF	Gets the first record of the current type. A GF command starts the search at the beginning of the data set.
GF*xxx*	Gets the first record of type *xxx.* A GF*xxx* command starts the search at the beginning of the data set. For example, GF240 displays the first type 240 record in the data set and specifies the active record type to 240 for the GET command.
G00	Gets the next record, regardless of type, and nullifies the current record type.
GF00	Gets the first record in the data set, regardless of type. The current record type becomes null.
C	Continues the search for an additional number of records set from the SEARCH LIMIT field of the Peruse and Print CMF Records panel. If the search limit is reached before the desired record type is found, the search is stopped and this message appears on the screen: xxx RECORDS SEARCHED When you see this message, type **C** to continue your search for the next *xxx* records.
P	Generates and prints a SNAP dump of the current displayed record.

Chapter 3 Producing DSO Reports

Setting up the DSO Analyzer to generate reports is the second major task in reorganizing the volumes in your system. This chapter shows you how to run the Analyzer to process the Extractor records.

Figure 3-1 shows all of the reports available with the DSO Analyzer.Some reports are produced by default; others must be requested by setting parameters with Analyzer control statements. The default and optional parameters to produce the reports are listed on the left in Figure 3-1. The default parameters are underlined

Figure 3-1 DSO Reports

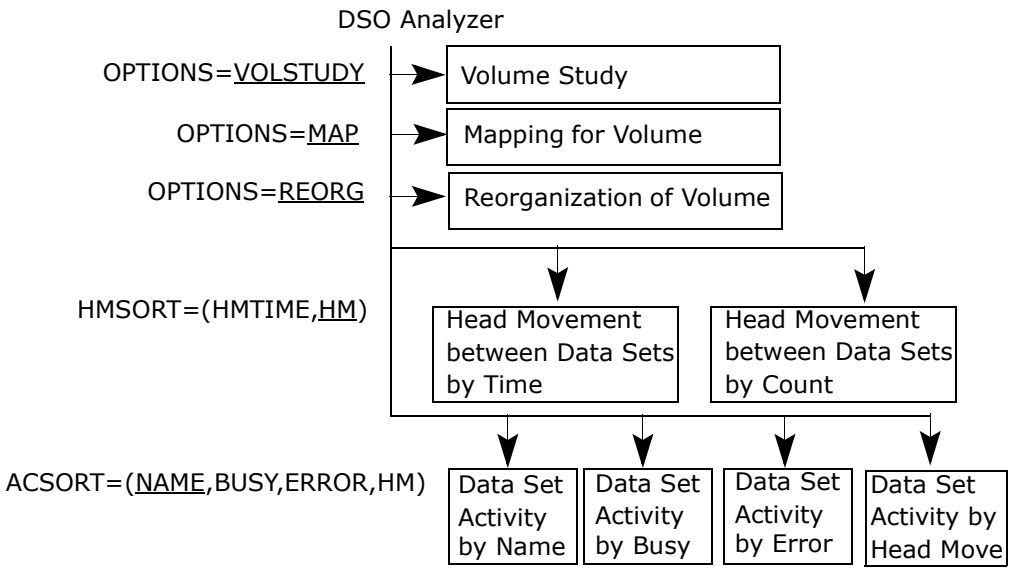

A brief summary of each DSO report is presented in Table 3-1. Chapter 4, “Evaluating DSO Reports,”, shows you how to use the reports to reorganize your volumes. Chapter 7, “DSO Reports,”, in the reference part of this book gives you a detailed description of each report.

Table 3-1 Summary of DSO Reports

Report title	Report summary
Study Report for Volume	Summarizes the activity that occurred on the volume during the sampling interval. Individual summary reports are created for every volume that is studied. The information given in the report is presented in three different categories: • Data set classes • Types of recorded samples • Head moves See “Study Report for Volume” on page 7-5 for more information about this report.
Mapping for Volume	Presents the physical mapping of all data held on the volume is presented in this report. The report lists the names of the data structures found between extent boundaries. The extents are listed in ascending cylinder and head order. The number of head moves to the extent during the sampling interval is also shown. See “Mapping for Volume Report” on page 7-7 for details.
Reorganization of Volume	Describes the recommended reorganization of the volume. The report lists the names of the data sets associated with a proposed extent. The extents are presented in ascending order by cylinder and head number. Data sets are placed on the volume according to the probability of head moves between adjacent extents. Refer to “Reorganization of Volume Report” on page 7-14 for a complete description.
Head Movement between Data Set Pairs by Time on Volume/by Count on Volume	Show head move activity between pairs of data sets with a default and optional report. The default report lists the data set pairs by the number of head moves between them in the sampling interval. The optional report lists the same data set pairs by their percentage share of the total time for all of the head moves recorded during the sampling interval. See “Head Movement Between Pairs of Data Sets by Time on Volume Report” on page 7-10 and “Head Movement Between Pairs of Data Sets by Count on Volume Report” on page 7-13 for details.
Data Set Activity Report by Name on Volume Data Set/by Busy on Volume Data Set/by Error Recovery on Volume Data Set/by Head Movement on Volume	List head move activity for each data set on the volume according to a particular head move statistic that is shown in the report. The default report lists the data sets by name in alphabetic order. The three optional reports list the data sets by the number of head moves, percentage of busy time, and percentage of total error time. “Data Set Activity Reports” on page 7-19 describes the fields that appear in all four Data Set Activity reports.

Setting Up the JCL for the DSO Analyzer

The DSO Analyzer reads Extractor records from an SMF or a CMF data set and reformats them onto a direct access work file managed by the BMC Software DMSS access method. Selected volume mapping and head move records are sorted and transferred to the DMSS file according to a date and time range specified by an Analyzer control statement. Calculations are performed on the records to produce statistics that appear in a series of reports about the data set organization of the sampled volumes. The Analyzer also produces utility control statements for automated volume reorganization using FDR COMPAKTOR or DFDSS.

A sort package compatible with DFSORT is required to run the DSO Analyzer. You must not presort the data. Only records required for the requested reports are inserted into the sort.

Figure 3-2 shows an example of a typical batch job to run the DSO Analyzer. The example is from a member named CMFANLD1 held in your BBPARM library for CMF MONITOR. Each statement in the example is described in Table 3-3 on page 3-13.

Figure 3-2 Sample DSO Analyzer JCL

```
//JOBCARD JOB
//*-------------------------------------------------------------------
//*                    SAMPLE JCL TO RUN THE DSO ANALYZER
//*   REVIEW THE JCL FOR APPLICABILITY TO YOUR INSTALLATION STANDARDS
//*   CHANGE ?HILEVEL TO THE HIGH-LEVEL QUALIFIER CHOSEN FOR THE CMF MONITOR LIBRARIES.
//*   CHANGE ?BBASMFID TO THE SMF ID (SYSTEM ID) OF THE TARGET SYSTEM.
//*   SUPPLY VALUES FOR ALL PARAMETERS MARKED WITH ?????
//*-------------------------------------------------------------------
//DSOCMFA EXEC PGM=DSOCMFA,REGION=4096K,
//        PARM='MAXVOL=255,MAXCC=8,RESERVE=2048K'
//STEPLIB  DD   DISP=SHR,                 - ANALYZER LOAD LIBRARY
//         DSN=?HILEVEL.BBLINK
//INPUT1   DD   DISP=SHR,                 - EXTRACTOR INPUT DATA
//         DSN=?HILEVEL.?BBASMFID.IPMOUT1
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SORTMSG  DD SYSOUT=*
//SYSOUT   DD SYSOUT=*
//INDEX    DD SYSOUT=*
//SYSIN    DD DUMMY
//DMSSMAIN DD UNIT=SYSDA,SPACE=(CYL,(5)),DCB=DSORG=DA
//DMSSALT  DD UNIT=SYSDA,SPACE=(CYL,(5,5))
//SORTLIB  DD DSN=SYS1.SORTLIB,DISP=SHR
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(10,10))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(10,10))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(10,10))
//*DFDSS    DD XXX....   <=== DFDSS JCL OUTPUT DATASET
//*COMPAKTR DD XXX....   <=== FDR COMPAKTOR CONTROL STATEMENT OUTPUT DATA SET
```

```
//DSOIN    DD *
  (DSO Analyzer control statements)
/*
//
```

Setting Up DSO Analyzer Control Statements

This section shows you how to prepare the Analyzer control statements to produce DSO reports. It presents some basic concepts that you should consider when you are preparing the job to run the DSO Analyzer. Refer to Chapter 6, “DSO Analyzer,”, for a complete description of each control statement.

Figure 3-3 is an example of a typical set of DSO Analyzer control statements. In the example, the control statements are placed within the job to run the Analyzer with the DSOIN DD statement. The parameters shown with the control statements produce all nine DSO reports and create an output job to reorganize volume ABC123 with DFDSS. The control statements are shown in Figure 3-3.

Only four of the possible six control statements are shown in Figure 3-3. DATETIME and HMTIME are the two other Analyzer control statements. Both are used typically under specific conditions:

- DATETIME selects specific records from your Extractor data set by date and time.
- HMTIME allows you to substitute an alternative set of equations to calculate seek time from the head move records.

Both control statements are discussed in Chapter 6, “DSO Analyzer,”. “Restricting Input Data with the DATETIME Control Statement” on page 3-8 shows how to use DATETIME with the DATA control statement to set local and global date/time ranges for selecting records from the Extractor data set.

Figure 3-3 Example Control Statements for DSO Analyzer

```
//DSOIN    DD *
 ANUM TST1
 DATA MONITOR=IPM,SYSTEM=SYSB,(INPUT1)
 DEFAULT OPTIONS=(DFDSS,REORG,HM,MAP,VOLSTUDY,ACTIVE),
         ACSORT=(HM,BUSY,NAME,ERROR),HMSORT=(HM,HMTIME)
 VOLUME NAME=ABC123
```

The control statements shown in Figure 3-3 are described in the following list:

ANUM — Optional DSO control statement that specifies a label for the DSO reports. "ANUM" on page 6-3 describes two other parameters that you can use with the control statement.

- TST1 — Positional keyword parameter that identifies the particular Analyzer job. In this example, the word TST1 is the analysis run number that appears in the header area at the top of page 2 of the reports. This parameter can be up to four characters long.

DATA — Regulates the type of Extractor records that are processed by the Analyzer to produce DSO reports. Normally, the three parameters shown with the DATA control statement in the example are those you use most often. "This example limits DSO reports to the first 25 volumes found in the Extractor data set. Any errors that occur as the reports are built are ignored if their condition code is less than or equal to 08. The analysis run number, TST1, appears at the upper right corner of the Extractor Data Sets report shown in Figure 6-1" on page 6-4 describes every parameter that can be declared with the control statement.

- MONITOR=IPM — IPM is the recommended sampling mode to collect the Extractor records used by DSO. The MONITOR= parameter restricts the input records processed by the Analyzer to those collected under IPM sampling. The MONITOR= parameter should always be specified with the DATA control statement if your Extractor data set holds a mixture of records that are collected with both CPM and IPM sampling.

- SYSTEM=SYSB — Identifies the SMF system ID whose records are stored on the Extractor data set. Up to four different systems can be declared with the SYSTEM= parameter.
  This parameter should always be specified with the DATA control statement if you collect your Extractor records with a system ID. If the SYSTEM= parameter is not declared, the input records are restricted to those without a system ID.

- (INPUT1) — Identifies the ddname that declares the input data set(s) holding Extractor records. This parameter must always be used with the DATA control statement.

DEFAULT — Sets global processing options for every volume that has Extractor records. In the example shown in Figure 3-3 on page 3-4 that DEFAULT precedes the VOLUME control statement in the job stream. You must always code the two control statements in this manner: DEFAULT before VOLUME.

There are 16 parameters associated with the DEFAULT control statement. “DEFAULT” on page 6-11 describes each parameter and subparameter in more detail.

OPTIONS= Declares the type of reports and volume reorganization produced by the Analyzer from the Extractor records. In the example shown in Figure 3-3 on page 3-4, the subparameters that appear with the OPTIONS= parameter specify that every DSO report is produced. The DFDSS parameter specifies that the Analyzer generate a job to reorganize a volume with the DFDSS facility.

ACSORT= Declares the types of Data Set Activity reports that are produced. This parameter must be declared with the DEFAULT or VOLUME control statements if you want to generate the optional Data Set Activity reports.

HMSORT= Declares the Headmove Between Data Set Pairs reports. This parameter must be declared if you want to generate an optional report that lists data set pairs by total head move time.

VOLUME Sets report and processing options for individual volumes. In the example shown in Figure 3-3 on page 3-4, all of the parameters declared with the other Analyzer control statements apply to the reports and processing options for the volume named ABC123.

NAME= Declares the name of the volume whose records are processed by the Analyzer. The volume is declared by volume serial name.

Setting Data Thresholds with Analyzer Control Statements

The Data Set Activity and Head Move by Data Set Pairs reports can become large if there are large numbers of data sets on the volumes that are studied. The reports become very too large if every data set has to be listed. You can limit the number of data sets that appear in the reports by setting thresholds with the VOLUME or DEFAULT control statements. Data sets that do not meet your threshold criteria are eliminated from the reports. This is a particularly effective way of keeping your reports manageable.

DEFAULT and VOLUME control statements provide six threshold parameters. Table 3-2 shows the six threshold parameters that are paired with the control statement subparameters that produce a particular DSO report. Unnecessary data sets are excluded from the reports by comparing a data set statistic against a threshold you set with these parameters. If the test criterion does not meet your threshold value, the data set does not appear in the report.

Table 3-2 DSO Analyzer Threshold Parameters

Report parameter	Report subparameter	Threshold parameter	Threshold description
ACSORT	NAME	NAMETH	Sets a threshold for the Data Set Activity report sorted by data set name. Data sets are included in the report if the data set name is true when compared to an alphanumeric string.
	BUSY	BUSYTH	Sets a threshold for the Data Set Activity report sorted by the percentage of total busy time. Data sets are included in the report if the test condition is true when compared to a percentage threshold.
	ERROR	ERRORTH	Sets a threshold for the Data Set Activity report sorted by the error recovery percentage of each data set. Data sets are included in the report if the test is true when compared to a percentage threshold.
	HM	HEADMVTH	Sets a threshold for the Data Set Activity report sorted by the total head move percentage. Data sets are included in the report if the test condition is true when compared to a percentage threshold.
HMSORT	HMTIME	HMTIMETH	Sets a threshold for the Head Movement by Data Set Pairs report sorted by the percentage of total head move time. Data sets are included in the report if the test condition is true when compared to a percentage threshold.
	HM	HMMOVETH	Sets a threshold for the Head Movement by Data Set Pairs report sorted by the percentage of total head moves. Data sets are included in the report if the test condition is true when compared to a percentage threshold.

The threshold parameters consist of a comparison operator and the threshold value expressed in the following form:

(comparison operator,threshold value)

Comparisons are made on the basis of the test criterion being less than (LT), equal to (EQ), or greater than (GT) the threshold value. The threshold values can be either a report percentage with two significant figures, *xx.xx*, or an alphanumeric string of five characters or less. Data sets are included in reports if the tested criterion is true when compared to the threshold value.

In the example of Analyzer control statements shown in Figure 3-3 on page 3-4, all four Data Set Activity reports are requested by coding the parameter ACSORT=(HM,BUSY,NAME,ERROR) with the DEFAULT control statement. As an example, if you want to limit the number of data sets that appears in the Data Set Activity report that sorts by head move percentage, you add the HEADMVTH parameter to the DEFAULT control statement. In the example of the DEFAULT control statement shown below, the highlighted HEADMVTH parameter limits the number of data sets that appear in the report to those with a percentage share of the total head moves greater than 12 percent.

```
DEFAULT OPTIONS=(DFDSS,REORG,HM,MAP,VOLSTUDY,ACTIVE),
        ACSORT=(HM,BUSY,NAME,ERROR),HEADMVTH=(GT,12.00),HMSORT=(HM,HMTIME)
```

Restricting Input Data with the DATETIME Control Statement

The DATETIME control statement sets a date and time range for the records held in the Extractor data set. You can be selectively include or exclude input records from processing by using the DATETIME control statement. A record is included in DSO reports if the date and time of the recording interval occurs within the specified DATETIME range.

You can selectively restrict the Extractor records you include in your DSO reports by using a combination of global and local DATETIME control statements. A global DATETIME statement restricts all input records to the DSO Analyzer to a specific date/time range when records were collected. A global DATETIME statement precedes all DATA control statements in the DSOIN data stream.

Each DATA control statement can also have an associated local DATETIME statement that sets a second date/time range. The local DATETIME range applies only to the records read from the ddname(s) listed on the preceding DATA statement.

If global and local DATETIME statements are used together, the resulting range from which records are selected is the logical intersection of both ranges. You can create an invalid range by specifying non-intersecting local and global DATETIME ranges. No input records are processed and no DSO reports are created if the two ranges do not intersect. Typically, the global DATETIME range is expressed as a Julian date and time, while the local DATETIME range is from the SELECT= parameter.

Figure 3-4 is an example of restricting input data to a narrow daily range by using a combination of global and local DATETIME control statements.

Figure 3-4 Local and Global DATETIME Control Statements

```
//DSOIN DD
     DEFAULT  VTOCSIZE=1234
     DATETIME JULIAN=(03032:000000,03059:235959)   1
     VOLUME   NAME=VOLSER
     DATA     (CMFIN1),(CMFIN2),SYSTEM=CPUA,
              RECTYPE=240,MONITOR=CPM
     DATETIME SELECT=(090000,153000,WKDAY)         2
     DATA     (CMFIN3),MONITOR=IPM,
              SYSTEM=CPUA,RECTYPE=240
```

The global DATETIME statement, 1, restricts all input records to the Analyzer to those created in February 2003. The JULIAN range applies to all data sets because the DATETIME control statement precedes all DATA statements. The local DATETIME statement, 2, further restricts the records from the CMFIN1 and ddname CMFIN2 data sets to those generated on the weekdays, Monday through Friday, between 9:00am and 3:30pm. The SELECT= range applies only to data read from the ddnames listed on the DATA statement that precedes the DATETIME control statement at 2.

The local range does not extend to the records read from the CMFIN3 data set, because the DATA statement does not have a corresponding local DATETIME statement. The input records are still restricted, however, to February 2003 because of the global DATETIME statement.

Running the DSO Analyzer

This section presents information about running the Analyzer under specific conditions and provides diagnostic information for Analyzer failures. Abend U1531 is a characteristic Analyzer failure that occurs when vital records are missing from the input data. This section shows you how to troubleshoot the abend to narrow down the source of the failure.

Running the Analyzer with Large Amounts of Data

Extractor records collected over a long period of time are seldom used by DSO because of the cost of processing massive amounts of data and the poor reorganization plans produced by the Analyzer. Long-term head move records tend to swamp out the characteristic pattern of head moves observable with short-term sampling. Volume reorganizations have limited benefit because data set placement is based on gross averages of head seeks rather than the more important transient activity.

There is a special case, however, when analyzing long-term head move records can be very helpful in producing a valid volume reorganization plan. This special case applies to reorganizing large-capacity volumes that are used by multiple applications. A typical example of a mixed-use volume is one that holds data for concurrent TSO and data base applications. Head seeks are not entirely random, but dependent upon the execution of specific jobs over a daily cycle.

It is often undesirable to make a reorganization plan for a mixed-use volume based upon sampling records obtained from short-term transient head moves. Short-term sampling often misses daily periods of peak activity as jobs are scheduled over the work day. Long-term sampling tends to average out these peak periods. A reorganization plan for a mixed-used volume is a *best fit* approximation to improve average daily performance rather than over a specific, limited period of time.

The sort facility preprocesses input records from Extractor data sets that hold a large number of head move records. Expect to use a significant amount of CPU resources and to see large numbers of sort record counts when the Analyzer reads data sets containing long-term trend records. The cost of processing these records can be reduced by eliminating all volume mapping records except the oldest and the most recent. Valid head move records can be used over the length of the sampling interval without the cost of processing duplicate volume mapping records.

The NOVOLMAP subparameter of the DATA control statement can be used to eliminate redundant VTOC records. NOVOLMAP reduces the amount of record processing by the Analyzer to organize the head move records.

Warning! NOVOLMAP must be used carefully; vital records can be lost if it is declared incorrectly in your Analyzer control statements. “DATA” on page 6-5 gives some tips about using the NOVOLMAP subparameter with the DATA control statement.

Combining Extractor Data from Multiple Systems

Valid mount records must precede head move records in the Extractor output. The mount records establish the relationship between device addresses and volume serials at the time the head move records are created. If the record order is disturbed, the DSO Analyzer does not process any head move records until a mount record defines the relationship between the volume and the rest of the storage subsystem. If records are processed from multiple extractions, the volume mapping records must be consistent over the entire length of the sampling interval for all extractions. If the volume serial to device address relationships vary across the extractions, the head move records are misordered and attributed to the wrong volume, resulting in invalid DSO reports.

Adhere to the following rules when combining data from multiple extractions:

- Preserve the time sequence and physical integrity of DSO Extractor output data.
  The time sequence is usually voided (accidentally) when Extractor data is output to multiple data sets and allowed, or forced, to switch from one data set to the other during the extraction. This causes the Extractor data to span two data sets. If the data sets are then fed to the Analyzer in reverse order, the positional dependencies between the mount and the head move records are violated. It is important to understand that many of the head move records are excluded from processing in this case.
  To avoid this problem, immediately archive those Extractor output data sets to a sequential data set on a medium such as tape. Use the DATETIME control statement to select the extractions to be analyzed from this sequential archive of DSO output.

