

MAINVIEW[®] Batch Optimizer Data Optimizer Reference Manual

Version 2.3

July 31, 2003



Copyright 2003 BMC Software, Inc. All rights reserved.

BMC Software, the BMC Software logos, and all other BMC Software product or service names are registered trademarks or trademarks of BMC Software, Inc. All other trademarks belong to their respective companies.

BMC Software considers information included in this documentation to be proprietary and confidential. Your use of this information is subject to the terms and conditions of the applicable End User License Agreement for the product and the proprietary and restricted rights notices included in this documentation.

Restricted Rights Legend

U.S. Government Restricted Rights to Computer Software. UNPUBLISHED -- RIGHTS RESERVED UNDER THE COPYRIGHT LAWS OF THE UNITED STATES. Use, duplication, or disclosure of any data and computer software by the U.S. Government is subject to restrictions, as applicable, set forth in FAR Section 52.227-14, DFARS 252.227-7013, DFARS 252.227-7014, DFARS 252.227-7015, and DFARS 252.227-7025, as amended from time to time. Contractor/Manufacturer is BMC Software, Inc., 2101 CityWest Blvd., Houston, TX 77042-2827, USA. Any contract notices should be sent to this address.

Contacting BMC Software

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

United States and Canada

Address BMC Software, Inc.
2101 CityWest Blvd.
Houston TX 77042-2827

Telephone 713 918 8800 or
800 841 2031

Fax 713 918 8000

Outside United States and Canada

Telephone (01) 713 918 8800

Fax (01) 713 918 8000

Customer Support

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

Support Web Site

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

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

Support by Telephone or E-mail

In the United States and Canada, if you need technical support and do not have access to the Web, call 800 537 1813. Outside the United States and Canada, please contact your local support center for assistance. To find telephone and e-mail contact information for the BMC Software support center that services your location, refer to the Contact Customer Support section of the Support page on the BMC Software Web site at www.bmc.com/support.html.

Before Contacting BMC Software

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

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

Contents

About This Book	xv	
Chapter 1	Introduction	
	MAINVIEW Batch Optimizer Overview	1-2
	Data Optimizer	1-2
	Job Optimizer	1-3
	Job Optimizer Pipes	1-3
	Job Optimizer for DB2 and IMS	1-3
	MAINVIEW Batch Optimizer Common Components	1-3
	Data Optimizer Overview	1-5
	VSAM Optimization	1-6
	Non-VSAM Optimization	1-7
	Features	1-9
	Benefits	1-10
Chapter 2	Understanding Data Optimizer	
	Overview	2-2
	Selection Processing	2-2
	Data Policy Definitions	2-2
	Data Definition Override Statements	2-3
	User Control Exit	2-3
	Selection Criteria Order	2-4
	Optimization Processing	2-7
	User-Specified Options	2-7
	System Resource Utilization	2-7
	Historical Data	2-8
	Performance Options	2-8
	General Options	2-9
	Resource Utilization Options	2-11
	VSAM and Non-VSAM Options	2-11
	Reporting	2-12
	Performance Reports	2-12
	Informational Messages	2-13

Chapter 3	Implementing Data Optimizer	
	Overview	3-2
	Setting Goals	3-2
	Assessing Candidates	3-3
	Locating Candidates	3-5
	Moving Optimization Into Production	3-6
	Initial Strategy	3-6
	Intermediate Strategy	3-6
	Advanced Strategy	3-6
Chapter 4	Managing a Data Policy	
	Overview	4-2
	Before You Begin	4-2
	Specifying Selection Criteria for Data Policies	4-2
	Ordering Definitions within a Data Policy	4-3
	Setting Global Options for a Data Policy	4-4
	Panel Flow for a Data Policy Registration	4-4
	Creating or Editing a Data Policy Definition	4-5
	Editing or Viewing Data Policy Global Options	4-15
	Activating a Data Policy	4-17
	Activating a Data Policy by Using a Command	4-18
	Activating a Data Policy by Using the Dialog	4-18
	Determining Which Data Policy is Active	4-20
	Determining which Subsystem Processes a Data Set	4-21
Chapter 5	Data Optimizer Options	
	Overview	5-1
	Option Description Format	5-1
	DD Statement Position	5-2
	Wildcard Characters	5-4
	Options	5-5
	General Options	5-5
	Resource Utilization Options	5-13
	Reporting Options	5-16
	Non-VSAM Options	5-22
	VSAM Options	5-30
	Summary of Data Optimizer Options, Keywords, and DD Overrides	5-53
Chapter 6	Data Optimizer Commands	
	Overview	6-2
	Operator Console Commands	6-2
	Starting the BMC Software Primary Subsystem	6-6
	Displaying the Status of the BMC Software Primary Subsystem	6-7
	Stopping the BMC Software Primary Subsystem	6-8
	Cancelling the BMC Software Primary Subsystem	6-9
	Starting the MAINVIEW Batch Optimizer Subsystem	6-10
	Displaying the Status of the MAINVIEW Batch Optimizer Subsystem	6-12

Stopping the MAINVIEW Batch Optimizer Subsystem	6-13
Cancelling the MAINVIEW Batch Optimizer Subsystem.	6-14
Displaying the Status of Data Optimizer	6-15
Enabling and Disabling Components.	6-16
Displaying the Status of a Data Policy.	6-17
Activating a Data Policy.	6-18
Displaying the Status of the User Control Exit	6-19
Activating the User Control Exit	6-20
Setting the History Key	6-21

Appendix A

Data Optimizer Reports

Overview	A-1
Non-VSAM Performance Report	A-2
Non-VSAM Performance Report Header	A-4
Non-VSAM Options Report.	A-6
Non-VSAM Statistics Report.	A-7
Non-VSAM Data Set Information Report	A-8
VSAM Performance Report	A-11
VSAM Performance Report Header	A-15
VSAM Options Report.	A-16
VSAM Statistics Report.	A-19
VSAM Data Set Information Report	A-20
VSAM Catalog Information Report.	A-22

Appendix B

Navigating the User Interface

Overview	B-2
ISPF Conventions	B-2
ISPF Profile	B-2
Panel Layout.	B-3
Accessing the User Interface	B-8
Modifying User Interface Processing	B-10
Dialog Options and Settings.	B-10
Modifying ISPF Settings	B-11
Modifying Key Lists	B-12
Using the User Interface	B-13
Navigation	B-14
Pop-up Panels.	B-15
Accessing the Action Bar.	B-16
Point-and-Shoot	B-17
Creating and Editing Control Data Set Members.	B-18
Object List Displays.	B-20
Online Help.	B-24
Panel Help	B-24
Field Help.	B-24
Message Help	B-24
Reference Phrase Help	B-25
Other User Interface Tasks	B-25
Product Version and Maintenance Level Information	B-25

	Environment Version and Release Information	B-26
	BMC Software Product Authorization User Interface	B-27
Appendix C	User Control Exit	
	Overview	C-2
	User Control Exit Decisions	C-4
	Status Change Call-Back Capability	C-4
	Testing the User Control Exit	C-5
	User Exit Statistics Report	C-6
Appendix D	Special Considerations	
	Overview	D-2
	Buffer Values Coded in the JCL	D-2
	Data Integrity with Non-VSAM Advanced Mode Processing	D-3
	EXCP and EXCPVR Processing Exclusions	D-4
	Excluded Programs	D-6
	Excluded DD Statements	D-9
	Excluded Data Sets	D-10
Appendix E	SMF Recording	
	Overview	E-1
	SMF Record Header	E-2
	Self-Defining Section	E-3
	Product Section	E-5
	Subsystem Identification Section	E-5
	Job Information Section	E-6
	Subsystem Processing Options Section	E-7
	General Data Set Information Section	E-11
	DCB Information Section	E-11
	ACB Information Section	E-13
	Non-VSAM I/O Statistics Section	E-14
	VSAM I/O Statistics Section	E-15
	VSAM LSR Statistics Section	E-15
	VSAM Catalog Information Section	E-16
Appendix F	HIPER-CACHE to Data Optimizer Conversion Utility	
	Overview	F-1
	Defaults, Definitions, and JCL Overrides	F-2
	Sample JCL	F-3
	Required DD Statements	F-4
	Optional DD Statements	F-5
	Reports	F-6
	Format	F-7
	Report Sample	F-9
	Conversion Procedure	F-17
	HIPER-CACHE to Data Optimizer Conversion Tables	F-18

Appendix G**Batch Utilities**

Overview	G-2
JOB Control Statements and Definitions	G-2
SYSIN Control Card Format	G-4
Keyword Specifications	G-5
SMF Report for VSAM Data Sets	G-6
SMF Report for Non-VSAM Data Sets	G-6

Index

Figures

Figure 1-1	MAINVIEW Batch Optimizer Components	1-5
Figure 1-2	VSAM Performance Component	1-7
Figure 1-3	Non-VSAM Performance Component	1-8
Figure 2-1	Data Optimizer Intercept Processing (Part 1 of 2)	2-5
Figure 4-2	Data Policy Definitions Panel	4-7
Figure 4-3	Data Policy Definition Panel	4-8
Figure 4-4	Data Policy Definition Panel–Default Values	4-10
Figure 4-5	Data Policy Selection Criterion Panel	4-11
Figure 4-6	Data Policy Selection Criterion Panel–Prompt Active	4-12
Figure 4-7	Data Policy Global Options Pop-up	4-16
Figure 4-9	Policy Activation BatchPlex Name Pop-up	4-19
Figure 4-10	Policy Activation Images List Pop-up	4-20
Figure B-1	MAINVIEW Batch Optimizer Objects List Panel	B-3
Figure B-2	Data Policy Definition Panel	B-4
Figure B-3	Data Policy Definition Selection Criterion Pop-up Panel	B-5
Figure B-4	Member Options Choice Entry Panel	B-6
Figure B-5	MAINVIEW Batch Optimizer Logon panel	B-9
Figure B-6	Dialog Options Panel	B-10
Figure B-7	ISPF Settings Panel	B-12
Figure B-8	Keylist Utility Panel	B-13
Figure B-9	Exit Confirmation Panel	B-14
Figure B-10	Cancel Confirmation Panel	B-15
Figure B-11	File Pull-Down Menu	B-17
Figure B-12	Data Policy Method Options Pop-up List Panel	B-18
Figure B-13	BatchPlex Definition Panel	B-21
Figure B-14	Job Optimizer Policy Selection Criteria and Action List	B-21
Figure B-15	Job Optimization Policy Multi-line Criteria Display Panel	B-23
Figure B-16	Components Maintenance Information Panel	B-26
Figure B-17	Version/Release Information Panel	B-26
Figure B-18	Product Authorization Primary Menu Panel	B-28
Figure F-4	Conversion Report Comment Block	F-9

Tables

Table C-1	User Exit Statistics Report Field Descriptions	C-6
Table D-1	Excluded CCW Op-codes	D-5
Table D-2	Excluded Programs - Non-VSAM Processing	D-6
Table D-3	Excluded Programs - VSAM Processing	D-8
Table D-4	Excluded DD Statements	D-9
Table F-1	HIPERCACHE to Data Optimizer Conversion Utility Required DD Statements	F-4
Table F-2	HIPERCACHE to Data Optimizer Conversion Utility Optional DD Statements	F-5
Table F-3	Conversion Utility Return Codes	F-6
Table F-4	Report Status Descriptions	F-8
Table G-1	JOB Control Statements and Definitions	G-2
Table G-2	Special Characters for SYSIN Control Card Keywords	G-5



About This Book

This book discusses the Data Optimizer functionality and provides information needed for implementing Data Optimizer in the most effective manner. This manual presumes that Product Name has previously been installed and customized in your site. For information on these topics, see *MAINVIEW Batch Optimizer Installation Manual*.

How This Book Is Organized

This book is organized as follows:

Chapter/Appendix	Description
Chapter 1, "Introduction"	explains the features and benefits of Data Optimizer
Chapter 2, "Understanding Data Optimizer"	explains what Data Optimizer is, what it does, and how it works
Chapter 3, "Implementing Data Optimizer"	describes the tasks necessary to implement Data Optimizer
Chapter 4, "Managing a Data Policy"	describes how to work with data policy definitions
Chapter 5, "Data Optimizer Options"	explains the concepts and options for the Data Optimizer component
Chapter 6, "Data Optimizer Commands"	provides task descriptions for Data Optimizer commands
Appendix A, "Data Optimizer Reports"	explains the reports that can be generated by Data Optimizer
Appendix B, "Navigating the User Interface"	explains how to work with the MAINVIEW Batch Optimizer user interface
Appendix C, "User Control Exit"	explains how to implement the User Control exit to intercept data sets for I/O optimization processing
Appendix D, "Special Considerations"	describes when, and how, to intervene manually in the normally automatic operation of Data Optimizer
Appendix E, "SMF Recording"	describes the SMF recording function of Data Optimizer and provides tables for interpreting SMF records

Chapter/Appendix	Description
Appendix F, "HIPER-CACHE to Data Optimizer Conversion Utility"	describes the function of the HIPER-CACHE to Data Optimizer conversion utility and explains how to use it
Appendix G, "Batch Utilities"	describes and provides examples of the batch utility programs provided with Data Optimizer

Related Documentation

BMC Software products are supported by several types of documentation:

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

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.html>. Log on and select a product to access the related documentation. (To log on, first-time users can request a user name and password by registering at the support page or by contacting a BMC Software sales representative.)

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.html>.

Online Help

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.html>.

Conventions

General Conventions

Item	Example
information that you are instructed to type	Type SEARCH DB in the designated field. Type search db in the designated field. (Unix)
specific (standard) keyboard key names	Press Enter .
field names, text on a panel	Type the appropriate entry in the Command field.
directories, file names, Web addresses	
nonspecific key names, option names	Use the HELP function key. KEEPDICTIONARY option

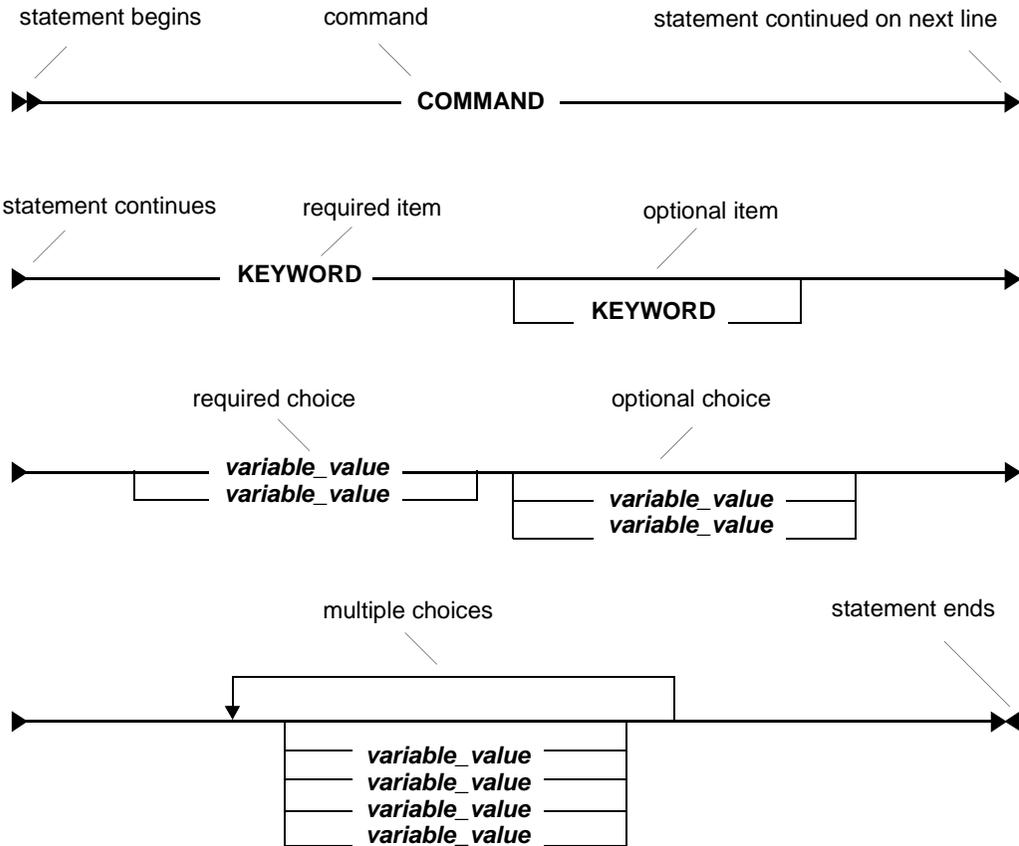
Item	Example
MVS calls, commands, control statements, keywords, parameters, reserved words	Use the .
code examples, syntax statements, system messages, screen text	//STEPLIB DD The table <i>table_name</i> is not available.
emphasized words, new terms, variables	
single-step procedures	»» To enable incremental backups, type y and press Enter at the next prompt.

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

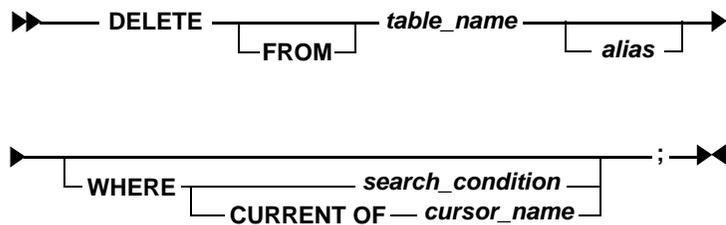
Syntax Statements

Syntax Diagrams

The following figure shows the standard format for syntax diagrams:



The following example illustrates the syntax for a DELETE statement. Because the FROM keyword, *alias* variable, and WHERE clause are optional, they appear below the main command line. In contrast, the *table_name* variable appears on the command line because the table name is required. If the statement includes a WHERE clause, the clause must contain either a search condition or a CURRENT OF clause. (The *search_condition* variable appears on the main line for the WHERE clause, indicating that this choice is required.)



The following guidelines provide additional information about syntax diagrams:

-
- A recursive (left-pointing) arrow above a stack indicates that you may choose more than one item in the stack.
 - An underlined item is a default option.
 - In general, MVS commands, keywords, clauses, and data types appear in uppercase. However, if an item can be shortened, the minimum portion of the MVS command or keyword may appear in uppercase with the remainder of the word in lowercase (for example, CANCEL).
 - The following conventions apply to variables in syntax diagrams:
 - Variables typically appear in lowercase and are always italicized.
 - If a variable is represented by two or more words, underscores connect the words (for example, *database_name* and *user_ID*).

Chapter 1 Introduction

This chapter explains what the Data Optimizer component is, what it does, and how it works. This chapter discusses the following topics:

MAINVIEW Batch Optimizer Overview	1-2
Data Optimizer	1-2
Job Optimizer	1-3
Job Optimizer Pipes	1-3
Job Optimizer for DB2 and IMS	1-3
MAINVIEW Batch Optimizer Common Components	1-3
Data Optimizer Overview	1-5
VSAM Optimization	1-6
Non-VSAM Optimization	1-7
Features	1-9
Benefits.	1-10

MAINVIEW Batch Optimizer Overview

The MAINVIEW® Batch Optimizer product is a specialized set of components designed to manage batch jobs, reduce elapsed processing times, and provide more efficient use of your available resources. Some MAINVIEW Batch Optimizer components provide common core functionality to the product, while others provide optimization for specific types of performance processing.

MAINVIEW Batch Optimizer is available in three product tiers, each which providing a different level of functionality to meet your batch processing optimization needs.

- MAINVIEW Batch Optimizer–Standard
- MAINVIEW Batch Optimizer–Advanced
- MAINVIEW Batch Optimizer–Enterprise

Table 1-1 lists the components included with each MAINVIEW Batch Optimizer product tier.

Table 1-1 Components of MAINVIEW Batch Optimizer by Product Tiers

Included component	MAINVIEW Batch Optimizer–Standard	MAINVIEW Batch Optimizer–Advanced	MAINVIEW Batch Optimizer–Enterprise
Data Optimizer	yes	yes	yes
Job Optimizer	no	yes	yes
Job Optimizer Pipes	no	yes	yes
Job Optimizer for DB2 and IMS	no	no	yes

The optimization components of MAINVIEW Batch Optimizer are described in more detail below.

Data Optimizer

Data Optimizer provides I/O performance benefits by applying a robust set of optimization techniques to VSAM and non-VSAM I/O processing.

Data Optimizer comprises MAINVIEW Batch Optimizer–Standard and is included with the Advanced and Enterprise product packages.

Job Optimizer

Job Optimizer provides performance benefits to batch jobs by running job steps concurrently and providing step-to-step piping to enable the movement of data between the split steps.

Job Optimizer is provided with MAINVIEW Batch Optimizer–Advanced and MAINVIEW Batch Optimizer–Enterprise.

Job Optimizer Pipes

Job Optimizer Pipes provides in-memory piping of application data between batch jobs and job steps. Job Optimizer Pipes allows two data-related applications to execute concurrently rather than sequentially, reducing the elapsed time that is required to process the jobs.

Job Optimizer Pipes is provided with MAINVIEW Batch Optimizer–Advanced and MAINVIEW Batch Optimizer–Enterprise.

Job Optimizer for DB2 and IMS

The Job Optimizer for DB2 and IMS components provides performance benefits to batch jobs that access DB2 and IMS databases. Performance is optimized by running job steps concurrently.

Job Optimizer for DB2 and IMS offers the same performance benefits as Job Optimizer; however, Job Optimizer for DB2 and IMS executes job steps in parallel that access DB2 or IMS databases.

Job Optimizer for DB2 and IMS is provided only with MAINVIEW Batch Optimizer–Enterprise.

MAINVIEW Batch Optimizer Common Components

The following components provide core functionality to the MAINVIEW Batch Optimizer product.

MAINVIEW Batch Optimizer User Interface

The MAINVIEW Batch Optimizer user interface (dialog) is an ISPF-based interactive dialog that provides access to the MAINVIEW Batch Optimizer control data set. You can use the ISPF Edit function to view or modify the control data set members, the MAINVIEW Batch Optimizer user interface affords a protected environment for this process. The dialog verifies that any control data set members you create or modify are syntactically correct and do not provide conflicting instructions to the optimization components.

The dialog provides control for all portions of the product and it is included with all tiers of MAINVIEW Batch Optimizer.

MAINVIEW Batch Optimizer Subsystem

MAINVIEW Batch Optimizer Subsystem (MBOS) is provided with all tiers of the product and is responsible for providing the following key services to the other components:

- console communication
- history data set access
- common service routines

The MBOS is included with all tiers of MAINVIEW Batch Optimizer.

BMC Software Primary Subsystem

BMC Software Primary Subsystem (BMCP) is responsible for establishing interception points so that MAINVIEW Batch Optimizer components get control at the required times. The BMCP is included with all tiers of MAINVIEW Batch Optimizer.

Figure 1-1 on page 1-5 shows the core and optimization components of MAINVIEW Batch Optimizer.

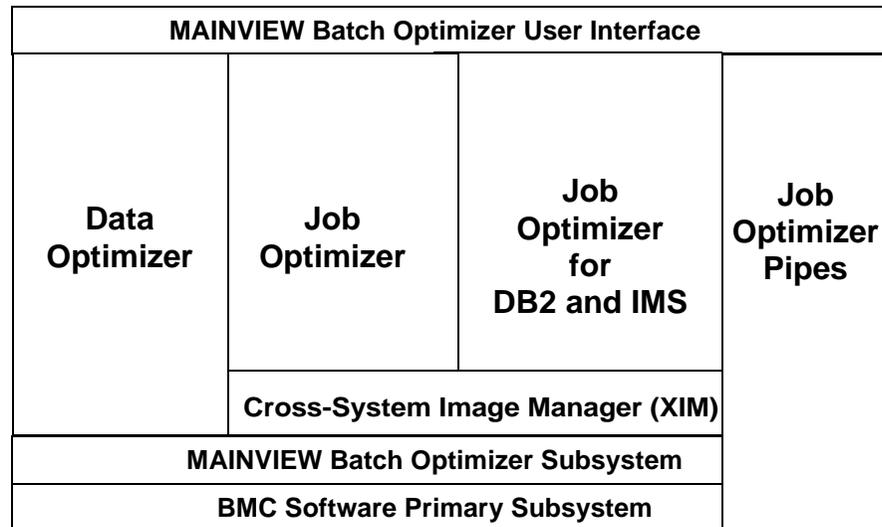
Cross-System Image Manager

Job Optimizer incorporates the BMC Software Cross-System Image Manager (XIM™) technology. XIM enables Job Optimizer to distribute and manage job steps across one or more MVS systems. XIM functions transparently within Job Optimizer, but it requires the presence of the IBM Cross-System Coupling Facility (XCF).

To permit the distribution of job steps across multiple operating system images, XCF must be executing in a multisystem environment.

Figure 1-1 shows the core components and optimization components of MAINVIEW Batch Optimizer.

Figure 1-1 MAINVIEW Batch Optimizer Components



This book describes functionality that is specific to the Data Optimizer component of the MAINVIEW Batch Optimizer.

Data Optimizer Overview

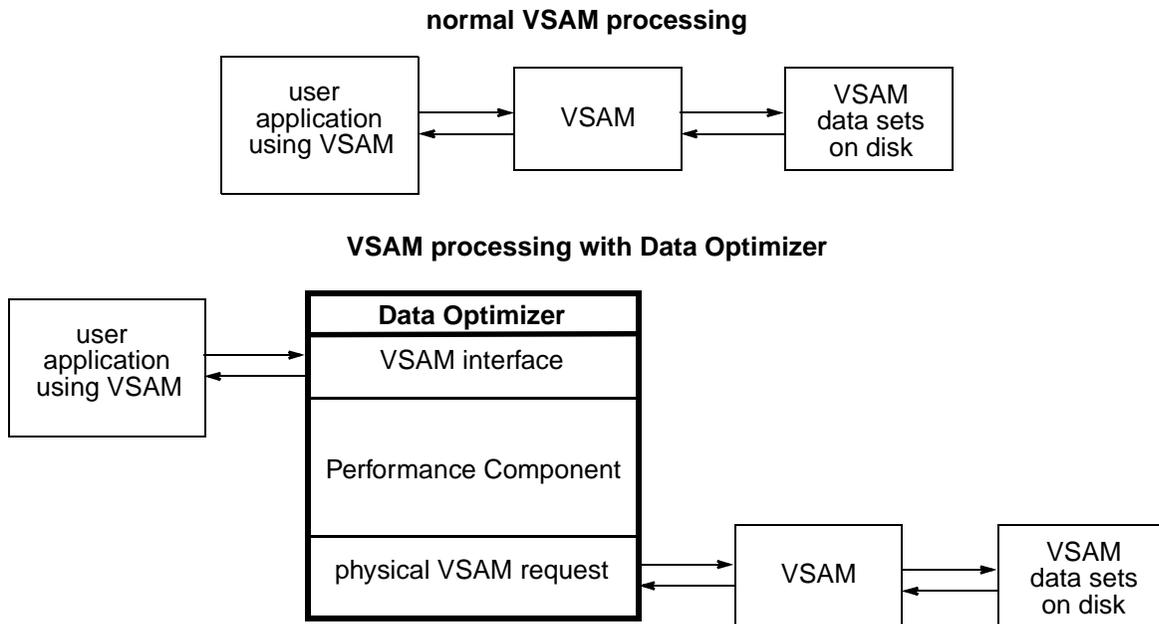
Data Optimizer is an I/O optimization product that provides performance benefits for VSAM and non-VSAM data sets. Using a variety of I/O optimization techniques, Data Optimizer reduces the number of accesses to disk and reduces or eliminates unnecessary wait times, which results in reduced elapsed run times. It performs automatic tuning to maintain optimum application performance without over-utilizing system resources and negatively affecting overall system performance. It provides performance benefits to almost any application that performs I/O without requiring program or JCL changes to do so.