- Avoid running sorts on DSO Extractor output records.
  The DSO Analyzer invokes a sort utility to separate head move records by volume and to consolidate volume mapping records from DSCB and VSAM catalogs. The most current volume mapping records are used to generate all DSO reports. It is crucial that the mount records be correctly ordered in the input stream.
  If possible, avoid sorting concatenated Extractor records. The records are easily misordered in a concatenated sort, leading to failures caused by missing records or invalid reports because obsolete volume mapping records were used.

- Specify the same system ID to a shared volume that the Extractor was running when collecting records from the volume.

When several systems share the same volumes, the system ID from which records were collected must agree with the system ID given with the SYSTEM= parameter of the DATA control statement. The volume mapping records must accurately represent the CPU and I/O configurations that were in effect when they were created. If they are not attributed to the correct system when they are processed by the Analyzer, invalid reports and reorganization plans are produced.

If DASD channel assignments are changed regularly, it might be convenient to associate each identifiable configuration with a different system ID. If data is tagged this way, it prevents possible errors if data is sorted incorrectly.

Concatenating Input Data Sets

The DSO Analyzer processes any number of input DD statements specified with the DATA control statement. Each DD statement refers to a single data set or multiple concatenated data sets.

If a spanned record is started in one data set of a concatenation, it can continue in the next data set in the concatenation, but records cannot span ddnames. The DSO Analyzer can avoid most concatenation errors, such as those caused by concatenating a large blocksize data set to a small blocksize data set.

If tape input is used, all tapes in a concatenation must be mounted simultaneously if they are assigned to separate drives.

Diagnosing Analyzer Abend U1531 Failures

An U1531 abend is a characteristic Analyzer failure. It occurs when the Analyzer is missing key records required to process other records. The job is aborted and no reports are created. This section lists some common errors that cause the abend. When you suffer an abend, you should carefully go through each item in Table 3-3 to confirm the coding in the Analyzer job.

Table 3-3 Diagnosing Analyzer Abend U1531

Cause of the abend	Diagnosis
Required Extractor records are missing.	Using the PERUSE option described in Chapter 2, “PERUSE Command Summary” on page 2-20, look at your Extractor data set to confirm the presence of the following records for the volume you are studying: • 240-00 SRM record (optional) • 240-11 GBLS record (optional) • 240-12 volume mount records (required) • 240-13 head move records (required) • 240-14 VTOC map records (required) • 240-22 VSAM records (optional)
The CX10HMOV or CM40HMOV extractor module aborted during CMF Extractor sampling.	Either module may abort if the REPORT control statement does not specify enough CSA storage. A secondary reason is that the VSAM=YES parameter is declared to collect VSAM records and the extraction process takes too long, resulting in an Extractor failure.
Different SMF record formats have been declared with the Extractor and Analyzer control statements.	Verify that the SMFRECID parameter of the Extractor REPORT control statement agrees with the RECTYPE parameter of the Analyzer DATA control statement.
The date and time range set with the DATETIME control statement excludes all Extractor records from processing.	Verify that the Extractor data set contains records that were created during the date and time range selected for Analyzer processing. See Register 12 to get the system ID from CMF records. See Register 14 to get the start date. See Register 15 to get the start time.
The SYSTEM= parameter was not coded with the the DATA control statement or it has an invalid system ID.	Look at the system ID of the records held in the Extractor data set for the studied volume. That ID should be coded with the SYSTEM= parameter. See Register 12 to get the system ID from CMF records. See Register 14 to get the start date. See Register 15 to get the start time.
No valid Extractor data set names were declared with the Analyzer JCL.	Confirm that the ddname(s) of the Extractor data set(s) agree between the Extractor JCL and the Analyzer DATA control statement.
The NOVOLMAP subparameter of the DATA control statement excluded all 240-14 VTOC mapping records from the volume requested for study.	Confirm that at least one valid 240-14 record is created for each volume that is studied.
The requested volume does not have any records in the Extractor data set.	Confirm that the volumes selected with the Extractor HEADMOVE parameter agree with the volumes chosen by equivalent parameters of the Analyzer DEFAULT or VOLUME control statements.

Table 3-3 Diagnosing Analyzer Abend U1531 (continued)

Cause of the abend	Diagnosis
An invalid type of DASD was declared with the TYPE= parameter of the DEFAULT or VOLUME control statement.	Confirm that the four-digit device type is supported by DSO. "DEFAULT" on page 6-11 lists valid DASD types in the discussion of the TYPE= parameter.
HEADMOVE records were concatenated in the wrong order.	Verify that the Extractor records are concatenated in the same order in which they were created. Normally, the four DSO records appear in the following order in the Extractor data set: 1. 240-12 volume mount records 2. 240-14 VTOC map records 3. 240-22 VSAM records 4. 240-13 head move records
The same volume was selected for study with two different VOLUME control statements. The Analyzer excludes multiple requests for the same volume.	Confirm that the volumes to be studied are specified only once in the JCL job to run the Analyzer.
The HEADMOVE control statement was missing from the Extractor job.	Confirm that a HEADMOVE control statement is included in the Extractor job that collects DSO data.
Extractor records were excluded by the Analyzer because the records were collected with a different sampling mode (CPM or IPM) than the type requested by the Analyzer.	Confirm that the sampling mode of the requested Extractor records agrees with the sampling mode declared with the MONITOR= parameter of the DATA control statement.

Usually, other DSO messages accompany the U1531 dump. These messages give more information about the cause of the failure. However, if the NOSTAE parameter is included with the EXEC statement in the JCL job, DSO messages are suppressed and are not printed.

Chapter 4 Evaluating DSO Reports

The third major task in reorganizing your volumes is evaluating the DSO reports. This chapter tells you how to analyze several DSO reports under specific conditions. Each section of the chapter describes a specific consideration that should be part of your evaluation of the DSO reports. Because of the large variation in modern storage subsystems, not every tuning condition can be addressed. This chapter presents the most common conditions.

When evaluating your reports, you should consider whether DSO is the appropriate tool to improve the performance of a specific DASD volume. DSO is designed to help you reduce seek time; it cannot solve other problems that afflict DASD performance.

If a volume is saturated with I/O requests, reorganizing the volume can alter the pattern of seek activity. Reorganizing the volume can possibly improve seek time or at least reveal the source of contention that is causing delay. However, if the DASD unit is running on an overloaded channel, the cause of poor performance is not at the volume level. Reorganizing the volume with DSO provides little benefit.

Always keep the following in mind:

- When a device has a low physical priority, the contention on the volume is commensurate with the service rate available to it; therefore, no improvement can be attained with DSO.

- When a device has a very high physical priority, the applications that use that device receive excessive time and resources while the devices with a lower physical priority are locked out. For example, a lower priority device cannot reconnect when it has data. The channel busy times and I/O queuing increase for reads and writes on the low priority volumes.

Unless the programs that use data stored on the volume can be rescheduled, the only solution to these problems is to redistribute the data sets to other volumes.

Another consideration when you conduct your DSO studies involves the depth of your interpretation of the reports. You must assess whether the benefit of improved DASD performance is worth the cost of your tuning efforts. Although the reports give you empirical information, they can be interpreted to suit your own needs. More importantly, the reports can be created under certain conditions to elicit the information that you need to systematically approach the problem of reducing seek time in your DASD subsystem.

As an example, the reorganization report enables you to reorganize the volumes in your system. If you decide that it is not worth the effort to completely reorganize the volume, you can selectively implement the reorganization recommendations. Splitting active data sets by moving one or more of them to other volumes is a relatively simple way to partially reduce the total seek time occurring on the volume without a complete reorganization. The Data Set Activity and the Head Movement by Pairs of Data Set reports list data sets by the number of head moves. The data sets at the top of either report are the best candidates to migrate to another volume.

Using DSO Reports

The remainder of this chapter gives you some tips about using the DSO reports when you reorganize your volumes.

Fixed-Head Studies

DSO can be used to simulate the conditions of a fixed-head volume. This is done by making two separate Analyzer runs on the same volume. The volume must be from 33*xxF* or 33*xx* DASD types.

The AUTOFIX or the FIXED parameter should be coded for the fixed-head study. For fixed-head type volumes, it is beneficial to make several runs of the Analyzer requesting only the Reorganization of Volume report while varying the AUTOFIX or FIXED parameters. Compare the reports to see which data sets were moved to the fixed-head cylinders.

VTOC Considerations

A traditional tenet of DASD tuning is to center the VTOC. The belief is that placing the VTOC at the midpoint of the data cylinders reduces the average time for seeks between the VTOC and all data sets distributed over the range of cylinders. Unfortunately, seeks between data sets and the VTOC are not distributed uniformly. In most circumstances, centering the VTOC does not significantly reduce the average seek time of a volume and provides little performance benefit.

DSO treats the VTOC the same as it does any data set and optimally places it near the most active data sets on the volume. DSO works best on volumes where VTOC activity is not an important consideration. If a volume has a lot of intentional VTOC activity, data sets continue to be allocated, opened, closed, or deleted on that volume.

DSO reports activity by data set or VSAM component name and is designed to optimize volumes that contain permanent data sets that experience a large number of I/O requests. Inter-data set and intra-data set head seeks minimize the importance of VTOC seeks as a significant consideration in reducing seek time.

If the VTOC is very active on a permanent database volume because it is mounted with the wrong attributes, change the mount attributes. It is not advantageous to reorganize a volume if the VTOC attributes are in error.

On volumes where the VTOC is moderately active, the best strategy is to let DSO treat the VTOC the same as it would any data set and optimize its location relative to the real activity on the volume.

If security for the VTOC is important, the VTOC should be placed as close as possible to the physical center of the volume. This can be important for older DASD with removable disk packs, such as the 3330-I and the 3330-II models.

DSO aligns the VTOC on a cylinder boundary automatically; therefore, when possible, the VTOC should be allocated in cylinders. BMC Software recommends cylinder allocation for data sets that have potential performance problems.

Cylinder allocation removes the requirement that IOS perform extent checking each time there is a logical head change. It has a substantial effect on CPU time of I/O-bound jobs.

Evaluating Volume Usage

Monitoring volume usage on a daily or weekly basis is easily done with the Mapping for Volume report. Data set usage is evaluated quickly by finding and eliminating fragmented space on each volume and removing unnecessary data sets. Furthermore, information is easily obtained regarding the amount of space being used by a given data set. This lets you judge whether space for the data set has been allocated properly.

The Data Set Activity report can also be used to evaluate volume usage. Seek times become a problem when a volume has several active data sets that are located far apart from each other. If the volume cannot be reorganized to bring the data sets closer together, you can split them up by moving them to different volumes. This follows the general DASD tuning philosophy of minimizing the demands on scant resources by spreading I/O activity to achieve a balance. Run Data Set Activity reports on the other volumes in your system to find the best candidate to receive the active data set.

Combining Data from Multiple Systems

Adhere to the following rules when you combine data from multiple systems for shared DASD evaluation:

- Use DATETIME statements to get a clean slice of data from each system. Do not rely on starting and stopping the Extractors in synchronization.
- Gather VTOC information from only one system.

- Restrict the sampling period to times when seek data is available from all systems that have access to the volume.

- Ignore the Reorganization of the Volume report ***unless***

  — Head seeks are caused by a single application

  — Head seeks are caused by a group of applications running simultaneously on all the systems, with a common pattern of access across all the systems

  For example, ignore the Reorganization of the Volume report if the database volume was accessed by two CICS applications running on loosely coupled CPUs.

DSO does not track the actual head seeks on a shared DASD volume. The fourth item in the above list represents the only case where valid results can be expected because of statistical averaging. For best results, compare Generalized Trace Facility (GTF) data from loosely coupled systems, but only if their time of day (TOD) clocks can be synchronized to the same master clock. Head seek tracking requires very close resolution to reconstruct the head moves accurately.

If the activity on the device created by two systems is not similar, analyze the data using the consolidated reorganization proposed by DSO. This type of analysis is done as follows:

1. Run a set of reports that includes volume reorganization for system 1 by using the SYSTEM operand of the DATA statement to eliminate data from system 2.

2. Run a set of reports from the same data that eliminates data from system 1 instead of system 2.

3. Run a third set of reports that lists the Data Set Activity by Busy reports and the Head Movement by Pairs of Data Sets reports, combining data from the two systems.

4. Compare the Data Set Activity by Busy reports from system 1 with the consolidated Data Set Activity by Busy report.

Long-Term Trending

The DSO Analyzer normally is not used for long-term trending except in the case of multiple-use volumes, discussed in “Running the Analyzer with Large Amounts of Data” on page 3-9. The accuracy of DSO reports is much more dependent upon the resolution of the HEADMOVE sampler than the number of recorded samples. Because of the enormous amount of data that can be generated by sampling DASD seek activity, most Extractor sampling is restricted to peak daily periods that last no more than several hours. The resulting reorganization plan optimizes the volume for I/O activity that occurs during these peak periods. This strategy may not be appropriate for volumes that experience several peak periods in a single day.

An alternative strategy is to tune your volumes with Extractor data that is collected over long-term sampling periods. Long-term sampling allows you to gather a complete daily history of I/O activity. DASD seeks can be sampled during the non-peak periods that are usually ignored.

Tuning your volumes with long-term trend data allows you to distribute I/O activity across your storage subsystem to balance your workload. This tuning strategy reduces the aggregate seek time of a volume over a period of a week or month. It does not optimize the seek performance of a volume for a relatively short daily period.

The focus of tuning with long-term trend data is to identify the most active data sets. These data sets are the best candidates to migrate to another volume. The overall performance of the volume is improved by reducing the number of seeks between these active data sets. Reorganization is not necessary or desirable for volumes that experience several daily periods of peak I/O activity.

The Data Set Activity reports identify the best candidates to relocate to another volume. Adjust the reporting thresholds for the Data Set Activity reports to list only the data sets whose busy percentage is greater than a predetermined figure (for example, 5 percent), and select the prime production shift by using the DATETIME control statement.

Head move samples can be extracted over a long period without incurring high system overhead if the intended use is to produce Data Set Activity reports. The Data Set Activity reports retain their validity when processing head move data from low-resolution samplings (that is, a greater than 50-millisecond sampling interval). This long-term trending data should be kept separate from the high-resolution samples gathered during peak daily periods. High-resolution samples do not improve the accuracy of long-term trending data for Data Set Activity reports. High-resolution samples skew the results heavily in favor of the data sets that were active during the brief periods of high-resolution sampling. An hour of high-resolution sampling data can contain as many samples as a day's worth of low-resolution sampling data. The DSO Analyzer does not differentiate between records that are created in either mode and the reports generated from both high- and low-resolution sampling records are distorted.

Do not use data from low-resolution sampling to reorganize your volumes. DSO needs an accurate history of all of the head moves that occurred on the volume to produce a valid reorganization plan. Low-resolution records lack information about relatively short-term head moves.

DSO and Dual-Logical Volume Drives

DASD tuning becomes more difficult with improvements in volume capacity because you cannot continue to dedicate a volume to one or two active data sets. This problem is aggravated with DASD that put multiple actuators in a head disk assembly to create dual-logical volumes.

Adhere to the following recommendations to reduce the problems created by these dual-logical volume drives:

- Avoid using these volumes with any application that attempts to optimize head moves dynamically. Because physical cylinder boundaries on these devices do not correspond to the cylinder boundaries of the emulated drive, reorganization algorithms are not used correctly and they decrease seek performance. The algorithm generates more head moves rather than less. Typical applications that do this are MVS paging, JES spooling, sorts, and in-house programs.

- Consider physical cylinder size when determining VSAM control area sizes; the drives can be ideal for VSAM applications because VSAM can be easily tailored to use them properly. VSAM must be forced to begin allocating at a point where a physical cylinder begins rather than at a point where a cylinder begins for the emulated device.

- Do not try to center the activity on the physical volume.

- Do not allow logical relationships between data sets to span the logical-volume boundary, unless a single large VSAM data set can span logical-volume boundaries with no ill effects.
- Ignore CMF measurements such as service time and queue depth for dual-logical volumes.
- Adjust the head move timetables associated with dual-logical volumes by doubling the lowest head move time.
- Use the double BUFNO/NCP/BUFNI/BUFND installation standard for data sets on these volumes to reduce the number of head moves.
- Mark at least one of the volumes PRIVATE.
- Do not back up both volumes during simultaneous runs.

The ideal uses for dual-logical volume drives are as follows:

Use	Reason
VSAM	CASIZE can be tailored to take advantage of an actual physical organization.
TSO work volumes	Activity is very random.
Sequential data sets for batch	String length can be easily altered for most programs by doubling BUFNO/NCP/BUFNI/BUFND.

Chapter 5 CMF Extractor

The CMF Extractor collects information from specific DASD volumes and formats the data into records written to SMF or CMF MONITOR data sets. For each volume, these records are a statistical representation of the head moves that occurred during the sampling interval and a composite map of the data stored on the volume. DSO creates reports from the head move and volume mapping records. This chapter describes how the CMF Extractor samples and collects data for DSO reports.

The Extractor has two different modes of sampling data: continuous performance monitoring and intermittent performance monitoring. Each mode controls a separate set of Extractor samplers that gather data and write records at different rates. Together, both sampling modes allow the Extractor to simultaneously collect data for long-term trending and for brief events that need a large number of samples over a short interval.

CPM	Continuous performance monitoring (CPM) mode. Designed for long-term system measurement. The samplers normally are run continuously at low sampling rates to collect trend data. The relatively slow sampling rate reduces the overhead of running the samplers.
IPM	Intermittent performance monitoring (IPM) mode. Provides short-term, high-resolution sampling. It is initiated at the discretion of the operator to investigate specific areas of activity or to monitor specific devices during peak periods of activity. IPM sampling is intended to monitor specific devices or events that need a large number of samples over a relatively short interval. Collecting head move data is an example of an event that needs IPM sampling.

The CPM and IPM sampling modes should be considered as separate submonitors of the Extractor. Both modes are independent of each other and can be run separately or together. Either sampling mode can be started, stopped, or changed without affecting the performance characteristics of the other.

The Extractor uses four sampling methods:

- Disabled interrupt exit (DIE)
- Service request block (SRB)
- System resources manager (SRM)
- Task mode (TCB)

Most samplers use the SRB sampling method. Under SRB, higher resolution sampling is provided but system interrupts are still honored. CMF keeps the use of the DIE sampling method to a minimum to avoid degrading system performance. The SRM sampling method is used for sampling functions that are driven directly by SYSEVENTS.

Extractor Data Sets

The Extractor can direct output to SMF data sets when you specify SMF=YES on the Extractor REPORT control statement.

The Extractor uses the IBM SMFEWTM macro to write to the SMF data set. If you want data recorded to SMF, you must specify the SYS and/or SUBSYS parameter of member SMFPRM*xx* in SYS1.PARMLIB so SMF type 70 through 78 records and the CMF user record type (as specified on the Extractor REPORT statement) are written.

Alternatively, the Extractor can write to physical sequential data sets. CPM and IPM output also can be sent to the same or separate data sets for non-SMF recording. Only allocate CPM and IPM data sets when you are not using an SMF data set.

DISP=MOD is forced on all output data sets unless

- The output is sent to tape
- The output data set flips to the next available output data set you specified on DD names CMFCDS*xx* or CMFIDS*xx*, as a data set is filled
- The Extractor flips to the next available output data set when you use the MODIFY command; for example:

  F CMF,FLIP={CPM|IPM}

- You code the DISP=NEW parameter on the REPORT control statement

A TCLOSE is issued after each Extractor record is written, so that data written prior to a system failure is still accessible.

Under certain conditions, the TCLOSE does not protect the data set. For example, if a four-block VBS record is being written and the system fails as the third block is being written out, there is no end-of-file marker. Although data set damage rarely occurs in these circumstances, data sets can be recovered by copying the damaged data set to a new one, using a utility such as IEBGENER or CMF COPY VBS.

The CMF Extractor uses a technique called *alternate data set support* to continue writing records to the next available data set after each data set is filled with sampling data.

You can use the DSN= and ALTDSN= parameters of the REPORT control statement to specify one primary and one alternate output data set, or use DSNLIST= to specify up to 101 data sets to hold all Extractor sampling data.

The primary and alternate data sets must store similar records. If the primary data set collects IPM records, the alternate data set also must collect IPM records.

You can use the DSN= and ALTDSN= parameters of the HEADMOVE control statement to specify one primary and one alternate output data set, or use DSNLIST= to specify up to 101 data sets to hold just the Extractor head move data.

You can define from 1 to 101 data sets to hold Extractor sampling data by using the CMF*x*PM*xx* DD names in the JCL, where

x	Contains C (CPM mode) or I (IPM mode)
xx	Contains one or two legal characters (alpha, numeric, or a combination of both) you specify for 1–101 DD names

The Extractor automatically switches to the next available data set, following the order in which you specify them, to continue writing data and avoid an abnormal end. The Extractor, however, stops sampling after all data sets become full.

Note: When the Extractor is running, you can clear a data set not currently in use by executing the CMFOCLR procedure in batch mode or by executing the procedure at the console. Cleared data sets are reused by the Extractor.

For more information on CMFOCLR, see the *CMF MONITOR Customization Guide*.

You can define from 1 to 101 data sets to hold Extractor head move data by using the CMF*x*DS*xx* DD names in the JCL, where

x	Contains C (CPM mode) or I (IPM mode)
xx	Contains one or two legal characters (alpha, numeric, or a combination of both) you specify for 1to 101 DD names

Extractor Records

The CMF Extractor writes one record to the CMF data set at every recording interval specified by the INTERVAL= parameter of the REPORT control statement.

Extractor samplers store a copy of the VTOC map from each sampled volume and their head move data to an output data set. The head move data consists of a series of timestamped records that describe the position of the read/write heads on the recording surface.

The Extractor writes head move records to the output data set at regular sampling intervals for the duration of the Extractor run. The Extractor writes the mount records and the VSAM catalog from each volume to the output data set during the Extractor initialization phase. After the initialization phase is complete, the Extractor scans VTOC maps at a paced interval in the sampling phase. Head move sampling begins after all of the VTOC records are scanned.

Table 5-1 shows the six possible records that can be collected for DSO. The two control statements (HEADMOVE and REPORT) that specify the operating characteristics. The Extractor you must use to collect these records are described in “CMF Extractor Control Statements” on page 5-5.

Table 5-1 DSO Extractor Records

Control statement	Sampler	Extractor record	Record description
HEADMOVE	HMOV	240-12	Describes the volumes being sampled by device address, device type, and volume serial number. The record is generated when the head move sampling begins and when a new volume is encountered during a sampling period.
		240-13	DSO makes an entry in the record whenever a change in seek address or device status occurs on a volume sampled for head moves. The record entry includes device address, seek address, device status, and the sample count.
		240-14	Contains VTOC information for each volume sampled for head moves. This record is created during the Extractor sampling phase if the VTOC=YES parameter is specified with the HEADMOVE control statement.
		240-22	Contains the VSAM catalog description segments of the volumes being sampled for head moves. This record is created during the Extractor initialization phase if the VSAM=YES parameter is specified with the HEADMOVE control statement.
REPORT	RECD	240-00	Contains SRM constants, installation performance specifications, Extractor control statements, and sampling rates. This record is written once every sampling interval.
	GBLS	240-11	Consists of a CPU section followed by a variable number of global data sections. The CPU section of the record describes the CPU configuration of the system. The global sections describe the status of the CPU, the executing program, and a variable length channel bit map. This record is created if the default GBLS=YES parameter is unchanged with the REPORT control statement. This record is written as needed and, at a minimum, once every recording interval.

CMF Extractor Control Statements

DSO requires two Extractor control statements to run the samplers that collect the records used to produce reports. Each Extractor control statement has associated parameters that control the manner in which a particular sampler functions. In the case of DSO, the REPORT control statement parameters set global operating conditions for the Extractor sampling interval. The HEADMOVE control statement parameters regulate the collection of volume mapping and head move records on selected DASD.

The remainder of this chapter discusses the REPORT and HEADMOVE control statements. Table 5-2 gives a page reference for further information about each control statement.

Table 5-2 CMF Extractor Control Statements Used with DSO

Control statement	Control statement function	Further information
HEADMOVE	Samples DASD head move activity	See “HEADMOVE” on page 5-6
REPORT	Defines the Extractor recording mode and other operating characteristics	See “REPORT” on page 5-12

HEADMOVE

The HEADMOVE control statement tells the CMF Extractor to sample and collect head move data from selected volumes within the measured system. The HEADMOVE control statement is required with DSO. Head move data from a maximum of 4096 volumes can be collected from a single system.

You can specify one HEADMOVE control statement for each Extractor sampling mode (CPM or IPM). However, the performance burden of the HEADMOVE sampler is quite large. With DSO, you should run the Extractor in IPM mode only at a recommended sampling interval between 25 and 33 milliseconds.

Note: A new parameter, DSNLIST=, enables you to define from 1 to 101 data sets to hold head move data. You can continue to use DSN=, ALTDSN=, and DDNAME= if you choose not to make HEADMOVE control statement changes.

You can specify up to four characters in device address ranges you assign to the RANGE= and EXCEPT= parameters.

HEADMOVE Syntax

```
HEADMOVE  [ALL]
[,CHANNEL=(c1,c2,...c15)]
[,NUMBER=(n1,n2...n15)]
[,RANGE=(xxxx[:xxxx]1,xxxx[:xxxx]2,...xxxx[:xxxx]32)]
[,EXCEPT=(xxxx[:xxxx]1,xxxx[,xxxx]2,...xxxx[:xxxx]32)]
[,VTOC={NO|YES}}
[,VSAM={NO|YES}]
[,DSNLIST=(xxx1,...,xxx101)]
[,DDNAME=ddname]
[,DSN='dsn']
[,ALTDSN='dsn']
[,SAMPLE={33|nnnn}]
[,OFFLINE={NO|YES}]
[,PACE={nn}]
```

Control Statement Parameters

ALL	Positional keyword parameter that specifies that all online volumes in the system are sampled for head move data. It supersedes all other HEADMOVE parameters that declare the volumes to be sampled for head move activity. The CHANNEL=, NUMBER=, and RANGE parameters are ignored if ALL is specified with the HEADMOVE control statement. Note: You can exclude specific volumes from sampling by coding the EXCEPT parameter with the ALL parameter.
CHANNEL=	Specifies the measurement of all on-line volumes that use a particular channel as their primary path. The CHANNEL= parameter is valid only in an MVS/370 environment. You can declare up to 16 channels using a unique hexadecimal number (0-F) for each channel. If you specify more than one channel, separate channel numbers by commas and enclose them with parentheses. Note: The CHANNEL parameter is ignored if the positional parameter ALL is included with the HEADMOVE control statement.
NUMBER=	Specifies the measurement of all on-line volumes with the same first digit as the three-digit DASD device address. For example, if NUMBER=2 is declared with the HEADMOVE control statement, all online volumes with a 2*xx* unit address are measured in the system. The NUMBER= parameter is valid only in MVS/XA and MVS/ESA environments. You can declare up to 16 numbers using a unique hexadecimal number (0-F) for each unit address range. If you specify more than one number, separate numbers by commas and enclose them within parentheses. Note: The NUMBER parameter is ignored if the positional parameter ALL is included with the HEADMOVE control statement.

RANGE=	Specifies a range of volumes for the Extractor to measure by device address. The range can consist of a list of single device addresses, a list of device address ranges (*xxxx*:*xxxx*), or a combination of both. You can use a zero prefix for three-digit device addresses and you can specify three or four digits in any combination; for example: • RANGE=(234:23F) • RANGE=(1234:124F) • RANGE=(234:12FF) • RANGE=(0200:2DD) You can designate a maximum of 32 volumes or device ranges for sampling with the RANGE= parameter. Note: The Extractor ignores the RANGE parameter when you include the positional parameter ALL in the HEADMOVE control statement.
EXCEPT=	Excludes a range of volumes from sampling by the HEADMOVE control statement. You can exclude a maximum of 32 volumes or device ranges. The range can consist of a list of single device addresses, a list of device address ranges (*xxxx*:*xxxx*), or a combination of both. You can use a zero prefix for three-digit device addresses and you can specify three or four digits in any combination; for example: • EXCEPT=(234:23F) • EXCEPT=(1234:124F) • EXCEPT=(234:12FF) • EXCEPT=(0200:2DD)
VTOC=	Declares whether VTOC data is extracted from the volumes being sampled for head move. The CMF Extractor default is NO. VTOC=YES must be explicitly specified with the HEADMOVE control statement to produce the head move and VTOC data required for DSO analysis. VTOC=YES must be specified on at least one Extractor run to provide the required VTOC maps. If records are collected from the same volumes with multiple extractions, only one run needs to collect VTOC maps. The remaining runs can omit redundant VTOC records by specifying VTOC=NO with their HEADMOVE control statement.
VSAM=	Declares whether VSAM cluster allocation information is extracted from the sampled volumes. The Extractor default is NO. The primary reason for the VSAM=YES parameter is to obtain cluster allocation information for VSAM clusters that are defined within a VSAM data space. The secondary reason for the VSAM=YES option is to provide the cataloged name of the cluster for the component name. Since the advent of DF/EF VSAM, multiple cluster data spaces no longer exist. Thus, coding VSAM=YES provides little benefit to DF/EF VSAM users. Coding VSAM=YES can add an excessive amount of time to Extractor processing because DSO scans through each VSAM or ICF catalog associated with a given volume. This is not a serious problem with a small number of volumes. The additional processing time becomes more serious with a larger number of volumes with VSAM data sets. Another problem associated with the VSAM=YES parameter occurs when VSAM catalogs are not kept current. Many data sets can be cataloged on volumes that no longer exist. As a result, catalog management requests that a volume be varied online that no longer exists in the system. You must manually terminate the request.
DSNLIST=	Specifies from 1 to 101 previously allocated, cataloged CMF Extractor data set names to hold head move data. The Extractor uses the CMF*x*DS*xx* names in the JCL. If the CMF*x*DS*xx* DD names are already allocated, the Extractor disables and reallocates the CMF*x*DS*xx* DD names for its use. Data is written to data sets in the order in which you specify them. When the last data set is full, the Extractor writes to the first and subsequent data sets again, writing over any existing data. Do not use DDNAME= or DSN= parameters in HEADMOVE when using DSNLIST=.

DDNAME=	Specifies the DD name of a previously allocated, cataloged CMF Extractor data set that holds head move data. You do not have to modify the JCL if the data set on the CMF*x*DS*xx* DD name in the JCL matches the data set name on the DDNAME= parameter. You must allocate the data set before running the CMF Extractor. The DSN=, DSNLIST=, and DDNAME= parameters are mutually exclusive. Use only one with the HEADMOVE control statement.
DSN=	Specifies the name of a previously allocated, cataloged CMF Extractor data set that holds head move data. The DSN=, DSNLIST=, and DDNAME= parameters are mutually exclusive. Use only one with the HEADMOVE control statement. If you use DSN=, you do not have to change the JCL if the data set on the CMF*x*DS*xx* DD name in the JCL matches the data set name on the DSN= parameter. You must allocate the data set in advance.
ALTDSN=	Specifies the name of a previously allocated, cataloged CMF Extractor data set that holds head move data if the primary data set becomes full. This is an alternate data set. Only use ALTDSN= when using DSN= to define a primary data set. If you use ALTDSN=, you do not have to change the JCL if the data set on the CMF*x*DS*xx* DD name in the JCL matches the data set name on the ALTDSN= parameter. You must allocate the data set in advance.
SAMPLE=	Declares the number of milliseconds between data gathering cycles. At the end of a sampling interval, head move data is written to the CMF Extractor data set. The default is 33. The sampling interval cannot be less than 20 nor greater than 9999 milliseconds.
OFFLINE=	OFFLINE=YES parameter reserves entries in the sample tables for selected volumes that are offline when the Extractor begins sampling. If these volumes are placed online during an Extractor run, head move data is collected from them for the remainder of the sampling period. The default, OFFLINE=NO, eliminates unwanted dummy device entries.
PACE=	Regulates the length of the interval between each volume's VTOC scan requests. The optional PACE= parameter is an integer between 0 and 64 that is multiplied by the SAMPLE= parameter. The product of the two parameter values is the resulting pacing interval between the VTOC scan of each volume. HEADMOVE VTOC scans occur during the recording phase of CMF Extractor sampling. VTOC data from a single volume is scanned during the data collection cycle of the sampling interval set by the SAMPLE= parameter. Scan processing continues until VTOC data from every sampled volume has been collected. If there is a large number of VTOC records, other CMF Extractor functions can be delayed and failures result. By lengthening the interval between VTOC scans to allow other processing, some CMF Extractor failures can be avoided. By default, VTOC scan requests occur at an approximate 2-second interval. If the PACE= parameter is not declared, the VTOC scan interval remains at 2 seconds. VTOC interval pacing is ignored if the PACE= parameter is greater than 64 or if the SAMPLE= parameter is greater than 2000.

Examples

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

The CMF Extractor collects data from all online volumes at the default 33-millisecond interval except volumes 142, 145, 146, and 147. Entries are reserved in the sample tables for offline volumes. Mandatory VTOC data is collected for DSO.

```
HEADMOVE CHANNEL=(4,5),RANGE=(201,204:305),VTOC=YES,SAMPLE=25
```

This example is for DSO running in an MVS/370 environment. The CMF Extractor samples the volumes that use channels 4 and 5 as their primary path. The volumes sampled from these two channels must have a device address 201 or in the range between 204 and 305. Head move data is collected every 25 milliseconds during the Extractor sampling period. Mandatory VTOC data is collected for DSO.

```
HEADMOVE NUMBER=(304,306),VTOC=YES,VSAM=YES,
           DSN=HIGH.LEVEL.CMFIDSO1,ALTDSN=HIGH.LEVEL.CMFIDSO2
```

VSAM cluster allocation data is collected from volumes 304 and 306 in addition to the other required DSO head move data. HEADMOVE records are stored in the primary CMFIDSO1 data set until it becomes full. The records are stored in the secondary CMFIDSO2 data set after the primary data set becomes full. Head move data is collected at the default 33-millisecond interval.

```
HEADMOVE NUMBER=(2,3),VTOC=YES,DSN=HIGH.LEVEL.CMFIDS01,
         SAMPLE=50,PACE=60
```

Head move records are stored in the HIGH.LEVEL.CMFIDSO1 data set from all volumes with a unit address that begins with 2 or 3. Head move records are collected every 50 milliseconds over the Extractor sampling period. VTOC records are paced to a 3000-millisecond interval by taking the product of the PACE=60 and SAMPLE=50 parameters.

```
HEADMOVE NUMBER=(2,3),VTOC=YES,DSNLIST=(USER.HEADMOVE.COLL001,
USER.HEADMOVE.COLL002,USER.HEADMOVE.COLL003),
SAMPLE=50,PACE=60
```

Head move records are stored in USER.HEADMOVE.COLL001 until it fills up; then records are stored in USER.HEADMOVE.COLL002 until it fills up; and then records are stored in USER.HEADMOVE.COLL003. When the USER.HEADMOVE.COLL003 data set is filled, the Extractor writes head move data to USER.HEADMOVE.COLL001 and the subsequent data sets, writing over the previous data.

Data is collected from all volumes with a unit address that begins with 2 or 3. Head move records are collected every 50 milliseconds over the Extractor sampling period. VTOC records are paced to a 3000-millisecond interval by taking the product of the PACE=60 and SAMPLE=50 parameters.

HEADMOVE Data Set Definition Summary

When you do not want to send head move records and all sampling data to the CMF Extractor data sets defined on the CMF*x*PM*xx* DD statements, you can define up to 101 data sets to hold head move records.

Table 5-3 summarizes the four different parameters of the HEADMOVE control statement you use to define data sets to hold head move data.

Table 5-3 HEADMOVE Data Set Definitions

data set specified on	Used by CMF MONITOR Extractor when
DSNLIST	• You specify from 1 to 101 previously allocated, cataloged CMF Extractor data sets on the DSNLIST= keyword. • You do not use the DDNAME= or DSN= parameters of HEADMOVE. Example: **In HEADMOVE:** DSNLIST=(USER.HEADMOVE.COLL001, USER.HEADMOVE.COLL002, USER.HEADMOVE.COLL003) You ***must*** allocate the data sets in advance. Note: You do not need to modify the JCL, because the Extractor ignores any existing CMF*x*DS*xx* DD names and assigns its own DD names. If the CMF*x*DS*xx* are already allocated, the Extractor disables and reallocates the CMF*x*DS*xx* DD names for its use.
DDNAME	• You specify the DDNAME= keyword in HEADMOVE and do not use DSN= or DSNLIST= parameters. Example: **In HEADMOVE:** DDNAME=USERDD You ***must*** allocate the data set in advance. Note: You do not need to modify the JCL if the JCL DD name matches the DD name you assigned on the DDNAME= keyword of HEADMOVE.
DSN	You specify the DSN= keyword in HEADMOVE and do not use any DDNAME= or DSNLIST= parameters. Example: **In HEADMOVE:** DSN=(USER.HEADMOVE.COLL001) You ***must*** allocate the data set in advance. Note: You do not need to modify the CMF*x*DS*xx* DD statement in the JCL if the JCL data set name in the JCL matches the data set name assigned to the DSN= keyword of HEADMOVE.
ALTDSN	You use the ALTDSN= keyword and the DSN= keyword to define an alternate data set. Example: **In HEADMOVE:** ALTDSN=(USER.HEADMOVE.ALTDSN) DSN=(USER.HEADMOVE.COLL001) You ***must*** allocate the data set in advance. Note: You do not need to modify the CMF*x*DS*xx* DD statement in the JCL if the JCL data set name in the JCL matches the data set name assigned to the ALTDSN= keyword HEADMOVE.

DD Statements for HEADMOVE Data Sets

Instead of using HEADMOVE parameters to define data sets to hold head move data, you can code DD statements in the JCL using the CMF*x*DS*xx* format, where

x Contains C (CPM mode) or I (IPM mode)

xx Contains one or two legal characters (alpha, numeric, or a combination of both) you specify for 1 to 101 DD names

For example, the following DD statements specify two data sets:

```
CMFCDSAA DD DSN=USER.HEADMOVE.COLL001,DISP=SHR
CMFCDSBB DD DSN=USER.HEADMOVE.COLL002,DISP=SHR
```

You ***must*** allocate the data sets in advance.

Do not use the DDNAME=, DSN=, or DSNLIST= parameters of the HEADMOVE control statement when using this type of DD statement.

If there are any CMF*x*DS*xx* DD names in the JCL and you specify the DSNLIST= parameter, the Extractor ignores the DD names in the JCL.

Default CMF Data Set

The CMF Extractor sends head move data and all sampling data to the CMF data sets specified on the CMF*x*PM*xx* DD statements when either of the following is true:

- You do not use the DDNAME=, DSN=, or DSNLIST= parameters of the HEADMOVE control statement.
- You do not use any CMF*x*DS*xx* DD statements in the JCL.

REPORT

The REPORT control statement sets the global operating conditions that regulate how data is collected from the Extractor over the course of an entire sampling period. REPORT must be the first control statement in the JCL job to run the Extractor. Only one REPORT statement can be defined in a single job.

REPORT has parameters that regulate the

- Sampling mode of the Extractor
- Length of time that the Extractor collects data
- Recording interval between data collection cycles
- Type of records created from Extractor data
- Data sets that hold Extractor data

Note: A new parameter (DSNLIST=) has been added to enable you to define from 1 to 101 data sets to hold Extractor sampling data.

You can continue to use DSN= and ALTDSN= if you choose not to make REPORT control statement changes.

Another way you can define output data sets is by using the CMF*x*PM*xx* DD names in the JCL (instead of using the REPORT control statement for this purpose), where

x Contains C (CPM mode) or I (IPM mode)

xx Contains one or two legal characters (alpha, numeric, or a combination of both) you specify for 1–101 DD names

REPORT Syntax

```
REPORT  {CPM|IPM}
[,INTERVAL={30|nn|HOUR|HALF|QTR|SMF}]
[,RUNTIME={1440|nnnn}]
[,SYNCH={00|nn|HOUR|HALF|QTR|SMF}]
[,CSA={80|nnn}]
[,SMFRECID={240|nnn}]
[,DSNLIST=(xxx1,...,xxx101)]
[,DSN= ‘dsn’]
[,ALTDSN=‘dsn’]
[,SMF={NO|YES}]
[,DISP=NEW]
[,GBLS={YES|NO|1000....|9000}]
```

REPORT Parameters

CPM or IPM Specifies whether the control statements associated with an Extractor job are for CPM or IPM sampling. One of the positional keyword parameters, CPM or IPM, must be specified with every REPORT control statement.

INTERVAL= Specifies the length of time between recording intervals in minutes. The recording interval is the period when the CMF Extractor gathers data collected by the samplers and writes it to a data set.

The default recording interval is 30 and it can range from a minimum of 2 to a maximum of 60. The special keywords HOUR, HALF, and QTR force the recording interval on the hour, half hour, or quarter hour respectively.

The length of the recording interval can be synchronized with the SMF data set by specifying INTERVAL=SMF.

RUNTIME= Specifies the number of minutes that the Extractor sampling mode is active. The RUNTIME= parameter is intended primarily for the IPM sampling mode. The CPM sampling mode is meant to be run as long as possible to establish long-term trend data.

RUNTIME=1440 is the default length that an Extractor sampling mode remains active. Values can range from a minimum of 2 to a maximum default of 1440. If the specified interval is greater than the maximum, the runtime defaults to 1440. This time value specifies elapsed time, not central processing unit (CPU) time.

SYNCH= Specifies the starting minute of an hour that the CMF Extractor recording interval begins. The values associated with the SYNCH= parameter range from a minimum of 00 to a maximum of 60.

The special keywords HOUR, HALF, and QTR force the recording interval on the hour, half hour, or quarter hour respectively. The starting minute of the Extractor recording interval can be synchronized with the SMF data set by specifying SYNCH=SMF.

SYNCH= overrides the recording synchronization set by the INTERVAL= parameter.

CSA= Specifies the amount of extended common storage area (CSA) allocated to the CMF Extractor in kilobytes (K). The default is 80K and the maximum is 512K. The recommended amount of CSA storage for DSO is 180K.

SMFRECID= Identifies the SMF record format of the records written by the CMF Extractor. This parameter must be declared if CMF records are written to an SMF data set.

The default is 240. The CMF Extractor verifies that the identification number is between 128 and 255. This number must be unique from all other SMF record IDs used by other software within the sampled system.

Note: The SMF=YES parameter must be declared with the SMFRECID= parameter.

DSNLIST= Specifies from 1 to 101 previously allocated, cataloged CMF Extractor data set names to hold sampling data from the CMF Extractor.

Data is written to data sets in the order in which you specify them. When the last data set is full, the Extractor attempts to write to the first and subsequent data sets again. When all data sets are full, the writing of output is suspended.

You can clear a data set by executing the CMFOCLR procedure in batch mode or at the console.For more information on CMFOCLR, see the *CMF MONITOR Customization Guide*.

Do not use the CMF*x*PM*xx* DD names in the CMF Extractor JCL because the job fails if DD names are found.

Do not use the DSN= and ALTDSN= parameters of REPORT when using DSNLIST=.