VSAM Optimization

For accesses to VSAM data sets, Data Optimizer optimizes buffer processing and selects appropriate VSAM options to improve performance. For sequential accesses, Data Optimizer optimizes index and data buffer values. For random accesses, Data Optimizer dynamically builds and uses LSR buffer pools without changes to the application. For accesses that indicate random and sequential access, Data Optimizer records history to learn how to best optimize future processing by the application. When Data Optimizer has selected LSR processing, if the application proceeds to perform significant sequential accesses, Data Optimizer can perform its own read-ahead to continue providing sequential performance benefits. For NSR and LSR optimization, Data Optimizer moves buffers and control blocks above the 16 MB line (when possible) to aid in virtual storage constraint relief.

Figure 1-2 shows the VSAM performance component of Data Optimizer.

Figure 1-2 VSAM Performance Component



Non-VSAM Optimization

For non-VSAM data sets, Data Optimizer replaces all low-level I/O functions with its own, providing it with complete control of buffer management and of the physical I/O requests. All I/O requests by the application are satisfied logically and transparently by the Data Optimizer internal buffer manager.

For sequential processing, Data Optimizer reads large amounts of data with each I/O and overlaps the I/Os to maximize performance, regardless of physical blocking characteristics.

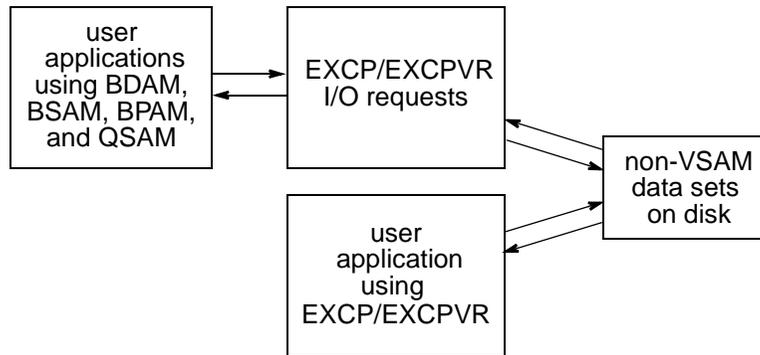
For random accesses, Data Optimizer can retain up to 255 tracks of data in memory to maximize cache-hit performance benefits. It also allows below-the-line buffers to be decreased to assist in virtual storage constraint relief while still providing performance benefits.

For tape data sets, PDS members, and striped data sets, Data Optimizer optimizes data management buffer values.

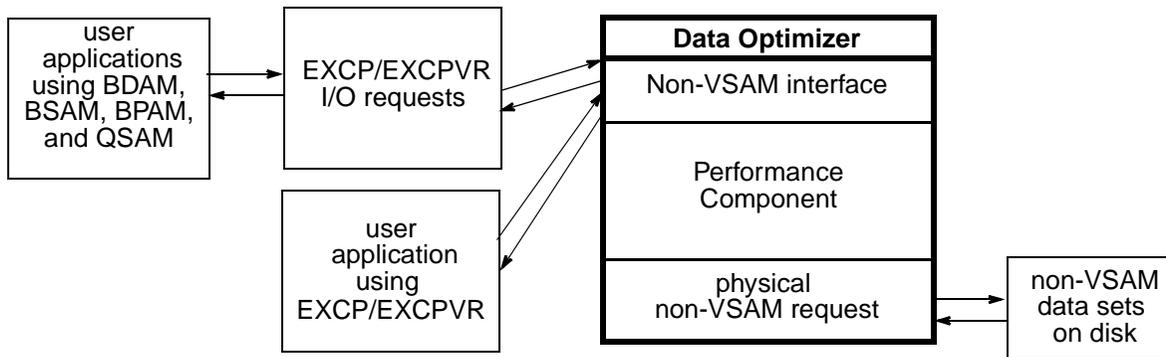
Figure 1-3 shows the non-VSAM performance component of Data Optimizer.

Figure 1-3 Non-VSAM Performance Component

normal non-VSAM processing



non-VSAM processing with Data Optimizer (advanced level)



Features

Data Optimizer provides the following features:

- **robust tuning techniques**

Data Optimizer provides a variety of sophisticated tuning techniques which provide the best possible performance benefits for each type of processing and file structure. Techniques range from simple buffer value changes to more advanced techniques such as dynamically caching buffers, overlapping I/O, building and using LSR pools, and complete rebuilding of channel programs.

- **dynamic tuning based on resource availability**

When a data set that is being opened has been selected for optimization, Data Optimizer checks the availability of a variety of system resources to ensure that capacity is available and that tuning values are chosen that will not negatively impact overall system performance.

- **dynamic region adjustment**

Data Optimizer can dynamically modify the region values in an address space to compensate for any additional I/O-related storage areas obtained on the application's behalf. This feature eliminates the need to increase the region value on the JOB or EXEC JCL statements to prevent S878 and S80A abends. It also prevents the application from seeing a decrease in the amount of *getmain-able* storage when optimization is used.

- **flexible implementation**

Data Optimizer provides a variety of methods for defining what is eligible for optimization benefits. The MAINVIEW Batch Optimizer dialog allows the creation of a data policy which provides a robust set of selection criteria including JOBNAME, DSNAME, SMSCLASS, and RACF associations. In addition, JCL DD override statements are provided for testing and ad-hoc optimization specifications at the job step or DD statement level. A user control exit interface is provided for users that have selection algorithms that are too complex to be accomplished with the data policy.

- **detailed informational messages and statistics reports**

Data Optimizer provides easy-to-understand messages that indicate when optimization has been performed. In addition, Data Optimizer indicates what source has requested optimization or exclusion and provides specific reasons when optimization is bypassed for special circumstances. Data Optimizer also provides detailed statistics reports that indicate information about the kind of processing performed and the amount of savings achieved. JOBNAME, STEPNAME, DDNAME, and DSNNAME information that is related to the data set being processed is also provided.

Benefits

Data Optimizer provides the following benefits:

- **larger data transfers per I/O operation**

Data Optimizer increases the amount of data movement per I/O request (when appropriate). The increased data movement decreases the number of I/O requests and reduces the elapsed time attributed to I/O.

- **improved I/O efficiency**

Data Optimizer learns and records an application's data access patterns and restructures I/O requests to take advantage of this information. It restructures the requests by using advanced I/O improvement techniques, which include improved buffering, overlapping I/O requests, and decreasing I/O path lengths.

- **reduced elapsed times**

Because Data Optimizer reduces the number of accesses to disk, it reduces or eliminates unnecessary wait times and significantly reduces the number of EXCPs. It automatically performs tuning to maintain optimum application performance without overusing system resources and negatively affecting overall system performance.

- **improved buffering and caching**

Data Optimizer automatically determines and sets the optimum buffering values for the utility or application. As I/O access patterns change for the utility or application, Data Optimizer adjusts these values.

- **virtual storage constraint relief**

Data Optimizer moves data areas above the 16 MB line whenever possible to maximize virtual storage constraint relief. This applies to both VSAM and non-VSAM processing.

Chapter 2 Understanding Data Optimizer

This chapter describes how Data Optimizer works and the features and functionality that it provides. This chapter discusses the following topics:

Overview	2-2
Selection Processing	2-2
Data Policy Definitions	2-2
Data Definition Override Statements	2-3
User Control Exit	2-3
Selection Criteria Order	2-4
Optimization Processing	2-7
User-Specified Options	2-7
System Resource Utilization	2-7
Historical Data	2-8
Performance Options	2-8
General Options	2-9
Resource Utilization Options	2-11
VSAM and Non-VSAM Options	2-11
Reporting	2-12
Performance Reports	2-12
Informational Messages	2-13

Overview

Data Optimizer is a powerful I/O optimization tool that can provide your site with significant performance benefits. This product contains many features and functions that enhance throughput performance by using dynamic and automatic I/O improvements. It also offers many user controls which provide flexibility for customizing it to your site's needs and environment. To realize maximum benefits from Data Optimizer, it is important to first understand how it works and the features and functions it provides.

Selection Processing

When the Data Optimizer component is enabled, it determines whether to provide performance benefits by getting control during every open request. After it identifies the data set and open requester, it determines whether any matching selection criteria exist. If Data Optimizer finds no matching selection criteria, optimization is bypassed. If it finds matching criteria, Data Optimizer continues processing by next evaluating the performance options associated with the selection criteria. The first option that is evaluated is the ACTION option. The ACTION option instructs Data Optimizer to intercept the open request for optimization (INCLUDE) or to bypass the request (EXCLUDE).

Selection criteria and their corresponding performance options can be specified to Data Optimizer in three ways:

- data policy definitions
- DD override statements
- user control exit

Data Policy Definitions

Data policy definitions are the most common means for specifying selection criteria and performance options. They are a powerful tool for managing optimization for a large number entities in one centralized location without requiring JCL changes. Data policy definitions are composed of two parts. The first part defines one or more selection criteria. Data Optimizer supports a robust set of selection criteria:

- file type (VSAM, non-VSAM, or both)
- data set name
- job name
- program name

- step name
- procedure step name
- data definition name (DD name)
- SMS data class name
- SMS storage class name
- SMS management class name
- security product group ID
- security product user ID
- execution job class
- time of day

These criteria can be specified independently or in combination to most accurately define the entity you wish to optimize. Wildcards are supported for the specified values.

The second part of the data policy definition is the performance options that you want to apply to the selected open request. For a complete discussion of creating and maintaining data policies, see Chapter 4, “Managing a Data Policy.”

Data Definition Override Statements

Data Definition (DD) override statements provide a simple means for specifying inclusion or exclusion and/or setting options in the JCL without having to create a data policy definition. They are particularly useful for short-term testing and ad-hoc changes. The DD override statements apply at the job step level and are position sensitive. If they appear as the first statement(s) after the EXEC statement (or STEPLIB), they apply to all DDs within the job step. If they appear after a particular DD, they apply to that DD only.

User Control Exit

Data Optimizer provides a user control exit interface which allows you to implement selection algorithms that are too complex for the data policy definition. An example would be if you wanted to include jobs that have jobnames whose fifth character position is a numeric between two and five. Although the user control exit interface is infrequently required, it can be a powerful tool when needed.

Selection Criteria Order

When Data Optimizer evaluates an open request, it must determine how to process it based on the three selection sources previously described. In some cases, the specifications of one selection source may override those of another. Data Optimizer uses a selection hierarchy to determine which selection criteria and performance options apply. DD override statements have the highest precedence in the hierarchy, followed by the user control exit, and finally the data policy definition. You can take advantage of this hierarchy to supply selection criteria and performance options from several sources. For example, a data policy definition could indicate inclusion for all data sets in a job along with various options and a DD override statement could be supplied to exclude performance processing for a specific DD statement.

Figure 2-1 illustrates the selection criteria hierarchy.

Figure 2-1 Data Optimizer Intercept Processing (Part 1 of 2)

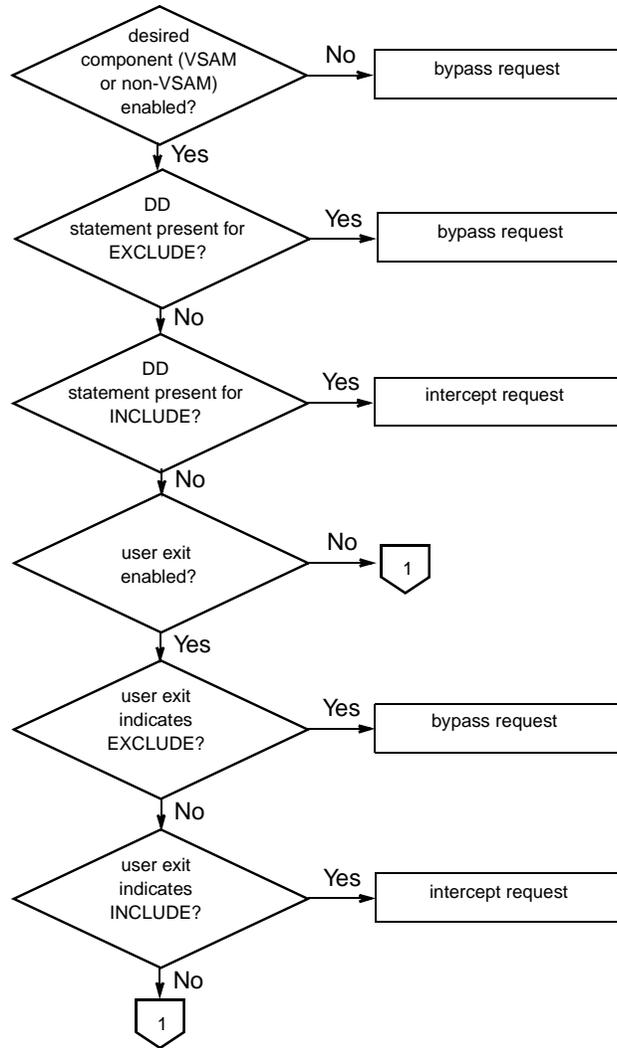
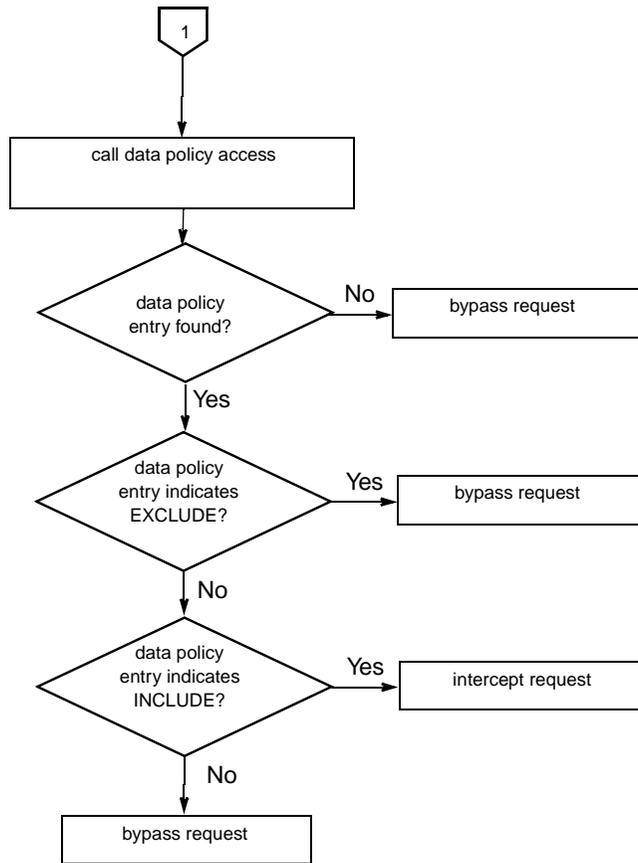


Figure 2-1 Data Optimizer Intercept Processing Flow (Part 2 of 2)



Optimization Processing

Once Data Optimizer has determined that an open request is to be included for optimization, it performs two functions. First, it determines if there are any circumstances that would prevent it from carrying out the request. This would include such things as unsupported file types or access methods. See Chapter 3, “Implementing Data Optimizer” and Appendix D, “Special Considerations” for a description of unsupported items. After Data Optimizer has determined that it can continue with the optimization process, it determines what type of optimization to perform based on the following criteria:

- user-specified options
- resource utilization
- historical data

User-Specified Options

The most common source for optimization criteria is user-specified options. These options can be specified explicitly using data policy definitions, DD override statements, the user exit, or can be specified implicitly by defaults which are also defined in the data policy. The user-specifiable options are summarized in “Performance Options” on page 2-8.

System Resource Utilization

When Data Optimizer is allowed to determine tuning values, it will take more into consideration than CI sizes, blk sizes, and so on. At the time of the open request, it will evaluate the current utilizations of key system resources such as CPU busy, paging rate, and storage availability and will take them into consideration when making tuning decisions. This results in highly tuned application processes that do not negatively impact overall system performance. You can change the thresholds for these values.

Historical Data

Certain I/O processes are impossible to fully evaluate at the time of open. For example, when a VSAM ACB specifies both sequential and direct processing, there is no way to know which type of access will predominate or whether one will be used at all. For this reason, Data Optimizer collects historical information during close processing, when appropriate, to use it for future reference. The next time an open occurs, Data Optimizer uses the information to make more informed tuning decisions. In this way, Data Optimizer learns the application access patterns to better provide performance benefits.

Data Optimizer has two methods to relate historical data to the file being opened. The first method, DSNPGM, uses the program name and data set name combination to access the data. The second method, FULL, uses jobname, stepname, procstepname, program name and DDNAME to access the historical data. The second method is the default method. Data Optimizer provides the HISTKEY command to set which method to use when creating the history. For more information on the HISTKEY command, see Chapter 6, “Data Optimizer Commands.”

If history has been previously recorded using the DSNPGM method then that data will be used to prime the decision-making process for the FULL method. The priming of the historical data will ensure that the buffering decision will be consistent.

Performance Options

Data Optimizer provides a large number of performance options that allow you to tune it to your applications and environment. These options are grouped into the following categories by function:

- general options
- resource utilization options
- reporting options
- non-VSAM options
- VSAM options

The more significant performance options are highlighted below. For a complete list of the performance options and their descriptions, see Chapter 5, “Data Optimizer Options.”

General Options

The general options allow you to control the general-purpose features and functions of Data Optimizer.

Action Option

The Action option is the most significant of all the performance options. It tells Data Optimizer whether to include or exclude the selected entity for optimization processing.

Optimization Mode Option

The Optimization Mode option specifies at a high level which performance techniques are used. For both VSAM and non-VSAM processing, Data Optimizer has the following optimization modes:

- ADVANCED
- BASIC
- REPORTONLY

The ADVANCED mode provides the most significant performance benefits by using the most advanced I/O optimization techniques such as dynamically building and using LSR buffer pools for VSAM processing and completely replacing the access method's I/O drivers to maximize data transfer rates for non-VSAM processing. These techniques maximize both read-ahead processing for sequential accesses and cache-hit performance for random accesses.

When the ADVANCED mode is not appropriate, BASIC mode can be used to provide benefits by optimizing buffer value specifications. For VSAM processing, optimum BUFNI and BUFND values are determined and implemented. Correspondingly, BUFNO values are optimized for QSAM processing and NCP values are optimized for BSAM, BPAM, and BDAM processing.

Data Optimizer also offers a REPORTONLY optimization mode which does not provide performance benefits, but rather provides you with a Data Set Information report which can give you an understanding of the type of processing and data management options that an application is using for a given data set access.

Dynamic Region Adjustment Option

When buffer values are increased as an I/O performance technique, two potential problems can arise. One is that this increased storage usage results in decreased storage availability for the application, which can occasionally cause the application's performance to degrade or cause the application to stop working altogether. The second problem is that increased buffer value specifications can lead to storage related abends, such as S80As, S878s, and so on.

To minimize these problems when buffers are increased, Data Optimizer will dynamically alter the region value for the step to compensate for the additional storage used. For example, if the region size specified on the EXEC statement is 512 KB and Data Optimizer increases the buffer usage by a total of 512 KB, then the Dynamic Region Adjustment feature will increase the internally represented region size to 1024 KB. This will make the apparent available storage to the application the same as it was without Data Optimizer and will minimize the risk of storage related abends. For additional information, see "Dynamic Region Adjustment Option" on page 5-8.

SMF Record Type Option

To aid in performance analysis, you can direct Data Optimizer to generate an SMF record during close processing for each optimized data set. These records can then be gathered and processed to provide analysis reports of I/O processing and benefits received. For additional information, see Appendix E, "SMF Recording," and Appendix G, "Batch Utilities."

Optimize STC Access and Optimize TSO Access Options

Because started tasks and TSO sessions typically access data sets which are not good candidates because of their small size, Data Optimizer bypasses optimization for started task and TSO address spaces. However, an application process is occasionally created in a started task or TSO environment that processes large data sets. For these situations, the Optimize STC Access and Optimize TSO Access options allow you to specify that optimization processing is to be provided for them. For additional information, see "Optimize STC Access Option" on page 5-8 and "Optimize TSO Access Option" on page 5-9.

Override User Values Option

Even before you install Data Optimizer, you may have already specified buffer values to the JCL of production jobs in an attempt to boost performance for them. In some cases, these values may not have been determined with a great deal of consideration for each specific application or they may have become outdated over time. In other cases, the tuning values may have been derived by very meticulous evaluation and represent the best-possible values for the particular application and should not be changed. For this reason, Data Optimizer provides the Override User Values option to allow you to determine whether you want to override the JCL-supplied buffer values. This option applies only if the values chosen by Data Optimizer are lower than the values specified in the JCL or access method control block (ACB)/data control block (DCB). For more information, see “Override User Values Option” on page 5-10.

Resource Utilization Options

As mentioned in “System Resource Utilization” on page 2-7, when Data Optimizer is allowed to select tuning values it will take into consideration the current use of several key system resources. Because a site may wish to artificially provide high-water marks for these resources to further limit resource consumption, Data Optimizer provides several Resource Utilization options. They allow you to specify high-water marks for CPU busy, paging rate, and private storage usage both above and below the 16 MB line.

VSAM and Non-VSAM Options

When the optimization mode is specified without other tuning options, Data Optimizer will dynamically evaluate a variety of criteria and will select what it perceives to be optimum techniques and values. Most of the time, allowing Data Optimizer to make these decisions will yield the maximum possible benefits. There are occasions however, when users may want to specify values other than those chosen by Data Optimizer. For this reason, Data Optimizer offers a wide variety of tuning controls that allow you to customize Data Optimizer processing to your specific applications and environment. These controls can be used to alter buffer value specifications, set maximum buffer values, etc. for both VSAM and non-VSAM processing.

Reporting

Data Optimizer provides a variety of messages and reports to make it easy to understand the type of processing performed and the amount of benefits received.

Performance Reports

Performance reports provide comprehensive status and processing information and share a common header that identifies the open request being processed. Data Optimizer provides the following performance reports:

- **Options Report**

The Options report describes the options that are supplied by the data policy global statement, data policy define statement, or DD override statement(s).

- **Statistics Report**

The Statistics report provides I/O statistics information, such as records counts and I/O savings that are provided.

- **Data Set Information Report**

The Data Set Information report provides a variety of information about the data set being processed, including unit type, volume serial, data set organization, open type, and so on.

- **Catalog Report (VSAM only)**

The Catalog report provides VSAM data set statistics that are maintained in the catalog, such as high-used RBA, number of CI-splits, and so on.

These reports can be requested individually or together, and will appear in a single printout (one each for VSAM and non-VSAM processing). The Data Set Information and Catalog reports can also be requested without requesting optimization using the REPORTONLY optimization mode.

The Reporting options provide control over the types of information that will be generated about data sets and messages. For information on how to request reports and control they types of information that will be generated, see Chapter 5, “Data Optimizer Options.” For a complete description of the contents of these reports, see Appendix A, “Data Optimizer Reports.”

Informational Messages

Data Optimizer provides informational messages to indicate whether an open request was selected for inclusion or exclusion and if so, what the source of the selection was (data policy, DD override statement, or user control exit). If the open request is selected for inclusion, a second message is also provided to indicate whether interception was actually performed or bypassed, and if bypassed, what the reason was.

To specify the detail of messages desired, Data Optimizer provides three different general-purpose message levels.

- Message level 0 indicates that no messages are to be generated.
- Message level 1 indicates that messages are to be generated only when the request has been explicitly excluded or was bypassed for some reason.
- Message level 2 indicates that messages are always to be generated.

A fourth, special-purpose message level is also provided. Message level 3 is a special message level that is used for diagnostic purposes only. It should be used only when directed by BMC Software Product Support. Depending on the circumstances, it can generate a large volume of messages to the operator's console.

For a complete description of the message levels, see “Message Level Option” on page 5-20.

Chapter 3 Implementing Data Optimizer

This chapter describes the tasks that must be performed to implement Data Optimizer. This chapter discusses the following topics:

Overview	3-2
Setting Goals	3-2
Assessing Candidates	3-3
Locating Candidates	3-5
Moving Optimization Into Production	3-6
Initial Strategy	3-6
Intermediate Strategy	3-6
Advanced Strategy	3-6

Overview

In order to receive the greatest possible benefits from Data Optimizer, it is important to develop a well thought out implementation plan. The implementation plan should consist of three important activities:

- setting goals
- choosing candidates
- moving optimization into production

Setting Goals

To successfully implement, you must first determine what type of benefits you want to achieve with the product. For example, is your primary objective to provide maximum performance benefits for certain mission-critical applications, to provide overall throughput improvement for your batch window, or to provide performance improvements for certain heavily used data sets when accessed by any application? When making these decisions, it is important to take into consideration the amount of system resources that you have available. Providing performance benefits to a large number jobs requires a greater amount of available resources. If these resources are not available, performance benefits will diminish.

After you decide what you wish to accomplish, you should determine the criteria that best describes the entities you wish to optimize. To do this, you must have a thorough understanding of the production standards and conventions used in your environment. In some environments, applications are designated by the first three characters of the jobname, such as *PAY*. In others, all hot-batch is run in jobclass *A*, *L*, or *Q*. Further, application data sets may be differentiated by a data set name prefix, such as *ACCTRCV*. The criteria that your installation uses for defining its workloads will most likely determine the selection criteria that you use in data policy definitions for Data Optimizer. The objectives you have will also influence the selection criteria you use. Using the objectives described in the paragraph above, here are some possible guidelines for choosing selection criteria:

- When attempting to provide performance benefits for certain applications, you would probably want to use selection criteria such as jobname (or jobname prefix), jobclass, or data set name (or data set name prefix).
- When attempting to provide more general throughput improvements, you would probably want to use selection criteria such as jobclass, SMS class, or possibly time-of-day.

- When attempting to provide performance improvements for specific data sets or groups of data sets, you would probably use data set name (or data set name prefix) or SMS class.

Because you may have several objectives you wish to accomplish, you can mix the types of selection criteria you use. Some data policy definitions may use jobnames for their selection criteria while others use data set name. These selection criteria can be combined to further qualify optimization recipients. For example, you could create a data policy definition that would provide optimization for data set *PAY01.GL.DATA* only when it is accessed by jobs whose names are prefixed with *PAYR*.

Assessing Candidates

The next important step in the implementation process is to determine what to optimize. As mentioned previously, it is important to keep in mind that attempting to optimize everything can unnecessarily consume valuable system resources. Although the criteria that determines a good candidate varies from environment to environment, there are three common and significant qualities:

- important to business/mission critical
- good potential for savings
- supported by the product

Identifying the work which is important to your company's business is a crucial part of planning I/O optimization. The best way to maximize value from Data Optimizer is to selectively optimize the work that is considered mission-critical to your business or heavily impacts your batch window. By doing so, you spend your system resources as wisely and efficiently as possible. There is little value in providing performance benefits for work that is insignificant to your company's business mission.

The next quality to consider when selecting performance candidates is the potential for realizing I/O and elapsed time savings. Although many factors can contribute to this, the following items can serve as a guideline in making this determination:

- programs that are I/O intensive
- data sets of at least 10 cylinders
- batch address spaces (typically not TSO and STC)
- programs that are not already highly-tuned including, but not limited to utilities such as DF/SORT, SYNC SORT, DF/DSS, and ICKDSF.

The last quality you must consider when determining performance candidates is whether the program, access method, file type, and so on are supported by Data Optimizer. Data Optimizer supports a wide variety of accesses and file types, but certain items are not supported for various reasons.

Table 3-1 shows the items that are eligible for performance benefits and the items that are not eligible.