DSN= Specifies the name of a previously allocated, cataloged data set that holds sampling data from the CMF Extractor. A default CMF Extractor data set is dynamically allocated if you do not specify a data set with the DSN= parameter.

If you use DSN=, do not include the default DD names CMFCPM*xx* or CMFIPM*xx* in the CMF Extractor JCL and do not use the DSNLIST= parameter.

When you do not use the DSN= or DSNLIST= parameters, the CMF Extractor sends sampling data to data sets specified on the CMF*x*PM*xx* DD statements.

ALTDSN= Specifies the name of an alternate, previously allocated and cataloged data set that holds sampling data written by the CMF Extractor if the data set specified by DSN= becomes full. If you use this option, the DD names CMFCPM*xx* and CMFIPM*xx* must not be included in the CMF Extractor JCL.

Note: The CMF Extractor automatically writes to the default DD names CMFCPM*xx* or CMFIPM*xx* when they are in the JCL.

SMF= Specifies whether you want data written to SMF data sets. If you specify SMF=YES, you must specify the SMFRECID parameter. SMF=NO is the default.

Note: The SMF=YES parameter is mutually exclusive from any data sets you specify with the DSN= or ALTDSN parameters. Do not use DSN= and ALTDSN= when you use the SMF=YES parameter.

DISP=NEW Specifies whether data is written to the beginning of an output data set or from the end of existing records when the Extractor opens the data set during the initialization phase. This parameter is not required when sampling data is written to SMF data sets.

The DISP=NEW parameter applies only during the initialization phase of the Extractor. After the primary data set becomes full and subsequent data is written to the alternate data set, the disposition automatically switches to MOD and DISP=NEW is no longer in effect.

GBLS= Specifies whether the global sampler collects data to construct the 240-11 records. The default is YES.

The sample rate for the global sampler can be controlled by specifying the number of milliseconds in multiples of 1000. The default sampling rate is 1000 expressed in the form GBLS=1000.

If NO is declared with the GBLS= parameter, the 240-11 records are not created. However, the global sampler continues to provide some sampling function under control of the DIE. Data collected under this condition is supplied to other Extractor samplers.

Examples

```
REPORT IPM,INTERVAL=QTR,RUNTIME=30,SMFRECID=187,
DSNLIST=(CMF.IPM.OUTPUT1,
         CMF.IPM.OUTPUT2,
         CMF.IPM.OUTPUT3)
```

Specifies intermittent sampling mode. SMF record 187 data is collected on the quarter hour over the 30-minute sampling period.

Data is written to CMF.IPM.OUTPUT1 until the data set is filled; then data is written to CMF.IPM.OUTPUT2 until the data set is filled; and then data is written to CMF.IPM.OUTPUT3 until the data set is filled.

When the last data set is filled, the Extractor checks the first and subsequent data sets in the list to see if any are available. When no more data sets are available, the writing of output is suspended. Clearing a previously used data set enables the Extractor to resume recording and reuse the cleared data set.

```
REPORT IPM,INTERVAL=QTR,RUNTIME=30,SMFRECID=187,
DSN='CMF.IPM.OUTPUT',ALTDSN='CMF.IPM.ALT.OUTPUT'
```

Specifies intermittent sampling mode. SMF record 187 data is collected on the quarter hour over the 30-minute sampling period. Data is written to ddname CMF.IPM.OUTPUT until the data set is completely filled and then is written to an alternate data set CMF.IPM.ALT.OUTPUT.

```
REPORT CPM,INTERVAL=30,SMFRECID=188
```

SMF 188 records are sampled continuously for the default length of 1440 minutes. The data is collected every 30 minutes and written to the default data set specified in the JCL on the CMF*x*PM*xx* DD statements. CSA storage is allocated for the default 80K.

```
REPORT IPM,RUNTIME=120,SMFRECID=199,SMF=YES
```

Intermittent sampling over 120 minutes creates SMF 199 records that are written to the SMF data sets. Data is collected at 30-minute intervals with 80K of the CSA area allocated.

Chapter 6 DSO Analyzer

The DSO Analyzer produces batch reports from extracted data. Up to nine different reports can be produced from the volume mapping and head move records stored on the Extractor data set. The Analyzer produces five of the reports by default. The remaining reports are requested when optional control statement parameters are specified in the JCL job to run the Analyzer. Chapter 7, “DSO Reports,”, discusses each report.

In addition, the Analyzer produces separate batch jobs containing JCL and control statements to reorganize volumes. These batch jobs contain the input code used by COMPAKTOR or DFDSS reorganization facilities. Both of these products make their volume reorganization in accordance with user-defined recommendations coded into the Analyzer input job. The recommendations that appear in the input jobs are specified with control statements in the batch job to run the Analyzer. Chapter 8, “DSO Output JCL,”, has separate sections that discuss the volume reorganization jobs produced for DFDSS and COMPAKTOR.

This chapter discusses the DSO Analyzer. An emphasis is placed on the six control statements included with the Analyzer job to produce batch reports and volume reorganization plans. Table 6-1 identifies the control statements and gives page references for further information. Each control statement is discussed in the remainder of this chapter.

Table 6-1 DSO Analyzer Control Statements

Control statement	Control statement function	Further information
ANUM	Specifies the run number that appears on the first page of an Analyzer report.	See “ANUM” on page 6-3
DATA	Specifies the ddnames of all data sets used as input for DSO reports when concatenation does not work.	See “DATA” on page 6-5
DATETIME	Declares the date and time range for the data used for head move analysis.	See “DATETIME” on page 6-8

Table 6-1 DSO Analyzer Control Statements

DEFAULT	Selects default analyzer control statement options and reports for all volumes.	See “DEFAULT” on page 6-11
HMTIME	Declares alternative head move formulas for specific types of DASD.	See “HMTIME” on page 6-18
VOLUME	Associates Analyzer processing options to a particular volume or device address range.	See “VOLUME” on page 6-21

DSO Analyzer Control Statements

You request DSO reports and reorganization plans by setting corresponding control statement parameters. Generally, each report or reorganization plan is associated with a control statement parameter on a one-to-one basis. Each control statement directs a program to run that selects, reads, and formats records for specific DSO output. The DSO Analyzer continues to run even when an individual program fails. If a program fails, a snap dump is taken, a diagnostic message is written, and the failing program is disabled for the remainder of the batch job.

DSO Analyzer reads Extractor records and reformats them on to a random access work file that is managed by the BMC Software access method, Data Management Subsystem (DMSS). The Analyzer selects records from the Extractor data set according to a user-defined date and time range. The records are chosen from the Extractor data set as they are read and formatted on to the DMSS work file.

During the selection and formatting process, the DSO Analyzer sorts the volume mapping and head move records for the requested reports. A sort package compatible with DFSORT is required to run the DSO Analyzer. Input records cannot be presorted, and only the records required to produce the reports are read, formatted, and sorted.

ANUM

ANUM is an optional DSO Analyzer control statement. It specifies the analysis run number, the number of volumes that appear in the reports, and an upper threshold value for any condition codes created by errors during processing.

Only one ANUM control statement is permitted in a set of DSO Analyzer control statements that generate reports for a particular batch job.

ANUM Syntax

```
ANUM xxxx
[,MAXVOL=nnn]
[,MAXCC={04 |nn}]
```

ANUM Parameters

xxxx Specifies the analysis run number for a set of reports. The analysis run number is a one- to four-character alphanumeric string that appears on the upper right corner of the Extractor Data Sets report. An example of an analysis run number is shown in Figure 6-1 on page 6-4.

MAXVOL= Limits Analyzer reports to the first *nnn* volumes held in the Extractor data set. Up to 255 volumes can be declared with the MAXVOL= parameter.

The MAXVOL= parameter should not be used with the VOLUME control statement or the MAXVOL parameter of the PARM field of a JCL EXEC statement. The MAXVOL= parameter is superseded by any parameters of the VOLUME control statement or EXEC PARM fields that declare the volumes for DSO reports.

MAXCC= Sets an upper threshold for condition codes that result from errors during Analyzer processing. Severity errors are ignored if their condition codes are less than or equal to the MAXCC= parameter. The default is 04 and the maximum condition code threshold is 15.

The MAXCC= parameter applies only if the similar MAXCC= parameter of the PARM field of the JCL EXEC statement is not coded.

Note: The MAXCC= parameter of the ANUM control statement is interpreted after the PARM field version of the same parameter. It cannot suppress the abnormal end to Analyzer processing caused by errors that occur before ANUM is interpreted.

Example

```
ANUM  TST1,MAXVOL=25,MAXCC=08
```

This example limits DSO reports to the first 25 volumes found in the Extractor data set. Any errors that occur as the reports are built are ignored if their condition code is less than or equal to 08. The analysis run number, TST1, appears at the upper right corner of the Extractor Data Sets report shown in Figure 6-1

Figure 6-1 Example of an Analysis Run Number with the ANUM Control Statement

```
                                D A T A   S E T   O P T I M I Z E R
PRODUCED BY DSO R2.0.0                                                                  PAGE 2
BMC SOFTWARE, INC.                                                    REPORT DATE   6 JUN 03
                                                                      REPORT TIME   14.45.55

                                                                    ANALYSIS RUN NUMBER-TST1
```

DATA

The DATA control statement specifies the ddnames of all Extractor data sets when concatenation does not work because of a variety of recording or screening requirements that apply to the various data sets. The DATA control statement requests input Extractor records by SMF record types, sampling mode (CPM and IPM), and system ID(s).

The DATA control statement can also specify that data from certain ddnames be used only to obtain volume mapping or head move data. Each name is any valid ddname of one to eight characters, except for any ddnames reserved for other purposes by the DSO Analyzer.

The DATA control statement can associate a ddname or list of ddnames with a unique DATETIME control statement. DATETIME immediately follows the DATA statement in the JCL job stream when they are associated. The date-time range expressed in the DATETIME statement applies only to the records read from ddnames listed in the preceding DATA statement. An example of the DATA and DATETIME control statements is shown in “Restricting Input Data with the DATETIME Control Statement” on page 3-8.

DATA Syntax

```
DATA
[RECTYPE={ 240 |nnn}]
[,MONITOR={IPM|CPM|ALL}]
[,SYSTEM={ NOID |(system1_ID,...system4_ID)}]
[,(name1,{VOLMAP|NOVOLMAP},...name25,{VOLMAP|NOVOLMAP})]
```

DATA Parameters

RECTYPE= Specifies the SMF record number of input Extractor records accepted from the ddname or ddname list. The default record number is 240, which matches the default record number of the CMF Extractor.

MONITOR= Restricts input data to the records collected by a particular Extractor sampling mode (CPM or IPM). If ALL is specified, both CPM and IPM records are accepted for analysis.

SYSTEM= Declares the one- to four-character name of the system whose data is held in an Extractor data set. The system ID is the same as the SMF ID. Records from four systems can be used to generate DSO reports.

If records are held in the Extractor data set without a system ID, the operand NOID should be used to select them. If SYSTEM is not specified with the DATA control statement, NOID is the default and only records without a system ID are used for analysis.

name1,... Specifies the ddname that identifies the input data set holding Extractor records. The default ddname of CMFIN is selected if no ddname is specified with the DATA control statement. Up to 25 ddnames can be declared with the DATA control statement.

The RECTYPE, MONITOR, and SYSTEM parameters apply to every ddname in the list. Individual VOLMAP or NOVOLMAP operands can be specified with each ddname to restrict the type of records processed by the Analyzer.

By default, volume mapping and head move records are accepted from the data sets specified by the ddname. If there is only one input ddname, the VOLMAP and NOVOLMAP operands are not applicable. The data set(s) for that ddname must contain volume mapping and head move records.

If VOLMAP is specified with the ddname, only the VTOC and VSAM records are used from the data set. Head move records are ignored.

If NOVOLMAP is specified with the ddname, only head move records are used from the data set. VTOC and VSAM records are ignored. NOVOLMAP can reduce the processing time required to correlate all of the volume mapping by omitting redundant records.

Note: Use the NOVOLMAP operand carefully. If it is not associated with the correct ddname, required records can be omitted, causing invalid reports or an Analyzer failure.

Adhere to the following rules when using the NOVOLMAP operand with the ddnames specified with the DATA control statement:

- The most recent volume mapping records should be used for any volume that is going to be reorganized. Do not code NOVOLMAP with the ddname containing the most recent volume mapping records.

- The oldest and newest volume mapping records should be used to create any DSO reports. Head move records are accepted by the Analyzer only if the data set extent locations are recognized in the current volume mapping records. The Analyzer must have the oldest map of the volume to count all the head moves to existing data set extents if the records are split across several data sets.

When concatenating records from several extractions, DSO Analyzer processing time can be reduced by suppressing all volume maps except the oldest and most recent. This technique produces valid results if the contents of the volume remain fairly constant.

If a volume contains a few permanent data sets with a larger number of transient data sets, this technique can filter out the activity from the short-lived data sets. The volume reorganization considers head moves that occur between the permanent data sets and ignores the temporary data sets.

Examples

```
DATA INPUT1,MONITOR=IPM,SYSTEM=SYSB
```

Extractor records from data sets specified by ddname INPUT1 are processed by the Analyzer. The selected records are created by IPM sampling mode from the system with the ID of SYSB. The records are in the default SMF 240 format.

```
DATA
RECTYPE=240,MONITOR=IPM,SYSTEM=(SYSA,SYSB),(CMFIN1,CMF
IN2)
```

Extractor records from data sets specified by ddnames CMFIN1 and CMFIN2 are processed by the DSO Analyzer. The selected records are created by IPM sampling mode from systems with the IDs of SYSA and SYSB. Only records with a SMF 240 record number are acceptable for input to the Analyzer.

```
DATA
MONITOR=IPM,SYSTEM=SYSA,(INPUT1,VOLMAP,INPUT2,NOVOLM
AP)
```

Extractor IPM records are processed from system SYSA. Head move records are ignored from the input data sets declared with the INPUT1 ddname. Volume mapping records are ignored from the data sets declared with the INPUT2 ddname.

DATETIME

DATETIME is an optional DSO control statement that specifies the overall date and time range of data included in DSO reports. A beginning date and time is paired with an ending date and time to define a range when Extractor records are created. An Extractor record is selected from the data set if the date and time of its recording interval falls within the DATETIME range.

When the DATETIME statement is interpreted, the DSO Analyzer initially looks for records that fall within the specified range. If the DATETIME statement is omitted, the Analyzer reads the entire input data set. In either case, the Analyzer always reads to the end of file on the input data set.

The DATETIME range is an overall range and should be defined to encompass the entire time period for any requested DSO reports. The time range for a specific report is not required to span the entire DATETIME range, but the records for DSO reports must be collected within the range.

The date and time ranges declared with the DATETIME control statement are expressed with JULIAN and SELECT parameters. If the JULIAN and SELECT parameters are specified together, the resulting date/time range is the logical intersection of both of the declared ranges. If the two ranges do not intersect, no records are accepted for Analyzer processing.

DATETIME Syntax

```
DATETIME
[JULIAN=(sdate_1:stime_1,edate_1:etime_1,...sdate_n:stime_n,edate_n:etime_n)]
[,SELECT=(stime_1,etime1,day1,...stime_n,etime_n,day_n)]
```

DATETIME Parameters

JULIAN= Specifies the beginning and ending date-time pairs in Julian format. If multiple pairs are specified with the JULIAN= parameter, they must appear in ascending date-time sequence and cannot overlap.

Julian date and time is expressed in the following manner:

JULIAN=yyddd:hhmmss

where

yy	is the year; represented by the last two digits of any year in this century
ddd	is the day; expressed as the number of the day in the year and must be represented as a three-digit number from 001 to 366
hh	is the hours; expressed as a two-digit number from 00 to 23
mm	is the minutes; expressed as a two-digit number from 00 to 59
ss	is the seconds; expressed as a two-digit number from 00 to 59

SELECT= Specifies a time interval within a day or a range of days. The interval is defined by a starting time, ending time, and day. More than one period can be declared from a single SELECT= parameter. The time and day for each period can overlap with another period and they can appear in any order within the SELECT= parameter.

Time and day are expressed in the form, hhmmss,hhmmss,day, where

hh	Are the hours; expressed as a two-digit number from 00 to 23
mm	Is the minutes; expressed as a two-digit number from 00 to 59
ss	Are the seconds; expressed as a two-digit number from 00 to 59
day	Is the day; expressed as a three or four-letter contraction of the name of the day. Similar contractions describe the weekend (WKEND) and the work days (WKDAY). The Boolean NOT form of the day is expressed with a NO prefix to the day name as follows:

MON	NOMON	TUE	NOTUE
WED	NOWED	THUR	NOTHUR
FRI	NOFRI	SAT	NOSAT
SUN	NOSUN	WKDAY	WKEND

Examples

```
DATETIME  JULIAN=(03032:000000,03038:235959)
```

Extractor records are accepted from any hour of the first seven days of February 2003.

```
DATETIME  JULIAN=(03001:000000,03031:235959),
SELECT=(113000,133000,MON,090000,130000,NOMON)
```

This DATETIME statement accepts records from the first 31 days of 2003. Records generated on Monday should be accepted if they are recorded between 11:30am and 1:30pm. Records are accepted from any other day if recorded between 9:00am and 1:00pm.

```
DATETIME SELECT=(100000,113000,WKDAY,133000,150000,WKDAY)
```

This DATETIME statement accepts records from any date if they are created on a weekday between 10:00am and 11:30am or between 1:30pm and 3:00pm.

DEFAULT

The DEFAULT control statement selects Analyzer processing options and reports for all sampled volumes. The VOLUME control statement sets processing options for individual volumes or a range of volumes by unit address. Only one DEFAULT control statement is permitted in the JCL job stream and it must precede all VOLUME control statements.

The parameters of the DEFAULT control statement have the same format and meaning as the corresponding parameters of the VOLUME control statement. VOLUME parameters supersede equivalent DEFAULT parameters if both control statements are included in the Analyzer JCL job stream. The DEFAULT parameters are in effect for all volumes unless explicitly overridden by VOLUME parameters. If a parameter is not declared with either control statement, the default value of the parameter is in effect.

DEFAULT Syntax

```
DEFAULT [AUTOFIX]
 [,TYPE=device type]
 [,OPTIONS=([{ REORG |NOREORG}][,{ HM |NOHM}]
          [,{ NOCOMPAK |COMPAK}][,{ MAP |NOMAP}]
          [,{ NODFDSS |DFDSS}]
          [,{ VOLSTUDY |NOVOLSTUDY}]
          [,{ ACTIVE |NOACTIVE}]
          [,{ CENTER |NOCENTER}][,{ MOVE |NOMOVE})]
 [,ISAM={MOVE |NOMOVE}]
 [,VSAM={MOVE |NOMOVE}]
 [,VVDS={MOVE |NOMOVE}]
 [,VTOCLOC=cccchh|FIXED]
 [,VTOCSIZE=yyy]
 [,XTNTPRT={HEX |DEC}]
 [,ACSORT=([NAME ][,BUSY][,ERROR][,HM])]
 [,HMSORT=([HM |[,HMTIME])]
 [,NAMETH=({LT|GT|EQ},nnnnn)]
 [,ERRORTH=({LT|GT|EQ},xx.xx)]
 [,HEADMVTH=({LT|GT|EQ},xx.xx)]
 [,BUSYTH=({LT|GT|EQ},xx.xx)]
 [,HMTIMETH=({LT|GT|EQ},xx.xx)]
 [,HDMOVETH=({LT|GT|EQ},xx.xx)]
```

DEFAULT Parameters

AUTOFIX Specifies that the Analyzer can automatically place the most active data sets on fixed-head cylinders in the proposed volume reorganization. If specified, AUTOFIX is a positional keyword parameter that must follow immediately after the DEFAULT control statement.

TYPE= Four-digit number that identifies the DASD type according to IBM device nomenclature. An alphanumeric suffix identifies particular models of a DASD type. DSO automatically furnishes the correct device type if the TYPE= parameter is not specified with the DEFAULT control statement. DSO supports the following DASD types:

3330-1	3330-11	3340-35
3340-70	3340-70F	3344
3344F	3350	3350F
3375	3380	3380D
3380E	3380J	3380K
3390-1	3390-2	3390-3
3390-1C	3390-2C	3390-3C

Note: The ending character F identifies a fixed-head model. The ending character C identifies 3390 DASD models designed to run in 3380 track compatibility mode.

DSO results are not reliable if the DASD associated with the TYPE= parameter is different from the actual DASD type of the volume being studied. The 3340 and 3350 DASD types are exceptions to this rule. Fixed-heads can be specified with the TYPE= parameter for the models of both types with movable heads.

OPTIONS= Declares the type of reports and volume reorganization produced by the Analyzer from the Extractor records. The following reports and volume reorganization options are declared with the OPTIONS= parameter. The underlined items in the two lists are the default values for the OPTIONS= parameter.

Analyzer report options	
REORG|NOREORG	Option default produces the Reorganization of Volume report.
HM|NOHM	Option default produces the Head Movements by Data Set Pairs reports.
MAP|NOMAP	Option default produces the Mapping for Volume report.
VOLSTUDY|NOVOLSTUDY	Option default produces the Volume Study report.
ACTIVE|NOACTIVE	Option default produces the Data Set Activity reports.

Analyzer reorganization options	
NOCOMPAK| COMPAK	Option default prohibits COMPAKTOR control statements from being produced. A COMPAKTR DD statement must be included in the Analyzer JCL if OPTIONS=COMPAK is declared. If OPTIONS=COMPAK is declared, the DSORG=U, ABSTR, and ISAM data sets become unmovable in the COMPAKTOR reorganization plan by forcing the OPTIONS=NOMOVE and ISAM=NOMOVE parameters. Note : Use COMPAK or NOCOMPAK for FDR COMPAKTOR Release 5.0 and earlier.
NOCOMPAK50| COMPAK50	Option default prohibits COMPAKTOR control statements from being produced. A COMPAKTR DD statement must be included in the Analyzer JCL if OPTIONS=COMPAK50 is declared. If OPTIONS=COMPAK50 is declared, the DSORG=U, ABSTR, and ISAM data sets become unmovable in the COMPAKTOR reorganization plan by forcing the OPTIONS=NOMOVE and ISAM=NOMOVE parameters. Note: Use COMPAK50 or NOCOMPAK50 for FDR COMPAKTOR Release 5.2 or later.
NODFDSS| DFDSS	Option default prohibits DFDSS control statements from being produced. A DFDSS DD statement must be included in the Analyzer JCL if OPTIONS=DFDSS is declared. If OPTIONS=DFDSS is declared, the DSORG=U, ABSTR, and ISAM data sets become unmovable in the DFDSS reorganization plan by forcing the OPTIONS=NOMOVE and ISAM=NOMOVE parameters.
CENTER| NOCENTER	Option default centers the VTOC in the reorganization plan of the volume.
MOVE| NOMOVE	Option default moves the DSORG=U and ABSTR data sets to the optimal position on the volume in the reorganization plan.

ISAM= Controls the repositioning of indexed sequential access method (ISAM) data sets in the proposed reorganization.

MOVE Causes the Analyzer to assign optimal placement to ISAM data sets.

NOMOVE Forces the Analyzer to treat ISAM data sets as unmovable.

VSAM= Controls the repositioning of VSAM data sets in the proposed reorganization.

MOVE Permits the Analyzer to move VSAM data sets to the optimal cylinder position in the reorganization plan.

NOMOVE Forces the Analyzer to treat VSAM data sets as unmovable.

VVDS= Controls the repositioning of VVDS data sets in the proposed reorganization.

MOVE Permits the Analyzer to move VVDS data sets to the optimal cylinder position in the reorganization plan.

NOMOVE Forces the Analyzer to treat VVDS data sets as unmovable.

VTOCLOC= Controls the repositioning of the VTOC in the proposed reorganization.

ccchh Forces the proposed placement of the VTOC to a specific cylinder and track address (decimal).

FIXED Forces the Analyzer to treat the VTOC as unmovable.

Warning! If VTOCLOC=FIXED is specified, it can significantly limit any performance improvements resulting from a volume reorganization. The reduction in potential performance improvement is attributable to the fact that the reorganization algorithms do not consider fixed data sets in the reorganization plan. Performance improvement is limited to the reduction in seek times by placing high activity movable data sets near immovable data sets.
If non-ICF VSAM data spaces with nonzero secondary allocation quantities exist on the volume, the COMPAKTR control cards must be manually edited. COMPAKTR does not honor changes to the position or size of the VTOC in the reorganization job.

VTOCSIZE= Allocates the number of tracks for the VTOC in the proposed reorganization of the volume. No warning message is produced if the track size is insufficient to hold the contents of the VTOC.

XTNTPRT= Specifies the numbering system to express data set extent boundaries in all DSO reports. The default, XTNTPRT=HEX, shows data set extent boundaries with hexadecimal numbers. Data set boundaries can be expressed with decimal numbers with the XTNPRT=DEC option.

ACSORT= Declares the type of Data Set Activity reports produced by the Analyzer. Up to four different types of Data Set Activity reports can be produced from the same data. Each report lists the data sets on the volume by one of the following parameters:

NAME Data sets are listed in alphanumeric order. This is the default report.

BUSY Data sets are listed by their percentage of the total busy time of the volume in the recording interval.

ERROR Data sets are listed by their percentage of the total time spent by the volume in error recovery.

HM Data sets are listed by their percentage of the total recorded head moves.

Thresholds can be set with the ACSORT= parameter to filter out unwanted data sets from the Data Set Activity reports. The Data Set Activity reports can be kept to a manageable size by restricting the number of significant data sets with a threshold value. Refer to “Threshold Parameters” for the parameters that are associated with the ACSORT= parameter.

HMSORT= Declares the type of Head Movement by Data Set Pairs reports. Two types of reports can be produced.

HM Data set pairs are listed by their percentage of the total number of head moves recorded in the sampling interval. This is the default report.

HMTIME Data set pairs are listed by their percentage of the total time spent moving heads during the sampling interval.

Thresholds can be set with the HMSORT= parameter to filter out unwanted data sets from the Head Movement by Data Set Pairs reports. Refer to “Threshold Parameters” for more information about setting associated parameters with the HMSORT= parameter.

Threshold Parameters

The DEFAULT and VOLUME control statements provide six parameters to regulate the number of data sets that appear in the Data Set Activity and Head Movement by Data Set Pairs reports. These six parameters set thresholds when paired with the ACSORT= and HMSORT= parameters. Data Sets that do not meet the threshold criteria are excluded from the reports.

The parameters consist of a comparison operator and the threshold value. Comparisons are made on the basis of the test criterion being less than (LT), equal to (EQ), or greater than (GT) the threshold. Data sets are included in reports if the tested criterion is true when compared to the threshold value.

The parameters are expressed in the form (comparison operator,threshold value), where

comparison operator	Is LT, EQ, or GT
threshold value	Is a percentage with two significant figures xx.xx Is an alphanumeric string of five or less characters

The following parameters set thresholds when they are associated with the ACSORT and HMSORT parameters:

NAMETH= Sets a threshold for the Data Set Activity report sorted by data set name. Data sets are included in the report if the data set name is true when compared to an alphanumeric string.

ERRORTH= Sets a threshold for the Data Set Activity report sorted by the error recovery percentage of each data set. Data sets are included in the report if the test is true when compared to a percentage threshold.

HEADMVTH= Sets a threshold for the Data Set Activity report sorted by the total head move percentage. Data sets are included in the report if the test condition is true when compared to a percentage threshold.

BUSYTH= Sets a threshold for the Data Set Activity report sorted by the percentage of total busy time. Data sets are included in the report if the test condition is true when compared to a percentage threshold.

HMTIMETH= Sets a threshold for the Head Movement by Data Set Pairs report sorted by the percentage of total head move time. Data sets are included in the report if the test condition is true when compared to a percentage threshold.

HDMOVETH= Sets a threshold for the Head Movement by Data Set Pairs report sorted by the percentage of total head moves. Data sets are included in the report if the test condition is true when compared to a percentage threshold.

Examples

```
DEFAULT   AUTOFIX,TYPE=3350F,OPTIONS=(DFDSS,NOCENTER)
```

The default Analyzer reports are created for 3350 DASD with fixed-heads. In addition, a DFDSS reorganization job is created with a requirement that the VTOC not be centered and that the most active data sets be placed on the fixed-head cylinders.

```
DEFAULT  OPTIONS=(NOMAP,REORG,NOACTIVE,COMPAK,NOCENTER),
          HMSORT=(HMTIME),HMTIMETH=(GT,10.00)
```

All default Analyzer reports are created except the Mapping for Volume and Data Set Activity reports.

A COMPAKTOR output job is created with a requirement that volume reorganization does not center the VTOC. By default, the ISAM, DSORG=U, and ABSTR data sets remain unmovable in the COMPAKTOR reorganization job.

A Head Movement by Data Set Pairs report is created that is sorted by the percentage of total head move time. It includes data sets with head move counts greater than 10 percent. All other data sets are omitted from the report.

```
DEFAULT   TYPE=3390-3,OPTIONS=(NOMAP,ACTIVE),
ACSORT=(NAME,BUSY),BUSYTH=(GT,00.05),
HMSORT=(HM),HMTIMETH=(GT,01.00)
```

Data Set Activity reports are created for Model 3 3390 DASD. The Mapping for Volume report is excluded. The Data Set Activity reports are sorted by name and total busy time. Only data sets with a busy percentage greater than 0.05 percent are included in the report that sorts data sets by total busy time.

A Head Movement by Data Set Pairs report includes only data sets with head move counts greater than 1.00 percent of the total.

HMTIME

HMTIME is an optional control statement that explicitly overrides the default head move formulas used to calculate seek time. The HMTIME control statement is most commonly used for plug-compatible DASD with different seek characteristics than the default IBM DASD types. The DSO Analyzer uses the seek time ranges given with the HMTIME control statement to calculate all head move values that appear in the reports.

The CMF Extractor does not make direct measurements of DASD head move times. Seek addresses are recorded that can be interpreted for the total distance of the head move based upon the difference between the destination cylinder and the current cylinder location.

The DSO Analyzer uses a series of formulas to calculate seek time based upon the cylinder distance of the head move and the type of DASD. The default formulas are based upon the seek performance of IBM DASD types. The formulas include the time to move one cylinder and the time to move the maximum number of cylinders for that DASD type. The seek times for all cylinder distances that are not explicitly stated are determined by linear interpolation.

Table 6-2 shows the default head move formulas for each type of IBM DASD. The formulas can be used to compute device seek times for any of the listed DASD types or any new devices defined with an Extractor DEVICE control statement. If the HMTIME control statement is omitted from the Analyzer job, seek times are computed with the formulas shown in Table 6-2.

Table 6-2 Formulas to Compute DASD Seek Times

IBM device type	Default formulas
3330-1	HMTIME TYPE=3330-1,(1,8),(17,16),(102,26),(240,36),(410,48)
3330 model 11	HMTIME TYPE=3330-11,(1,8),(34,16),(204,26),(480,36),(814,48)
3340 model 35	HMTIME TYPE=3340-35,(1,10),(116,25),(347,50)
3340 model 70 or 70F	HMTIME TYPE=3340-70(F),(1,10),(232,25),(695,50)
3344 or 3344F	HMTIME TYPE=3344(F),(1,10),(928,25),(2783,50)
3350 or 3350F	HMTIME TYPE=3350(F),(1,10),(200,25),(558,50)
3375	HMTIME TYPE=3375,(1,4),(480,19),(959,38)
3380A or B	HMTIME TYPE=3380,(1,3),(295,16),(885,30)
3380D	HMTIME TYPE=3380D,(1,3),(295,15),(885,28)
3380E	HMTIME TYPE=3380E,(1,3),(590,17),(1770,31)
3380J	HMTIME TYPE=3380J,(1,2),(295,12),(885,21)
3380K	HMTIME TYPE=3380K,(1,2),(885,16),(2655,29)
3390-1	HMTIME TYPE=3390-1,(1,2),(371,10),(1113,18)
3390-1C	HMTIME TYPE=3390-1C,(1,2),(371,10),(1113,18)
3390-2	HMTIME TYPE=3390-2,(1,2),(742,13),(2226,23)
3390-2C	HMTIME TYPE=3390-2C,(1,2),(742,13),(2226,23)

Table 6-2 Formulas to Compute DASD Seek Times

IBM device type	Default formulas
3390-3	HMTIME TYPE=3390-3,(1,2),(1113,15),(3339,33)
3390-3C	HMTIME TYPE=3390-3C,(1,2),(1113,15),(3339,33)
9345-1	HMTIME TYPE=9345-1,(1,2),(480,10),(1440,16)
9345-2	HMTIME TYPE=9345-2,(1,2),(719,11),(2155,20)

Note: In 3390 DASD formulas, the minimum and average seek times have been rounded up by 0.5 milliseconds due to DSO restrictions on floating point numbers.

HMTIME Syntax

HMTIME TYPE= device_type[,(x1,y1)] [,(xn,yn)]

HMTIME Parameters

TYPE= Identifies the type of DASD by a four-digit number and a suffix that identifies a particular model of the DASD type. The four-digit number agrees with IBM DASD nomenclature. Supported DASD types are shown in Table 6-2 on page 6-18.

(x1,y1) Represents the time required to move one cylinder in milliseconds. The value x1 defaults to 1 and the y1 value is the minimum seek time for the DASD type.

(xn,yn) Represents the time required to move the maximum cylinder distance in milliseconds. The value xn represents the maximum number of cylinders for the DASD type. The value yn represents the time to move the heads over the maximum cylinder distance in milliseconds.

The (x1,y1) and (xn,yn) pairs represent the end point coordinates of a distribution of cylinder distances to seek times. Additional pairs can be specified between the end points. The (x,y) pairs must be placed in ascending order from the smallest cylinder distance to the largest with the HMTIME control statement.

It is not necessary to include the alternate cylinders in the maximum cylinder (x,y) pair for 33xx DASD. Alternate cylinders are not observed by the CMF ExtractorCMF Extractor.

Examples

```
HMTIME TYPE=3380K,(1,2),(882,14),(2655,24)
```

The default head move formula is overridden for 3380K volumes. The minimum seek time is 2 milliseconds and the maximum seek time is 24 milliseconds. An intermediate distance of 882 cylinders is included in the formula with a seek time of 14 milliseconds.

```
HMTIME TYPE=3390-2,(1,2),(1000,11),(1500,18),(2226,27)
```

The default head move formula is overridden for Model 2 3390 volumes. The minimum seek time is 2 milliseconds and the maximum seek time is 27 milliseconds. Intermediate seek times are included for 1000 and 1500 cylinder head moves.

```
HMTIME TYPE=3380K,(1,2),(1000,11),(1500,16),(2655,23)
```

The seek characteristics of Amdahl 6380 model K DASD have been substituted into the formula for the equivalent IBM 3380K default DASD type. DSO recognizes the 6380K as a 3380K device, but computes seek times with the correct values given in the substituted formula.

VOLUME

The VOLUME control statement associates Analyzer processing options with a specific volume or unit address range. The parameters specified with the VOLUME control statement override similar parameters specified with the DEFAULT control statement. All VOLUME control statements must follow the DEFAULT control statement in the Analyzer JCL job stream.

If no VOLUME control statements are present, default parameters are generated for the first *nnn* volumes found in the input data set. Up to 255 VOLUME control statements can be included with each Analyzer batch job.

The syntax of the VOLUME control statement is the same as the DEFAULT control statement with many equivalent parameters. Only parameters that are unique to the VOLUME control statement are described in this section. The parameters that are shared between the two control statements are described in “DEFAULT” on page 6-11.

VOLUME Syntax

```
VOLUME [AUTOFIX]
[,TYPE=device type]
[,DEVTYPE=xxxxxxxx]
{[,NAME=volume]|[,UNIT=(cuu1[,cuu2])]}
[,FIXED=(dsn1,dsn2,...dsnn)]
[,OPTIONS=([{REORG |NOREORG}][,{HM|NOHM}]
                [,{NOCOMPAK|COMPAK}][,{MAP|NOMAP}]
                [,{NODFDSS|DFDSS}]
                [,{VOLSTUDY|NOVOLSTUDY}]
                [,{ACTIVE|NOACTIVE}]
                [,{CENTER|NOCENTER}][,{ MOVE|NOMOVE})]
 [,ISAM={MOVE|NOMOVE}]
 [,VSAM={MOVE|NOMOVE}]
 [,VVDS={MOVE|NOMOVE}]
 [,VTOCLOC=cccchh|FIXED]
 [,VTOCSIZE=yyy]
 [,XTNTPRT={HEX|DEC}]
 [,ACSORT=([NAME][,BUSY][,ERROR][,HM])]
 [,HMSORT=([HM|[,HMTIME])]
 [,NAMETH=({LT|GT|EQ},nnnnn)]
 [,ERRORTH=({LT|GT|EQ},xx.xx)]
 [,HEADMVTH=({LT|GT|EQ},xx.xx)]
 [,BUSYTH=({LT|GT|EQ},xx.xx)]
 [,HMTIMETH=({LT|GT|EQ},xx.xx)]
 [,HDMOVETH=({LT|GT|EQ},xx.xx)]
```

VOLUME Parameters

DEVTYPE= Declares nonstandard UCB device types in an installation. The alternate path bit is ignored when matching this value to the contents of the CMF Extractor output file. It defaults to the DASD device associated with the TYPE= parameter.

NAME= Declares the volume whose records are processed by the Analyzer. The volume is identified by volume serial name.

UNIT= Specifies a volume for study by a three-digit unit address or a sequence of volumes by a range of unit addresses (*xxx*,*xxx*). Analyzer reports are generated for the first volume encountered in the input data set with the specified unit address. The UNIT= parameter should be used carefully with input data containing extractions from several systems. Invalid reports are created if a volume is represented by a different device address on different systems.

FIXED= Requests a specific data set be placed on fixed-head cylinders in the planned volume reorganization. A warning message is issued if the volume does not have fixed-heads or there is insufficient cylinder space. No embedded blanks are allowed within this parameter.

Table 6-3 shows the use of the FIXED= and AUTOFIX parameters in reorganizing a fixed-head volume.

Table 6-3 Analyzer Action of the FIXED and AUTOFIX Parameters

Parameters	Analyzer action
FIXED	Data sets specified with the FIXED parameter are placed on fixed- head cylinders. The entire data set must fit within a single fixed-head cylinder.
AUTOFIX	The Analyzer selects data sets with the highest number of head moves to reside on fixed-head cylinders if the volume has fixed-heads.
FIXED and AUTOFIX	After allocating the data sets specified by the FIXED= parameter to fixed-head cylinders, the Analyzer selects high activity data sets to fill any remaining free space on fixed-head cylinders.
Neither FIXED nor AUTOFIX	The original data sets remain on fixed-head cylinders if individual data sets can fit within a single cylinder.

Examples

```
VOLUME NAME=ABC012
```

A set of default Analyzer reports are produced for volume ABC012. The default reports are shown with underlined parameters in the Selection Option column of Table 7-1 on page 7-1.

```
VOLUME NAME=SYSRES,ACSORT=(HM,NAME,BUSY,ERROR),HM=(HM,HMTIME)
```

A complete set of Analyzer reports are produced for volume SYSRES. For the complete set of reports, see Table 7-1 on page 7-1.

```
VOLUME UNIT=12A,ISAM=MOVE,VSAM=MOVE,VVDS=MOVE,XTNPRT=DEC
```

The default reports are produced for the volumes with unit address 12A. The reports show data set extent boundaries with decimal numbers. The ISAM, VSAM, and VVDS catalogs can be moved in the proposed volume reorganization.

```
VOLUME UNIT=(12A,12F),ISAM=MOVE,VSAM=MOVE,VVDS=MOVE,XTNPRT=DEC
```

This example produces the same reports and reorganization plan as the previous example for all volumes in the unit address range of 12A to 12F.

Chapter 7 DSO Reports

The DSO Analyzer produces a series of reports from the head move and volume mapping records stored in the Extractor or SMF data sets. Each report presents a unique perspective of the head move activity that occurred during the sampling period. These reports can be printed after the successful completion of the batch job to run the DSO Analyzer.

This chapter describes each DSO report and concludes with instructions for accessing online Messages and Codes. Several preliminary reports are produced with the main DSO reports. These preliminary reports appear together in “Preliminary DSO Reports” on page 7-3 and list the input data sets and Analyzer control statements used to generate the main DSO reports.

Table 7-1 shows the main DSO reports. The nine reports are listed in the order in which they are printed. The report number shown in the left column of Table 7-1 appears with each report and in an index printed with a series of reports if the INDEX DD statement is coded in the Analyzer JCL job. The index lists the page number for every report by title and report number.

Table 7-1 DSO Reports

Report number	Report title	Selection option	Further information
1	Study Report for Volume	OPTIONS=VOLSTUDY	See “Study Report for Volume” on page 7-5
2	Mapping for Volume	OPTIONS=MAP	See “Mapping for Volume Report” on page 7-7
3	Head Movement Between Pairs of Data Sets by Time on Volume	HMSORT=HMTIME	See “Head Movement Between Pairs of Data Sets by Time on Volume Report” on page 7-10
4	Head Movement Between Pairs of Data Sets by Count on Volume	HMSORT=HM	See “Head Movement Between Pairs of Data Sets by Count on Volume Report” on page 7-13
5	Reorganization of Volume	OPTIONS=REORG	See “Reorganization of Volume Report” on page 7-14

Table 7-1 DSO Reports

Report number	Report title	Selection option	Further information
6	Data Set Activity Report by Name on Volume	ACSORT=NAME	See “Data Set Activity Reports” on page 7-19
7	Data Set Activity Report by Busy on Volume	ACSORT=BUSY	See “Data Set Activity Reports” on page 7-19 for an example of an equivalent report
8	Data Set Activity Report by Error Recovery on Volume	ACSORT=ERROR	See “Data Set Activity Reports” on page 7-19 for an example of an equivalent report
9	Data Set Activity Report by Head Movement on Volume	ACSORT=HM	See “Data Set Activity Reports” on page 7-19 for an example of an equivalent report

Some reports are produced automatically; others are requested when you code specific Analyzer control statement parameters. The reports with the underlined parameters in the Selection Option column of Table 7-1 on page 7-1 are produced by default.

The remaining reports must have the parameters shown in the Selection Option column coded in the Analyzer control statements before they are produced. Chapter 6, “DSO Analyzer,”, lists the control statements and their associated parameters that produce the reports described in this chapter.

DSO Report Headers

Every page of a DSO report has a header area. The title, DATA SET OPTIMIZER, is centered at the top of each page. Beneath the title are two columns of information at the left and right margins of the header area. Two different kinds of report headings are produced, depending upon the type of report.

Primary Heading

A primary heading is produced for both preliminary reports described in “Preliminary DSO Reports” on page 7-3. The column at the left margin contains the release number of DSO and the BMC Software product ownership notice. The column at the right margin shows the report date and time, the DSO patent number, and the analysis run number that is specified with the ANUM control statement.

Figure 7-1 is an example of a primary DSO report heading.