Table 3-1 Summary of Data Optimizer Support (Part 1 of 2)

Data Set Type Basic	Basic Processing	Advanced Processing
Non-VSAM		
Basic sequential access method (BSAM)	X	X
Queued sequential access method (QSAM)	X	X
Basic direct access method (BDAM) for non-keyed, physical sequential data sets	X	X
Hiperbatch access to QSAM data sets	X	—
Basic partitioned access method (BPAM)	X	—
Data sets used by started tasks, see “Optimize STC Access Option” on page 5-8.	X	—
Data sets allocated by TSO users, see “Optimize TSO Access Option” on page 5-9.	X	—
Extended partitioned data sets (PDS/Es)	X	—
Striped data sets	X	—
Tape data sets	X	—
VIO data sets	—	—
Execute channel program (EXCP) macro and execute channel program virtual real (EXCPVR) macro	—	X
Keyed data sets, such as indexed sequential access method (ISAM) data sets	—	—
Keyed access to basic direct access method (BDAM) data sets	—	—
Data sets on any DASD device older than a 3380	—	—
EXCP and EXCPVR access to tape data sets	—	—
Data sets that are opened as update-in-place (UPDAT)	—	—
Utilities that are already tuned (such as BMC Software Database Utilities, SYNCSORT, DFSORT, DFDSS, ICKDSF, and IAM data sets)	—	—
Data sets accessed by a restricted program. See “Excluded Programs” on page D-6 for a list of restricted programs.	—	—
Data sets referenced by a restricted DD statement. See “Excluded DD Statements” on page D-9.	—	—
Restricted data set names, see “Excluded Data Sets” on page D-10.	—	—
VSAM		

Table 3-1 Summary of Data Optimizer Support (Part 2 of 2)

Data Set Type Basic	Basic Processing	Advanced Processing
Key-sequenced data sets (KSDSs)	X	X
Entry-sequenced data sets (ESDSs)	X	X
Fixed-length relative record data sets (RRDSs)	X	X
Variable-length relative record data sets (VRRDSs)	X	X
Applications that use: <ul style="list-style-type: none"> • Global shared resource (GSR) processing • Improved control interval (ICI) processing 	—	—
Files and databases accessed by DB2 regions and DB2 utilities	—	—
Databases accessed by IMS regions and IMS utilities	—	—
Files accessed by CICS regions	—	—
Data sets that were allocated with cross-region or cross-system SHAREOPTIONS(4); for example, SHAREOPTIONS(4 x) or SHAREOPTIONS(x 4)	—	—
Hiperbatch processing of VSAM data sets	—	—
Utilities that are already tuned (such as BMC Software Database Utilities)	—	—
Data sets accessed by a restricted program. See “Excluded Programs” on page D-6 for a list of restricted programs.	—	—
Data sets referenced by a restricted DD statement. See “Excluded DD Statements” on page D-9 for a list of restricted DD statements.	—	—
Restricted data set names. See “Excluded Data Sets” on page D-10.	—	—

Locating Candidates

After you have determined what constitutes a good candidate in your environment, you must proceed by locating the entities that match those criteria. There are two primary ways of gathering this information. The first is MAINVIEW Batch Optimizer's Candidate Utility. Using SMF data, this utility provides estimated elapsed time savings for jobs or data sets that you deem as important. The second method for locating candidates is to use an in-house or OEM utility that can process SMF data such as SAS or MXG. By setting up such a utility to filter and report entities that match your criteria for good candidates, you can create a valuable tool that is customized to your site's specific needs.

Moving Optimization Into Production

No matter what your objectives and candidates are for Data Optimizer, gradually phasing in optimization is usually the safest and easiest way to ensure a non-disruptive implementation. Using this approach, you can learn how the product behaves in your environment as you go and can minimize the impact of unforeseen product or environmental incompatibilities. The following strategies indicate how Data Optimizer can be moved into production in using a gradual approach.

Initial Strategy

If you have a separate test system, start by bringing Data Optimizer up on it to become familiar with the product's structure and behavior. Create and run test jobs. Copies of production jobs work well for this as they allow you observe how your applications will behave in a non-impactive situation. For initial testing, the DD override statements provide the quickest way to control the product. Try the options, review the reports, and observe the product's behavior to gain familiarity.

Intermediate Strategy

After you have become familiar with using Data Optimizer in a test environment, bring it up on your production system(s). Proceed by selecting one or more less critical production jobs and placing them under control of Data Optimizer. At some point you may want to begin using the MAINVIEW Batch Optimizer dialog to create a data policy for defining what will be optimized and what options will be used.

Advanced Strategy

Once you become comfortable with Data Optimizer in a production environment, you can continue to place more jobs under Data Optimizer's control and to include more mission-critical jobs and applications. For jobs requiring maximum benefits, you can experiment with the many tuning controls available to fine-tune Data Optimizer for specific programs, file types, and environmental conditions.

Chapter 4 Managing a Data Policy

This chapter describes how to view or modify data policy definitions, activate policies, and determine which policy will be used, including information about the panels and dialog actions that you use to accomplish these tasks. This chapter discusses the following topics:

Overview	4-2
Before You Begin	4-2
Specifying Selection Criteria for Data Policies	4-2
Ordering Definitions within a Data Policy	4-3
Setting Global Options for a Data Policy	4-4
Panel Flow for a Data Policy Registration	4-4
Creating or Editing a Data Policy Definition	4-5
Editing or Viewing Data Policy Global Options	4-15
Activating a Data Policy	4-17
Activating a Data Policy by Using a Command	4-18
Activating a Data Policy by Using the Dialog	4-18
Determining Which Data Policy is Active	4-20
Determining which Subsystem Processes a Data Set	4-21

Overview

Data policies allow you to define when and how optimization processing is to occur from a single centralized repository. Data policy definitions associate one or more selection criteria with a set of performance options. When an open request matches the selection criteria, Data Optimizer uses the performance options specified in the data policy definition to determine how to provide performance improvements. This chapter provides information for using the MAINVIEW Batch Optimizer dialog to establish data policy definitions.

Before You Begin

If you have not already read Chapter 1, “Introduction,” and Chapter 2, “Understanding Data Optimizer,” please do so before continuing with this chapter.

This chapter describes how you can use the MAINVIEW Batch Optimizer dialog to perform the following data policy definition tasks:

- review or modify data policy global defaults in a control data set
- review or modify a data policy definition in a control data set

The MAINVIEW Batch Optimizer customization process creates a sample data policy definition in the control data set. The panels and pop-ups shown in this chapter reflect that sample data policy.

The data policy registration task descriptions in this chapter assume that you are familiar with the dialog’s initial panel sequence and basic navigation techniques. For a complete description of this facility, see Appendix B, “Navigating the User Interface.”

Note: Help is available online for all panels, pop-ups, and fields while registering BatchPlex and policy definitions. For details about obtaining help online, see “Online Help” on page B-24.

Specifying Selection Criteria for Data Policies

Data Optimizer provides selection criteria that allow you to control when I/O performance processing occurs. These criteria include the following items:

- type of file to be used (VSAM, non-VSAM, or both)
- data definition name (ddname)
- data set name (dsname)

- job name
- program name
- step name
- procedure step name
- SMS data class name
- SMS storage class name
- SMS management class name
- security product user ID
- security product group ID
- job class
- time of day

You can use one or all of these items when specifying the criteria.

For example, if you want the data set I/O of a specific SMS data class to be processed as quickly as possible, you can create a definition in a data policy that identifies the SMS data class name and the I/O performance options that should be used with that SMS data class name. A more complex selection criteria Data Optimizer could intercept all VSAM data sets for performance processing that have a specific high-level qualifier and that are being processed during a specific time range by a specific job.

Ordering Definitions within a Data Policy

As Data Optimizer searches through the definitions in a data policy, the first match that Data Optimizer encounters is the definition that provides I/O performance processing. Therefore, the order in which you place definitions in a data policy is extremely important.

The following example shows order with a definition.

Example

Basic Example of Definition Ordering for a Data Policy

Suppose that you have two definitions. One definition instructs Data Optimizer to exclude all jobs from I/O performance processing. The other definition instructs Data Optimizer to include selected jobs for I/O performance processing. When you specify the definitions in the data policy, you must list the “include specific jobs” definition before the “exclude all jobs” definition.

Setting Global Options for a Data Policy

When you are creating definitions for data policies, you can leave any of the I/O performance options blank. This instructs Data Optimizer to use the global options value, which are set using the Options pull-down from the Data Policy Definition panel.

Global options provide two benefits:

- easier data entry

The values that you select for some I/O performance options will remain constant. For example, your installation will probably always choose Y or N (yes or no) for the Dynamic Region Adjustment option. You can set the Global option to YES or NO, and allow the dialog to take the default option by leaving the field blank in the data policy definitions.

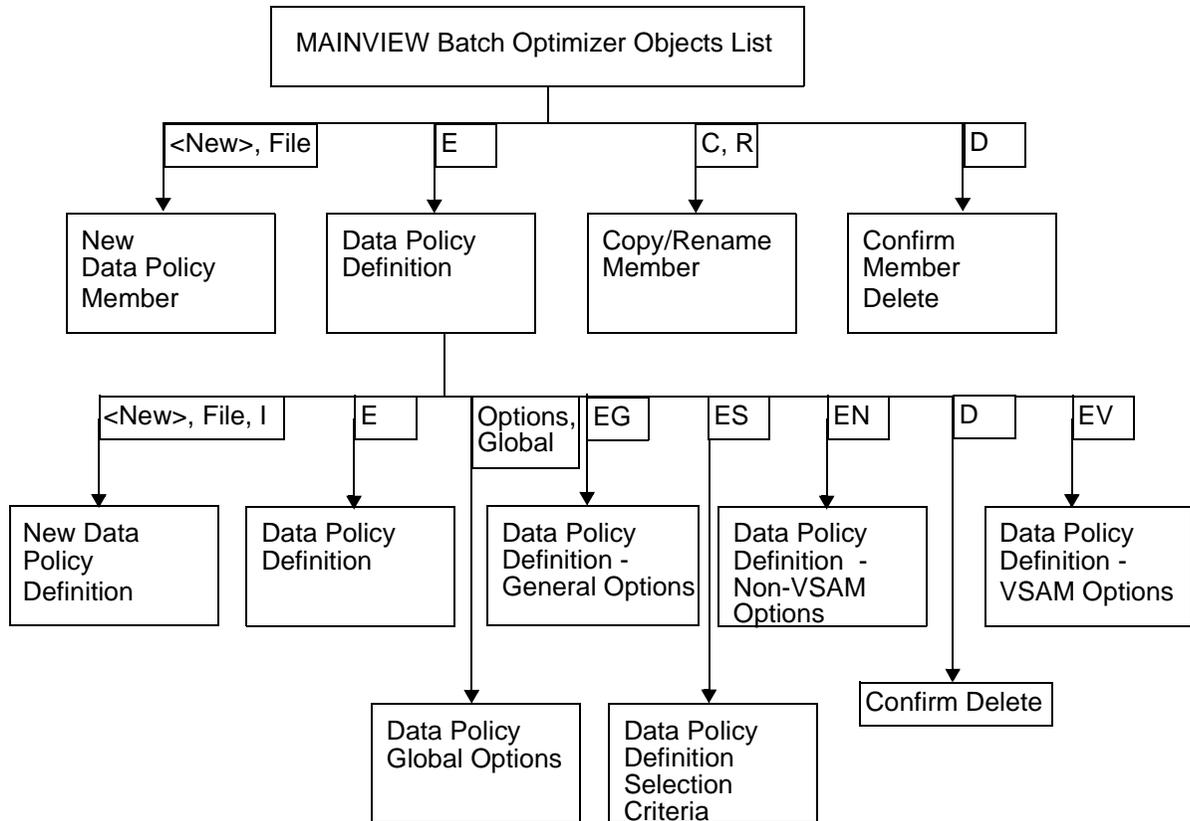
- global toggling

When you leave an I/O performance option value blank, you can globally change a value without accessing each definition and changing the value. For example, you might leave the actual I/O count option blank in all definitions and globally toggle the option on and off for different batch runs.

Panel Flow for a Data Policy Registration

Figure 4-1 shows the basic panel flow for the data policy registration area of the MAINVIEW Batch Optimizer dialog. Figure 4-1 indicates the action codes (E, ES, EG, EN, EV, I, C, R, D), commands (New, Globals), or the pull-down option (File, Options) that you select to move from one panel to another.

Figure 4-1 Data Policy Registration–Dialog Panel Flow



Creating or Editing a Data Policy Definition

You can review or modify the information that is saved for a data policy definition associated with the policy.

To edit or view information for a data policy in the control data set, complete the following steps:

- Step 1** Access the MAINVIEW Batch Optimizer Objects List panel.
- Step 2** Type **E (Edit)** in the action entry field to the left of your choice of data policy and press **Enter**.
- Step 3** Review the Selection Criteria

The Data Policy Definitions panel (Figure 4-2 on page 4-7) is displayed. The Data Policy Definitions panel lists the definitions in the data policy. The list reflects your choice of definition order.

The Selection Criteria column contains an abbreviated version of the selection information for each definition statement in the Data policy member. If more selection criterion exists than can be displayed in the allotted space for a definition, a plus (+) sign is displayed at the end of the affected definition. You may view a complete list of selection criterion by positioning your cursor beneath the “+” and pressing **Enter** or by using the **ES** line command to select this definition. You can view the following abbreviations for selection criterion by positioning your cursor beneath the Selection Criteria column and pressing **F1** for Help.

- DSN—data set name
- JBN—job name
- PGM—program name
- UID—RACF user ID
- DTY—data set type (VSAM, NONVSAM, and ALLFILES)
- JBC—job class
- STP—step name
- PRC—procedure step name
- DDN—data definition (DD) name
- DCL—SMS data class
- SCL—SMS storage class
- MCL—SMS management class
- GRP—RACF group
- TIM—time range (*hhmm-hhmm*)

Step 4 Review the **Actions** and **Method** values and modify if needed.

For each definition, the list indicates the associated action and method to employ (the **Actions** and **Method** fields). These two fields indicate the options that you explicitly specified or they indicate implied behavior based upon your selections. When creating or modifying a statement, if you indicate that Action and/or Method should use the default value, these columns display the result of that choice. To view the options that are explicitly specified for a particular statement, select the definition statement by using one of the Edit line commands.

Figure 4-2 is an example of the Data Policy Definitions panel. The customization process created this sample data policy. You can use this data policy for practice as you learn to manipulate data policies by using the MAINVIEW Batch Optimizer dialog. You can edit this data policy and tailor it to suit the needs of your site, or you can delete this data policy and create a new one.

Figure 4-2 Data Policy Definitions Panel

```

File   View   Display   Applications   Options   Help

                                Data Policy Definitions           Row 1 to 17 of 17
Command ==> _____ SCROLL ==> CSR_

Data Policy Information                               System: SYSL
Name . . : DATPOL00                                  SMF ID: SYSL
Comment . Sample data policy                         Date : 2000/06/11
                                                    Time : 14:49:26

Type an action code. Then press Enter.

Edit line commands:                                Statement line commands:
E=All fields          EG=Common fields              C=Copy D=Delete I=Insert
EN=NonVSAM fields    EV=VSAM Fields                 X=Cut   P=Paste
ES=Selection criterion only

      Selection Criteria                                Action Mode Response
.. <New>
.. PGM(EASY*) DTY(VSAM)                               Exclude Advanced
.. JBN(APPRODJ*)                                       Exclude Advanced
.. PGM(EZT*) DTY(VSAM)                                Exclude Advanced
.. PGM(IEBDG) UID(MEB*)                               Include Advanced
.. PGM(IEBCOMPR) UID(MEB*)                            Include Advanced
.. DSN(PROD.AP.VSAM.BIAS01.** ) DTY(VSAM)              Include Advanced
.. DSN(PROD.AP.VSAM.BIAS03.** ) DTY(VSAM)              Include Advanced
.. DSN(PROD.AP.VSAM.BIAS05.** ) DTY(VSAM)              Include Advanced
.. DSN(PROD.AP.VSAM.BIAS10.** ) DTY(VSAM)              Include Advanced

```

The Data Policy Definitions panel is a table display and includes a **Row** field. Figure 4-2 shows this panel as if it is displayed on a screen large enough to display all table rows.

Step 5 *Optional.* Change any input field on this screen by positioning the cursor in the field and typing a new value. Clear any characters that remain from the previously displayed value.

Step 6 *Optional.* Change the Data Policy Definitions by typing one of the following edit line commands to the left of the policy that you want to change:

- **E**—displays the complete definition
- **EG**—displays optimization options that affect all data sets
- **EN**—displays optimization options that affect non-VSAM data sets only
- **EV**—displays optimization options that affect VSAM data sets only
- **ES**—displays selection criteria only

The individual edit line commands allow you to reference only the data that is of interest to you. The commands are useful when you want to view or modify a subset of the optimization options for a specific definition.

If you select options *E*, *EG*, *EN*, or *EV*, the Data Policy Definition panel (or a variant of this panel with the appropriate subset information) is displayed.

Note: The Data Policy Definitions panel contains two scrollable areas; one for the selection criteria, and one for the data optimization options. To scroll a specific area, place the cursor within the area and enter a scroll command. The active selection criteria are displayed at the top of the list; blank specifications are displayed at the bottom. If you see one or more blank selection criteria, there is no need to scroll down for more selection criteria.

Figure 4-3 is an example of the Data Policy Definition panel. The customization process defines these sample data policy options.

Note: The Data Policy Definition panel has two scroll fields that are designated by the More specification; one for the selection criterion, and one for the optimization options.

Figure 4-3 Data Policy Definition Panel

```

File  View  Display  Applications  Options  Help
-----
BSSP00MA                      Data Policy Definition
Command ===> _____ SCROLL ===> CSR_

Data policy name . . . . . : DATPOL00
Definition comment . . . . . _____

Selection Criterion          Value                                     More:  +
-----
Data set name . . . . . PROD.AP.VSAM.BIAS01.** _____
Data set type . . . . . VSAM_____ +
Job name . . . . . _____
Program name . . . . . _____

General options              Value          Valid settings          More:  +
-----
Action . . . . . INCLUDE          + Include Exclude
Optimization mode . . . . . _____ + Basic Advanced Reportonly
Dynamic region adjustment . . _          Y=Yes N=No
Resource usage bias . . . . . 01          1-10 99
Honor RUB value . . . . . _          Y=Yes N=No
Optimize STC access . . . . . _          Y=Yes N=No
Optimize TSO access . . . . . _          Y=Yes N=No
Optimize TAPE data sets . . . . . _          Y=Yes N=No
Optimize PO data sets . . . . . _          Y=Yes N=No
Optimize EXCP access . . . . . _          Y=Yes N=No
Optimize DA data sets . . . . . _          Y=Yes N=No
SMF record type . . . . . _____          0, 128-255
Override user values . . . . . _          Y=Yes N=No
Resource Utilization options  Valid settings
Page fault rate threshold . . _____          0-99999 CPU utilization threshold
    
```

Step 7 *Optional.* Change any input field as necessary.

You can make changes to any input field on this panel by positioning the cursor in the field and typing a new value. Clear any characters that remain from the previously displayed value. Fields followed by a plus (+) sign can be prompted for a list of valid values by positioning the cursor to the field and pressing **F4** (Prompt) or by positioning the cursor beneath the “+” and pressing **Enter**.

Note: The column to the right of the input fields describes the valid settings for those fields.

Step 8 *Optional.* Change the display if you wish to modify a different subset of information or to see the comprehensive list of options.

You can use **F11** to toggle through the various displays or use the Display pull-down option and directly select a subset. If you display your function keys on the panel, the dialog sets the **F11** label to the name of the next subset of data.

Step 9 *Optional.* Determining the current default (global) values. If desired, use **F6** or the View pull-down option to toggle the column to the right of the optimization value input fields between the valid settings and the default values display. This is useful in determining the resulting behavior for optimization values that are blank, implying the global default.

If you press **F6** (Defaults), the default values are displayed (Figure 4-4).

Figure 4-4 Data Policy Definition Panel–Default Values

```

File   View   Display   Applications   Options   Help

                                Data Policy Definition
Command ==>> _____ SCROLL ==>> CSR_

Data policy name . . . . . : DATPOL00
Definition comment . . . . . _____

Selection Criterion          Value

                                More:   +
Data set name . . . . . PROD.AP.VSAM.BIAS01.** _____
Data set type . . . . . VSAM _____ +
Job name . . . . . _____
Program name . . . . . _____
                                More:   +

General options              Value (or D)  Default values
Action . . . . . INCLUDE_  + EXCLUDE
Optimization mode . . . . .                + ADVANCED
Dynamic region adjustment . .                Y
Resource usage bias . . . . . 01           99
Honor RUB value . . . . .                N
Optimize STC access . . . . .                N
Optimize TSO access . . . . .                N
SMF record type . . . . .                0
Override user values . . . . .                N
Resource Utilization options                Default values
Page fault rate threshold . .                0
    
```

Note: The column to the right of the input fields now describes the default values that are specified in the Defaults (Global) panel.

Step 10 *Optional.* Select option ES to display the Data Policy Definition Selection Criterion pop-up panel.

In this panel, you can modify the selection criterion for a data policy definition. You can make changes to any input field on this screen by positioning the cursor in the field and typing a new value. Clear any characters that remain from the previously displayed value. Fields followed by a plus (+) sign can be prompted for a list of valid values by position the cursor to the field and enter F4 (Prompt) or by positioning the cursor beneath the “+” and pressing **Enter**.

Figure 4-5 is an example of the Data Policy Selection Criterion panel. This panel displays the complete list of available selection criterion in a scrollable area (designated by the More specification).

Figure 4-5 Data Policy Selection Criterion Panel

```

Data Policy Definition Selection Criterion

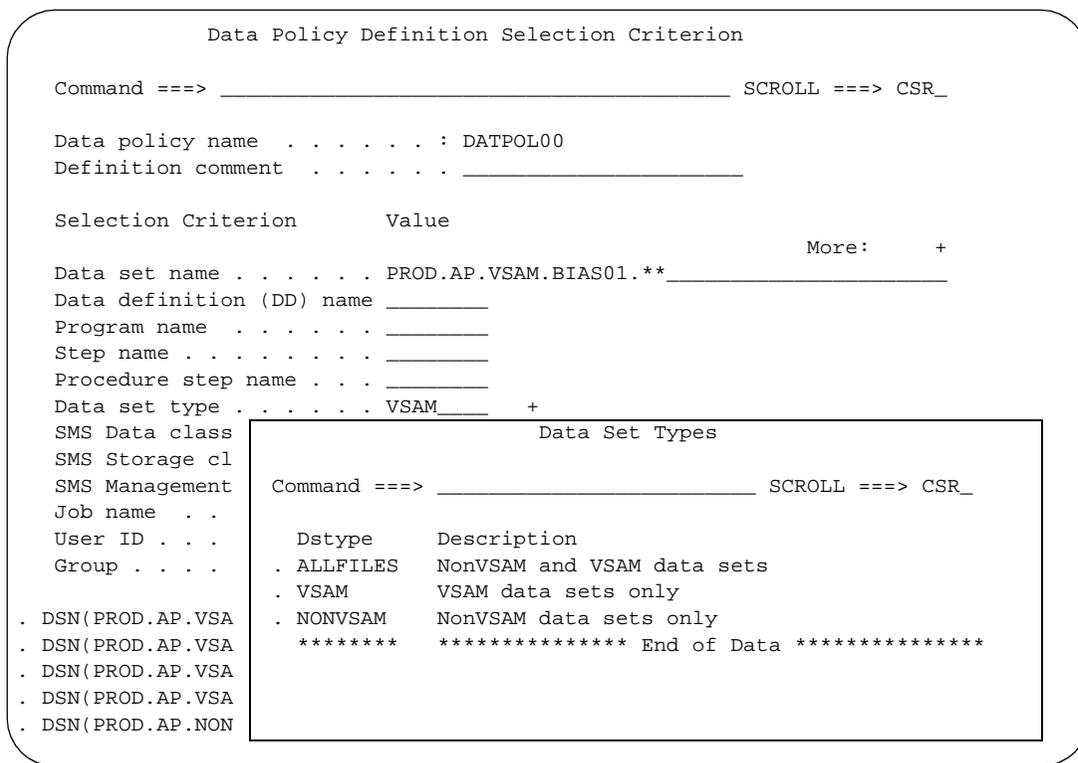
Command ==> _____ SCROLL ==> CSR_

Data policy name . . . . . : DATPOL00
Definition comment . . . . . _____

Selection Criterion      Value
Data set name . . . . . PROD.AP.VSAM.BIAS01.** _____ More:  +
Data definition (DD) name _____
Program name . . . . . _____
Step name . . . . . _____
Procedure step name . . . _____
Data set type . . . . . VSAM_____ +
SMS Data class . . . . . _____
SMS Storage class . . . . . _____
SMS Management class . . . _____
Job name . . . . . _____
User ID . . . . . _____
Group . . . . . _____
    
```

Figure 4-6 is an example of the Data Policy Selection Criterion panel after using the Prompt command to request a list of acceptable data set types. Select an appropriate data set type by placing the cursor next to the desired value and pressing **Enter**.

Figure 4-6 Data Policy Selection Criterion Panel—Prompt Active



Note: On some smaller displays stations (24 lines by 80 columns), the Prompt pop-up may obscure fields that you need to see to aid you with your decision. Use the IBM ISPF Window command to move the pop-up display.

Step 11 *Optional.* Make changes to the list of definitions by performing the administrative tasks described in Table 4-1 on page 4-13.

Note: The order of data policy definitions affects Data Optimizer operation.

Table 4-1 Data Policy Definition Administrative Acts (Part 1 of 2)

Objective	Action	Result
<p>add a new definition to the top of the list</p>	<ul style="list-style-type: none"> • select 1 (New) from the File pull-down option and type E next to the <New> line –or– • enter New on the command line Press Enter to process your selection 	<p>The dialog displays the New Data Policy Definition panel.</p> <p>Make changes to any selection criteria and/or optimization options input fields and press Enter to return to the Data Policy Definitions panel (Figure 4-2 on page 4-7).</p> <p>The added definition will appear with the word “Added” beneath the Response column next to the definition.</p>
<p>insert a new definition, (similar to add, but you choose the position in the list)</p>	<ul style="list-style-type: none"> • decide where in the list you wish to place the new definition • position the cursor to the left of the statement immediately preceding this position • type I (for Insert), and press Enter 	<p>The New Data Policy Definition panel is displayed with a message describing where this definition is to reside in the list.</p> <p>Make changes to any selection criteria and/or optimization options input fields and press Enter to return to the Data Policy Definitions panel (Figure 4-2 on page 4-7).</p> <p>The inserted definition will appear at the selected point in the list and with the word “Inserted” beneath the Response column next to the definition.</p>

Table 4-1 Data Policy Definition Administrative Acts (Part 2 of 2)

Objective	Action	Result
<p>copy (and paste) a definition</p>	<ul style="list-style-type: none"> • position the cursor in the action entry field to the left of your choice of definition • type C (for Copy), and press Enter <p>The dialog response depends on the clipboard contents:</p> <ul style="list-style-type: none"> • If the clipboard is empty, the dialog issues a message and places Copied beneath the Response column next to the definition to indicate that the definition has been copied. • If the clipboard contains a definition, the dialog displays a Confirm Clipboard Overwrite pop-up. <p>To empty the clipboard, press Enter. The dialog copies the cut definition from the list to the emptied clipboard. To cancel the operation, press F12 or enter Cancel.</p> <p>Decide where in the list you wish to place the copied definition. Position the cursor to the left of the statement immediately preceding this position, type the letter P (for Paste), and press Enter.</p>	<p>The dialog moves the definition from the clipboard to the list at the selected position. If you want to edit the copied definition, see "Creating or Editing a Data Policy Definition" on page 4-5</p>
<p>move (cut and paste) a definition</p>	<p>position the cursor in the action entry field to the left of your choice of definition. Type X (for Cut), and press Enter. The dialog response depends on the clipboard contents:</p> <ul style="list-style-type: none"> • If the clipboard is empty, the dialog moves the definition from the list to the clipboard and issues a message. • If the clipboard contains a definition, the dialog displays a Confirm Clipboard Overwrite pop-up. To empty the clipboard, press Enter. The dialog moves the cut definition from the list to the emptied clipboard. To cancel the operation, press F12 or enter Cancel. <p>Decide where in the list you wish to place the cut definition. Position the cursor to the left of the statement immediately preceding this position, type the letter P (for Paste), and press Enter.</p>	<p>The dialog moves the definition from the clipboard to the list. If you want to edit the moved definition, see "Creating or Editing a Data Policy Definition" on page 4-5.</p> <p>Warning: If you save a Data Policy member after you have "Cut" a definition but before you "Paste" it, the "Cut" (moved) definition will not be saved in the Data Policy member.</p>

Step 12 *Optional.* Delete a definition. Position the cursor in the action entry field to the left of your choice of definition.

12.A Type **D** (for Delete).

12.B Press **Enter**.

The Confirm Delete pop-up is displayed, which allows you to confirm the delete request before the dialog removes the definition from the data policy.

Step 13 Save your changes by using one of the following methods:

13.A Select the File pull-down choice **2** (Save policy).

The dialog saves the changed data policy and redisplay the same panel.

13.B Type **Save** on the command line.

The dialog saves the changed data policy and redisplay the same panel.

Press **F3** (Exit). The dialog displays a Confirm Exit pop-up. Type **1** (Save changes to disk and exit), and press **Enter**. The dialog saves the changed data policy, and the MAINVIEW Batch Optimizer Objects List panel is displayed with “Edited” displayed beneath the Response column, next to the changed policy.

Editing or Viewing Data Policy Global Options

You can review or modify the global option values saved for a data policy.

To edit or view global options for a data policy in the control data set, complete the following steps:

Step 1 Access the Data Policy Global Options pop-up.

1.A Position the cursor on the action bar Options option, and press **Enter**.

The Options pull-down is displayed.

1.B Type **5** (Data policy global options), and press **Enter**.

Note: You can also access this option by entering **Globals** on the command line.

Step 2 The dialog Data Policy Global Options pop-up is displayed. Specify values that serve as the default optimization values.

These values are used in cases where an optimization value on a data policy was left blank. You can specify or change the same information on the Data Policy Global Options pop-up as you did on the Data Policy Definition panel (Figure 4-3 on page 4-8).

Figure 4-7 is an example of the Data Policy Global Options pop-up. The customization process defines these sample global options. If you leave the value for an optimization option blank on the Data Policy Definition panel, Data Optimizer uses the values that you specified on the Data Policy Global Options pop-up.

Figure 4-7 Data Policy Global Options Pop-up

```

----- Data Policy Global Options -----
|BSSP00MG
|Command ==>> _____
|
|Type values in fields below. Press Enter to continue.
|
|Data policy name . . . . . : DATPOL02
|
|General options          Value          Valid settings          More:  +
| Action . . . . . EXCLUDE          + Include Exclude
| Optimization mode . . . . . ADVANCED__ + Basic Advanced Reportonly
| Dynamic region adjustment . . Y          Y=Yes N=No
| Resource usage bias . . . . . 99          1-10 99
| Honor RUB option . . . . . N          Y=Yes N=No
| Optimize STC access . . . . . N          Y=Yes N=No
| Optimize TSO access . . . . . N          Y=Yes N=No
| Optimize TAPE data sets . . . Y          Y=Yes N=No
| Optimize PO data sets . . . . Y          Y=Yes N=No
| Optimize EXCP access . . . . Y          Y=Yes N=No
| Optimize DA data sets . . . . Y          Y=Yes N=No
| SMF record type . . . . . 0__          0, 128-255
|
-----

```

Step 3 *Optional.* Change any input fields, if desired.

3.A Position the cursor in the field and typing a new value. Clear any characters that remain from the previously displayed value.

3.B Repeat this step for each field you want to edit.

Step 4 *Optional.* Save or cancel changes.

4.A Press **Enter** to retain your changes within the dialog.

4.B Press **F12** (Cancel) to discard your changes.

The dialog returns to the Data Policy Definition panel, where you can commit dialog changes to the control data set.

Step 5 *Optional.* Commit changes.

5.A Commit the changes to the control data set by selecting the File pull-down choice **2** (Save Policy).

–or–

5.B Type **Save** on the command line.

The dialog saves the changed data policy and redisplay the Data Policy Definition panel.

Activating a Data Policy

The MAINVIEW Batch Optimizer customization process created an initial data policy for your site. It becomes active when you start the batch performance subsystem on the images in the BatchPlex. The BatchPlex definition defines the name of the data policy.

Note: For a new data policy, you must activate the policy on each MVS image in the BatchPlex.

You can use a MAINVIEW Batch Optimizer command or the MAINVIEW Batch Optimizer dialog to activate a data policy. However, using the command or dialog activates the policy for the current activation of the BatchPlex only. When you restart the BMC Software Primary subsystem, it activates the data policy that is referenced by the BatchPlex definition.

To make a data policy change permanent, you must change the data policy name in the BatchPlex definition.

When you activate a data policy, using the command or dialog, the BMC Software Primary subsystem responds with messages such as the following to indicate whether the activation was successful:

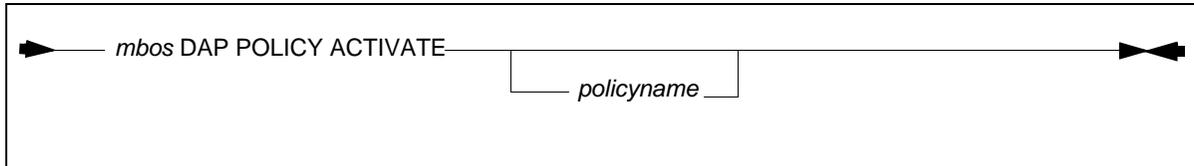
- Policy was activated successfully.
- Policy was not found in the control data set.
- Policy contains an error. The subsystem does not activate the policy.

Activating a Data Policy by Using a Command

To activate a data policy by using a command, enter the DAP POLICY ACTIVATE command from the MVS console.

Figure 4-8 shows the syntax of the DAP POLICY ACTIVATE command:

Figure 4-8 Syntax for ACTIVATE Command



mbos is the ID of the MAINVIEW Batch Optimizer subsystem (MBOS) that is active on the MVS image, and *policyname* is the name of an existing data policy definition in the control data set.

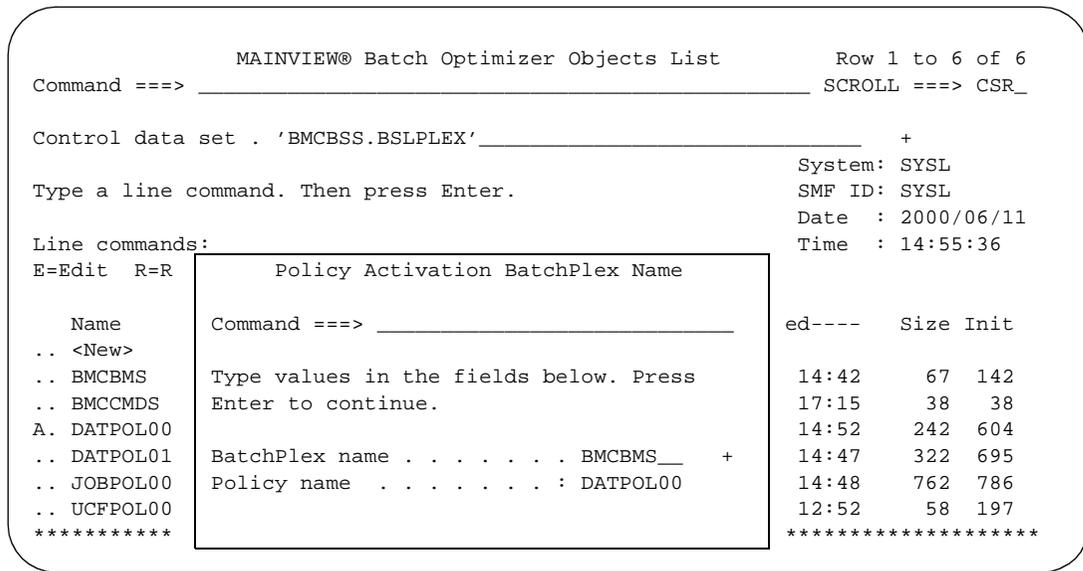
Activating a Data Policy by Using the Dialog

To activate a data policy using the dialog, complete the following steps:

- Step 1** Access the MAINVIEW Batch Optimizer Objects List panel.
- Step 2** Type **A** (Activate) to the left of your choice of data policy, and press **Enter**.

The Policy Activation BatchPlex Name pop-up dialog is displayed (Figure 4-9).

Figure 4-9 Policy Activation BatchPlex Name Pop-up



Step 3 Type the name of the active BatchPlex member that is the target of your activation request, and press **Enter**.

A plus sign (+) follows the BatchPlex name input field, indicating that you may prompt for a list of BatchPlex member from which to choose. To obtain a list, place your cursor in the BatchPlex member field and press **F4** (Prompt) or position your cursor beneath the “+” and press **Enter**. Select the BatchPlex member that defines your active BatchPlex by positioning the cursor to the left of a name and pressing **Enter**. The dialog Policy Activation BatchPlex Name pop-up is displayed (Figure 4-10) with your choice of BatchPlex in the **BatchPlex name** field. If you do not want to select an item type from the list, press **F12** (Cancel). The dialog Policy Activation BatchPlex Name pop-up is displayed.

Note: Figure 4-10 on page 4-20 displays only if Job Optimizer is operational. Otherwise, the dialog issues the **ACTIVATE** command only on the current MVS image. To route the activation command to multiple images through the dialog, you must have defined a BatchPlex and started Job Optimizer.

Figure 4-10 Policy Activation Images List Pop-up

```

Policy Activation Images List
Row 1 to 1 of 1
Command ==> _____ SCROLL ==> _____

BatchPlex name . . . . . : BMCBMS      System: SYSL
Policy name . . . . . : DATPOL00     SMF ID: SYSL
Policy type . . . . . : DATAPOL      Date : 2000/06/09
                                           Time : 09:29:06

Select one or more images for activation and Press Enter.
/=Activate Policy on Image

----MVS Image----- --Subsystems--- Number of
Name      Status      MBOS BMPC   XIM Initiators Response
.. SYSL   Active      BMCP MBOS   XIM      40
***** Bottom of data *****

```

Step 4 Type / next to each MVS Image in the BatchPlex to activate the policy, and press **Enter**.

The Response field for each of the selected MVS images indicates whether the policy was successfully activated.

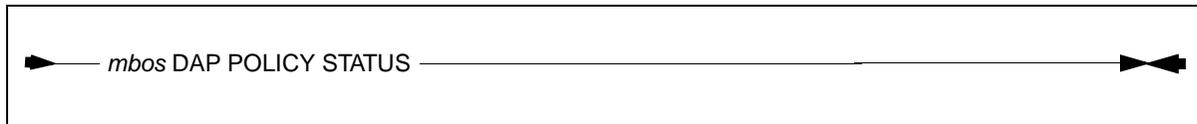
Step 5 Press **F12** (Cancel) or **F3** (Exit) to return to the MAINVIEW Batch Optimizer Objects List panel.

Determining Which Data Policy is Active

Because a new data policy can be activated at any time, the policy that is identified in the BatchPlex definition might not be the policy in effect. Before making changes to a data policy based on current activities at your site, you should determine which definitions are active.

To determine which data policy definition is active, enter the DAP POLICY STATUS command from the MVS console.

Figure 4-11 shows the syntax for the DAP POLICY STATUS command:

Figure 4-11 Syntax for STATUS Command

mbos is the ID of the MBOS that is active on the MVS image.

The MBOS responds with messages that indicate the name of the active data policy. The active BatchPlex definition is displayed.

Determining which Subsystem Processes a Data Set

Within a BatchPlex, multiple batch performance subsystems can be active. Each subsystem can have a different data policy active. In this situation, a data set might meet the selection criteria in more than one data policy.

Data Optimizer determines which data policy will apply in the following manner:

- The data policy on the subsystem that was started last examines the data set to see if it meets its criteria. If the data set meets the data policy criteria, the data policy action definition will be used for the data set.
- If the data set does not meet the criteria of the last subsystem, the subsystem that was started just before the last one examines the data set.
- If the data set does not meet the criteria of the second-to-last subsystem, the subsystem started before that one examines the data set. This continues until the data set meets the criteria of a data policy. When there are no more data policies left to examine the data set, the data set is excluded from I/O performance processing.

Chapter 5 Data Optimizer Options

This chapter explains the options for Data Optimizer. This chapter discusses the following topics:

Overview	5-1
Option Description Format	5-1
DD Statement Position	5-2
Wildcard Characters	5-4
Options	5-5
General Options	5-5
Resource Utilization Options	5-13
Reporting Options	5-16
Non-VSAM Options	5-22
VSAM Options	5-30
Summary of Data Optimizer Options, Keywords, and DD Overrides . . .	5-53

Overview

Data Optimizer provides options for determining how and when to provide performance improvements. This section describes the format used in this manual for the options, summarizes the selection criteria for the options, and summarizes the Data Optimizer options.

Option Description Format

The description of each option provides the following information:

Values This item identifies the values that you can specify for the option. If multiple field names or keywords exist for the option, each is explained. This is an example of how values are displayed:

Include Includes the data set in Data Optimizer's I/O performance processing.

Exclude Excludes the data set from I/O performance processing.

Default This item identifies the internally set default value for the option.

DD Overrides This item identifies the DD statements that you can use at the job step level to override a previously specified option. This is an example of a DD override:

```
//DAP@LSR DD DUMMY
```

If a DD override does not exist for the option, the following text is displayed:

```
This option does not provide a DD override.
```

Data Policy Keywords This item identifies the keywords that you can specify for the option. This is an example of how keywords are displayed:

```
INCLUDE  
EXCLUDE
```

DD Statement Position

The DAP@ overrides are position sensitive. If the override follows the JCL EXEC statement or other special DD statements and precedes any other non-Data Optimizer DD statement, the option is considered to be global in scope. It applies to all data sets intercepted by Data Optimizer in the step (including any that are dynamically allocated). The special DD statements that are treated like the JCL EXEC statement include:

- JOBLIB
- JOBCAT
- STEPLIB
- STEPCAT
- VVHC\$MSG
- VVHC\$M\$\$
- VVHC\$IOS
- DAPMESGS
- DAPRPTS

- DAPSTATS
- DAPDSINF
- UCC11NR
- DAP\$...

If the DAP@ override follows any other non-Data Optimizer DD statement, it applies only to the DD statement that it follows.

Data Optimizer also supports an older version of DD override statements that are prefixed with DAP\$. These statements are global in nature regardless of position.

For example, in the following JCL, the LSR and VPO options apply to all VSAM data sets accessed in the step, but the VDW option applies only to the data set represented by ddname DD2. The DRA options applies to all VSAM and non-VSAM data sets accessed in the step.

```
//STEP1 EXEC PGM=TESTPGM
//DAP@LSR DD DUMMY
//DAP@VPO DD DUMMY
//DD1 DD DSN=.....
//DD2 DD DSN=.....
//DAP$DRA DD DUMMY
//DAP@VDW DD DUMMY
//DD3 DD DSN=.....
```

Local level JCL overrides supersede global overrides of the same type. For example, if you place a DAP@VPO DD at the global level signifying that you want all VSAM data sets intercepted, you can still place a DAP@NVPO DD after a DD statement to indicate that you do not want that particular data set intercepted. In the following example, DD1 and DD3 are intercepted, but DD2 is not be intercepted:

```
//STEP1 EXEC PGM=TESTPGM
//DAP@VPO DD DUMMY
//DD1 DD DSN=.....
//DD2 DD DSN=.....
//DAP@NVPO DD DUMMY
//DD3 DD DSN=.....
```

Typically, JCL overrides are specified with the DUMMY keyword. Except as noted in the appropriate override, the presence of the DD is noted and any keywords are ignored.

Wildcard Characters

When you are specifying selection criteria in a data policy definition, you can specify a value string that includes a wildcard character. For example, you can specify `JOBNAME=ABC*` to indicate that all jobs whose names start with “ABC” will be processed by this definition. Wildcard characters have the following meaning:

- % A percent sign represents any single character.
- * A single asterisk represents zero or more characters (see `DSNAME EXCEPTIONS`)
- ** A double asterisk represents zero or more characters (see `DSNAME EXCEPTIONS`)

DSNAME EXCEPTIONS

- * A single asterisk represents a partial or entire qualifier. You can use one or two asterisks in a qualifier. If you use one asterisk, it can be in any position within the qualifier.

If you use two asterisks, the first one must be in the first position of the qualifier and the second one must be in the last position of the qualifier.

- ** A double asterisk represents zero or more entire qualifiers.

If you do not use any other wildcards in the qualifier, a double asterisk is assumed at the end of the qualifier.

- Note:** For `FILETYPE` or `TIMERANGE` selection criteria, you cannot specify wildcard characters in the **Value** field.

Options

The section heading for each option describe the field name as it appears appear on the MAINVIEW Batch Optimizer dialog Data Policy Definition panel. For details on how to access this dialog, see Appendix B, “Navigating the User Interface.”

General Options

General options are valid for VSAM and non-VSAM performance components. General options comprise the following categories:

- Action option
- Optimization Mode option
- Optimization Percentage option
- Dynamic Region Adjustment option
- Optimize STC Access option
- Optimize TSO Access
- SMF Record Type option
- Override User Values option
- Honor Resource Usage Bias
- Resource Usage Bias

Action Option

The Action option dictates whether Data Optimizer should intercept or bypass I/O optimization processing for a data set.

Values

This option has the following values:

INCLUDE Data Optimizer will intercept the data set and process it according to the Optimization Mode option.

EXCLUDE Data Optimizer will not intercept and process the data set.

Default

EXCLUDE

DD Overrides

`//DAP@NPO DD DUMMY`

Data Optimizer will intercept and process non-VSAM data sets.

`//DAP@NNPO DD DUMMY`

Data Optimizer will not intercept or process non-VSAM data sets.

`//DAP@VPO DD DUMMY`

Data Optimizer will intercept and process VSAM data sets.

```
//DAP@NVPO DD DUMMY
```

Data Optimizer will not intercept or process VSAM data sets.

Data Policy Keywords INCLUDE
EXCLUDE

Optimization Mode Option

Optimization Mode consists of three mutually exclusive options, ADVANCED, BASIC, and REPORTONLY.

Values This option has the following values:

ADVANCED	The ADVANCED option states that for non-VSAM data sets, Data Optimizer will use its own internal buffering mechanisms where possible and will not change the BUFNO/NCP values unless Data Optimizer determines that it cannot use its own internal buffering (perhaps due to storage constraints). For VSAM data sets, ADVANCED tells Data Optimizer to perform LSR or NSR processing and to choose its own buffer values, regardless of which buffering methodology is chosen.
BASIC	The BASIC option states that for non-VSAM data sets, only the BUFNO/NCP values will be optimized and that for VSAM data sets, only NSR BUFND and BUFNI values will be optimized.
REPORTONLY	The REPORTONLY option indicates that Data Optimizer will not perform performance optimization but will intercept processing of the data set to the degree necessary to gather data set and catalog information.

Default ADVANCED

DD Overrides //DAP@ADV DD DUMMY
This statement sets the Optimization Mode to ADVANCED for VSAM and non-VSAM performance processing.

```
//DAP@BASC DD DUMMY
```

This statement sets the Optimization Mode to BASIC for VSAM and non-VSAM performance processing.

```
//DAP@RPTO DD DUMMY
```

This statement sets the Optimization Mode to REPORTONLY for VSAM and non-VSAM performance processing.

Data Policy Keywords BASIC
ADVANCED
REPORTONLY

Optimization Percentage Option

The Optimization Percentage option allows you to define a percentage to apply to Data Optimizer's calculated performance optimization buffering values in order to determine the final buffering value that you will use. The Optimization Percentage option is used to scale up or scale down Data Optimizer's performance decisions.

Example

Assuming Data Optimizer calculated that a BUFNO of 12 should be used for a non-VSAM data set, the value will be modified as follows for the various listed NONVSAMOPTPCT specifications:

NONVSAMOPTPCT	Original BUFNO	New BUFNO
100	5	12
50	5	6
200	5	24

Values

This option has the following values:

0 to 999 Specify a value between 0 and 999; 0 indicates that no value is specified is equivalent to a specification of 100 percent.

Note: For VSAM data sets, this value will not also be applied to buffer specifications associated with hiperspaces. Previously, this percentage that are only applied for non-hiperspace, Data Optimizer-calculated buffer values.

Default

0

DD Overrides

```
//DAP@Vxxx DD DUMMY
```

xxx is a value from 0 to 999; specifying 0 is the same as 100; defines the VSAM optimization percentage.

```
//DAP@Nxxx DD DUMMY
```

xxx is a value from 0 to 999; specifying 0 is the same as 100; defines the non-VSAM optimization percentage.

Data Policy Keywords VSAMOPTPCT
NONVSAMOPTPCT

Dynamic Region Adjustment Option

The Dynamic Region Adjustment option allows you to specify whether Data Optimizer should modify region specifications automatically to compensate for additional I/O-related storage areas obtained on an application's behalf. For example, if you set the REGION parameter to 800 KB and Data Optimizer requires 400 KB for I/O performance processing, the REGION parameter can be dynamically adjusted to 1200 KB.

Values

This option has the following values:

Y Modify region specifications as needed.

N Do not modify region specifications.

Default

Y

DD Overrides

```
//DAP@DRA DD DUMMY
```

This statement sets the Dynamic Region Adjustment option to "Yes" for VSAM and non-VSAM performance processing.

```
//DAP@NDRA DD DUMMY
```

This statement sets the Dynamic Region Adjustment option to "No" for VSAM and non-VSAM performance processing.

Data Policy Keywords DRA
NODRA

Optimize STC Access Option

The Optimize STC Access option lets you specify whether Data Optimizer will attempt to optimize a data set when it is accessed by a started task (STC) address space.

Values

This option has the following values:

Y Data Optimizer will attempt to optimize I/O when a data set is accessed by an STC address space and it meets all other selection criteria.

N Data Optimizer will not attempt to optimize I/O when a data set is accessed by a STC address space.

Default

N

DD Overrides

This option does not provide a DD override.

Data Policy Keywords STC
NOSTC

Optimize TSO Access Option

The Optimize TSO Access option lets you specify whether Data Optimizer will attempt to optimize a data set when it is accessed by a time sharing (TSO) address space.

Values This option has the following values:

- Y Data Optimizer will attempt to optimize I/O when a data set is accessed by a TSO address space and it meets all other selection criteria.
- N Data Optimizer will not attempt to optimize I/O when a data set is accessed by a TSO address space.

Default N

DD Overrides This option does not provide a DD override.

Data Policy Keywords TSO
NOTSO

SMF Record Type Option

The SMF Record Type option lets you specify whether Data Optimizer writes a System Management Facilities (SMF) record and indicates the record type to which Data Optimizer should output informational, performance, and statistical data.

Note: The dialog limits the valid values to those defined by IBM as user-defined values (128 to 255).

Values This option has the following values:

- 0 Do not write SMF records.
- 128 to 255 Write SMF records using the indicated record type.

Default 0

DD Overrides //DAP@NSMF DD DUMMY
This statement disables SMF recording for VSAM and non-VSAM performance processing.

Data Policy Keywords **SMFRECTYPE**

Note: Once the SMF record type has been chosen and specified to Data Optimizer via the SMFRECTYPE option, it must added to SMFPRM00 in SYS1.PARMLIB to be collected. Data Optimizer checks to ensure that the specified record type is being collected and will only create the SMF record if it is. For additional information, see Appendix E, “SMF Recording.”

Override User Values Option

The Override User Values option lets you specify whether Data Optimizer should override a user’s non-VSAM DCB BUFNO/NCP values or VSAM ACB BUFND/BUFNI values even if the Data Optimizer value is lower. By default, if Data Optimizer calculates a buffer value that is lower than that already associated with a DCB or ACB, the current user value will be honored.

Values This option has the following values:

- Y Override the user’s buffer values.
- N Do not override the user’s buffer values.

Note: In some cases, Data Optimizer may issue a freepool during close processing when this option is specified. If the application issuing the close cannot tolerate the buffers being freed, this option should not be specified.

Default N

DD Overrides //DAP@OVU DD DUMMY
This statement sets the override user values option to “Yes” for VSAM and non-VSAM performance processing.

//DAP@NOVU DD DUMMY
This statement sets the override user values option to “No” for VSAM and non-VSAM performance processing.

Data Policy Keywords **OVERRIDEUSER**
NOOVERRIDEUSER

Honor Resource Usage Bias Option

The Honor Resource Usage Bias (RUB) option is provided for backward compatibility with the SmartBatch for OS/390 product. With this option, you can specify how Data Optimizer should handle an existing RUB value. The RUBHONOR value specifies to use the SmartBatch for OS/390 method of specifying performance parameters by defining a RUB value.

Values

This option has the following values:

- Y Honor any user-specified RUB value as defined in the methods that are available in SmartBatch for OS/390 or choose a RUB value by using SmartBatch for OS/390 logic.
- N Do not honor any user-specified RUB value, and do not choose RUB values.

Default

N

DD Overrides

```
//DAP@RBH DD DUMMY
```

This statement sets the RUBHONOR option to “Yes” for VSAM and non-VSAM performance processing.

```
//DAP@NRBH DD DUMMY
```

This statement sets the RUBHONOR option to “No” for VSAM and non-VSAM performance processing.

Data Policy Keywords RUBHONOR NORUBHONOR

Resource Usage Bias Option

Note: This option is provided for backward compatibility with the SmartBatch for OS/390 product. BMC Software recommends that you use the Optimization Mode and Optimization Percentage options instead.

The Resource Usage Bias option lets you indicate how Data Optimizer uses resources when providing performance improvements. For the non-VSAM Performance component, Data Optimizer manipulates CPU and main storage to obtain performance improvements. For the VSAM Performance component, Data Optimizer manipulates the size of the buffer pool to obtain performance improvements.

You may have certain applications that you want to run as quickly as possible. A smaller elapsed time, however, usually means greater CPU usage or a larger buffer pool. You may have other applications where virtual storage constraints—especially 24-bit storage—are an issue. Use this option to indicate how resources are to be used when providing performance improvements.

A special RUB value, 99, is provided to indicate that Data Optimizer should dynamically select the optimum RUB value based on the following criteria:

- type of processing being performed
- system resource availability at the time of processing

For most situations, allowing Data Optimizer to select the RUB value will provide the best performance benefits without over-committing system resources. Under some circumstances, however, you might want to override the Data Optimizer choice to favor CPU or virtual storage for a given application or environment.

Generally, for non-VSAM data sets, CPU usage increases and application storage below 16 MB decreases as you move from a low RUB value to a high RUB value. For VSAM data sets, the number of buffers increases as you move from a low RUB value to a high RUB value.