Figure 7-1 Primary DSO Report Heading

```
PRODUCED BY DSO R2.0.0                                                    PAGE 2
 BMC SOFTWARE, INC.                                                REPORT DATE   6 JUN 03
                                                                   REPORT TIME   14.45.55

                                                                ANALYSIS RUN NUMBER-TST1
```

Secondary Heading

Every main DSO report has a secondary heading at the top of each page. The secondary heading is identical to the primary heading but does not display the DSO patent number or the run number.

Figure 7-2 shows an example of a secondary DSO report heading that appears with all of the main reports.

Figure 7-2 Secondary DSO Report Heading

```
1                                D A T A   S E T   O P T I M I Z E R
 PRODUCED BY DSO R2.0.0                                                   PAGE 3
 BMC SOFTWARE, INC.                                                REPORT DATE   6 JUN 03
```

Preliminary DSO Reports

DSO generates two preliminary reports about the Analyzer batch job that creates the main reports shown in Table 7-1 on page 7-1. These preliminary reports are concerned with the input data sets that hold DSO data and the Analyzer control statements used to produce the main reports. Both reports are printed on the first two pages of the report listing with a primary report heading.

Control Card Edit Listing Report

The Control Card Edit Listing report, shown in Figure 7-3, shows the control statements from the batch JCL job to run the Analyzer. The control statements regulate the performance of the Analyzer as the Extractor records are processed. They specify what DSO reports are created and how the data is interpreted for volume reorganization plans.

Any error messages that appear during processing are listed in the report with the text of the message and message number. An explanation of DSO messages and abend codes are available online. For more information concerning messages and codes, see “” on page 7-22.

Figure 7-3 Control Card Edit Listing

```
                                    D A T A   S E T   O P T I M I Z E R
PRODUCED BY DSO R2.0.0                                                                                  PAGE 1
BMC SOFTWARE, INC.                                                                        REPORT DATE   6 JUN 03
                                                                                          REPORT TIME   14.45.55



                     *** CONTROL CARD EDIT LISTING ***

  ANUM TST1,MAXVOL=255,MAXCC=08                                        00420007
  DATA INPUT1,MONITOR=IPM,SYSTEM=(SYSA),RECTYPE=240                    00442007
 DSO 4222 FIRST 255 VOLUMES ENCOUNTERED WILL BE PROCESSED
```

Extractor Data Sets Report

The Extractor Data Sets report lists the input data sets that provide Extractor records for analysis. The report includes all time intervals in which Extractor data was collected and combines this information into one report.

Each input data set appears on a separate line in the report. These data sets are listed in the order in which they were processed by the Analyzer. The data set at the top of the list contains the initial records used for DSO reports.

Figure 7-4 shows an example of the Extractor Data Sets report. If there is no 240-00 record, no Extractor Data Sets report is generated. Each column heading of the report is described in Table 7-2.

Figure 7-4 Extractor Data Sets Report

```
                                    D A T A   S E T   O P T I M I Z E R
PRODUCED BY DSO R2.0.0                                                                                  PAGE 2
BMC SOFTWARE, INC.                                                                        REPORT DATE   6 JUN 03
                                                                                          REPORT TIME   14.45.55

                                                                                        ANALYSIS RUN
NUMBER-TST1


-EXTRACTOR DATA SETS REPORT
0DDNAME        START    START     END      END     EXTRACTOR  SYSTEM    SYS ID/   SAMP     DIE SAMPLE  REC# WITHIN
               DATE     TIME      DATE     TIME     RELEASE  RELEASE    SAMPLER   RATE       COUNT       DATASET
0INPUT1      00.307  08:04:15  00.307  09:04:15  CMF 5.5  MVS       SYSA/IPM     33      107,615          965
```

Table 7-2 Extractor Data Sets Report Column Descriptions

Column heading	Description
DDNAME	ddname in the Analyzer JCL that specifies the CMF Extractor input data set. If the ddname defines a concatenation, the concatenation number appears beside the ddname.
START DATE	Starting date of the Extractor job run that collected data held in the input data set. The date is expressed in Julian format (*yy.ddd*), where *yy* represents the last two digits of the year and *ddd* the day number of that year.
START TIME	Starting time of the Extractor job run that collected data held in the input data set. Time is expressed in military hours, minutes, and seconds (*hh.mm.ss*).
END DATE	Ending date of the Extractor job run that collected data held in the input data set. The date is expressed in Julian format.
END TIME	Ending time of the Extractor job run that collected data held in the input data set.
EXTRACTOR RELEASE	Eight-character code that identifies the CMF Extractor release number.
SYSTEM RELEASE	Release number of the MVS operating system from which data was extracted.
SYS ID/ SAMPLER	Four-character SMF system identifier, followed by the sampling mode (CPM or IPM) used for sampling.
SAMP RATE	Sample interval (in milliseconds) specified by the SAMPLE= parameter of the HEADMOVE control statement.
DIE SAMPLE COUNT	Number of samples that the CMF Extractor collected under the disabled interrupt exit (DIE).
REC# WITHIN DATA SET	Physical record number of the CMF Extractor startup record, relative to the beginning of the data set.

Study Report for Volume

The Study Report for Volume is the first main DSO report. The report summarizes the characteristics of the DASD, the data set statistics on the volume, and the total counts for various head move records.

By default, a separate Study Report is written for each volume that has records in the input Extractor data set and is requested for analysis by the Analyzer control statements. It can be omitted from the DSO report listing by specifying OPTIONS=NOVOLSTUDY with the DEFAULT or VOLUME control statements.

Figure 7-5 is an example of the Study Report for Volume. Each column heading of the report is described in Table 7-3.

Figure 7-5 Study Report for Volume

```
1                                          D A T A   S E T   O P T I M I Z E R
-PRODUCED BY DSO R2.0.0                                                                    PAGE 3
 BMC SOFTWARE, INC.                                                                REPORT DATE   6 JUN
03
                                                                                   REPORT TIME

14.45.55
0  1. STUDY REPORT FOR VOLUME WC9014

0

      VOLUME INFORMATION
0         DEVICE TYPE                          3390-1
            UCB DEVTYPE                      X'3010200F'
          AVAILABLE FORMAT 0 DSCBS              7,349
          AVAILABLE ALTERNATE TRACKS                0
          DATA SETS WITH NO EXTENTS                 1
          O/S DATASETS                            104
          VSAM DSCBS                                9
          VSAM CATALOGED ENTRIES                    0
          TOTAL SAMPLES                        77,784
             SAMPLES BEFORE DATASET EXISTED         0
             SAMPLES TO UNDEFINED EXTENTS           0
             SAMPLES ACCEPTED                  77,784
          PHYSICAL HEAD MOVEMENTS                   2
          LOGICAL HEAD MOVEMENTS                    1
          VOLUME BUSY/ERROR COUNT                   4
          MOVEMENTS TO ALTERNATE CYLINDERS          0
0     VOLUME MAPPING INFORMATION
0         DDNAME USED FOR MAPPING              INPUT1
          DATE AND TIME USED FOR MAPPING       02 NOV 00  08.04.21
```

Table 7-3 Study Report for Volume Column Descriptions

Column heading	Description
DEVICE TYPE	Specifies a four-digit number that identifies the type of DASD that the studied volume is resident. The device type is identified according to IBM DASD naming convention. An alphanumeric suffix follows the four-digit number to identify a particular model of the device type. Supported DASD types are shown in Table 6-2 on page 6-18.
UCB DEVTYPE	An eight-digit hexadecimal number that identifies the device type in the MVS unit control block (UCB).
AVAILABLE FORMAT 0 DSCBS	Number of format 0 data set control blocks (DSCBs) in the VTOC that are not allocated. When the number of free DSCBs approaches zero, it is unlikely that more data sets can be added to the volume.
AVAILABLE ALTERNATE TRACKS	Number of available alternate tracks remaining on the volume.
DATA SETS WITH NO EXTENTS	Number of data sets with format 1 DSCB entries in the VTOC not allocated to an extent. The DSCBs take up space in the VTOC but are not allocated to any storage space on the volume.
O/S DATA SETS	Number of recorded OS data sets on the volume.
VSAM DSCBS	Number of recorded VSAM data sets on the volume.
VSAM CATALOGED ENTRIES	Number of recorded VSAM cataloged entries on the volume.

Table 7-3 Study Report for Volume Column Descriptions (continued)

Column heading	Description
TOTAL SAMPLES	Total number of head moves recorded from the volume during the sample interval. Head moves are categorized as either valid or invalid samples: • SAMPLES BEFORE DATASETS EXISTED list the number of head moves to a free extent before a data set is allocated to it. The head moves are considered invalid and are not used for analysis. • SAMPLES TO UNDEFINED EXTENTS list the number of head moves to an undefined extent. The head moves are considered invalid and are not used for analysis. • SAMPLES ACCEPTED are the number of valid head move records. It represents the difference after the sum of all invalid records is subtracted from the total number of recorded head moves.
PHYSICAL HEAD MOVEMENTS	Total number of recorded head moves that are the result of the actuator moving to another physical cylinder.
LOGICAL HEAD MOVEMENTS	Total number of recorded head moves that are the result of changing tracks or logical cylinders without moving the actuator.
VOLUME BUSY/ERROR COUNT	Number of recorded head moves that occur when the volume is busy or in error recovery.
MOVEMENTS TO ALTERNATE CYLINDERS	Total number of head moves on the combined data sets to alternate cylinders. Alternate cylinder head moves are ignored in the data set head move and reorganization reports. For 33*xx* device types, moves to alternate cylinders are not observed by the CMF Extractor.
VOLUME MAPPING INFORMATION	Lists the ddname of the Extractor input data set and the date/time period that data is collected. The data supplied to the Analyzer can include more than one mapping of the same volume from different sampling periods. The most recent mapping is used by the Analyzer and its ddname appears on this line of the report.

Mapping for Volume Report

The Mapping for Volume report shows the current location of each data set on the volume. It includes information regarding each data set extent, each VTOC extent, and the remaining areas of free space left on the volume.

Each data set or other information is listed in a separate line of the report. Line entries are presented in the order of the physical mapping of the information on the volume. The report lists the data sets on each recording surface by

- Ascending cylinder number (CC)
- Ascending head order (HH)

If the NOVOLMAP parameter is specified with the DATA control statement, only head move records are read from the data set. VTOC and VSAM records are ignored. This forces the Analyzer to use existing mapping records for the volume. An invalid report can result when the Analyzer processes obsolete mapping records for the Mapping for Volume report.

Figure 7-6 shows an example of the Mapping for Volume report. Each column heading of the report is described in Table 7-4 on page 7-8.

Figure 7-6 Mapping for Volume Report

```
1                                        D A T A   S E T   O P T I M I Z E R

0PRODUCED BY DSO R2.0.0                                                                                    PAGE       4
 BMC SOFTWARE, INC.                                                                                 REPORT DATE   6 JUN 03
                                                                                                    REPORT TIME   14.45.55


0  2. MAPPING FOR VOLUME WC9014

0 DATA SET NAME                                    EXT  EXTENT BOUNDARIES   EXTENT  MOVEMENTS  LAST RCD   DSORG  RECORD   BLKSIZE
                                                          (HEX)            LENGTH  TO EXTENT  WRITTEN          FORMAT
                                                       CC  HH   CC  HH    CYL TRK              TT   R

  *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*        1      01    9 0E     9  14          1                UNK    UNK        0
  SYS1.VTOCIX.WC9014                                        A 00    B 0E     2              0     1D 15      PS      F     2,048
  SDBF.ISPF.PROFTAB                                         C 00    C 01         2         1      1 14      PO     FB    27,920
  SDWT1.ISPF.PROFTAB                                        C 02    C 0B        10         0      9 0E      PO     FB    27,920
  SDWT.DDIR.I                                               C 0C    C 0C         1         0                VSM    UNK    4,096
  SDWT.OS39024.SCPPEENU.INDEX                               C 0D    C 0D         1         0                VSM    UNK    4,096
  SDWT.OS39024.SCPPHENU.INDEX                               C 0E    C 0E         1         0                VSM    UNK    4,096
  SDWT1.ISPF.ISPPROF                                        D 00    D 0E     1             0      9 0E      PO     FB    27,920
  SDWT.DDIR.D                                               E 00    F 0E     2             0                VSM    UNK    4,096
  SYS1.VVDS.VWC9014                                        10 00   10 09        10         0                VSM    UNK    4,096
  SDWT.OS39024.SCPPVENU.INDEX                              10 0A   10 0A         1         0                VSM    UNK    4,096
  SDWT.BBSDEF                                              10 0B   10 0B         1         0        06      PO     FB    27,920
  SDWT.BBVDEF                                              10 0C   10 0C         1         0        06      PO     FB    27,920
  SDWT.DITPROF                                             10 0D   10 0D         1         0        01      PS     FB    27,920
  SDWT.DUTPROF                                             10 0E   10 0E         1         0        01      PS     FB    27,920
  SDWT.OS39024.SCPPEENU.DATA                               11 00   11 0E     1             0                VSM    UNK    4,096
  SDWT.OS39024.SCPPHENU.DATA                               12 00   14 0E     3             0                VSM    UNK    4,096
  SDWT.OS39024.SCPPVENU.DATA                               15 00   18 0E     4             0                VSM    UNK    4,096
  SDWT.COB2.CNTL                                           19 00   19 0E     1             0      5 05      PO     FB    27,920
  SDWT.CPPTEMP1.SCPPWORK                                   1A 00   1A 0E     1             0        01      PS     FB    27,920
  SDWT.CPPTEMP2.SCPPWORK                                   1B 00   1B 0E     1             0        01      PS     FB    27,920
  SDWT.DIALOG.JCL                                          1C 00   2C 0E    17             0     F1 04      PO     FB    27,920
  SDWT.EXEC.RACF.CLIST                                     2D 00   2D 00         1         0        01      PS     VB    27,998
  SDWT.FIN.FRONTALL                                        2D 01   2D 01         1         0        01      PS     FB    27,920
  SDWT.FIN.FRONTASK                                        2D 02   2D 02         1         0        01      PS     FB    27,920
  SDWT.FIN.KAREN                                           2D 03   2D 03         1         0        01      PS     FB    27,920
```

Table 7-4 Mapping for Volume Report Column Descriptions

Column heading	Description
DATA SET NAME	Name of the data set in DSCB 1 format, the VTOC, or other data physically located in the defined storage area of the volume.
EXT	Extent number that holds the data shown in the DATA SET NAME column of the report.
EXTENT BOUNDARIES	Location of an extent on a disk recording surface. The extent location is a range of storage between the spatial boundaries defined by two cylinder (CC) and head (HH) pairs.
EXTENT LENGTH	Length of the extent by cylinders (CYL) and tracks (TRK).
MOVEMENTS TO EXTENT	Number of head moves to an extent from any other extent on the volume. Cylinder-to-cylinder head moves within an extent are not reported. The number of head moves to an extent on a fixed-head volume is always zero.
LAST RCD WRITTEN	Address of the last extent of all sequential or partitioned data sets that points to the end of file mark. The address is given by track (TT) and record number (R).

Table 7-4 Mapping for Volume Report Column Descriptions

Column heading	Description
DSORG	Data set organization code located in the first extent of each data set. The data set organization code is one of the following values: IS Index sequential organization PS Physical sequential organization DA Direct organization PO Partitioned organization U Unmovable; the data contains location-dependent information UNK Unknown VI VSAM INDEX component VD VSAM DATA component (KSDS or ESDS) VA VSAM alternate index component VSM VSAM data space DSCB The combinations ISU, POU, PSU, and DAU are also possible.
RECORD FORMAT	Record format of the data set located in the first extent. The record format is one of the following values: F Fixed length V Variable length U Undefined length FB Fixed block FBA Fixed block with ANSI carriage control FBM Fixed block with machine carriage control VB Variable block VBS Variable block spanned VT Variable length with track overflow UT Undefined with track overflow UNK Not specified (VSAM is always UNK)
BLKSIZE	Maximum number of bytes in the physical records of the data set as specified in the data set label or the VSAM catalog.

Head Movement Between Pairs of Data Sets by Time on Volume Report

The Head Movement Between Pairs of Data Sets by Time on Volume Report describes head move activity between two different volume locations. Although the report is titled Head Movement Between Pairs of Data Sets, head moves to the VTOC are shown, as well as other non-data set locations on the volume.

This report is optional. You produce it by specifying the HMSORT=HMTIME parameter with the DEFAULT or VOLUME control statements. Data set pairs are listed by their percentage share of the total head move time recorded for the volume over the sampling interval.

Note: A companion report lists the same data set pairs by their percentage of the total number of head moves. The Head Movement Between Pairs of Data Sets by Count on Volume is a default DSO report. It is described in “Head Movement Between Pairs of Data Sets by Count on Volume Report” on page 7-13.

This report includes logical head moves between data set pairs in addition to physical head moves. Including logical head moves is particularly important for reorganizing shared volumes. The volume reorganization is based partly on the historical pattern of I/O requests causing head moves. Different systems have different I/O patterns that affect the proposed volume reorganization.

When a volume is shared between systems, automated reorganization is complicated by the difficulty of attributing responsibility for the head move to the correct system. A complete record of all head move activity is important to discern the source of I/O requests causing the head moves between data set pairs. Generally, a logical head move is caused by one of the following reasons:

- The seek occurs between data set extents that occupy the same cylinder.
- One or both of the seek locations are on a fixed-head cylinder.

One important exception to logical head move counts is the seek between the VSAM index and data components. Both components frequently share the same cylinder because the IMBED parameter is specified with the VSAM DEFINE command. Logical seeks between an embedded sequence set and the alternate index data component on the same cylinder are not included in the head move counts because

- Controlling the placement of the sequence set and alternate index data component is done as part of VSAM and is not part of the fundamental volume organization

- Physical head moves between non-shared index extents and data extents are more important in the volume reorganization

- If logical and physical head moves are reported together, their counts are indistinguishable because they are merged in the sequence-set-to-data head move pairs

The number of data set pairs that appear in the report can be restricted by specifying a threshold value for total head move time with the HMTIMETH= parameter. Data set pairs that do not meet the threshold value for total head move time are excluded from the report. When one or more entries are omitted from the report because they do not meet the threshold criterion, a message indicates the value of the threshold and the number of deleted entries.

Figure 7-7 shows an example of the Head Movement Between Pairs of Data Sets By Time report. Each column heading of the report is described in Table 7-5.

Figure 7-7 Head Movement Between Pairs of Data Sets by Time on Volume Report

```
1                                        D A T A   S E T   O P T I M I Z E R

0PRODUCED BY DSO R2.0.0                                                                                    PAGE      13
 BMC SOFTWARE, INC.                                                                                 REPORT DATE   5 JUL 01
                                                                                                    REPORT TIME   17.10.26

0  3. HEAD MOVEMENT BETWEEN PAIRS OF DATA SETS BY TIME ON VOLUME SYSDA2

0     DATA SET PAIRS                                HEAD MVMNT BETWEEN DATA SETS      TOTAL HM TIME       AVERAGE      CUMULATIVE
                                                          COUNT   PERCENT           AMOUNT  PERCENT      MVMNT TIME     PERCENT
                                                                                      (MS)                   (MS)

0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            6     17.14            90      22.39          15           22.39
      TADF3.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            3      8.57            54      13.43          18           35.82
      PLDP2.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            3      8.57            45      11.19          15           47.01
      BCKD.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            2      5.71            32       7.96          16           54.98
      SDKK.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            2      5.71            30       7.46          15           62.44
      BCVC.ISPF.PROFTAB
0     HMKJ.ISPF.PROFTAB                                        1      2.86            19       4.73          19           67.16
      WW.PAY.CYBORG.GL.PRINT2
0     SYS1.VTOCIX.SYSDA2                                       1      2.86            19       4.73          19           71.89
      WW.PAY.CYBORG.GL.PRINT2
0     BI.PAY.CYBORG.PAYXTR10.CNTL                              1      2.86            18       4.48          18           76.37
      TADF3.ISPF.PROFTAB
0     BI.PAY.CYBORG.PRINTU                                     1      2.86            18       4.48          18           80.85
      BCKD.ISPF.PROFTAB
0     HMKJ.ISPF.PROFTAB                                        1      2.86            16       3.98          16           84.83
      SDKK.ISPF.PROFTAB
0     BCKD.ISPF.PROFTAB                                        3      8.57             9       2.24           3           87.06
      SDKK.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            1      2.86             8       1.99           8           89.05
      BI.PAY.CYBORG.PRINTU
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            1      2.86             8       1.99           8           91.04
      BI.PAY.CYBORG.PAYXTR10.CNTL
0     BCKD.ISPF.PROFTAB                                        1      2.86             8       1.99           8           93.03
      PLDP2.ISPF.PROFTAB


0     BI.PAY.CYBORG.PRINTU                                     1      2.86             8       1.99           8           95.02
      HMKJ.ISPF.PROFTAB
0     BI.PAY.CYBORG.PRINTU                                     1      2.86             8       1.99           8           97.01
      SYS1.VTOCIX.SYSDA2
0     BCKD.ISPF.PROFTAB                                        3      8.57             6       1.49           2           98.51
      TADF3.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            2      5.71             4       1.00           2           99.50
      SYS1.VTOCIX.SYSDA2
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*            1      2.86             2        .50           2          100.00
      HMKJ.ISPF.PROFTAB

0     TOTALS                                                  35    100.00           402     100.00          11          100.00
```

Table 7-5 Head Movement Between Pairs of Data Sets Report Column Descriptions

Column heading	Description
DATA SET PAIRS	Names of both data sets with head moves between the extents where both data sets are located.
HEAD MVMNT BETWEEN DATA SETS	Head move activity between a pair of data sets measured by the number of head moves and the percentage of the total that the head move count represents. COUNT Total number of head moves between the respective locations of a pair of data sets. It includes head-to-head seeks that cross the boundaries between the data sets. PERCENT Percentage of the total head moves recorded for the volume by head moves between the data set pair.