You may need to experiment with several values before you determine the optimum value for your applications. Use the global commands to set default values for your site for VSAM and non-VSAM performance improvements. Then use data policy definitions for exceptions to the default value.

For example, if you have an application that requires the best elapsed time, regardless of CPU usage, create a data policy definitions for the application, and set the RUB option to 9. If you are concerned about CPU usage, but still want performance improvements for an application, create a data policy definitions for the application, and set the RUB option to 1 or 2.

Values

This option has the following values:

- 0 For data policy definitions and DD overrides, this value indicates the absence of a selection. Continue through the hierarchy of intercept processing to determine the resource usage bias value.
- 1 to 9 For the non-VSAM Performance component, as you move from 1 to 9, CPU usage increases and 24-bit application storage decreases. For the VSAM Performance component, as you move from 1 to 10 the size of the buffer pool increases.
- 10 For the non-VSAM Performance component, minimize all storage constraints, especially 24-bit storage.

99 For data policy definitions and DD overrides, this value indicates that Data Optimizer should select an appropriate value dynamically for your installation.

Default 99

DD Overrides

//DAP@NRxx DD DUMMY

This statement sets the RUB option value to *xx* for non-VSAM performance processing (*xx* is the 01 to 10 or 99).

//DAP@VRxx DD DUMMY

This statement sets the RUB option value to *xx* for VSAM performance processing (*xx* is 01 to 10 or 99).

Data Policy Keywords RUB

Resource Utilization Options

Resource Utilization options allow you to set the thresholds for determining whether a system is storage-constrained and to establish limits on Data Optimizer storage utilization. Resource utilization comprise the following categories:

- Page Fault Rate Threshold option
- CPU Utilization Threshold option
- Below 16 Megabyte Reserve option
- Above 16 Megabyte Reserve option
- Below 16 Megabyte Maximum option
- Above 16 Megabyte Maximum option

Page Fault Rate Threshold Option

The Page Fault Rate Threshold option lets you specify the value to be used by Data Optimizer to determine whether a system is storage-constrained at OPEN. Storage constraint is one criteria used by Data Optimizer in calculating optimal VSAM and non-VSAM buffer values. The maximum value permitted is 9999. If no value is specified, the RCCPTRTH value is used.

Values

This option has the following values:

0 to 9999 Specify a value between 0 and 9999; 0 indicates that no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords PAGINGLIMIT

CPU Utilization Threshold Option

The CPU Utilization Threshold option lets you specify the value to be used by Data Optimizer to determine if a system is CPU constrained at OPEN. CPU constraint is one criteria used by Data Optimizer in calculating optimal VSAM and non-VSAM buffer values. The maximum value permitted is 999.

Values This option has the following values:

0 to 999 Specify a value between 0 and 999; 0 indicates that no value is specified. When no value is specified, 101% is used as the CPU constrained threshold value.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords CPULIMIT

Below 16 Megabyte Reserve Option

The Below 16 Megabyte Reserve option lets you specify the amount of storage below the 16 MB line that you want to be left unallocated after storage is allocated for control blocks and data buffers. Data Optimizer will not alter buffer values if the storage needed to back the specified value would cause less than this amount of storage to be left unallocated in the user private region below 16 MB.

Values This option has the following values:

0 to 16777215 0 means no value is specified.

16383K *K* denotes kilobytes.

15M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords STG24RESERVE

Above 16 Megabyte Reserve Option

The Above 16 Megabyte Reserve option lets you specify the amount of storage above the 16 MB line that you want to be left unallocated after storage is allocated for control blocks and data buffers. Data Optimizer will not alter the buffer values if the storage needed to back the specified value would cause less than this amount of storage to be left unallocated in the user private region above 16 MB.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords STG31RESERVE

Below 16 Megabyte Maximum Option

The Below 16 Megabyte Maximum option lets you specify the maximum amount of additional storage Data Optimizer is to use within the user private area below 16 MB because of its allocation of storage control blocks and data buffers.

Values This option has the following values:

0 to 16777215 0 means no value is specified.

16383K *K* denotes kilobytes.

15M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords STG24MAX

Above 16 Megabyte Maximum Option

The Above 16 Megabyte Maximum option lets you specify the maximum amount of additional storage Data Optimizer is to use within the user private area above 16 MB because of its allocation of storage control blocks and data buffers.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords STG31MAX

Reporting Options

The Reporting options provide control over the types of information that will be generated about data sets and messages. Reporting options comprise the following categories:

- Report Destination option
- Statistics Report option
- Report SYSOUT Class option
- Data Set Information Report option
- Catalog Information Report option
- Options Report option
- Message Level option
- Message Destination option
- Message SYSOUT Class option

Report Destination Option

The Report Destination option lets you control where Data Optimizer will output Data Optimizer reports.

Values This option has the following values:

SYMSMSG Include reports as part of the JES SYMSMSG log.

SYSOUT Spin off reports as a separate SYSOUT data set to the class defined in the report SYSOUT class field.

Default SYMSMSG

DD Overrides

//DAPRPTS DD DUMMY

Send all reports output to JES SYMSMSG log.

//DAPRPTS DD SYSOUT=*n*

Send all reports output to separate SYSOUT data set class *n*. *n* is any valid sysout class.

//DAPRPTS DD DSN=*data.set.name*

Send all reports to data set.

Note: DAPSTATS can be used in place of DAPRPTS in the DD statements listed above to support JCL changes you may have made for SmartBatch. BMC Software recommends using DAPRPTS when possible.

Data Policy Keywords RPTDEST

Statistics Report Option

The Statistics Report option lets you specify whether Data Optimizer produces a performance statistics report.

Values This option has the following values:

Y Produce a statistics report.

N Do not produce a statistics report.

Default N

DD Overrides

//DAP@STA DD DUMMY

Produce a statistics report.

//DAP@NSTA DD DUMMY

Do not produce a statistics report.

Data Policy Keywords STATRPT
NOSTATRPT

Report SYSOUT Class Option

The Report SYSOUT Class option lets you specify which SYSOUT class Data Optimizer should use if the Report Destination option is defined as “SYSOUT.” The value in this field defines the JES SYSOUT class for reports.

Values The valid values for this field are any valid JES SYSOUT classes for your environment.

Default *

DD Override //DAPRPTS DD SYSOUT=*n*
Send all reports output to separate SYSOUT data set class *n*. *n* is all valid sysout class.

Note: DAPSTATS can be used in place of DAPRPTS in the DD statements listed above to support JCL changes you may have made for SmartBatch. BMC Software recommends using DAPRPTS when possible.

Data Policy Keywords RPTCLASS

Data Set Information Report Option

The Data Set Information Report option lets you specify whether Data Optimizer produces a Data Set Information report. This report lists DCB/ACB and data set attributes.

Values This option has the following values:

Y Produce a Data Set Information Report.

N Do not produce a Data Set Information Report.

Default N

DD Overrides //DAP@DSI DD DUMMY
Produce a Data Set Information Report.

//DAP@NDSI DD DUMMY
Do not produce a Data Set Information Report.

To produce a Data Set Information Report without optimizing the data set, include the following DD statement:

//DAP@RPTO DD DUMMY

Data Policy Keywords DSINFORPT
NODSINFORPT

Catalog Information Report Option

The Catalog Information Report provides VSAM data set statistics that are maintained in the catalog, such as high-used RBA, number of CI-splits, etc.

Values This option has the following values:

Y Produce a Catalog Information Report.

N Do not produce a Catalog Information Report.

Default N

DD Overrides //DAP@CAT DD DUMMY
Produce a Catalog Information Report.

//DAP@NCAT DD DUMMY
Do not produce a Catalog Information Report.

To produce a Catalog Information Report without optimizing the data set, include the following DD statement:

//DAP@RPTO DD DUMMY

Data Policy Keywords CATRPT
NOCATRPT

Options Report Option

The Options Report describes the options that are supplied by the data policy global statement, data policy define statement, or DD override statement(s).

Values This option has the following values:

Y Produce an Options Report.

N Does not produce an Options Report.

Default N

DD Overrides //DAP@OPT DD DUMMY
Produce an Options Report.

```
//DAP@NOPT DD DUMMY  
Do not produce an Option Report.
```

Data Policy Keywords **OPTIONRPT**
NOOPTIONRPT

Message Level Option

The Message Level option lets you indicate what types of informational messages Data Optimizer will write.

Values This option has the following values:

- 0 Do not write informational messages.
- 1 Only write exception informational messages. See the example below. This includes rejection by data policy or ddname trigger, as well as interception that is bypassed.
- 2 Write all informational messages.
- 3 This level is a special message level that is used for diagnostic purposes only. It should be used only when directed by BMC Software Customer Support. Depending on the circumstances, it can generate a large volume of messages to the operator's console.

Default 0

DD Overrides

```
//DAP@MSG0 DD DUMMY
```


Sets Message Level option to 0 (do not write informational messages).

```
//DAP@MSG1 DD DUMMY
```


Sets Message Level option to 1 (write exception informational messages).

```
//DAP@MSG2 DD DUMMY
```


Sets Message Level option to 2 (write all informational messages).

```
//DAP@MSG3 DD DUMMY
```


Sets Message Level option to 3 (use only when directed by BMC Software Customer Support).

Data Policy Keywords **MSGLEVEL0**
MSGLEVEL1
MSGLEVEL2
MSGLEVEL3

Depending on whether the open request is to be included or excluded, Data Optimizer will write one or two informational messages for each open that occurs. The messages are written to the destination defined using the Message Destination option.

The first message is a REQUESTED/REJECTED message. The message informs you that the data set being opened has been specifically requested or rejected by an entry in the data policy or by a DDname trigger. The message describes Data Optimizer's intent and names the source of the intent.

If the data set is requested rather than rejected, a second message is written. The second message is an INTERCEPTION/BYPASS message. This message indicates that one of the following actions will occur:

- I/O performance processing
- basic-only performance processing
- bypassed for performance processing

If the processing is set to BASIC or is bypassed, the reason for this action is also indicated.

Message Destination Option

The Message Destination option lets you specify where Data Optimizer is to output the Data Optimizer informational messages if requested.

Values

This option has the following values:

SYSMSG

Include the report as part of the JES SYSMSG log.

SYSOUT

Spin off the report as a separate SYSOUT data set to the class defined in the message SYSOUT class field.

Default

SYSMSG

DD Overrides

```
//DAPMESGS DD DUMMY
```

Send message output to JES SYSMSG log.

```
//DAPMESGS DD SYSOUT=n
```

Send message output to separate SYSOUT data set class *n*. *n* is any valid sysout class.

```
//DAPMESGS DD DSN=data.set.name
```

Send message output to a data set.

Data Policy Keywords MSGDEST

Message SYSOUT Class Option

The Message SYSOUT Class option lets you specify which SYSOUT class Data Optimizer should use if the Message Destination option is defined as “SYSOUT.” The value in this field defines the JES SYSOUT class for the diagnostic messages.

Values	The valid values for this field are any valid JES SYSOUT classes for your environment.
Default	*
DD Override	//DAPMESGS DD SYSOUT= <i>n</i> Send statistics report output to separate SYSOUT data set class <i>n</i> . <i>n</i> is any valid sysout class.

Data Policy Keywords MSGCLASS

Non-VSAM Options

Non-VSAM options allow you to control processing for non-VSAM data sets and processes. Non-VSAM options comprise the following categories:

- Actual I/O Counts option
- Allow Shared Writes option
- BUFNO Value option
- BUFNO Storage option
- Maximum BUFNO Value option
- Maximum BUFNO Storage option
- NCP Value option
- NCP Storage option
- Maximum NCP Value option
- Maximum NCP Storage option
- Optimize Tape Data Sets Option
- Optimize PO Data Sets Option
- Optimize DA Data Sets Option
- Optimize EXCP Access Option

Actual I/O Counts Option

The Actual I/O Counts option lets you specify how Data Optimizer should report to SMF the number of I/O requests for non-VSAM data sets.

Values This option has the following values:

Y Report the actual number of I/O requests made by Data Optimizer. The I/O count indicates the number of physical I/O requests that Data Optimizer issued to perform the data transfer. You may want to specify this value for capacity planning or I/O subsystem tuning purposes.

N Report the number of logical EXCP requests made by the application. For application access methods, the I/O count indicates the number of blocks transferred. You may want to specify this value for billing and accounting purposes.

Default N

DD Overrides //DAP@AIC DD DUMMY
Sets Actual I/O Counts option to “Yes” for non-VSAM performance processing.

//DAP@NAIC DD DUMMY
Sets Actual I/O Counts option to “No” for non-VSAM performance processing.

Data Policy Keywords AIC
NOAIC

Allow Shared Writes Option

The Allow Shared Writes option lets you specify whether Data Optimizer will allow shared write access to a data set.

Values This option has the following values:

Y Data Optimizer allows shared write access to a data set that has been selected for I/O optimization. Select this option only in cases when you are confident that no data integrity exposure exists.

N Data Optimizer does not allow shared write access to a data set that has been selected for I/O optimization. If Data Optimizer detects a shared write access attempt, the opener receives message BMC102615E and abends with system abend U2210-7709.

Default N

DD Overrides This option does not provide a DD override.

Data Policy Keywords SHAREDWRITE
NOSHAREDWRITE

BUFNO Value Option

The BUFNO Value option defines a value from 0 to 255 to be used as the BUFNO value within any intercepted DCB; this value honored for BASIC and ADVANCED forms of interception.

Values This option has the following values:
0 to 255 0 indicates that no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords BUFNO

BUFNO Storage Option

The BUFNO Storage option lets you specify the amount of storage to be allocated to buffers with the BUFNO DCB parameter.

A BUFNO value is calculated for the associated DCB, based on this value and the block size of the data set. The calculated value for BUFNO will not exceed 255. This value is honored for BASIC and ADVANCED optimization methods.

Values This option has the following values:
0 to 2147483647
0 indicates that no value is specified.
2097151K *K* denotes kilobytes.
2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords BUFNOSTG

Maximum BUFNO Value Option

The Maximum BUFNO Value option lets you specify the maximum permitted BUFNO value in a DCB. The BUFNO value as user-specified or calculated by Data Optimizer will not be permitted to exceed this value. This value is honored if the user specified a BUFNO value, a BUFNO Storage value, or if Data Optimizer is using the BASIC Optimization Mode.

Values	This option has the following values: 0 to 255 0 indicates that no value is specified.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords MAXBUFNO**Maximum BUFNO Storage Option**

The Maximum BUFNO Storage option lets you specify the maximum permitted amount of storage allocated to buffers defined by the BUFNO DCB parameter. The amount of storage allocated to BUFNO buffers as specified by the user or calculated by Data Optimizer is not allowed to exceed this value. The BUFNO value is adjusted accordingly to ensure that the maximum is not exceeded.

This value is honored if the user specified a BUFNO value, a BUFNO storage value, or if Data Optimizer is using the BASIC optimization methods. Any value higher than 255 times the block size is ineffective.

Values	This option has the following values: 0 to 2147483647 0 indicates that no value is specified. 2097151K <i>K</i> denotes kilobytes. 2047M <i>M</i> denotes megabytes.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords MAXBUFNOSTG

NCP Value Option

The NCP Value option lets you specify a value from 0 to 255 to be used as the NCP value within an intercepted DCB. This value is honored for both BASIC and ADVANCED optimization methods.

Note: The effectiveness of this value depends upon the ability of the user program to recognize and use the additional concurrent requests that are permitted by the new NCP value.

Values	This option has the following values: 0 to 255 0 indicates that no value is specified.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords NCP**NCP Storage Option**

The NCP Storage option lets you specify the amount of storage to be allocated for NCP processing. Although specifying the NCP value does not allocate buffer storage directly, Data Optimizer assumes that usage of the additional concurrent requests will be accommodated by the program via the acquisition of additional buffer storage equal to the block size times the additional NCP value.

An NCP value is calculated for the associated DCB, based on the NCP storage value and the block size of the data set. The calculated value will not exceed 255. This value is honored for BASIC and ADVANCED optimization methods. Any value higher than 255 times the block size is ineffective.

Values	This option has the following values: 0 to 2147483647 0 indicates that no value is specified. 2097151K <i>K</i> denotes kilobytes. 2047M <i>M</i> denotes megabytes.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords NCPSTG

Maximum NCP Value Option

The Maximum NCP Value option lets you specify the maximum permitted NCP value in a DCB. The NCP value as user-specified or calculated by Data Optimizer will not be permitted to exceed this value. This value is honored if the user specified a NCP value, a NCP storage value, or if Data Optimizer is using the BASIC Optimization Mode.

Values	This option has the following values: 0 to 255 0 indicates that no value is specified.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords MAXNCP

Maximum NCP Storage Option

The Maximum NCP Storage option lets you specify the maximum permitted amount of storage allocated to buffers as implied by the DCB NCP parameter. While specifying the NCP value itself does not allocate buffer storage directly, Data Optimizer will assume that usage of the additional concurrent requests will be accommodated by the program via the acquisition of additional buffer storage equal to the block size times the additional NCP value. The NCP value will be adjusted accordingly to ensure the maximum is not exceeded.

This value will be honored if the user specified a NCP value, a NCP storage value, or if Data Optimizer is using the BASIC optimization methods. Any value higher than 255 times the block size is ineffective.

Values	This option has the following values: 0 to 2147483647 0 indicates that no value is specified. 2097151K <i>K</i> denotes kilobytes. 2047M <i>M</i> denotes megabytes.
Default	0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXNCPSTG

Optimize Tape Data Sets Option

This option lets you specify whether Data Optimizer should try to optimize tape data sets. If a concatenation of data sets contains at least one tape data set, this option will affect the processing of the entire concatenation.

Values This option has the following values:

Y Allow optimization of tape data sets or concatenations containing at least one tape data set.

N Do not optimize tape data sets or concatenations containing at least one tape data set.

Default Y

DD Overrides //DAP@TAP DD DUMMY
Allow optimization of tape data sets or concatenations containing at least one tape data set.

//DAP@NTAP DD DUMMY
Do not optimize tape data sets or concatenations containing at least one tape data set.

Data Policy Keywords TAPE
NOTAPE

Optimize PO Data Sets Option

This option lets you specify whether Data Optimizer should optimize data sets whose organization is partitioned (PO). If a concatenation of data sets contains at least one PO data set, this option will affect the processing of the entire concatenation.

Values This option has the following values:

Y Allow optimization of PO data sets or concatenations containing at least one PO data set.

N Do not optimize PO data sets or concatenations containing at least one PO data set.

Default Y

DD Overrides //DAP@PDS DD DUMMY
Allow optimization of PO data sets or concatenations containing at least one PO data set.

//DAP@NPDS DD DUMMY
Do not optimize PO data sets or concatenations containing at least one PO data set.

Data Policy Keywords PDS
NOPDS

Optimize DA Data Sets Option

This option lets you specify whether Data Optimizer should try to optimize data sets whose organization is direct access (DA). If a concatenation of data sets contains at least one DA organization data set, this option will affect the processing of the entire concatenation.

Values This option has the following values:

Y Allow optimization of DA data sets or concatenations containing at least one DA data set.

N Do not optimize DA data sets or concatenations containing at least one DA data set.

Default Y

DD Overrides //DAP@DRC DD DUMMY
Allow optimization of DA data sets or concatenations containing at least one DA data set.

//DAP@NDRC DD DUMMY
Do not optimize DA data sets or concatenations containing at least one DA data set.

Data Policy Keywords DIRECT
NODIRECT

Optimize EXCP Access Option

This option lets you specify whether Data Optimizer should try to optimize data sets being processed using the EXCP access method.

Values	This option has the following values:
Y	Allow optimization of data sets that are being processed which are using the EXCP access method.
N	Do not optimize data sets that are being processed which are using the EXCP access method.
Default	Y
DD Overrides	<pre>//DAP@XCP DD DUMMY</pre> Allow optimization of data sets that are being processed which are using the EXCP access method. <pre>//DAP@NXCP DD DUMMY</pre> Do not optimize data sets that are being processed which are using the EXCP access method.
Data Policy Keywords	EXCP NOEXCP

VSAM Options

VSAM options let you control processing for VSAM data sets and processes. VSAM options comprise the following categories:

- LSR option
- Defer Writes option
- RMODE31 Control Blocks option
- SHAREOPTIONS(3,3) option
- NSR Index Buffers option
- NSR Index Storage option
- Maximum NSR Index Buffers option
- Maximum NSR Index Storage option
- NSR Data Buffers option
- NSR Data Storage option
- Maximum NSR Data Buffers option
- Maximum NSR Data Storage option
- LSR Index Buffers option
- LSR Index Storage option
- Percent of Index to Buffer option
- Maximum LSR Index Buffers option
- Maximum LSR Index Storage option
- LSR Data Buffers option
- LSR Data Storage option
- Maximum LSR Data Buffers option

- Maximum LSR Data Storage option
- Hiperspace Support option
- Hiperspace Support Percentage option
- Hiperspace Storage Range option
- Maximum Expanded LSR Buffers Storage option
- LSR Maximum System Expanded Storage option
- Maximum System Expanded Storage Percent option
- Page-Fix Buffers option
- Sequential Insert Strategy option
- Define LSR Pool ID option
- Override LSR Pool Specification option
- LSR Alternate Index option
- Multistring Protection option
- Disable Batch LSR option
- Detect Exclusive Control Conflicts option

LSR Option

The LSR option lets specify the degree to which LSR processing is used.

Values

This option has the following values:

Y (Yes) Use LSR processing when appropriate. Data Optimizer uses captured history data to determine whether LSR or NSR processing will best benefit I/O processing. Data Optimizer does not enable LSR where history data indicate LSR use is unacceptable or NSR yields better performance. You may want to specify this value to improve direct access index processing.

Warning! This value can cause problems for applications that cannot tolerate LSR processing.

N (No) Do not enable LSR processing for files selecting this option. You may want to specify this value for applications that cannot tolerate VSAM LSR processing for specific files.

P (Prefer) Enable LSR processing for VSAM files selecting this option if no other restrictions preclude the use of this option. Data Optimizer does not use history data to aid in this decision. Data Optimizer uses NSR to optimize I/O if use of LSR is not possible. You may want to specify this value where you have intimate knowledge of a file and its associated applications. With this option, Data Optimizer enables LSR processing.

R (Require) Optimize I/O for files using this option only if LSR processing is enabled. If Data Optimizer cannot enable LSR, it bypasses any further attempts to optimize I/O for this execution.

F (Force) Force Data Optimizer to place a VSAM data set into an LSR pool even if the ACB DIR bit is off. Without this option, VSAM data sets are only eligible to be placed into an LSR pool if the ACB DIR bit is on. All other restrictions that would disallow placement into an LSR pool will still apply.

Default Y

DD Overrides

/DAP@FLSR DD DUMMY

Force Data Optimizer to place a VSAM data set into an LSR pool.

//DAP@LSR DD DUMMY

Use LSR processing when appropriate.

//DAP@NLSR DD DUMMY

Do not enable LSR processing for files selecting this option.

//DAP@PLSR DD DUMMY

Enable LSR processing for VSAM files selecting this option if no other restrictions preclude the use of this option.

//DAP@RLSR DD DUMMY

Optimize I/O for files using this option only if LSR processing is enabled.

Data Policy Keywords FORCELSR
LSR
NOLSR
PREFERLSR
REQUIRELSR

Note: The LSR, NOLSR, PREFERLSR, REQUIRELSR, FORCELSR keywords are only honored if you are working in Advanced Optimization mode.

Defer Writes Option

During normal VSAM processing, VSAM writes a CI to DASD when an application issues a direct put. Enabling the VSAM Defer Writes option reduces physical I/O because the CI is not written to DASD until the file is closed or until the buffer is needed by another request.

Values This option has the following values:

Y Write requests in a buffer until the file is closed or the buffer is needed by another request. You may want to specify this value to reduce physical I/O.

Warning! This value can cause problems for applications that share VSAM data sets. First, because the CIs are not written immediately, an application can access back-leveled CIs. Second, for recovery/restart programs that assume that all writes are on DASD, a recovery or restart may fail. Updates are not guaranteed until a file is closed.

N Write requests to DASD as they occur, normally when an application issues a direct put. You may want to specify this value for applications that share VSAM data sets.

Default N

DD Overrides

//DAP@VDW DD DUMMY

Sets Defer Writes option to “Yes” for VSAM performance processing.

//DAP@NVDW DD DUMMY

Sets Defer Writes option to “No” for VSAM performance processing.

Data Policy Keywords VDW
NOVDW

RMODE31 Control Blocks Option

The RMODE31 Control Blocks option lets you specify whether Data Optimizer will manipulate the ACB RMODE31 option at VSAM open time.

Values This option has the following values:

Y Data Optimizer modifies the RMODE31 ACB option at open time. This will cause the open process to place VSAM control blocks in 31-bit storage. If this option is already enabled in the ACB, it is not changed.

N Data Optimizer does not modify the RMODE31 ACB option.

Default N

DD Overrides

//DAP@DVC DD DUMMY

Modify the RMODE31 ACB option to allow VSAM control blocks to reside in 31-bit storage.

//DAP@NDVC DD DUMMY
Do not modify the RMODE31 ACB option.

Data Policy Keywords VSAMCB31
NOVSAMCB31

SHAREOPTIONS(3,3) Option

The SHAREOPTIONS(3,3) option lets you select whether Data Optimizer provides performance improvements to VSAM data sets that are allocated as SHAREOPTIONS(3,3).

Values This option has the following values:

Y Provide performance improvements for otherwise qualified VSAM SHAREOPTIONS(3,3) data sets.

N Do not intercept VSAM SHAREOPTIONS(3,3) data sets.

Default N

DD Overrides //DAP@SHR DD DUMMY
Provide performance improvements for otherwise qualified VSAM SHAREOPTIONS(3,3) data sets.

//DAP@NSHR DD DUMMY
Do not intercept VSAM SHAREOPTIONS(3,3) data sets.

Data Policy Keywords SHROPT33
NOSHROPT33

NSR Index Buffers Option

The NSR Index Buffers option lets you specify a value for the number of VSAM NSR index buffers that Data Optimizer provides in an intercepted ACB. This number is limited by any NSR Maximum index buffer specifications.

If you also specify the NSR Index buffer storage value, Data Optimizer uses the lesser of the actual (NSR Index buffer) value and the calculated (NSR Index buffer storage) value.

Values This option has the following values:

0 to 255 For a value of 0, Data Optimizer does not provide a value for VSAM NSR index buffers.

Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords NSRIBUF**NSR Index Storage Option**

The NSR Index Storage option lets you specify a storage value that is used to calculate the number of VSAM NSR index buffers that Data Optimizer provides in an intercepted ACB. The calculated value is limited by any NSR Maximum Index Buffer specifications.

If you also specify the NSR Index buffer value, Data Optimizer uses the lesser of the actual (NSR Index buffer) value and the calculated (NSR Index buffer storage) value.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords NSRISTG**Maximum NSR Index Buffers Option**

The Maximum NSR Index Buffers option lets you specify a value for the maximum number of VSAM NSR index buffers that Data Optimizer provides in an intercepted ACB. This limit is applied to the actual (NSR Index Buffer) value and the calculated (NSR Index Buffer Storage) value.

If you also specify the NSR Maximum index buffer storage value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

Values This option has the following values:

0 to 255 0 indicates no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXNSRIBUF

Maximum NSR Index Storage Option

The Maximum Index Storage option lets you specify a storage value that is used to calculate the maximum number of VSAM NSR index buffers that Data Optimizer provides in an intercepted ACB. The calculated value will be used to limit the actual (NSR Index buffer) value and the calculated (NSR Index buffer storage) value.

If you also specify the NSR Maximum index buffers value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXNSRISTG

NSR Data Buffers Option

The NSR Data Buffers option lets you specify a value for the number of VSAM NSR data buffers that Data Optimizer provides in an intercepted ACB. This number is limited by any NSR Maximum data buffer specifications.

If you also specify the NSR Data buffer storage value, Data Optimizer uses the lesser of the actual (NSR Data buffer) value and the calculated (NSR Data buffer storage) value.

Values	This option has the following values: 0 to 255 0 indicates no value is specified.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords NSRDBUF**NSR Data Storage Option**

The NSR Data Storage option lets you specify a storage value that is used to calculate the number of VSAM NSR data buffers that Data Optimizer provides in an intercepted ACB. The calculated value is limited by any NSR Maximum data buffer specifications.

If you also specify the NSR Data buffer value, Data Optimizer uses the lesser of the actual (NSR Data buffer) value and the calculated (NSR Data buffer storage) value.

Values	This option has the following values: 0 to 2147483647 0 indicates that no value is specified. 2097151K <i>K</i> denotes kilobytes. 2047M <i>M</i> denotes megabytes.
Default	0
DD Overrides	This option does not provide a DD override.

Data Policy Keywords NSRDSTG**Maximum NSR Data Buffers Option**

The Maximum NSR Data Buffers option lets you specify a value for the maximum number of VSAM NSR data buffers that Data Optimizer provides in an intercepted ACB. This limit is applied to the actual (NSR Data buffer) value and the calculated (NSR Data buffer storage) value.

If you also specify the NSR Maximum data buffer storage value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

- Values** This option has the following values:
- 0 to 255 0 indicates no value is specified.
- Default** 0
- DD Overrides** This option does not provide a DD override.

Data Policy Keywords MAXNSRDBUF

Maximum NSR Data Storage Option

The Maximum NSR Data Storage option lets you specify a storage value that is used to calculate the maximum number of VSAM NSR data buffers that Data Optimizer provides in an intercepted ACB. The calculated value is used to limit both the actual (NSR Data buffer) value and the calculated (NSR Data buffer storage) value.

If you also specify the NSR Maximum data buffers value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

- Values** This option has the following values:
- 0 to 2147483647
- 0 indicates that no value is specified.
- Default** 0
- DD Overrides** This option does not provide a DD override.

Data Policy Keywords MAXNSRDSTG

LSR Index Buffers Option

The LSR Index Buffers option lets you specify a value for the number of VSAM LSR index buffers that Data Optimizer provides in an intercepted ACB. This number is limited by any Maximum LSR Index Buffer specifications.

If you also specify the LSR Index Buffer Storage value, Data Optimizer uses the lesser of the actual (LSR Index Buffer) value and the calculated (LSR Index Buffer Storage) value.

Note: The Percent of Index to Buffer option overrides this option.

Values This option has the following values:

0 to 65535 0 indicates that no value is specified.

23321K *K* denotes kilobytes.

1034M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords LSRIBUF

LSR Index Storage Option

The LSR Index Storage option lets you specify a storage value that is used to calculate the number of VSAM LSR index buffers that Data Optimizer provides in an intercepted ACB. The calculated value is limited by any LSR Maximum index buffer specifications.

If you also specify the LSR Index buffer value, Data Optimizer uses the lesser of the actual (LSR Index buffer) value and the calculated (LSR Index buffer storage) value.

Note: The Percent of Index to Buffer option overrides this option.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords LSRISTG**Percent of Index to Buffer Option**

The Percent of Index to Buffer option lets you specify a value for the number of buffers to use for VSAM LSR index buffers as a percentage of the number of index records which exist for the data set being opened. The number of index records is the total of the number of index set and sequence set records.

This value overrides any LSR Index buffer storage value or LSR Index buffers value that is specified.

Values This option has the following values:

0 to 100 0 indicates that no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords PCTLSRIBUF**Maximum LSR Index Buffers Option**

The Maximum Index Buffers option lets you specify a value for the maximum number of VSAM LSR index buffers that Data Optimizer provides in an intercepted ACB. This limit is applied to the actual (LSR Index Buffer) value and the calculated (LSR Index Buffer Storage) value.

If you also specify the Maximum LSR Index Buffer Storage value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

Values This option has the following values:

0 to 65535 0 indicates that no value is specified.

23321K *K* denotes kilobytes.

1034M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXLSRIBUF

Maximum LSR Index Storage Option

The Maximum LSR Index Storage option lets you specify a storage value that is used to calculate the maximum number of VSAM LSR index buffers that Data Optimizer provides in an intercepted ACB. The calculated value is used to limit both the actual (LSR Index buffer) value and the calculated (LSR Index buffer storage) value.

If you also specify the LSR Maximum index buffers value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXLSRISTG

LSR Data Buffers Option

The LSR Data Buffers option lets you specify a value for the number of VSAM LSR data buffers that Data Optimizer provides in an intercepted ACB. This number is limited by any LSR Maximum data buffer specifications.

If you also specify the LSR Data buffer storage value, Data Optimizer will use the lesser of the actual (LSR Data buffer) value and the calculated (LSR Data buffer storage) value.

Values This option has the following values:

0 to 65535 0 indicates that no value is specified.

23321K *K* denotes kilobytes.

1034M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords LSRDBUF**LSR Data Storage Option**

The LSR Data Storage option lets you specify a storage value that is used to calculate the number of VSAM LSR data buffers that Data Optimizer provides in an intercepted ACB. The calculated value is limited by any LSR Maximum data buffer specifications.

If you also specify the LSR Data buffer value, Data Optimizer uses the lesser of the actual (LSR Data Buffer) value and the calculated (LSR Data Buffer Storage) value.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords LSRDSTG**Maximum LSR Data Buffers Option**

The Maximum LSR Data Buffers option lets you specify a value for the maximum number of VSAM LSR data buffers that Data Optimizer provides in an intercepted ACB. This limit is applied to both the actual (LSR Data buffer) value and the calculated (LSR Data buffer storage) value.

If you also specify the LSR Maximum Data Buffer Storage value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

Values This option has the following values:

0 to 65535 0 indicates that no value is specified.

23321K *K* denotes kilobytes.

1034M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXLSRDBUF

Maximum LSR Data Storage Option

The Maximum Data Storage option lets you specify a storage value that is used to calculate the maximum number of VSAM LSR data buffers that Data Optimizer will provide in an intercepted ACB. The calculated value is used to limit both the actual (LSR Data buffer) value and the calculated (LSR Data buffer storage) value.

If you also specify the LSR Maximum Data Buffers value, Data Optimizer uses the lesser of the actual maximum value and the calculated maximum value.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXLSRDSTG

Hiperspace Support Option

The Hiperspace Support option lets you include the definition of hiperspace buffers when creating a VSAM LSR pool.

Values This option has the following values:

Y Enable hiperspace usage.

N Disable hiperspace usage.

Default N

DD Overrides //DAP@DVH DD DUMMY
Enable hiperspace usage; available at both the step and DDNAME levels.

//DAP@NDVH DD DUMMY
Disable hiperspace usage; available at both the step and DDNAME levels.

//DAP@VHSY DD DUMMY
Enable hiperspace usage at the step level; HIPER-CACHE compatibility DDNAME.

//DAP@VHSN DD DUMMY
Disable hiperspace usage at the step level; HIPER-CACHE compatibility DDNAME.

Data Policy Keywords HIPERSPACE
NOHIPERSPACE

Hiperspace Support Percentage Option

The Hiperspace Support Percentage option, if the Hiperspace option is specified, defines the number of hiperspace buffers to allocate as a percentage of the number of used data control intervals in the data set.

Values This option has the following values:

0 to 100 0 indicates that no value is specified.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords HIPERPCT

Hiperspace Storage Range Option

If the Hiperspace option is specified, this option lets you define the number of hiperspace buffers as the number that would fit in the amount of storage specified by this keyword.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords HIPERSTG

Maximum Expanded LSR Buffers Storage Option

The Maximum Expanded LSR Buffers Storage option lets you specify the maximum amount of expanded storage to allow Data Optimizer to allocate to hiperspace LSR buffers for an individual data set.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Note: This value now represents the maximum amount of storage, as defined by the number of 4KB pages, to allow Data Optimizer to allocate to hiperspace LSR buffers for an individual data set; previously, this field was specified as a number of bytes rather than a number of 4KB pages.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords MAXESTG

LSR Maximum System Expanded Storage Option

The LSR Maximum System Expanded Storage option specifies the maximum amount of expanded storage to permit Data Optimizer to allocate to hiperspace LSR buffers for all data sets that are intercepted within the associated system image.

Values This option has the following values:

0 to 2147483647

0 indicates that no value is specified.

2097151K *K* denotes kilobytes.

2047M *M* denotes megabytes.

Note: This value now represents the maximum amount of storage, as defined by the number of 4KB pages, to allow Data Optimizer to allocate to hiperspace LSR buffers for all data sets intercepted within the associated system image; previously, this field was specified as a number of bytes rather than a number of 4KB pages.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords STGEMAX

LSR Maximum System Expanded Storage Percent Option

The LSR Maximum System Expanded Storage Percent option specifies the maximum amount of expanded storage to permit Data Optimizer to allocate to hiperspace LSR buffers for all data sets that are intercepted within the associated system image as a percentage of the available expanded storage on the system.

Values This option has the following values:

0 to 100 0 indicates that no value is specified.

Note: For operating system levels that support expanded storage, the percentage value is calculated against the amount of online, available expanded storage. For OS390 R10 ESAME mode and z/OS systems, the percentage value is applied to the amount of online, available real storage; in either case, the resulting value is used to limit the amount of storage that is allocated to hiperspace LSR buffers for all data sets which are intercepted on the associated system image.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords STGEMAXPCT

Page Fix Option

The Page Fix option indicates whether or not to page-fix Data Optimizer LSR I/O control blocks.

Values This option has the following values:

Y Page-fix LSR I/O control blocks.

N Do not page-fix LSR I/O control blocks.

Default N

DD Overrides //DAP@FIX DD DUMMY
Page-fix LSR I/O control blocks; available at both step and DDNAME levels.

//DAP@NFIX DD DUMMY
Do not page-fix LSR I/O control blocks; available at both step and DDNAME levels.

//VVHCFIXY DD DUMMY
HIPER-CACHE compatibility DDNAME; page-fix LSR I/O control blocks; only available at the step level.

//VVHCFIXN DD DUMMY
HIPER-CACHE compatibility DDNAME; do not page fix LSR I/O control blocks; only available at the step level.

Data Policy Keywords PAGEFIX
NOPAGEFIX

Sequential Insert Strategy Option

The Sequential Insert Strategy option indicates to set the VSAM Sequential Insert Strategy option in the ACB.

Values This option has the following values:

Y Use VSAM Sequential Insert Strategy.

N Do not use VSAM Sequential Insert Strategy.

Default N

DD Overrides //DAP@SIS DD DUMMY
Use VSAM Sequential Insert Strategy; available at both step and DDNAME levels.

```
//DAP@NSIS DD DUMMY
```

Do not use the VSAM Sequential Insert Strategy; available at both step and DDNAME levels.

```
//VVHCSISY DD DUMMY
```

HIPER-CACHE compatibility DDNAME; use VSAM Sequential Insert Strategy; only available at the step level.

```
//VVHCSISN DD DUMMY
```

HIPER-CACHE compatibility DDNAME; do not use VSAM Sequential Insert Strategy; only available at the step level.

```
//VVHC$SIS DD DUMMY
```

HIPER-CACHE compatibility DDNAME; use VSAM Sequential Insert Strategy; only available at the DDNAME level

Data Policy Keywords SIS
NOSIS

Define LSR Pool ID Option

Defines a specific LSR pool number for Data Optimizer to use.

Note: Use this option with extreme caution and only if you are very familiar with the usage and interactions between data sets within a program. Improper use can compromise data integrity.

Note: This option is only available if neither DDNAME nor Data Set Name VSAM control block sharing is active.

Values This option has the following values:

0 to 255 0 indicates Data Optimizer chooses the LSR pool number.

Default 0

DD Overrides This option does not provide a DD override.

Data Policy Keywords LSRPOOLID

Override LSR Pool Specification Option

The Override LSR Pool Specification option indicates to replace an existing user LSR pool specification with Data Optimizer's own LSR pool definition. NOLSROVERRIDE is the default.

Values	This option has the following values:
	Y Override user's existing LSR specifications.
	N Do not override user's existing LSR specifications.
Default	N
DD Overrides	//DAP@OVL DD DUMMY Override user's existing LSR specifications; available at both step and DDNAME levels.
	//DAP@NOVL DD DUMMY Do not override user's existing LSR specifications; available at both step and DDNAME levels.
Data Policy Keywords	LSROVERRIDE NOLSROVERRIDE

LSR Alternate Index Option

The LSR Alternate Index option indicates to utilize LSR buffer processing for a path defined over an alternate index or for a cluster itself which contains at least one alternate index.

In previous releases of Data Optimizer, use of LSR buffering for alternate indexes and base clusters with alternate indexes was based upon various ACB options which resulted in not all alternate indexes being eligible for placement in an LSR pool. Now almost all alternate indexes processed through a path and their base clusters can be processed using LSR. With the addition of the new AIXLSR/NOAIXLSR option, and considering that earlier versions of the product did not provide the ability to control the buffering methodology for alternate indexes, the more conservative default of NOAIXLSR was chosen. Consequently, some clusters which were intercepted before may not be intercepted now if the default value of NOAIXLSR is used.

Note: If DDNAME or Data Set Name sharing is in effect for a group of data sets being intercepted by Data Optimizer, the first opener creates the LSR pool and subsequent openers participating in the sharing of the pool using the values established by the first opener.

Values This option has the following values:

- Y Allow processing of alternate index data sets.
- Note:** When AIXLSR is specified and LSR processing is used for an alternate index, buffer values calculated by Data Optimizer or specified by the user are for the base cluster associated with the path. Data Optimizer also assigns extra buffers for processing the alternate index associated with path under which the ACB was opened as well as for any members of the upgrade set.
- N Do not allow processing of alternate index data sets.
- Note:** When NOAIXLSR is specified and NSR processing is used for an alternate index, the buffer values calculated by Data Optimizer or specified by the user apply only to the AIX associated with the path. VSAM automatically assigns a minimal number of buffers for accessing the base cluster and any members of the upgrade set.

Default N

DD Overrides

//DAP@AIX DD DUMMY
Allow processing of alternate index data sets; available at both step and DDNAME levels.

//DAP@NAIX DD DUMMY
Do not allow processing of alternate index data sets; available at both step and DDNAME levels.

//VVHCAIXY DD DUMMY
HIPER-CACHE compatibility DDNAME; allow processing of alternate index data sets; only available at the step level.

//VVHCAIXN DD DUMMY
HIPER-CACHE compatibility DDNAME; do not allow processing of alternate index data sets; only available at the step level.

Data Policy Keywords AIXLSR
NOAIXLSR

Multistring Protection Option

The Multistring Protection option indicates to place VSAM data sets in an LSR pool if the original number of strings defined in the ACB is greater than one.

Values This option has the following values:

- Y Allow multi-string ACBs to use LSR buffering.
- N Do not allow multi-string ACBs to use LSR buffering.

Default N

DD Overrides //DAP@MST DD DUMMY
Allow multi-string ACBs to use LSR buffering; available at both step and DDNAME levels.

//DAP@NMST DD DUMMY
Do not allow multi-string ACBs to use LSR buffering; available at both step and DDNAME levels.

Data Policy Keywords MULTISTRING
NOMULTISTRING

Disable Batch LSR Interception Option

The Disable Batch LSR Interception option specifies whether or not to honor any IBM Batch Local Shared Resource specifications made in JCL for a data set that is eligible to be intercepted by Data Optimizer.

Values This option has the following values:

- Y Disable any Batch LSR subsystem specification and Data Optimizer decides the best buffering methodology and specifications. The Batch LSR subsystem name to be overridden is assumed to be BLSR.
- N Do not disable any Batch LSR subsystem specification.
- SSN Disable any Batch LSR subsystem specification and Data Optimizer decides the best buffering methodology and specifications. SSN is the four-character subsystem name. Any other value is assumed to be the subsystem name of a Batch LSR subsystem that Data Optimizer should disable.

Default N

DD Overrides //DAP@BLS DD DUMMY
Disable the Batch LSR subsystem with the subsystem name of BLSR; available at both step and DDNAME levels.

//DAP@NBLS DD DUMMY
Do not disable the Batch LSR subsystem with the subsystem name of BLSR; available at both step and DDNAME levels.

//VVHCBLRY DD DUMMY
 Disable the Batch LSR subsystem with the subsystem name of BLSR;
 HIPER-CACHE compatibility available at both step and DDNAME levels.

//VVHCBLRN DD DUMMY
 Hiper-Cache compatibility DDNAME; do not disable any Batch LSR
 specifications; available only at the step level

Data Policy Keywords DISABLEBATCHLSR=NO
 DISABLEBATCHLSR=YES
 DISABLEBATCHLSR=SSNM (*SSNM* is the four-character subsystem name)

Detect Exclusive Control Conflict Option

When a program opens a VSAM cluster for update purposes, the program requires exclusive control of the cluster. If the program also opens paths that are defined over an alternate index belonging to the same base cluster, and the program does not release its exclusive control prior to processing the other associated cluster, an exclusive-control conflict will occur if LSR mode processing is also in effect.

Furthermore, Data Optimizer often converts the buffering methodology for a VSAM data set from NSR to LSR. If the application that is using the data set cannot tolerate LSR processing (as a result of performing multiple concurrent requests to the same control interval, at least one of which requires exclusive control), a deadlock will occur.

The Detect Exclusive Control Conflict option lets you specify whether or not Data Optimizer should detect exclusive-control conflicts and issue messages BMC103687 and BMC103688 identifying the conflicting data sets.

Values This option has the following values:

Y Detect exclusive-control conflicts and issue messages identifying the conflicting data sets.

N Do not detect exclusive-control conflicts.

Default Y

DD Overrides //DAP@DCF DD DUMMY
 Data Optimizer should detect exclusive-control conflicts and issue messages identifying the conflicting data sets.

//DAP@NDCF DD DUMMY
 Data Optimizer should not detect exclusive-control conflicts

Data Policy Keywords DETECTCONFLICT
 NODETECTCONFLICT

Summary of Data Optimizer Options, Keywords, and DD Overrides

Table 5-1 provides a reference for Data Optimizer options, keywords, and DD overrides.

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 1 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
General Options				
• Action	EXCLUDE	EXCLUDE	EXCLUDE	DAP@NVPO, DAP@NNPO
	INCLUDE		INCLUDE	DAP@VPO, DAP@NPO
• Optimization Mode	ADVANCED	ADVANCED	ADVANCED	DAP@ADV
	BASIC		BASIC	DAP@BASC
	REPORTONLY		REPORTONLY	DAP@RPTO
• Optimization Percentage	0,1–999	0	VSAMOPTPCT NONVSAMOPTPCT	—
• Dynamic Region Adjustment	Y	N	DRA	DAP@DRA
	N		NODRA	DAP@NDRA
	N		NORUBHONOR	DAP@NRBH
• Optimize STC Access	Y	N	STC	—
	N		NOSTC	—
• Optimize TSO Access	Y	N	TSO	—
	N		NOTSO	—
• SMF Record Type	0, 128–255	0	SMFRECTYPE	—
• Override User Values	Y	N	OVERRIDEUSER	DAP@OVU
	N		NOOVERRIDEUSER	DAP@NOVU
• Honor RUB Value	Y	N	RUBHONOR	DAP@RBH

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 2 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
<ul style="list-style-type: none"> • Resource Usage Bias 	1–10, 99	99	RUB	DAP@VRxx, DAP@NRxx

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 3 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
Resource Utilization				
• Page Fault Rate Threshold	0–9999	0	PAGINGLIMIT	—
• CPU Utilization Threshold	0–999	0	CPULIMIT	—
• Below 16M Reserve	0–16777215	0	STG24RESERVE	—
• Above 16M Reserve	0–2147483647	0	STG31RESERVE	—
• Maximum Below 16M Reserve	0–16777215	0	STG24MAX	—
• Maximum Above 16M Reserve	0–2147483647	0	STG31MAX	—
Reporting Options				
• Report Destination	SYSMSG SYSOUT	SYSMSG	RPTDEST	DAPRPTS
• Statistics Report	Y N	N	STATRPT NOSTATRPT	DAP@STA DAP@NSTA
• Report SYSOUT Class	JES SYSOUT	*	RPTCLASS	DAPRPTS
• Data Set Information Report	Y N	N	DSINFORPT NODSINFORPT	DAP@DSI DAP@NDSI
• Catalog Information Report	Y N	N	CATRPT NOCATRPT	DAP@CAT DAP@NCAT
• Options Report	Y N	N	OPTIONRPT NOOPTIONRPT	DAP@OPT DAP@NOPT
• Message Level	0 1 2 3	0	MSGLEVEL0 MSGLEVEL1 MSGLEVEL2 MSGLEVEL3	DAP@MSG0 DAP@MSG1 DAP@MSG2 DAP@MSG3
• Message Destination	SYSMSG SYSOUT	SYSMSG	MSGDEST	DAPMESGS —
• Message SYSOUT Class	JES SYSOUT	*	MSGCLASS	DAPMESGS

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 4 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
Non-VSAM Options				
• Actual IO Counts	Y	N	AIC	DAP@AIC
	N		NOAIC	DAP@NAIC
• Allow Shared Writes	Y	N	SHAREDWRITE	DAP@SHW
	N		NOSHAREDWRITE	DAP@NSHW
• BUFNO Value	0–255	0	BUFNO	—
• BUFNO Storage	0–2147483647	0	BUFNOSTG	—
• BUFNO Maximum Value	0–255	0	MAXBUFNO	—
• BUFNO Maximum storage	0–2147483647	0	MAXBUFNOSTG	—
• NCP Value	0–255	0	NCP	—
• NCP Storage	0–2147483647	0	NCPSTG	—
• NCP Maximum value	0–255	0	MAXNCP	—
• NCP Maximum storage	0–2147483647	0	MAXNCPSTG	—
• Optimize Tape	Y	Y	TAPE	DAP@TAP DD DUMMY
	N		NOTAPE	DAP@NTAP DD DUMMY
• Optimize PO Data Sets Option	Y	Y	PDS	DAP@PDS DD DUMMY
	N		NOPDS	DAP@NPDS DD DUMMY
• Optimize DA Data Sets Option	Y	Y	DIRECT	DAP@DRC DD DUMMY
	N		NODIRECT	DAP@NDRC DD DUMMY
• Optimize EXCP Access Option	Y	Y	EXCP	DAP@XCP DD DUMMY
	N		NOEXCP	DAP@NXCP DD DUMMY

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 5 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
VSAM Options				
• LSR Option	F	Y	FORCELSR	DAP@FLSR
	Y		LSR	DAP@LSR
	N		NOLSR	DAP@NLSR
	P		PREFERLSR	DAP@PLSR
	R		REQUIRELSR	DAP@RLSR
• Defer Writes	Y	N	VDW	DAP@VDW
	N		NOVDW	DAP@NVDW
• RMODE31 Control Blocks	Y	N	VSAMCB31	DAP@DVC
	N		NOVSAMCB31	DAP@NDVC
• SHAREOPTIONS (3 3) Support	Y	N	SHROPT33	DAP@SHR
	N		NOSHROPT33	DAP@NSHR
• NSR Index Buffers	0–255	0	NSRIBUF	—
• NSR Index Buffer Storage	0–2147483647	0	NSRISTG	—
• NSR Max Index Buffers	0–255	0	MAXNSRIBUF	—
• NSR Max Index Buffer Storage	0–2147483647	0	MAXNSRISTG	—
• NSR Data Buffers	0–255	0	NSRDBUF	—
• NSR Data Buffer Storage	0–2147483647	0	NSRDSTG	—
• NSR Max Data Buffers	0–255	0	MAXNSRDBUF	—
• NSR Max Data Buffer Storage	0–2147483647	0	MAXNSRDSTG	—
• LSR Index Buffers	0–65535	0	LSRIBUF	—
• LSR Index Buffer Storage	0–2147483647	0	LSRISTG	—
• LSR Max Index Buffers	0–65535	0	MAXLSRIBUF	—
• LSR Max Index Buffer Storage	0–2147483647	0	MAXLSRISTG	—
• LSR Data Buffers	0–65535	0	LSRDBUF	—

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 6 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
VSAM Options (continued)				
• LSR Data Buffer Storage	0–2147483647	0	LSRDSTG	—
• LSR Max Data Buffers	0–65535	0	MAXLSRDBUF	—
• LSR Max Data Buffer Storage	0–2147483647	0	MAXLSRDSTG	—
• Percent Index to Buffer	0–100	0	PCTLSRIBUF	—
• User LSR Hiperspace Buffers	Y	N	HIPERSPACE	DAP@DVH
	N		NOHIPERSPACE	DAP@NDVH
• Percent Hiperspace Buffers	0–100	0	HIPERPCT	—
• Hiperspace Buffer Storage	0–2147483647	0	HIPERSTG	—
• LSR Max Expanded Storage	0–2147483647	0	MAXESTG	—
• LSR Max Sys Expanded Storage	0–2147483647	0	STGEMAX	—
• Percent Max Sys Expanded Storage	0–100	0	STGEMAXPCT	—
• Page Fix	Y	N	PAGEFIX	DAP@FIX
	N		NOPAGEFIX	DAP@NFIX
• Sequential Insert Strategy	Y	N	SIS	DAP@SIS
	N		NOSIS	DAP@NSIS
• Multistring	Y	N	MULTISTRING	DAP@MST
	N		NOMULTISTRING	DAP@NMST
• Override LSR	Y	N	LSROVERRIDE	DAP@OVL
	N		NOLSR OVERRIDE	DAP@NOVL
• LSR Pool ID	1–255	0	LSRPOOLID	—
• LSR Alternate Index	Y	N	AIXLSR	DAP@AIX
	N		NOAIXLSR	DAP@AIXN

Table 5-1 Reference Table for Options, Keywords, and DD Overrides (Part 7 of 7)

ISPF Dialog Field Names	Dialog Values	Default	Report Keyword	DD Override
VSAM Options (continued)				
<ul style="list-style-type: none"> • Disable Batch LSR 	Y	N	DISABLEBATCHLSR= YES	DAP@BLS
	N		DISABLEBATCHLSR= NO	DAP@NBLS
	SSNM (four-character subsystem name)		DISABLEBATCHLSR= SSNM	—
<ul style="list-style-type: none"> • Detect Exclusive Control Conflicts 	Y	Y	DETECTCONFLICT	DAP@DCF
	N		NODETECTCONFLICT	DAP@NDCF

Chapter 6 Data Optimizer Commands

This chapter describes the MAINVIEW[®] Batch Optimizer and Data Optimizer commands necessary to control I/O optimization processing. This chapter discusses the following topics:

Overview	6-2
Starting the BMC Software Primary Subsystem	6-6
Displaying the Status of the BMC Software Primary Subsystem	6-7
Stopping the BMC Software Primary Subsystem	6-8
Cancelling the BMC Software Primary Subsystem	6-9
Starting the MAINVIEW Batch Optimizer Subsystem	6-10
Displaying the Status of the MAINVIEW Batch Optimizer Subsystem	6-12
Stopping the MAINVIEW Batch Optimizer Subsystem	6-13
Cancelling the MAINVIEW Batch Optimizer Subsystem	6-14
Displaying the Status of Data Optimizer	6-15
Enabling and Disabling Components	6-16
Displaying the Status of a Data Policy	6-17
Activating a Data Policy	6-18
Displaying the Status of the User Control Exit	6-19
Activating the User Control Exit	6-20

Overview

I/O optimization requires processing from both the MAINVIEW Batch Optimizer subsystem and Data Optimizer components. Commands provide control to these components and allow you to display status information about them. You can enter any of these commands from the operator console or from any facility that emulates an operator console. Status changes made by commands entered from the operator console will be active only until the next IPL. You can also place any of the MAINVIEW Batch Optimizer subsystem or Data Optimizer commands in the MAINVIEW Batch Optimizer subsystem commands member BCSCMD.xx so that they will be issued each time the MAINVIEW Batch Optimizer Subsystem is started.