Table 7-5 Head Movement Between Pairs of Data Sets Report Column Descriptions (continued)

Column heading	Description
TOTAL HM TIME	Cumulative time for all head moves between a pair of data sets. Total head move time is measured in milliseconds and the percentage of the total time for all head moves on the volume. AMOUNT Total time (in milliseconds) spent moving the head between a data set pair. Head-to-head seeks are not considered because there is no physical actuator move. PERCENT Percentage of the total head move time for all head moves on a volume represented by the cumulative time of all head moves between a data set pair.
AVERAGE MVMNT TIME (MS)	Average time spent moving the head between a pair of data sets. The average is derived by dividing the total head move time between the data set pair by the number of head moves. To use this report properly, use the total head move time with the average head move time. If either data set has extents under fixed-heads and extents under movable heads, this average is only a theoretical number.
CUMULATIVE PERCENT	Cumulative percentage of total head move time by all data set pairs preceding this line of the report.

Head Movement Between Pairs of Data Sets by Count on Volume Report

This report is similar to the Head Movement Between Pairs of Data Sets by Time shown in Figure 7-7 on page 7-12. Both reports provide identical information using the same columnar format. The only difference in this report is that the data set pairs are listed by the number of head moves rather than by their share of total head move time.

Diagnostic information can be derived from the two reports by comparing the order of the data set pairs between the two reports. A data set pair with a relatively small number of head moves ranked highly in the report that organizes data sets by time indicates that the average head move time for the data set pair is larger than the volume average. The data set pair should be moved closer together in a volume reorganization to reduce the average head move time, particularly if the head move time for the data set pair represents a large percentage of the total head move time for the volume.

Figure 7-8 shows an example of the report. The columns are the same as the companion report discussed in “Head Movement Between Pairs of Data Sets by Time on Volume Report” on page 7-10. Refer to Table 7-5 on page 7-12 for a description of the columns.

Figure 7-8 Head Movement Between Pairs of Data Sets by Count on Volume Report

```
1                                        D A T A   S E T   O P T I M I Z E R

0PRODUCED BY DSO R2.0.0                                                                                    PAGE      73
 BMC SOFTWARE, INC.                                                                                 REPORT DATE   5 JUL 01
                                                                                                    REPORT TIME   17.10.26

0  4. HEAD MOVEMENT BETWEEN PAIRS OF DATA SETS BY COUNT ON VOLUME SYSDA2

0     DATA SET PAIRS                                  HEAD MVMNT BETWEEN DATA SETS      TOTAL HM TIME       AVERAGE      CUMULATIVE
                                                            COUNT   PERCENT           AMOUNT  PERCENT      MVMNT TIME     PERCENT
                                                                                        (MS)                   (MS)
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                6     17.14           90      22.39          15           17.14
      TADF3.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                3      8.57           45      11.19          15           25.71
      BCKD.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                3      8.57           54      13.43          18           34.29
      PLDP2.ISPF.PROFTAB
0     BCKD.ISPF.PROFTAB                                            3      8.57            6       1.49           2           42.86
      TADF3.ISPF.PROFTAB
0     BCKD.ISPF.PROFTAB                                            3      8.57            9       2.24           3           51.43
      SDKK.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                2      5.71            4       1.00           2           57.14
      SYS1.VTOCIX.SYSDA2
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                2      5.71           30       7.46          15           62.86
      BCVC.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                2      5.71           32       7.96          16           68.57
      SDKK.ISPF.PROFTAB
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                1      2.86            8       1.99           8           71.43
      BI.PAY.CYBORG.PRINTU
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                1      2.86            8       1.99           8           74.29
      BI.PAY.CYBORG.PAYXTR10.CNTL
0     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*                1      2.86            2        .50           2           77.14
      HMKJ.ISPF.PROFTAB
0     BCKD.ISPF.PROFTAB                                            1      2.86            8       1.99           8           80.00
      PLDP2.ISPF.PROFTAB
0     BI.PAY.CYBORG.PAYXTR10.CNTL                                  1      2.86           18       4.48          18           82.86
      TADF3.ISPF.PROFTAB
0     BI.PAY.CYBORG.PRINTU                                         1      2.86            8       1.99           8           85.71
      HMKJ.ISPF.PROFTAB
0     BI.PAY.CYBORG.PRINTU                                         1      2.86            8       1.99           8           88.57
      SYS1.VTOCIX.SYSDA2
0     BI.PAY.CYBORG.PRINTU                                         1      2.86           18       4.48          18           91.43
      BCKD.ISPF.PROFTAB
0     HMKJ.ISPF.PROFTAB                                            1      2.86           16       3.98          16           94.29
      SDKK.ISPF.PROFTAB
0     HMKJ.ISPF.PROFTAB                                            1      2.86           19       4.73          19           97.14
      WW.PAY.CYBORG.GL.PRINT2
0     SYS1.VTOCIX.SYSDA2                                           1      2.86           19       4.73          19          100.00
      WW.PAY.CYBORG.GL.PRINT2

0     TOTALS                                                      35    100.00          402     100.00          11          100.00
```

Reorganization of Volume Report

The Reorganization of Volume report describes the recommended reorganization of the volume based on the probability of head moves between data set pairs. The report concludes with a summary area that projects the percentage reduction in total seek time after the volume is reorganized.

The CMF Extractor head move records include the head number and cylinder numbers. Head moves are assigned to specific data sets, even when data sets share cylinders. Data set groups are not forced to remain at contiguous extents because they originally shared cylinders. However, the various extents of each data set are combined in the reorganization.

ISAM indexes and independent overflow areas are treated as separate data sets if they were originally allocated as separate entities. DSO issues a warning message when a split allocation is made for data sets. Split-allocated data sets are reorganized as a group.

VSAM components that share the same data space are treated as affiliated data sets. The reorganization algorithm weights them toward retaining contiguity within the data space. If the overall performance is not improved by keeping the data space intact, DSO splits VSAM components out of the data space. For this reason, the specified VSAM component should be allocated as a UNIQUE component in its own dedicated data space.

The reorganization process preserves the original location of unmovable data sets if the NOMOVE, VSAM=NOMOVE, or ISAM=NOMOVE parameters are specified with the DEFAULT or VOLUME control statements. If the NOMOVE parameter is declared for specific data sets, the data sets remain at their current location and no attempt is made to optimize their placement in the volume reorganization. Expect to see little or no projected improvement with a reorganization if the NOMOVE, VSAM=NOMOVE, or ISAM=NOMOVE parameters are selected for volumes with a considerable number of head moves to unmovable data sets.

The reorganization algorithm minimizes the number of data sets that must be split unless unmovable data sets are present. If COMPAKTR or DFDSS options are specified, the Analyzer does not split up data sets on the volume. However, processing for the volume is terminated when data sets cannot be reallocated contiguously.

Figure 7-9 on page 7-16 is a partial example of the Reorganization of Volume report. It shows a listing of the data sets on the volume, starting from the first extent. A summary area at the end of the report compares recorded head move times for the volumes to the projected times after reorganization. The calculated improvement is shown as the percentage reduction in total seek time Each column heading of the report is described in Table 7-6 on page 7-16.

Figure 7-9 Reorganization of Volume Report

```
1                                          D A T A   S E T   O P T I M I Z E R
-PRODUCED BY DSO R2.0.0                                                                                PAGE 1
 BMC SOFTWARE, INC.                                                                          REPORT DATE  11 JUL
01
                                                                                             REPORT TIME
16.00.40
0  5. REORGANIZATION OF VOLUME SYSDA2

0     DATA SET NAME                                    EXTENT BOUNDARIES      ABSTR    ABSTR        IMMOVABLE
                                                        CC  HH    CC  HH      START   LENGTH         EXTENT

      *** FREE SPACE ***                                   01     AB 0E          1     2,579
      WC.MEAP98.MR47                                     AC 00    1C4 0E      2,580     4,215
      BH.PAY.CYBORG.Y2000Q1                             1C5 00    261 0E      6,795     2,355
      ZA.PAY.CYBORG.P20.MASTER1.G0017V00                262 00    2DC 0E      9,150     1,845
      ZA.PAY.CYBORG.P20.MASTER1.G0014V00                2DD 00    355 0E     10,995     1,815
      CYBORG.CONTROL.FILE01.D102298.S                   356 00    3C1 0E     12,810     1,620
      CYBORG.CONTROL.FILE01.D100298.S2                  3C2 00    40F 0E     14,430     1,170
      HW.PAY.CYBORG.P20.MNTRUN                          410 00    45C 0E     15,600     1,155
      CYBORG.CONTROL.FILE01.D102198.S                   45D 00    498 0E     16,755       900
      SR.PAY.CYBORG.P20.MASTER1.G0070V00                499 00    4D1 0E     17,655       855
      BG.PAY.CYBORG.BATCHL                              4D2 00    507 0E     18,510       810
      DB.PAY.CYBORG.HH25XT.HFCC                         508 00    524 0E     19,320       435
      WC.TEMP.CONT                                      525 00    525 0E     19,755        15
      BI.PAY.CYBORG.PAYXTR10.CNTL                       526 00    527 0E     19,770        30
      BCKD.ISPF.PROFTAB                                 528 00    528 01     19,800         2
      SDKK.ISPF.PROFTAB                                 528 02    528 05     19,802         4
      TADF3.ISPF.PROFTAB                                528 06    528 07     19,806         2
      PLDP2.ISPF.PROFTAB                                528 08    528 09     19,808         2
      HMKJ.ISPF.PROFTAB                                 528 0A    528 0B     19,810         2
      BCVC.ISPF.PROFTAB                                 528 0C    528 0D     19,812         2
      LICC.ATHENA3                                      528 0E    528 0E     19,814         1
      *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*     529 00    52A 0E     19,815        30
      SYS1.VTOCIX.SYSDA2                                52B 00    52B 0D     19,845        14
      OL.OLSETCHR                                       52B 0E    52B 0E     19,859         1
      BI.PAY.CYBORG.PRINTU                              52C 00    52E 0E     19,860        45
      WW.PAY.CYBORG.GL.PRINT2                           52F 00    535 0E     19,905       105
      LA.PAY.CYBORG.P20.MASTER1.G0037V00                536 00    565 0E     20,010       720
      TE.PAY.CYBORG.PAY2OUT.DATA                        566 00    594 0E     20,730       705
      LV.PAY.CYBORG.P20.MASTER1.G0002V00                595 00    5BF 0E     21,435       645
      WCISC.INW2.Y94                                    5C0 00    5EA 0E     22,080       645
      GK.MEAP98.DATA58                                  5EB 00    610 0E     22,725       570
      CR.PAY.CYBORG.P20.PAYXTR                          611 00    632 0E     23,295       510
      DB.PAY.CYBORG.GG25XT.HFCC                         633 00    64F 0E     23,805       435
      CR.PAY.CYBORG.P20.MASTER1.G0027V00                650 00    66A 0E     24,240       405
      PE.PAY.CYBORG.P20.MNTRUN                          66B 00    683 0E     24,645       375
      FDRABR.QUARTER.REPORT.G0006V00                    684 00    69C 0E     25,020       375
      MA.PAY.CYBORG.TEMP04                              69D 00    6B5 0E     25,395       375
      GK.MEAP98.DATA47                                  6B6 00    6CD 0E     25,770       360
      ZB.PAY.CYBORG.P20.MNTRUN                          6CE 00    6E4 0E     26,130       345
      CI.PAY.CYBORG.P20.PAYXTR.YREND                    6E5 00    6FB 0E     26,475       345
.................
.................
      *** FREE SPACE ***                                9B2 09    A5E 0E     37,239     2,586
0
0          (1) TOTAL HEAD MOVEMENT TIME BEFORE REORGANIZATION                      402 MS
           (2)       IMMOVABLE/IMMOVABLE HEAD MOVEMENT TIME                          0 MS
           (3) TOTAL HEAD MOVEMENT TIME SUBJECT TO IMPROVEMENT                     402 MS
0          (4) APPROXIMATE HEAD MOVEMENT TIME AFTER REORGANIZATION                  54 MS
           (5)     REORGANIZED DATASET HEAD MOVEMENT TIME                           54 MS
           (6)     AVERAGE HEAD MOVEMENT TIME AFTER REORGANIZATION                   1 MS
0          (7) PERCENTAGE IMPROVEMENT RESULTING FROM REORGANIZATION               86.57 %
                   (((3)-(5))/(3))
0          (8) TOTAL PERCENTAGE IMPROVEMENT FOR VOLUME                            86.57 %
                   (((1)-(4))/(1))
0                        **OPTIONS IN EFFECT FOR REORGANIZATION
                                  **CENTER VTOC:     NO
                                  **VSAM IMMOVABLE:  NO
                                  **ISAM IMMOVABLE:  NO
                                  **O/S  IMMOVABLE:  NO
                                  **AUTO FIX:        NO
```

Table 7-6 Reorganization of Volume Report Column Descriptions (Part 1 of 2)

Column heading	Description
DATA SET NAME	Data set name for each line entry.
EXTENT BOUNDARIES	Boundaries of the data set are shown for each line entry in cylinders (CC) and tracks (HH).
ABSTR START	Absolute start location by track number.

Table 7-6 Reorganization of Volume Report Column Descriptions (Part 2 of 2)

Column heading	Description
ABSTR LENGTH	Absolute track length by the number of tracks.
IMMOVABLE EXTENT	Extents that are location-dependent (such as VSAM and ISAM extents).

Summary Area of the Reorganization of Volume Report

At the end of the Reorganization of Volume report, a summary area compares the recorded total head move time to the projected time after the volume is reorganized. Table 7-7 lists the fields are included in the summary area of the report.

Table 7-7 Reorganization of Volume Report, Summary Area Field Descriptions (Part 1 of 2)

Field	Description
TOTAL HEAD MOVEMENT TIME BEFORE REORGANIZATION	Total head move time measured for a volume before it is reorganized.
TOTAL HEAD MOVEMENT TIME SUBJECT TO IMPROVEMENT	Total head move time before reorganization minus the head move time associated with those data sets marked immovable (that is, the NOMOVE, the VSAM=NOMOVE, or the ISAM=NOMOVE control statement parameters were used). This measurement ignores data sets under fixed-heads, since there is no head move associated with them.
APPROXIMATE HEAD MOVEMENT TIME AFTER REORGANIZATION	Estimated total head move time after the proposed volume reorganization is complete. The approximation of the total head move time is based on the proposed distance between individual data sets after the volume is reorganized. Threshold specifications do not affect the calculation of estimated total head move time. Head move times are calculated differently before and after the volume is reorganized. The total head move time before reorganization is the sum of all valid head moves that are sampled on the volume. The approximation of total head move time after reorganization is the sum of the estimated average head move time between the extents of a data set pair. The total head move times can vary before and after reorganization, even if the placement of data sets is not changed by reorganizing the volume. The approximate head move time calculation is based in the following concepts: • If the extents of a data set pair were formerly on the same cylinder that allowed head-to-head seek moves, these head-to-head seeks are converted into real head moves and are included in the estimated total. Head move time is included for any data set pair whose extents have been moved to different cylinders after the volume is reorganized. • Head moves between a pair of data sets are subtracted from the estimated total if their extents are combined within a single cylinder or placed under a fixed-head volume. Head move time is subtracted for any data set pair whose extents placement permits an electronic head-to-head seek after the volume is reorganized. • The estimated head move time between a pair of data sets is reduced if the data set extents are closer together after the volume is reorganized. The estimate of the head move time is based upon the volume average for moving heads between the proposed distance of the data set extents.

Table 7-7 Reorganization of Volume Report, Summary Area Field Descriptions (Part 2 of 2)

Field	Description
REORGANIZED DATA SET HEAD MOVEMENT TIME	Estimated total head move time after the proposed reorganization of all movable data sets on the volume.
AVERAGE HEAD MOVEMENT TIME AFTER REORGANIZATION	Estimated average head move time after the volume is reorganized.
PERCENTAGE IMPROVEMENT RESULTING FROM REORGANIZATION	Estimated percentage reduction in head move time of the data sets on the volume that are subject to reorganization. The restrictions applied to the APPROXIMATE HEAD MOVEMENT TIME calculation also apply to this value. The estimate of the percentage reduction of total head move time for the data sets on the volume subject to reorganization is calculated with the following formula: (Total HM time subject to reorganization - Reorg data set HM time) X 100 --- Total HM time subject to improvement
TOTAL PERCENTAGE IMPROVEMENT FOR VOLUME	Estimated percentage reduction in total head move time for all data sets on the volume after reorganization. The restrictions noted under APPROXIMATE HEAD MOVEMENT time numbers also apply to this number. The estimate is derived using the following formula: (Total HM time before reorg - Approx. HM time after reorg) X 100 --- Total HM time before reorg

Reorganization Options for the Reorganization Volume Report

Volumes are reorganized according to several parameters set with the DEFAULT or VOLUME control statements. The positioning of the VTOC, VSAM data sets, and other data structures are examples of some of the possible reorganization options selected with the appropriate parameters of these control statements. Refer to Chapter 4, “Evaluating DSO Reports,” and Chapter 6, “DSO Analyzer,” for more information about declaring volume reorganization options with Analyzer control statements.

The end of the Reorganization of Volume report lists the possible reorganization options for the volume. A YES or NO response indicates whether the option was in effect as the report was created.

Table 7-8 lists the reorganization options that appear at the end of the Reorganization of Volume report.

Table 7-8 Reorganization Options for the Reorganization of Volume Report

****OPTIONS IN EFFECT FOR REORGANIZATION**	The options in effect for the Reorganization of Volume report. If a reorganization is impossible with the selected options, DSO overrides the options with default reorganization parameters.
CENTER VTOC	A YES response indicates that the VTOC is placed at the center cylinder, equidistant between the inner and outer boundaries of the recording surface of the disk that holds the data cylinders.
VSAM IMMOVABLE	A YES response indicates that the VSAM data sets are regarded as immovable during the reorganization and remain at their current location on the volume.
ISAM IMMOVABLE	A YES response indicates that the ISAM data sets are regarded as immovable during the reorganization and remain at their current location on the volume.
O/S IMMOVABLE	A YES response indicates that the O/S data sets are regarded as immovable during the reorganization and remain at their current location on the volume.
AUTOFIX	A YES response indicates that highly active data sets are placed on fixed-head cylinders.

Data Set Activity Reports

The Data Set Activity reports list every data set on the volume by a particular head move statistic gathered during the sampling interval. Each line of the report shows a separate data set.

DSO produces four different Data Set Activity reports based upon the ordering of the data sets on the volume. The same information appears in every report; only the listing of the data sets varies according to data set name, the percentage of total head moves, the percentage of time in error recovery, or the percentage of device busy time. Each report is selected as an option with the ACSORT= parameter of the DEFAULT or VOLUME control statements.

The method of ordering the data sets is shown from the complete titles of the four Data Set Activity reports:

Data Set Activity Report by Name on Volume

The standard Data Set Activity report that lists the data sets on the volume in alphabetical order. The report is produced automatically because ACSORT=NAME is the default parameter of the DEFAULT and VOLUME control statements.

Data Set Activity Report by Busy on Volume

Data sets are listed in order of their percentage share of the total device busy time over the measured interval. You specify the report by setting the ACSORT=BUSY parameter with the DEFAULT or VOLUME control statements.

Data Set Activity Report by Error Recovery on Volume

Data sets are listed in order of their percentage share of the total error recovery time spent by the volume over the measured interval. Data sets are listed alphabetically if the error recovery percentage is the same for different data sets. You specify the report by setting the ACSORT=ERROR parameter with the DEFAULT or VOLUME control statements.

Data Set Activity Report by Head Movement on Volume

Data sets are listed in order of their percentage share of the total head moves recorded for the volume. Data sets are listed alphabetically if the head move percentage is the same for different data sets. You specify the report by setting the ACSORT=HM parameter with the DEFAULT or VOLUME control statements.

If an ACTIVITY threshold is declared with an Analyzer control statement, data sets outside of the threshold are omitted from the report. You include or exclude data sets from Data Set Activity reports by comparing the threshold with the field used to sort the report entries. A message at the end of a Data Set Activity report indicates how many data sets are omitted.

Figure 7-10 is an example of the default Data Set Activity Report by Name on Volume. The three remaining Data Set Activity reports have the same report columns and are not shown. The report columns are described in Table 7-9.