The following tasks are discussed in this chapter:

- starting and stopping the subsystems
- displaying the status of the subsystems
- displaying the status of Data Optimizer
- displaying data policy status information
- enabling and disabling components
- enabling and disabling the user control exit

Operator Console Commands

MAINVIEW Batch Optimizer operator console commands require RACF access authority. Table 6-1 lists the commands for MAINVIEW Batch Optimizer components and subsystems and lists the RACF operator command-level authority that is required for entering the commands.

Table 6-1 MAINVIEW Batch Optimizer Component and Subsystem Commands (Part 1 of 4)

Command	Function	Required RACF Access
BMC Software Primary Subsystem		
START bmcpproc	starts the BMC Software Primary Subsystem	CONTROL
bmcps STATUS	provides the subsystem address space ID and indicates whether DEBUG mode is active or inactive	READ
bmcps SHUTDOWN	offers a nondisruptive way of stopping a BMC Software Primary subsystem, allowing jobs to finish before ending the subsystem	CONTROL
CANCEL bmcpproc	stops the BMC Software Primary Subsystem immediately, regardless of whether work is in progress	CONTROL
MAINVIEW Batch Optimizer Subsystem		

Table 6-1 MAINVIEW Batch Optimizer Component and Subsystem Commands (Part 2 of 4)

Command	Function	Required RACF Access
START mbosproc	starts the MAINVIEW Batch Optimizer Subsystem	CONTROL
mbos STATUS	provides the subsystem address space ID and indicates whether DEBUG mode is active or inactive	READ
mbos SHUTDOWN	stops the MAINVIEW Batch Optimizer Subsystem and stops all processing for Job Optimizer and Data Optimizer	CONTROL
CANCEL mbosproc	stops the MAINVIEW Batch Optimizer Subsystem immediately, regardless of whether work is in progress	CONTROL
Extended Job Execution Subsystem		
mbos BSL XJS STARTUP	starts the extended job execution subsystem and the MBOX initiators	CONTROL
mbos BSL XJS STATUS	provides information about the extended job execution subsystem and indicates whether the MBOX initiators are active or inactive	CONTROL
mbos BSL XJS QUIESCE	stops the MBOX initiators on a particular MVS image	CONTROL
mbos BSL XJS ACTIVATE	starts the MBOX initiators on a particular MVS image	CONTROL
mbos BSL XJS SHUTDOWN	stops the extended job execution subsystem and the MBOX initiators	CONTROL
mbos BSL XJS TRACE	provides trace information that BMC Software can use to follow the processing of the extended job execution subsystem	CONTROL

Table 6-1 MAINVIEW Batch Optimizer Component and Subsystem Commands (Part 3 of 4)

Command	Function	Required RACF Access
Job Optimizer Component		
mbos REINIT BSL	reinitializes Job Optimizer after an IPL or abend	CONTROL
mbos BSL STATUS	provides general information about the status of Job Optimizer	CONTROL
mbos BSL DISABLE	stops the Job Optimizer component from intercepting jobs for performance processing	CONTROL
mbos BSL ENABLE	starts the Job Optimizer component, which can begin intercepting jobs for performance processing	CONTROL
mbos BSL QUIESCE	stops a BatchPlex image from accepting Job Optimizer work	CONTROL
mbos BSL WAITS	displays the waiting status of job steps that are running in an MBOX initiator	CONTROL
mbos BSL POLICY ACTIVATE	activates a new or changed job policy	CONTROL
mbos BSL POLICY STATUS	displays the active job policy definition	CONTROL
mbos POLICY DISPLAY	is where Job Optimizer is to write the job policy action definition that is associated with a job upon termination of the job	CONTROL
mbos BSL READJTL	sets or displays the default action for the READJTL job policy action definition	CONTROL
mbos BSL UCF STATUS	displays the active UCF definition	CONTROL
mbos BSL UCF Activate	activates a new or changed UCF definition	CONTROL
mbos BSL GROUP	specifies which JES3 initiator groups MAINVIEW Batch Optimizer is to intercept and lists the initiator groups that MAINVIEW Batch Optimizer is intercepting	CONTROL
mbos BSL DUMP	produces various dumps that pertain to the specified batch performance subsystem	CONTROL
mbos BSL DEBUG	controls the capturing of dumps for the specified batch performance subsystem	CONTROL
Job Optimizer Pipes Component		
START mbop	starts Job Optimizer Pipes	CONTROL
F mbop, SHUT TYPE=NORMAL	stops Job Optimizer Pipes	CONTROL
Data Optimizer Component		
mbos REINIT DAP	reinitializes Data Optimizer after an IPL or abend	CONTROL
mbos DAP STATUS	provides general information about the status of Data Optimizer	CONTROL
mbos DAP DISABLE	stops the Data Optimizer component from providing performance processing for data set I/O requests	CONTROL
mbos DAP ENABLE	starts the Data Optimizer component, which begins providing performance processing for data set I/O requests	CONTROL

Table 6-1 MAINVIEW Batch Optimizer Component and Subsystem Commands (Part 4 of 4)

Command	Function	Required RACF Access
mbos DAP HISTKEY DSNPGM	creates the history key, where the history key is DNS and program name	CONTROL
mbos DAP HISTKEY FULL	creates the history key, where the history key is job, step, procstep, program and DDNAME	CONTROL
mbos DAP POLICY STATUS	displays the active data policy definition	CONTROL
mbos DAP POLICY ACTIVATE	activates a new or changed data policy	CONTROL

Starting the BMC Software Primary Subsystem

Summary: The BMC Software Primary subsystem (BMCP) is required by the MAINVIEW Batch Optimizer subsystem (MBOS). If you start the MBOS before starting the BMCP, the MBOS detects that the BMCP is not available. The MBOS then enters a command to start the BMCP.

To start the BMCP, complete the following steps:

Step 1 *Required.* Issue the following command:

```
START bmcpproc
```

Step 2 *Optional.* Review the system log to ensure that the command completed successfully.

The BMCP responds with messages indicating that it is available. Status messages are also displayed.

Displaying the Status of the BMC Software Primary Subsystem

Summary: Once you see message BMC10100I in the system log, you can display the status of the BMCP to determine whether the subsystem is active. Each MVS image can have only one BMCP active.

To display the status of the BMCP, complete the following steps:

Step 1 *Required.* Issue the following command:

```
bmcp STATUS
```

Step 2 *Optional.* Review the system log to ensure that the command completed successfully.

The BMCP responds with message BMC10700I indicating the status of the BMCP.

Stopping the BMC Software Primary Subsystem

Summary: BMC Software Product Support recommends that you leave the BMCP running continuously. You must stop it, however, as part of an orderly shutdown of MVS before an IPL. To stop the BMCP, issue the SHUTDOWN command.

To stop the BMCP, complete the following steps:

Step 1 *Required.* Issue the following command:

```
bmcpr  SHUTDOWN
```

Step 2 *Optional.* Review the system log to ensure the command successfully completed.

The BMCP responds with message BMC10106I indicating that termination was initiated and message BMC10228I indicating that termination completed. Status messages are also displayed.

Cancelling the BMC Software Primary Subsystem

Summary: To stop the BMCP immediately, use the MVS CANCEL command.

To cancel the BMCP, issue the following command:

```
CANCEL bmcpproc
```

Note: BMC Software recommends that you use the SHUTDOWN command to stop the BMCP in a non disruptive way. If you use the CANCEL command, unpredictable results might occur.

Warning! Use the MVS FORCE command only as a last resort when the CANCEL command fails to perform its function even after you have entered it several times. Using the FORCE command might require you to re-IPL the system.

Starting the MAINVIEW Batch Optimizer Subsystem

Summary: You must start the MBOS before any Data Optimizer processing can occur. Start the MBOS by entering the MVS START command.

To start the MBOS, complete the following steps:

Step 1 *Required.* Issue the following command:

```
START mbobcss ,AUTOSTRT=x
```

Include the AUTOSTRT=Y parameter to have the MBOS to use the autostart specifications for MVS images in the BatchPlex definition. MVS START commands are routed to all MVS images that are associated with the BatchPlex for which the **Automatically Start Subsystems** field is set to Y for the image.

Specify the AUTOSTRT=N parameter to have the MBOS ignore the autostart specifications for MVS images in the BatchPlex definition. The subsystems are not automatically started on other MVS images in the BatchPlex.

Step 2 *Optional.* Review the system log to ensure the command successfully completed.

The MBOS responds with message BMC11100I indicating that it is available. Status messages are also displayed.

When you start the MBOS, the following occur:

- The MBOS responds with messages indicating its availability.
- If the BMCP is not active, the MBOS starts it automatically.

If you have not performed an IPL since the last time you started the MBOS, the BMCP is not restarted. The BMCP is still in memory from the previous start.

The BatchPlex that is identified in the MBOS procedure is started, and the data policy that is identified in the BatchPlex definition is activated.

If you previously used the ACTIVATE command to start another data policy, this policy is not refreshed from memory when the MBOS is started. Only the data policy and that is identified in the BatchPlex definition is started. To ensure that the correct data policy is activated when the MBOS is started, you must change the BatchPlex definition.

Set Automatic Start of Subsystems for IPL

You can place the MVS START command for the MBOS in a COMMND_{xx} member of PARMLIB so that the subsystem starts automatically during an IPL. Do not include the AUTOSTRT parameter on START commands that are placed in a COMMND_{xx} member of PARMLIB.

Starting Under an Alternate JES Subsystem

If your site uses alternate JES subsystems, you can initialize MAINVIEW Batch Optimizer under an alternate JES subsystem by using the SUB parameter in the START command of the MBOS as in the following command

```
s mbobcss ,SUB=altj
```

where:

- *mbobcss* is the procedure name for the MBOS
- *altj* is the subsystem ID of the alternate JES subsystem

When MAINVIEW Batch Optimizer is initialized under the alternate JES subsystem, Data Optimizer intercepts only jobs that run under the alternate JES subsystem.

Displaying the Status of the MAINVIEW Batch Optimizer Subsystem

Summary: You can display the status of the MBOS to determine whether the subsystem is active. Display the status of the MBOS by entering the STATUS command.

To display the status of the MBOS, complete the following steps:

Step 1 *Required.* Issue the following command:

```
mbos STATUS
```

Step 2 *Optional.* Review the system log to ensure that the command successfully completed. The MBOS responds with messages indicating the status of the address space.

Stopping the MAINVIEW Batch Optimizer Subsystem

Summary: You can leave the MBOS running continuously. You must stop it, however, as part of an orderly shutdown of MVS before an IPL. To stop the MBOS address space, issue the SHUTDOWN command.

To stop the MBOS, complete the following steps:

Step 1 *Required.* Issue the following command to stop the MBOS address space in a nondestructive manner:

```
mbos SHUTDOWN
```

When you stop the MBOS, the following actions occur:

- The MBOS responds with message BMC11106I indicating that termination was initiated and message BMC11228I indicating that termination completed.
- The MBOS remains in memory until you perform an IPL.

Cancelling the MAINVIEW Batch Optimizer Subsystem

Summary: To stop the MBOS immediately, use the MVS CANCEL command.

To cancel the MBOS, complete the following step:

Step 1 *Required.* Issue the following command:

```
CANCEL mbobcss
```

Note: BMC Software Product Support recommends that you use the SHUTDOWN command to stop the MBOS in a non disruptive way. If you use the CANCEL command, unpredictable results might occur.

Warning! Use the MVS FORCE command only as a last resort when the CANCEL command fails to perform its function, even after you have entered it several times. Using the FORCE command might require you to IPL the system.

Displaying the Status of Data Optimizer

Summary: You can display status information about Data Optimizer and its components.

To display status information, complete the following steps:

Step 1 *Required.* Issue one of the following commands:

```
mbos DAP STATUS ALL
```

This command displays all status information about Data Optimizer, such as whether the VSAM or non-VSAM components are enabled, version and release information, and the date on which the Data Optimizer tape was generated.

```
mbos DAP STATUS VSAM
```

This command displays the general status information plus component-specific status information for the initialized VSAM components.

```
mbos DAP STATUS [NVSAM | NONVSAM]
```

This command displays the general status information plus component-specific status information for the initialized non-VSAM components.

Note: MBOS DAP STATUS is also valid and defaults to all.

Step 2 *Optional.* Review the system log to ensure that the command completed successfully.

Data Optimizer responds with messages indicating the status of the selected components.

Enabling and Disabling Components

Summary: You can enable and disable the Data Optimizer components on any Data Optimizer image. When a component is disabled, it will not intercept any requests—VSAM or non-VSAM, as appropriate.

To enable or disable Data Optimizer components, complete the following steps:

Step 1 *Required.* Issue any of the following commands:

```
mbos DAP [VSAM | NVSAM | NONVSAM] ENABLE
```

Include the VSAM or NONVSAM parameter to limit the request to the VSAM or non-VSAM components. Omit the VSAM and NONVSAM parameters to enable both the VSAM and non-VSAM components.

```
mbos DAP [VSAM | NVSAM | NONVSAM] DISABLE
```

Include the VSAM or NONVSAM parameter to limit the request to the VSAM or non-VSAM components. Omit the VSAM and NONVSAM parameters to disable both the VSAM and non-VSAM components.

Step 2 *Optional.* Review the system log to ensure that the command completed successfully. Data Optimizer responds with messages indicating which components have been enabled or disabled.

Displaying the Status of a Data Policy

Summary: You can display status information about Data Optimizer data policies.

To display the status of the active data policy, complete the following steps:

Step 1 *Required.* Issue the following command:

```
mbos DAP POLICY STATUS [policymembername]
```

This command displays general status information about the data policy, where *policymembername* is the name of the specific data policy. If *policymembername* is omitted, the status of the currently active data policy is displayed.

Step 2 *Optional.* Review the system log to ensure the command completed successfully.

MAINVIEW Batch Optimizer responds with messages indicating that the status of the selected components.

Activating a Data Policy

Summary: You can use this command to activate data policies.

To activate data policies, complete the following steps:

Step 1 *Required.* Issue the following command:

```
mbos DAP POLICY ACTIVATE [policymembername]
```

This command activates a specific data policy, where *policymembername* is the name of the specific data policy. If *policymembername* is omitted, the currently active data policy is re-activated.

Step 2 *Optional.* Review the system log to ensure the command completed successfully.

MAINVIEW Batch Optimizer responds with message BMC101111I indicating the activation of the selected policy, or message BMC101112E if the activation process encounters errors.

Displaying the Status of the User Control Exit

Summary: You can display the status of the user exit to determine which exit is currently active.

To enable or disable the user control exit, complete the following steps:

Step 1 *Optional.* Issue the following command to review the status of the user exit:

```
mbos DAP USEREXIT STATUS
```

When you issue this command, the last line of the display shows the user exit that is currently active.

Step 2 *Required.* Review the system log to ensure the command completed successfully.

Data Optimizer responds with message BMC101700I indicating the availability of the user control exit.

Activating the User Control Exit

Summary: You can use commands to enable and disable the user control exit at a global level.

To enable or disable the user control exit, complete the following steps:

Step 1 *Required.* Issue one of the following commands:

```
mbos DAP USEREXIT ENABLE [exitname]
```

This command enables the user control exit and lets it provide interception decisions. *Where exitname indicates a specific user control exit name.* If you do not specify a name, Data Optimizer uses DAPUEXIT as the default.

Note: If you have a user exit enabled and want to enable a new user exit, you must first disable any previously enabled user exits.

```
mbos DAP USEREXIT DISABLE [exitname]
```

This command disables the user control exit and keeps it from making interception decisions. *Where exitname indicates a specific user control exit name.* If you do not specify a name, Data Optimizer uses DAPUEXIT as the default.

Step 2 *Required.* Review the system log to ensure the command completed successfully.

Data Optimizer responds with message BMC101700I indicating the availability of the user control exit.

Setting the History Key

You can set which information should be used to create the history key. Data Optimizer will use the specified format to create and retrieve historical data.

To set history key information, complete the following steps:

Step 1 *Required.* Issue one of the following commands:

```
mbos DAP HISTKEY [FULL]
```

Where mbos is the default MAINVIEW Batch OPTimizer subsystem identifier.

This command indicates that the historical information should be based on the unique combination of job name, step name, procedure name, program name, and DD name. FULL is the default method.

```
mbos DAP HISTKEY [DSNPGM]
```

This command indicates that the historical information should be based on the unique combination of data set name and program name.

Step 2 *Optional.* Review the system log to ensure the command completed successfully. Data Optimizer responds with message indicating which components have been enabled or disabled.

Appendix A Data Optimizer Reports

This appendix explains the reports that can be generated by the Data Optimizer component of MAINVIEW® Batch Optimizer. This appendix discusses the following topics:

Overview	A-1
Non-VSAM Performance Report	A-2
Non-VSAM Performance Report Header	A-4
Non-VSAM Options Report	A-6
Non-VSAM Statistics Report	A-7
Non-VSAM Data Set Information Report	A-8
VSAM Performance Report	A-11
VSAM Performance Report Header	A-15
VSAM Options Report	A-16
VSAM Statistics Report	A-19
VSAM Data Set Information Report	A-20
VSAM Catalog Information Report	A-22

Overview

Data Optimizer generates reports that provide performance information for each data set access that is processed. The following types of reports are available:

- Options Report
- Statistics Report
- Data Set Information Report
- Catalog Information Report (for VSAM data sets only)

You can control where the reports are written using the Report Destination option. You can direct the output to the JES system message log (SYSMSG), a separate JES SYSOUT listing, or to a disk data set.

Non-VSAM Performance Report

The Non-VSAM Performance Report is actually a group of reports, each which you can request or exclude. All requested reports are produced on one printable output. You can customize the information generated by the Non-VSAM Performance Report by using the data policy options or DD override statements.

There are three types of information that can be included in this report:

- data policy options
- performance statistics
- data set information

Figure A-1 on page A-3 shows an example of the Non-VSAM Performance Report. Subsequent pages provide descriptions of the data fields included in each section of the report.

Figure A-1 Non-VSAM Performance Report

```

***** DATA OPTIMIZER PERFORMANCE
REPORT *****
* JOBNAME      : JJHT1          STEPNAME : STATTEST   PROCSTEPNAME :
*
* JOBCLASS    : A              USERID   : JJH2        GROUP        :
CC105413 *
* ADSPC TYPE  : JOB           PGMNAME  : S
*
* DDNAME      : DD01          DSTYPE   : NON-VSAM  SUBSYSTEM    : DAPR
*
* DATACLASS  : **NONE**      STORCLASS: **NONE**  MGMTCLASS    :
**NONE** *
* DSNAME      : SYS00217.T021202.RA000.JJHT1.DD01.H05
*
*
* INTERCEPTED BY DATAPOLICY(REGDATA) POLICYLINE(00011)
*
*
* OPTIMIZATION MODE..... ADVANCED..... CENT.....
100%*
* HONOR RUB VALUE..... NO OPTIMIZE STC ACCESS..... NO
*
* RESOURCE USAGE BIAS..... 0 OPTIMIZE TSO ACCESS..... NO
*
* DYNAMIC REGION ADJUSTMENT..... YES OVERRIDE USER VALUES..... NO
*
* OPTIMIZE TAPE DATA SETS..... YES OPTIMIZE PO DATA SETS..... YES
*
* OPTIMIZE EXCP ACCESS..... YES OPTIMIZE DA DATA SETS..... YES
*
* PAGE FAULT RATE THRESHOLD..... 0 BELOW 16M RESERVE..... 100K
*
* CPU UTILIZATION THRESHOLD..... 101% ABOVE 16M RESERVE..... 12M
*
* ACTUAL IO COUNTS..... NO MA
*
* ALLOW SHARED WRITES..... NO MAXIMUM ABOVE 16M ..... 0
*
* BUFNO VALUE..... 0 NCP VALUE..... 0
*
* BUFNO STORAGE..... 0 NCP STORAGE..... 0
*
* BUFNO MAXIMUM VALUE..... 0 NCP MAXIMUM VALUE..... 0
*
* BUFNO MAXIMUM STORAGE..... 0 NCP MAXIMUM STORAGE..... 0
*
* SMF RECORD TYPE..... 100%
*

```

Non-VSAM Performance Report Header

Non-VSAM Options Report

Non-VSAM Statistics Report

Non-VSAM Data Set Information Report

```

*
* I/O SAVINGS .....90.00%      ADDITIONAL BUFFERS.....
800 *
* LOGICAL EXCPS..... 800      PHYSICAL EXCPS.....
80 *
* LOGICAL
BLOCKS.....1290                      *
*
*
* DSORG..... PS      RECFM.....
F *
* LRECL..... 80      BLOCKSIZE.....
80 *
* VOLSER..... WORK0R      DEVICE TYPE.....
3390 *
* VOLUME COUNT..... 1      DCBE.....
NO *
* ACCESS METHOD..... QSAM      OPEN OPTION.....
INPUT *
* MACRO OPTION..... GM      DCBMACRF.....
5000 *
* BUFNO..... 5      NCP.....
N/A *
* DCBOPTCD..... 00      ERROR OPTION.....
NONE *
* EODAD..... NO      SYNAD.....
NO *
* DCB OPEN EXIT..... NO      DCB ABEND EXIT.....
NO *
* JFCB EXIT..... NO      JFCBE EXIT.....
NO *
* EOVS EXIT..... NO
*
*****
*****

```

Non-VSAM Performance Report Header

The first part of Figure A-1 on page A-3 shows the header information that is common to all requested non-VSAM reports. Table A-1 describes the fields that are included in this header.

Table A-1 Non-VSAM Performance Report Common Field Descriptions (Part 1 of 2)

Field	Description
JOBNAME	Name of the job in which the data set was opened
STEPNAME	Name of the step in which the data set was opened

Table A-1 Non-VSAM Performance Report Common Field Descriptions (Part 2 of 2)

Field	Description
PROCSTEPNAME	Name of the procedure step in which the data set was opened
JOBCLASS	JES job class of the job in which the data was opened
USERID	Security system userid associated with the job in which the data set was opened
GROUP	Security system group associated with the job in which the data set was opened
ADSPC TYPE	Type of address space in which the data set was opened. The possible values are: <ul style="list-style-type: none"> • JOBbatch job • TSUTSO user • STCstarted task
PGMNAME	Name of the job step program which opened the data set
OPEN TIME	Time the data set was opened in hour and minute (24-hour clock format)
DDNAME	The DD name of the data set which was opened
DSTYPE	The value <i>NON-VSAM</i> denotes a non-VSAM data set type.
SUBSYSTEM	Subsystem ID of the MAINVIEW Batch Optimizer subsystem which processed the data set
DATACLASS	SMS Data Class of the data set which was opened. If the data set is not under SMS control, the value "***NONE**" will appear.
STORCLASS	SMS Storage Class of the data set which was opened. If the data set is not under SMS control, the value "***NONE**" will appear.
MGMTCLASS	SMS Management Class of the data set which was opened. If the data set is not under SMS control, the value "***NONE**" will appear.
DSNAME	The name of the data set which was opened
INTERCEPTED BY (interception reason)	Indicates the reason the data set was intercepted by Data Optimizer. Possible values are: <ul style="list-style-type: none"> • INTERCEPTED BY DATA POLICY(<i>XXXXXXXX</i>) POLICYLINE(<i>NNNNN</i>) where <i>XXXXXXXX</i> is the Data Policy member name active when the data set was opened and <i>NNNNN</i> is the line number within the Data Policy which caused interception • INTERCEPTED BY DD STATEMENT OVERRIDE indicating that a JCL DD statement caused interception • INTERCEPTED BY USER EXIT indicating that the Data Optimizer user exit requested interception If the Report Only option was specified, the interception reason will be suffixed with "FOR REPORTING ONLY"

Non-VSAM Options Report

The Options Report will display the options in effect at the time of open according to the options hierarchy, DD override, user exit, then data policy. The second group of data items in Figure A-1 on page A-3 reflect the Non-VSAM Options Report information. Table A-2 shows fields that are included in the non-VSAM Options Report.

Note: The following values are user-defined, but are not necessarily the values actually used.

Table A-2 Non-VSAM Options Report Field Descriptions (Part 1 of 2)

Field	Description
OPTIMIZATION MODE	Defines the optimization technique. Possible values are ADVANCED, BASIC, or REPORTONLY.
OPTIMIZATION PERCENT	Defines the percentage to apply to Data Optimizer's calculated performance optimization buffering values
HONOR RUB VALUE	Set to YES or NO to indicate whether the Resource Usage Bias value should be used. A setting of YES indicates that Data Optimizer will run in compatibility mode with earlier versions of the product.
OPTIMIZE STC ACCESS	Set to YES or NO to indicate whether data sets opened within started tasks will be optimized
RESOURCE USAGE BIAS	Defines the Resource Usage Bias value in effect. Only has meaning if HONOR RUB VALUE is set to YES.
OPTIMIZE TSO ACCESS	Set to YES or NO to indicate whether data sets opened within Time Sharing Option. User address spaces will be optimized.
DYNAMIC REGION ADJUSTMENT	Set to YES or NO to indicate whether Dynamic Region Adjustment was active
OVERRIDE USER VALUES	Set to YES or NO to indicate whether to replace user-coded DCB values if they are higher than those chosen by Data Optimizer
PAGE FAULT RATE THRESHOLD	The page fault rate used to determine whether the system is paging constrained at the time of the open
BELOW 16M RESERVE	The amount of storage below the 16-megabyte line Data Optimizer is to leave unused and available for user processing. The suffixes K and M can be used to denote kilobytes and megabytes.
CPU UTILIZATION THRESHOLD	The CPU utilization percentage value used to determine whether the system is cpu constrained at the time of the open.
ABOVE 16M RESERVE	The amount of storage above the 16-megabyte line Data Optimizer is to leave unused and available for user processing. The suffixes K and M can be used to denote kilobytes and megabytes.
ACTUAL IO COUNTS	Set to YES or NO to indicate whether the Actual I/O Counts option was enabled at the time of the open
MAXIMUM BELOW 16M	The maximum amount of storage below the 16-megabyte line Data Optimizer is to allocate for performance processing over and above that which would have been used without Data Optimizer involvement. The suffixes K and M can be used to denote kilobytes and megabytes.