Figure 7-10 Data Set Activity by Name on Volume Report

```
                                    D A T A   S E T   O P T I M I Z E R

0PRODUCED BY DSO R2.0.0                                                                                    PAGE      28
 BMC SOFTWARE, INC.                                                                                 REPORT DATE  11 JUL 01
                                                                                                    REPORT TIME   16.00.40

0  6. DATA SET ACTIVITY REPORT BY NAME ON VOLUME SYSDA2

0                                                        ERROR RECOVERY                  HEAD MOVEMENTS    PERCENT   CUMULATIVE
     DATA SET NAME                                 PERCENT  PERCENT  PERCENT   BLOCKSIZE   TO THE     WITHIN   OF TOTAL    PERCENT
                                                     BUSY   DS IDLE  DS BUSY              DATA SET   DATA SET    (HM)

     *.*.*.*.*.*.*.*.____VTOC____.*.*.*.*.*.*.*.*   .0145   .0000    .0000          0         10          1      28.95      28.95
     AL.GBSCHE21.TEMP                                .0000   .0000    .0000     23,461          0          0        .00      28.95
     AL.HSPT1096.DATA                                .0000   .0000    .0000     30,500          0          0        .00      28.95
     AL.PAY.CYBORG.PRINTU                            .0000   .0000    .0000      1,330          0          0        .00      28.95
     AL.PAY.CYBORG.RUNCNTLG.G0002V00                 .0000   .0000    .0000         80          0          0        .00      28.95
     AL.STU.PIC021                                   .0000   .0000    .0000     23,432          0          0        .00      28.95
     AL.STU29.GR200                                  .0000   .0000    .0000      6,160          0          0        .00      28.95
     ALAL.ISPF.PROFTAB                               .0000   .0000    .0000     23,440          0          0        .00      28.95
     ALJP2.ISPF.PROFTAB                              .0000   .0000    .0000     23,440          0          0        .00      28.95
     AT.PAY.CYBORG.GL.PRINT2                         .0000   .0000    .0000      1,330          0          0        .00      28.95
     AT.PAY.CYBORG.P50OT5                            .0000   .0000    .0000      7,680          0          0        .00      28.95
     AT.PAY.CYBORG.RUNCNTLG.G0006V00                 .0000   .0000    .0000         80          0          0        .00      28.95
     ATRH.ISPF.PROFTAB                               .0000   .0000    .0000     23,440          0          0        .00      28.95
     ATST.ISPF.PROFTAB                               .0000   .0000    .0000     23,440          0          0        .00      28.95
     BC.STU.DOWN198                                  .0000   .0000    .0000     23,452          0          0        .00      28.95
     BC.STU.STAT160                                  .0000   .0000    .0000     23,460          0          0        .00      28.95
     BC.ST903I.TESTINPT                              .0000   .0000    .0000      1,503          0          0        .00      28.95
     BCAB2.ISPF.PROFTAB                              .0000   .0000    .0000     23,440          0          0        .00      28.95
     BCKD.ISPF.PROFTAB                               .0145   .0000    .0000     23,440          6          0      15.79      44.74
     BCPL2.ISPF.PROFTAB                              .0000   .0000    .0000     23,440          0          0        .00      44.74
     BCSB2.ISPF.PROFTAB                              .0000   .0000    .0000     23,440          0          0        .00      44.74
     BCVC.ISPF.PROFTAB                               .0019   .0000    .0000     23,440          1          0       2.63      47.37
     BETSO.ISPF.PROFTAB                              .0000   .0000    .0000     23,440          0          0        .00      47.37
     BG.BLDGCALL.DOWN006                             .0000   .0000    .0000      9,360          0          0        .00      47.37
     BG.BUYBACK                                      .0000   .0000    .0000      5,500          0          0        .00      47.37
     BG.COMP                                         .0000   .0000    .0000     23,408          0          0        .00      47.37
     BG.GEN5CDS                                      .0000   .0000    .0000         80          0          0        .00      47.37
     BG.PAY.CYBORG.BATCHL                            .0000   .0000    .0000      9,040          0          0        .00      47.37
     BG.PAY.CYBORG.GLFILE                            .0000   .0000    .0000      7,680          0          0        .00      47.37
     BG.PAY.CYBORG.NXTHRLYA                          .0000   .0000    .0000     23,440          0          0        .00      47.37
     BG.PAY.CYBORG.TEMP.C015                         .0000   .0000    .0000      1,330          0          0        .00      47.37
     BG.PAY.CYBORG.UPLOAD.PR234TCH                   .0000   .0000    .0000      6,233          0          0        .00      47.37
     BG.PAY.CYBORG.UPLOAD.PR234TCH.TEST              .0000   .0000    .0000      6,233          0          0        .00      47.37
     BG.PAY.LASER.DEDT                               .0000   .0000    .0000     23,125          0          0        .00      47.37
     BGCB.NAME                                       .0000   .0000    .0000      2,900          0          0        .00      47.37
     BGCB.TOM                                        .0000   .0000    .0000      1,830          0          0        .00      47.37
     BGCT.ISR0001.BACKUP                             .0000   .0000    .0000     13,680          0          0        .00      47.37
     BGCT.ISR0003.BACKUP                             .0000   .0000    .0000     13,680          0          0        .00      47.37
     BGDW.PAYOFF                                     .0000   .0000    .0000         80          0          0        .00      47.37
     BGDW.SRCHFOR.LIST                               .0000   .0000    .0000      3,325          0          0        .00      47.37
.............
.............
     WYKA.SPF1.LIST                                  .0000   .0000    .0000      3,146          0          0        .00     100.00
     ZA.PAY.CYBORG.P20.MASTER1.G0014V00              .0000   .0000    .0000      1,536          0          0        .00     100.00
     ZA.PAY.CYBORG.P20.MASTER1.G0017V00              .0000   .0000    .0000      1,536          0          0        .00     100.00
     ZA.PAY.CYBORG.RUNCNTLG.G0001V00                 .0000   .0000    .0000         80          0          0        .00     100.00
     ZB.PAY.CYBORG.P20.MNTRUN                        .0000   .0000    .0000      1,536          0          0        .00     100.00
     ZZ.FIN.NEWACT                                   .0000   .0000    .0000      1,600          0          0        .00     100.00
     ZZ.FIN.OPENPO                                   .0000   .0000    .0000      2,660          0          0        .00     100.00
     ZZ.PAY.PAYROLL.DATA                             .0000   .0000    .0000     23,440          0          0        .00     100.00

0     TOTALS                                           .11     .00      .00                    35          3     100.00     100.00
```

Table 7-9 Data Set Activity by Name on Volume Report Column Descriptions

Column heading	Description
DATA SET NAME	Name of the data set or the VTOC described in the report.
PERCENT BUSY	Percentage of time the recording head was positioned over the data set but the device was busy.

Table 7-9 Data Set Activity by Name on Volume Report Column Descriptions (continued)

Column heading	Description
ERROR RECOVERY	Percentage of total extraction time the device spent in error recovery with each data set. DSO detects that the device is in error recovery when the I/O Supervisor Block (IOSB) flag IOSERR becomes active. The IOSERR flag is set when the Error Recovery Procedure (ERP) of the device is in control or ERP retry is in progress. DS IDLE Percentage share of error recovery time that occurred during system operations. DS BUSY Percentage share of error recovery time that occurred during a period of I/O activity on the volume.
BLOCKSIZE	Length of the records in bytes that compose the storage allocations of data stored on DASD. Blocksize is the physical record size for VSAM data sets.
HEAD MOVEMENTS	Number of head moves associated with a data set. Data set head moves are differentiated between external head moves to the data set from another cylinder and internal head moves that originate within the data set. External head moves include head-to-head seeks. Internal head moves exclude all seeks that do not physically move the device actuator. TO THE DATA SET Number of external head moves to the data set. WITHIN DATA SET Number of internal head moves within the data set.
PERCENT OF TOTAL (HM)	Number of individual data set head moves as a percentage of the total head moves recorded for the volume. Data Set Activity reports sorted by name or the number of head moves express the percent of the total by head moves. Device Activity reports that sort data sets by busy or error time represent the percent of the total as the ratio between individual data set error or busy counts to the total recorded for the volume.
CUMULATIVE PERCENT	Cumulative percentage of the total head moves for all data sets preceding this line of the report.

Chapter 8 DSO Output JCL

This chapter briefly discusses COMPAKTOR, DFDSS, and some of the limitations imposed on volume reorganization. An example of the DSO reorganization job for each product is also shown.

Both DFDSS and COMPAKTOR use an output job created by DSO as their input to reorganize a volume. A DD statement declares the name of the output data set that holds the job created by the Analyzer. Either the DFDSS or COMPAKTR subparameter must be declared with the OPTIONS= parameter of the VOLUME or DEFAULT control statement to produce the output job. A reorganization job is produced for the studied volumes.

DFDSS Reorganizing JCL and Control Card Output

DFDSS is a member of the IBM Data Facility Family. This integrated family of products performs the key storage management functions in an MVS environment. DFDSS is the primary data mover in the family. It has facilities to back up, restore, and manage space. In addition, it converts data to and from the volumes that are part of a Storage Management Subsystem environment.

All DFDSS JCL and control cards created by DSO are held in the output data set declared from the batch job to run the Analyzer. There are many limitations regarding the file types that can be processed by DFDSS for logical data set name dumps or restores. Refer to DFDSS publications for information about the types of data set that are supported. The limitations listed below apply to certain data sets that are reorganized with plans developed by DSO:

- ISAM is not supported.

- Specific placement for VSAM is not supported for files that DFDSS can dump or restore.

- Multivolume data sets are not supported except for the first volume of the data set. DFDSS only supports movement of multivolume data sets using logical dump and logical restore. Since the DFDSS job generated by DSO performs physical dump and physical restore, only the first segment of the data set is restored properly. The remaining allocated segments (on other volumes) will not be restored because of a mismatch in volume sequence numbers.

Figure 8-1 on page 8-3 shows an example of the 10-step job generated by DSO to reorganize a volume with DFDSS.

Before submitting the reorganization jobs to DFDSS, edit and make any necessary changes to the following list:

- Perform a full volume backup before running the reorganization job. A volume backup is necessary when certain files cannot be restored with a data set restore (for example, ISAM).

- In STEP0002, mark the volume as PRIVATE before attempting a reorganization. Using the PRIVATE subparameter ensures that the volume is not accessed by another job during reorganization. As an example, the PRIVATE subparameter should be coded in the form VOL=(PRIVATE,SER=(ABC123)).

- If the volume is initialized as an indexed VTOC, it must be converted to OSVTOC before reorganization. The following JCL example should be inserted between STEP0001 and STEP0002 to make the conversion to OSVTOC. Refer to the appropriate IBM ICKDSF manual that discusses the BUILDIX command for more information.

```
//STEP001A EXEC PGM=ICKDSF
//         COND=(0,NE)
//DASD  DD DSN=SYS1.VTOCIX.ABC123,
//         UNIT=(3380),
//         VOL=(PRIVATE,SER=(SYS1.VTOCIX.ABC123)),
//         DISP=OLD
//SYSPRINT  DD SYSOUT=*
//SYSIN     DD *
//BUILDIX   DDNAME(DASD) OS NOPURGE
/*
```

Figure 8-1 Example DSO Job—Reorganizing a Volume with DFDSS

```
//DSODFDSS JOB (XXXX),'PGMRNAME',CLASS=A,MSGCLASS=A,REGION=2048K
//* ************************************************************
//*
//*  DFDSS REORG OF VOLUME SYS002
//*
//*  "HLTAPE"  ==> DSN PREFIX OF BACKUP TAPE
//*  "HILEVEL" ==> DSN PREFIX OF TEMPORARY DISK FILES
//*
//* ************************************************************
//*
//*  DUMP DATASETS ON VOLUME WITH DFDSS
//*
//STEP0001 EXEC PGM=ADRDSSU,COND=(0,NE)
//DASD     DD   UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=OLD
//TAPE     DD   DSN=HLTAPE.SYS002.BACKUP,
//         UNIT=TAPE,
//         VOL=(,,,10),
//         LABEL=(,SL),
//         DISP=(NEW,CATLG,DELETE)
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   *
  DUMP     FULL -
           ALLDATA(*) ALLEXCP -
           TOL(ENQF) INDD(DASD) OUTDD(TAPE)
/*
//*
//*  INIT VOLUME AND PLACE VTOC WITH DSF
//*
//STEP0002 EXEC PGM=ICKDSF,COND=(0,NE)
//DASD     DD   UNIT=3380,
//         VOL=(PRIVATE,SER=(SYS002)),
//         DISP=OLD
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   *
 INIT     DDNAME(DASD) VOLID(SYS002) -
          VERIFY(SYS002) -
          VTOC(0390,00,0015) PURGE

/*
//*
//*  HOLD FREESPACE FOLLOWING VTOC
//*
//STEP0003 EXEC PGM=IEFBR14,COND=(0,NE)
//FILE0001 DD   DSN='HILEVEL.VTOCFSPC',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(,KEEP,KEEP),
//         SPACE=(ABSTR,(33960,05865))
```

```
//*
//*  ALLOCATE THE FILES BEFORE VTOC
//*
//STEP0004 EXEC PGM=IEFBR14,COND=(0,NE)
//FILE0001 DD   DSN='HILEVEL.FSPC0001',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(,KEEP,KEEP),
//         SPACE=(ABSTR,(00388,00002))
//FILE0002 DD   DSN='TMS.AUDIT.OLD',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(,KEEP,KEEP),
//         SPACE=(CYL,(0006,0000))
           .
//STEP0005 EXEC PGM=IEFBR14,COND=(0,NE)
//FILE0001 DD   DSN='HILEVEL.VTOCFSPC',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(OLD,DELETE,DELETE)
//*
//*  ALLOCATE FILES AFTER VTOC
//*

//STEP0006 EXEC PGM=IEFBR14,COND=(0,NE)
//* SYS2ZSMP.PNCP.NCPTZN.CSI.INDEX                    *VSAM FILE*
//FILE0001 DD   DSN='HILEVEL.VSPC0019',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(,KEEP,KEEP),
//         SPACE=(ABSTR,(00015,05865))
//* SYS2ZDAV.PR100.DATABASE.INDEX                     *VSAM FILE*
//FILE0002 DD   DSN='HILEVEL.VSPC0020',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(,KEEP,KEEP),
//         SPACE=(ABSTR,(00010,05880))
//* SYS2PCIC.CPRD1.DFHRSD.INDEX                       *VSAM FILE*
           .

//STEP0007 EXEC PGM=IEFBR14,COND=(0,NE)
//FILE0001 DD   DSN='HILEVEL.VSPC0001',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(OLD,DELETE,DELETE)
//FILE0002 DD   DSN='HILEVEL.VSPC0002',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(OLD,DELETE,DELETE)
           .
```

```
//STEP0008 EXEC PGM=ADRDSSU,COND=(0,NE)
//DASD     DD   UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=OLD
//TAPE     DD   DSN=HLTAPE.SYS002.BACKUP,
//         DISP=OLD
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   *
  RESTORE  DATASET(INC(**)  -
           EXC(SYS1.VVDS        REPLACE TOL(ENQF) OUTDD(DASD) INDD(TAPE)

/*
//*
//*  RELEASE ALL FREESPACE ON VOLUME
//*
//STEP0009 EXEC PGM=IEFBR14,COND=(0,NE)
//FILE0001 DD   DSN='HILEVEL.FSPC0001',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(OLD,DELETE,DELETE)
//FILE0002 DD   DSN='HILEVEL.FSPC0002',
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=(OLD,DELETE,DELETE)
//*
//*  BUILD VOLUMES VTOC INDEX
//*
//STEP0010 EXEC PGM=ICKDSF,COND=(0,NE)
//DASD     DD   DSN=SYS1.VTOCIX.SYS002,
//         UNIT=3380,
//         VOL=SER=SYS002,
//         DISP=OLD
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   *
 BUILDIX   DDNAME(DASD) IXVTOC
/*
//
```

FDR COMPAKTOR Control Statement Output

All FDR COMPAKTOR control statements are held in the output data set declared from the batch job to run the Analyzer.

The center of the pack is described in these control statements (up to 255 statements are generated for each volume). The following rules apply for any volume that is reorganized:

- Immovable data sets are included only if option MOVE is specified.

Note: DSO does not ensure that COMPAKTOR control cards are valid for a given file. A simulated COMPAKTOR run should be done before reorganizing the volume.

- ISAM data sets should never be optimized if COMPAKTOR is used to reorganize. Space is reserved for them in even cylinders.

- Explicit positioning is avoided as much as possible. Any control statements produced by DSO contain positioning data for noncylinder-aligned data sets until the first cylinder-aligned data set is encountered; the remaining data sets are set up as a SEQUENCE group.

- Only data sets with activity are specifically placed on the volume.

Explicit positioning is avoided so the control statement group cannot be invalidated when one or more data sets expands between the time the volume is observed with DSO and the time COMPAKTOR is run.

Note: Even though space is reserved for VSAM and ISAM data sets, FDR may use up that space for the remaining nonsequenced data sets on the pack. Data sets may have to be relocated manually to clear that free space. This can be accomplished by deleting all data sets contaminating the free areas, allocating the VSAM or ISAM data sets, and restoring the data sets by their names that were deleted. It may also be necessary to fill the pack with dummy data sets to force allocation to fill the spaces reserved for VSAM or ISAM data sets.

Figure 8-2 is a representative example of COMPAKTOR reorganization JCL produced by DSO. It shows sample output.

Figure 8-2 Sample COMPAKTOR Output

```
*VOLSER=SYS002
 COMPAKT FROMDD=TAPE1,  **INPUT  DDN **
         TODD=DISK1,    **TARGET DDN **
         DUMP=YES,      **DUMP BEFORE**
         CPYVOLID=YES,  **KEEP VOLSER**
         TEMPS=KEEP,    **KEEP TEMPS **
*        EXTENTS=       **  REGROUP  **
*        PSRLSE=        ** P/S AS IS **
*        PORLSE=        ** P/O AS IS **
*        %FREE=         ** % FREE SPC**
         OBJECT=MAXFREE,   FREE SPACE**
         OVERRIDE=YES,  ** IGNORE ENQ**
         VTOC=COMPAKT   ** RELEASE F0**
 DSN SYS2ZCIC.CR17.MACLIB,POS=22550012
 SEQUENCE POS=22830000
 DSN SYS2PNTM.P.MONLOG
 DSN SYS2ZCIC.CR17.LOADLIB2
 DSN CICS161.SDF.ITFB.INDEX
 DSN SYS1.VTOCIX.SYS002
 DSN SYS1.VVDS.VSYS002
 DSN CICS161.SDF.TML.INDEX
 ENDSEQ
 SEQUENCE POS=23330000
 DSN SYS2ZCIC.CR17.COBLIB
 DSN SYS2PNTM.P.CHANGE.DATA
 DSN SYS2DCIC.CR17.LOADLIB1
 DSN CICS161.SDF.MSL.DATA
 ENDSEQ
 SEQUENCE POS=23450000
 DSN CICS161.SDF.TML.DATA
 DSN SYS2ZCIC.CR17.LOADLIB
 ENDSEQ
 SEQUENCE POS=23810000
 DSN ***VTOC,SIZE=0015
 DSN SYS2DSPT.P.ONLLOAD
 ENDSEQ
 SEQUENCE POS=24840000
 DSN SYS2PCIC.CPRD3.DFHDMPA
 DSN SYS2ZCIC.CR17.SOURCE
 ENDSEQ
```

Appendix A Report Calculations

The focus of DSO reports is the head move activity that occurs on a volume during a sampling interval. Head move activity is reported under a variety of conditions that existed on the volume when the sample was recorded. The first section in this appendix shows the variables used by the DSO Analyzer to calculate head move values that appear in the reports. The equations that calculate report values from the variables are shown in the second section.

DSO Report Variables

The following variables represent various aspects of head move activity recorded by the Extractor. The DSO Analyzer processes the raw Extractor input records to derive the head move variables shown below. The variables represent head move activity recorded on the volume under different conditions.

B(v)	Total number of recorded head moves that occurred when the volume was busy but not in error recovery.
E(v)T	Total number of recorded head moves that occurred on a volume during error recovery.
S(v)	Total number of DIE samples recorded from a valid online volume. The total may be less than the reported total shown in the DIE SAMPLE field of the CMF Extractor Data Sets report.
Teb(v) = B(v)+E(v)	Total number of recorded head moves that occurred when the volume was busy or in error recovery.
H(v)	Total number of recorded head moves on a volume. The Extractor sampler records a head move regardless of the busy or error recovery status of the volume during the head move. Logical head moves are part of the total.
L(v)	Total number of recorded head moves in which the cylinder/head location changed.

P(v)	Total number of recorded physical head moves.
P(d)	The total number of physical head moves to an extent from a location outside of the extent.
I(d)	Total number of head moves to an extent attributed to data set (*d*) from a location within the extent.
B(d)	Total number of head moves to an extent attributed to data set (*d*) when the status of the UCB is busy and the volume is not in error recovery.
Eb(d)	Total number of head moves to an extent attributed to data set (*d*) when the status of the UCB is busy and the volume is in error recovery.
Ei(d)	Total number of head moves to an extent attributed to data set (*d*) when the volume is idle and in error recovery.
A(d) = B(d)+Eb(d)+ Ei(d)X(d1,d2)	Total number of head moves from the extent holding data set *d1* to the extent holding data set *d2*

Report Equations

The equations that derive head move values are listed in this section. The equations appear with the DSO report that presents a head move value calculated from the particular equation. The titles of three DSO reports are shown below. Beneath the titles are the report field descriptors and equations. The field descriptors are the areas in the DSO reports that show the head move values calculated from the equations.

Study Report for Volume

PHYSICAL HEAD MOVEMENTS	= P(v)
LOGICAL HEAD MOVEMENTS	= L(v)
VOLUME BUSY/ERROR COUNT	= Teb(v)

Head Movement between Pairs of Data Sets

COUNT	= X(d1,d2)+X(d2,d1)
PERCENT	= ((X(d1,d2)+X(d2,d1))*100/P(v)

Note: These numbers include counts of head moves to fixed-head cylinders.

Data Set Activity Reports

PERCENT BUSY= (A(d)*100)/S(v)

ERROR RECOVERY

PERCENT DS IDLE= (Ei(d)*100)/S(v)

PERCENT DS BUSY= (Eb(d)*100)/S(v)

HEAD MOVEMENTS

TO THE DATA SET= P(d)

WITHIN DATA SET= I(d)

PERCENT OF TOTAL and CUMULATIVE PERCENT

(BUSY)= (A(d)*100)/Teb(v)

(ERROR)= ((Ei(d)+Eb(d))*100)/E(v)

(HM)= ((P(d)+I(d))*100)/H(v)

Index

Numerics

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

X

Notes

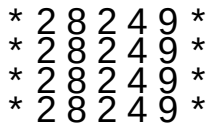

28249