Table A-2 Non-VSAM Options Report Field Descriptions (Part 2 of 2)

Field	Description
ALLOW SHARED WRITES	Set to YES or NO to indicate whether Data Optimizer is to permit multiple opens for output processing when it intercepts the data set
MAXIMUM ABOVE 16M	The maximum amount of above the 16-megabyte line Data Optimizer is to allocate for performance processing over and above that which would have been used without Data Optimizer involvement. The suffixes K and M can be used to denote kilobytes and megabytes.
BUFNO VALUE	Defines the DCB BUFNO value to use
NCP VALUE	Specifies the NCP value within an intercepted DCB
BUFNO STORAGE	Specifies the amount of storage to be allocated to buffers with the BUFNO DCB parameter
NCP STORAGE	Specifies the amount of storage to be allocated for NCP processing
BUFNO MAXIMUM VALUE	Specifies the maximum BUFNO value in a DCB
NCP MAXIMUM VALUE	Specifies the maximum NCP value in a DCB
BUFNO MAXIMUM STORAGE	Specifies the maximum amount of storage allocated to buffers defined by the BUFNO DCB parameter
NCP MAXIMUM STORAGE	Specifies the maximum amount of storage allocated to buffers by the DCB NCP parameter
SMF RECORD TYPE	Defines the SMF record type to be used when producing Data Optimizer SMF records. If set to a value of zero, no SMF records will be produced.
OPTIMIZE TAPE DATA SETS	Set to YES or NO to indicate whether optimization should be attempted for tape data sets
OPTIMIZE PO DATA SETS	Set to YES or NO to indicate whether optimization should be attempted for PO (partitioned) organization data sets
OPTIMIZE EXCP ACCESS	Set to YES or NO to indicate whether optimization should be attempted for data sets being processed using the EXCP access method
OPTIMIZE DA DATA SETS	Set to YES or NO to indicate whether optimization should be attempted for DA (direct access) organization data sets

Non-VSAM Statistics Report

The Statistics Report displays I/O statistics information for the processed data set. The third group of data items in Figure A-1 on page A-3 reflect the Non-VSAM Statistics Report information. Table A-3 shows fields that are included in the non-VSAM Figure A-1 on page A-3.

Table A-3 Non-VSAM Statistics Report Field Descriptions

Field	Description
I/O SAVINGS	Percentage savings of EXCPs realized by using Data Optimizer
ADDITIONAL BUFFERS	Amount that the BUFNO/NCP value (as coded in the DCB or JCL) was increased by Data Optimizer, either dynamically or by user request

Table A-3 Non-VSAM Statistics Report Field Descriptions

Field	Description
LOGICAL EXCPS	Number of EXCP requests issued by the application (or access method) and intercepted by Data Optimizer. This field will appear <i>only</i> if ADVANCED Data Optimization Mode is requested <i>and</i> can be accommodated by the product based on conditions at time of processing.
PHYSICAL EXCPS	Number of physical I/Os issued to the data set by Data Optimizer
LOGICAL BLOCKS	Number of blocks processed by the application (the value reported as "EXCP count" by SMF). This field will appear <i>only</i> if ADVANCED Data Optimization Mode is requested <i>and</i> can be accommodated by the product based on conditions at time of processing.

Non-VSAM Data Set Information Report

The Data Set Information Report provides detailed attribute information about the data set being processed. Values generally reflect those in effect after open processing has completed and before the first I/O.

For concatenations, this data represents only the first data set in the concatenation. If a particular value does not apply to a specific type of DCB or is unavailable, the value "N/A" is displayed.

The final group of data items in Figure A-1 on page A-3 reflect the non-VSAM Data Set Information Report information. Table A-4 shows fields that are included in the non-VSAM Data Set Information Report.

Table A-4 Non-VSAM Data Set Information Report Field Descriptions (Part 1 of 4)

Field	Description
DSORG	Data set organization of the data set. Possible values are: <ul style="list-style-type: none"> • PPhysical sequential • PPartitioned organization • DAdirect access • OTHERdata set organization other than those above
RECFM	Record format of the data set. Possible values are: <ul style="list-style-type: none"> • FFixed format • VVariable format • UUndefined format • BBlocked format • SSpanned/standard format • AASA control characters present • MMachine control characters present
LRECL	Logical record size of the data set
BLOCKSIZE	Block size of the data set
VOLSER	First or only volume serial number. If additional volumes exist, a plus sign (+) will follow the volume serial number.

Table A-4 Non-VSAM Data Set Information Report Field Descriptions (Part 2 of 4)

Field	Description
DEVICE TYPE	The model number or compatible model number of the device as defined in the Unit Control Block for the VOLSER specified
VOLUME COUNT	The number of volumes allocated to the data set. This value represents the volume count after close processing.
DCBE	Set to YES or NO to indicate whether a DCB Extension exists for this DCB
ACCESS METHOD	Access method used to process the data set. Possible values are: <ul style="list-style-type: none"> • BSAMBasic Sequential Access Method • BPAMBasic Partitioned Access Method • BDAMBasic Direct Access Method • QSAMQueued Sequential Access Method • EXCPExecute Channel Program Access Method • OTHERAccess method other than those listed above
OPEN OPTION	Open mode used to process this DCB. Possible values are: <ul style="list-style-type: none"> • INPUTData set will be read • OUTPUTData set will be written • UPDATEDData set will be both read and updated in place • INOUTData set will first be read, then written • OUTINData set will first be written, then read • EXTENDData set will be appended. • INOUTXData set will first be read, then appended. • RDBACKData set will be read backward.

Table A-4 Non-VSAM Data Set Information Report Field Descriptions (Part 3 of 4)

Field	Description
MACRO OPTION	<p>Processing options used in accessing the data set. Input and output type operations will be separated from each other with a comma. Format is similar to that used when specifying the MACRF option within the DCB macro. Possible values for each access method are:</p> <p>EXCP:</p> <ul style="list-style-type: none"> • AAppendages exist • BAccurate block count option specified • CCommon interface exists • PPage fix appendage exists • XFoundation extension exists • 5Five-word device interface • 3Three-word device interface • 1One-word device interface <p>QSAM:</p> <ul style="list-style-type: none"> • GGET support • PPUT support • CCNTRL support • DData mode • LLocate mode • MMove mode <p>BSAM:</p> <ul style="list-style-type: none"> • RREAD support • WWRITE support • CCNTRL support • LLoad mode • PPOINT support <p>BDAM:</p> <ul style="list-style-type: none"> • RREAD support • WWRITE support • CCHECK support • AAdd mode • IBlock ID search • KKey search • SDynamic buffering • XExclusive control required <p>BPAM:</p> <ul style="list-style-type: none"> • RREAD support • WWRITE support
DCBMACRF	A hexadecimal representation of the DCBMACR1 and DCBMACR2 fields of the DCB
BUFNO	The DCB BUFNO value present in the DCB
NCP	The DCB NCP value present in the DCB
DCBOPTCD	A hexadecimal representation of the DCBOPTCD field of the DCB
ERROR OPTION	The DCB error option chosen. Possible values are ACCEPT, SKIP, and ABEND corresponding to the equivalent DCB EROPT values of ACC, SKP, and ABE. If no EROPT value is specified in a qualifying DCB, NONE is displayed.
EODAD	Set to YES or NO to indicate the presence of a DCB EODAD routine

Table A-4 Non-VSAM Data Set Information Report Field Descriptions (Part 4 of 4)

Field	Description
SYNAD	Set to YES or NO to indicate the presence of a DCB SYNAD routine
DCB OPEN EXIT	Set to YES or NO to indicate the presence of a DCB OPEN EXIT
DCB ABEND EXIT	Set to YES or NO to indicate the presence of a DCB ABEND EXIT
JFCB EXIT	Set to YES or NO to indicate the presence of a DCB JFCB EXIT
JFCBE EXIT	Set to YES or NO to indicate the presence of a DCB JFCBE EXIT
EOV EXIT	Set to YES or NO to indicate the presence of a DCB EOV EXIT

VSAM Performance Report

The VSAM Performance Report is actually a group of reports, each which you can request or exclude. All requested reports are produced on one printable output. You can customize the information generated by the VSAM Performance Report by using the data policy options or DD override statements.

There are four types of information that can be included in this report:

- data policy options
- performance statistics
- data set information
- catalog information

Figure A-2 on page A-12 shows an example of the VSAM Performance Report. Subsequent pages provide descriptions of the data fields included in each section of the report.

Figure A-2 VSAM Performance Report

```

***** DATA OPTIMIZER PERFORMANCE REPORT
*****
* JOBNAME      : JJHT1          STEPNAME : STATTEST   PROCSTEPNAME:
*
* JOBCLASS    : A              USERID   : JJH2        GROUP       :
CC105413 *
* ADSPC TYPE  : JOB           PGMNAME  : STATTEST   OPEN TIME   : 02:16
*
* DDNAME      : KSDS1DD       DSTYPE   : VSAM        SUBSYSTEM  : PR
*
* DATACLASS  : **NONE**      STORCLASS: **NONE**   MGMTCLASS   :
**NONE** *
* DSNAME      :
BMCBSL.QA.COMPNONE.KSDS.STRIPED.PATH *
* BASE DSN    :
BMCBSL.QA.COMPNONE.KSDS.STRIPED *
*
* INTERCEPTED BY DATAPOLICY(REGDATA) POLICYLINE(00012)
*
*
* OPTIMIZATION MODE..... ADVANCED   OPTIMIZATION
PERCENT.....100% *
* HONOR RUB VALUE..... NO          OPTIMIZE STC ACCESS.....
NO *
* RESOURCE USAGE BIAS..... 6       OPTIMIZE TSO ACCESS.....
NO *
* DYNAMIC REGION ADJUSTMENT..... ..YES  OVERRIDE USER VALUES.....
NO *
* PAGE FAULT RATE THRESHOLD..... 0   BELOW 16M RESERVE.....
100K *
* CPU UTILIZATION THRESHOLD..... 0%  A
12M *
* LSR OPTION.....YES               MAXIMUM BELOW 16M .....
0 *
* RMODE31 CONTROL BLOCKS..... NO    MAXIMUM ABOVE 16M .....
0 *
* DEFER WRITES..... NO            SHAREOPTIONS(3 3) SUPPORT.....
NO *
* MULTIPLE STRINGS..... NO        SEQUENTIAL INSERT STRATEGY....
NO *
* AIX IN LSR..... NO              PAGE FIX.....
NO *
* NSR INDEX BUFFERS..... 0        NSR DATA BUFFERS.....
0 *
* NSR INDEX BUFFER STG..... 0     NSR DATA BUFFER STG.....
0 *
* NSR MAX INDEX BUFFERS..... 0    NSR MAX DATA BUFFERS.....
0 *

```

VSAM Performance Report Header

VSAM Options Report

VSAM Statistics Report

```

* NSR MAX INDEX BFR STG..... 0      NSR MAX DATA BFR STG.....
0 *
* LSR INDEX BUFFERS..... 0      LSR DATA BUFFERS.....
0 *
* LSR INDEX BUFFER STG..... 0      LSR DATA BUFFER STG.....
0 *
* LSR MAX INDEX BUFFERS..... 0      LSR MAX DATA BUFFERS.....
0 *
* LSR MAX INDEX BFR STG..... 128M *
* PCT INDEX TO BUFFER.....0%      SMF RECORD
TYPE.....211 *
* OVERRIDE USER LSR POOL ID.....NO  LSR POOL
ID..... 0 *
* USE HIPERSPACE LSR BUFFERS.....NO  PERCENT HIPERSPACE
BUFFERS.....50% *
* HIPERSPACE BUFFER STG.....0      LSR MAX HIPER
STG.....240K *
* LSR MAX SYS HIPER STG.....0      PERCENT MAX SYS HIPER
STG.....0% *
*
*
* RND NON-UPD READS.....          EADS.....
0 *
* SEQ NON-UPD READS..... 7,200     SEQ UPDATE READS.....
0 *
* RND INSERTS..... 0      RND UPDATES.....
0 *
* SEQ INSERTS..... 0      SEQ UPDATES.....
0 *
* SKIP SEQUENTIAL READS..... 0     POINTS.....
3,600 *
* ENDREQS..... 0      ERASES.....
0 *
*
*
* INDEX BUFFERS..... 16      DATA BUFFERS.....
1,536 *
* INDEX BUFFER READS..... 16      DATA BUFFER READS.....
1,851 *
* INDEX BUFFER HITS..... 9,049     DATA BUFFER HITS.....
3,599 *
* NON-USER WRITES..... 0      USER WRITES.....
0 *
* MAXIMUM STRINGS..... 8      HIPERSPACE BUFFERS.....
1,000 *
* HIPERSPACE READS..... 1,000     HIPERSPACE WRITES.....
0 *
* BATCH LSR INTERCEPTION
DISABLED-BLSR *
*
*

```

VSAM Data Set Information Report

VSAM Catalog Information Report

```

* INDEX PHYSICAL I/O'S..... 16      DATA PHYSICAL I/O'S.....
1,851 *
* ESTIMATED I/O SAVINGS..... 0.00%
*
*
* ACCESS..... KEY      RESET TO EMPTY.....
NO *
* INPUT..... YES      OUTPUT.....
YES *
* SEQUENTIAL..... YES  DIRECT.....
YES *
* SKIP SEQUENTIAL..... YES  RLS.....
YES *
* SEQUENTIAL INSERT STRATEGY..... NO  DEFER WRITES.....
NO *
* LSR..... YES      GSR.....
NO *
* LSR POOL ID..... 254  EXCLUSIVE CONTROL WAIT.....
YES *
* 31-BIT BUFFERS..... YES  31-BIT CONTROL BLOCKS.....
NO *
* USER BUFFERING..... NO  FIXED BUFFERS.....
NO *
* STRINGS..... 1      BASE STRINGS.....
0 *
* AIX..... NO      DSN SHARING.....
NO *
* OPENAD EXIT..... YES  EODAD EXIT.....
YES *
* JRNAD EXIT..... YES  UPAD EXIT.....
NO *
* LERAD EXIT..... YES  SYNAD EXIT.....
NO *
* ACBOFLGS..... 00      ACBMACR.....
00000000 *
* ACBINFLG..... 0000
ACBERFLG.....00 *
*
*
* ORGANIZATION..... KSDS  MAX LRECL.....
128 *
* INDEX VOLSER..... WORK0R+  DATA VOLSER.....
WORKAB+ *
* INDEX DEVICE TYPE..... 3390  DATA DEVICE TYPE.....
3390 *
* INDEX VOLUME COUNT..... 2  DATA VOLUME COUNT.....
2 *
* INDEX CI SIZE..... 8,192  DATA CI SIZE.....
1,536 *
* INDEX CI SPLITS..... 115  DATA CI SPLITS.....
0 *

```

```

* INDEX CA SPLITS..... 9      DATA CA SPLITS.....
0 *
* INDEX HI-ALLOC RBA... 99,999,999,999 DATA HI-ALLOC RBA.....
128,000 *
* INDEX HI-USED RBA..... 87,543,064 DATA HI-USED RBA.....
8,192 *
* INDEX EXTENTS..... 9      DATA EXTENTS.....
2 *
* INDEX READS..... 22,222 DATA READS.....
4 *
* INDEX INSERTS..... 0      DATA INSERTS.....
1 *
* INDEX UPDATES..... 46     DATA UPDATES.....
39 *
* INDEX DELETES..... 0      DATA DELETES.....
2 *
* INDEX RECORDS..... 5,000 DATA RECORDS.....
12 *
* INDEX LEVELS..... 2      SHAREOPTIONS.....
1,2 *
*****
****

```

VSAM Performance Report Header

The first part of Figure A-2 on page A-12 shows the header information that is common to all requested VSAM reports. Table A-1 on page A-4 describes the fields that are included in this header.

Table A-5 VSAM Performance Report Header Field Descriptions (Part 1 of 2)

Field	Description
JOBNAME	Name of the job in which the data set was opened
STEPNAME	Name of the step in which the data set was opened
PROCSTEPNAME	Name of the procedure step in which the data set was opened
JOBCLASS	JES job class of the job in which the data was opened
USERID	Security system Userid associated with the job in which the data set was opened
GROUP	Security system group associated with the job in which the data set was opened
ADSPC TYPE	Type of address space in which the data set was opened. Possible values are as follows: <ul style="list-style-type: none"> • JOB (a batch job) • TSU (TSO user) • STC (started task type of job)
PGMNAME	Name of the job step program which opened the data set
OPEN TIME	Time the data set was opened in hour and minute (24-hour clock format)

Table A-5 VSAM Performance Report Header Field Descriptions (Part 2 of 2)

Field	Description
DDNAME	The DD name of the data set which was opened
DSTYPE	The value VSAM denoting a non-VSAM data set
SUBSYSTEM	Subsystem ID of the MAINVIEW Batch Optimizer subsystem which processed the data set
DATACLASS	SMS Data Class of the data set which was opened. If the data set is not under SMS control, the value "***NONE**" will appear.
STORCLASS	SMS Storage Class of the data set which was opened. If the data set is not under SMS control, the value "***NONE**" will appear.
MGMTCLASS	SMS Management Class of the data set which was opened. If the data set is not under SMS control, the value "***NONE**" will appear.
DSNAME	The name of the data set which was opened
BASE DSN	Indicates the name of the base cluster, if a path is being processed.
INTERCEPTED BY (interception reason)	Indicates the reason the data set was intercepted by Data Optimizer. The possible values are: <ul style="list-style-type: none"> • INTERCEPTED BY DATA POLICY(XXXXXXXX) POLICYLINE(NNNNN) where XXXXXXXX is the Data Policy member name active when the data set was opened and NNNNN is the line number within the Data Policy which caused interception • INTERCEPTED BY DD STATEMENT OVERRIDE indicating that a JCL DD statement caused interception • INTERCEPTED BY USER EXIT indicating that the Data Optimizer user exit requested interception

VSAM Options Report

The Options Report will display either data policy options or override options, or both in effect at OPEN time. This section of the report appears only if the Options Report option is specified. The second group of data items in Figure A-2 on page A-12 reflect the VSAM Options Report information. Table A-6 shows fields that are included in the VSAM Options Report.

Note: The following values are user-defined, but are not necessarily the values actually used.

Table A-6 VSAM Options Report Field Descriptions (Part 1 of 4)

Field	Description
OPTIMIZATION MODE	Defines the optimization technique. Possible values are ADVANCED, BASIC, REPORTONLY.
OPTIMIZATION PERCENT	The percentage value to apply to Data Optimizer's calculated buffer amounts. Used to scale Data Optimizer's performance recommendations higher or lower.

Table A-6 VSAM Options Report Field Descriptions (Part 2 of 4)

Field	Description
HONOR RUB VALUE	Set to YES or NO to indicate whether the Resource Usage Bias value should be used. A setting of YES indicates that Data Optimizer will run in compatibility mode with earlier versions of the product.
OPTIMIZE STC ACCESS	Set to YES or NO to indicate whether data sets opened within started tasks will be optimized
RESOURCE USAGE BIAS	Defines the Resource Usage Bias value in effect. Only has meaning if HONOR RUB VALUE is set to YES.
OPTIMIZE TSO ACCESS	Set to YES or NO to indicate whether data sets opened within Time Sharing Option User address spaces will be optimized
DYNAMIC REGION ADJUSTMENT	Set to YES or NO to indicate whether Dynamic Region Adjustment was active
OVERRIDE USER VALUES	Set to YES or NO to indicate whether to replace user-coded DCB values if they are higher than those chosen by Data Optimizer
PAGE FAULT RATE THRESHOLD	The page fault rate used to determine whether the system is paging constrained at the time of the open
BELOW 16M RESERVE	The amount of storage below the 16-megabyte line Data Optimizer is to leave unused and available for user processing. The suffixes K and M can be used to denote kilobytes and megabytes.
CPU UTILIZATION THRESHOLD	The CPU utilization percentage value used to determine whether the system is CPU constrained at the time of the open
ABOVE 16M RESERVE	The amount of storage above the 16-megabyte line Data Optimizer is to leave unused and available for user processing. The suffixes K and M can be used to denote kilobytes and megabytes.
LSR OPTION	One of the following values: YES - Perform normal Data Optimizer LSR processing NO - Do not perform Data Optimizer LSR processing FORCE - Place a VSAM data set into an LSR pool even if the ACB DIR bit is off PREFER - Convert NSR processing to LSR if possible; do not use History information REQUIRE - Must convert NSR processing to LSR if possible; do not use History information; if conversion is not possible, bypass interception
MAXIMUM BELOW 16M	The maximum amount of storage below the 16-megabyte line Data Optimizer is to allocate for performance processing over and above that which would have been used without Data Optimizer involvement. The suffixes K and M can be used to denote kilobytes and megabytes.
RMODE31 CONTROL BLOCKS	Set to YES or NO to indicate whether or not to place VSAM control blocks above or below the 16-megabyte line
MAXIMUM ABOVE 16M	The maximum amount of above the 16-megabyte line Data Optimizer is to allocate for performance processing over and above that which would have been used without Data Optimizer involvement. The suffixes K and M can be used to denote kilobytes and megabytes.
DEFER WRITES	Set to YES or NO to indicate whether or not to defer buffer writes until necessary.
SHAREOPTIONS(3 3) SUPPORT	Set to YES or NO to indicate whether or not to intercept VSAM clusters with SHAREOPTIONS which are set to 3 for both intra-system and inter-system sharing
MULTIPLE STRINGS	Indicates the value of the Multi-string option.

Table A-6 VSAM Options Report Field Descriptions (Part 3 of 4)

Field	Description
SEQUENTIAL INSERT STRATEGY	Indicates the value of the Sequential Insert Strategy option.
AIX IN LSR	Indicates the value of the AIX LSR option.
PAGE FIX	Indicates the value of the Page Fix option.
NSR INDEX BUFFERS	The number of index buffers to allocate for NSR processing
NSR DATA BUFFERS	The number of data buffers to allocate for NSR processing
NSR INDEX BUFFER STG	The amount of storage to allocate to index buffers for NSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
NSR DATA BUFFER STG	The amount of storage to allocate to data buffers for NSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
NSR MAX INDEX BUFFERS	The maximum number of index buffers to allocate for NSR processing
NSR MAX DATA BUFFERS	The maximum number of data buffers to allocate for NSR processing
NSR MAX INDEX BFR STG	The maximum amount of storage to allocate to index buffers for NSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
NSR MAX DATA BFR STG	The maximum amount of storage to allocate to data buffers for NSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
LSR INDEX BUFFERS	The number of index buffers to allocate for LSR processing
LSR DATA BUFFERS	The number of data buffers to allocate for LSR processing
LSR INDEX BUFFER STG	The amount of storage to allocate to index buffers for LSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
LSR DATA BUFFER STG	The amount of storage to allocate to data buffers for LSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
LSR MAX INDEX BUFFERS	The maximum number of index buffers to allocate for LSR processing
LSR MAX DATA BUFFERS	The maximum number of data buffers to allocate for LSR processing
LSR MAX INDEX BFR STG	The maximum amount of storage to allocate to index buffers for LSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
LSR MAX DATA BFR STG	The maximum amount of storage to allocate to data buffers for LSR processing. The suffixes K and M can be used to denote kilobytes and megabytes.
PCT INDEX TO BUFFER	The percentage of the total number of index control intervals for which to provide buffers
SMF RECORD TYPE	Defines the SMF record type to be used when producing Data Optimizer SMF records. If set to a value of zero, no SMF records will be produced.
OVERRIDE USER LSR POOL ID	Indicates the value of the LSR Override option.
LSR POOL ID	Indicates the value of the LSR Pool ID option.
USER LSR HIPERSPACE BUFFERS	Indicates the value of the Hiperspace option.

Table A-6 VSAM Options Report Field Descriptions (Part 4 of 4)

Field	Description
PERCENT HIPERSPACE BUFFERS	Defines the number of hiperspace buffers to create as the percentage of used data control intervals.
HIPERSPACE BUFFER STG	Defines the number of hiperspace buffers as that amount that will fit into the storage amount specified for the keyword.
LSR MAX HIPER STG	Defines the maximum amount of hiperspace storage that a single cluster is permitted to use.
LSR MAX SYS HIPER STG	Defines the maximum amount of hiperspace storage that all open clusters which are being intercepted by Data Optimizer are permitted to use.
PERCENT MAX SYS HIPER STG	Defines the maximum amount of available online expands storage, as a percentage.

VSAM Statistics Report

The Statistics Report displays I/O statistics information for the data set processed. The third, fourth and fifth group of data items in Figure A-2 on page A-12 reflect the VSAM Statistics Report information. Table A-7 shows fields that are included in the VSAM Statistics Report.

Table A-7 VSAM Statistics Report Field Descriptions (Part 1 of 2)

Field	Description
RND NON-UPD READS	Defines the number of random non-update reads issued against the data set.
RND UPDATE READS	Defines the number of random update reads issued against the data set.
SEQ NON-UPD READS	Defines the number of sequential non-update reads issued against the data set.
SEQ UPDATE READS	Defines the number of sequential update reads issued against the data set.
RND INSERTS	Defines the number of random inserts issued against the data set.
RND UPDATES	Defines the number of random updates issued against the data set.
SEQ INSERTS	Defines the number of sequential inserts issued against the data set.
SEQ UPDATES	Defines the number of sequential updates issued against the data set.
SKIP SEQUENTIAL READS	Defines the number of skip sequential read requests issued against the data set.
POINTS	Defines the number of positioning requests issued against the data set.
ENDREQS	Defines the number of terminate requests issued against the data set.
ERASES	Defines the number of delete requests issued against the data set.
INDEX BUFFERS	Defines the number of NSR/LSR index buffers allocated to the data set.
DATA BUFFERS	Defines the number of NSR/LSR data buffers allocated to the data set.
INDEX BUFFER READS	Defines the number of times a control interval was physically read into an index buffer in support of a logical VSAM request.

Table A-7 VSAM Statistics Report Field Descriptions (Part 2 of 2)

Field	Description
DATA BUFFER READS	Defines the number of times a control interval was physically read into an LSR data buffer in support of a logical VSAM request.
INDEX BUFFER HITS	Defines the number of times a logical VSAM request was satisfied directly from an LSR index buffer without having to do physical I/O.
DATA BUFFER HITS	Defines the number of times a logical VSAM request was satisfied directly from an LSR data buffer without having to do physical I/O.
NON-USER WRITES	Defines the number of times VSAM flushed an LSR buffer to disk.
USER WRITES	Defines the number of times a user directly initiated the write of an LSR buffer to disk.
MAXIMUM STRINGS	Defines the maximum number of static strings allocated to the LSR buffer pool.
HIPERSPACE BUFFERS	Defines the number of buffers allocated in hiperspace for this cluster, if this cluster is a base cluster. Defines the number of buffers allocated in hiperspace for the base cluster, if the cluster is being processed is a path.
HIPERSPACE READS	Defines the number of reads from a hiperspace buffer into a virtual storage buffer including MOVE PAGE operations. This data only appears if hiperspace buffers have been allocated.
HIPERSPACE WRITES	Defines the number of writes into a hiperspace buffer from a virtual storage buffer including MOVE PAGE operations. This data only appears if hiperspace buffers have been allocated.
BATCH LSR INTERCEPTION DISABLED	Indicates that IBM's Batch Local Shared Resource product was found to be enabled for this data set and was disabled by Data Optimizer as requested by the DISABLEBATCHLSR option. The four-digit subsystem ID of the Batch LSR subsystem is displayed with this message.
INDEX PHYSICAL I/O's	Defines the number of physical I/Os issued against the index component of the cluster.
DATA PHYSICAL I/O's	Defines the number of physical I/Os issued against the data component of the cluster.
ESTIMATED I/O SAVINGS	Defines the estimated percentage of physical I/O's saved by Data Optimizer as compared to default VSAM buffer settings.

VSAM Data Set Information Report

The Data Set Information Report provides detailed attribute information about the data set being processed. The values reflect those in affect after OPEN processing has completed and before the first I/O.

The following tables contain information about logical and physical I/O counts after CLOSE of the data set.

The sixth group of data items in Figure A-2 on page A-12 reflect the VSAM Data Set Information Report information. Table A-8 shows fields that are included in the VSAM Data Set Information Report.

Table A-8 VSAM Data Set Information Report Field Descriptions (Part 1 of 2)

Field	Description
ACCESS	Indicates the type of access used to process the data set: KEY - keyed access ADR - RBA access CNV - control interval access
RESET TO EMPTY	Set to YES or NO to indicate whether or not to set the high-used RBA to zero at OPEN time
INPUT	Set to YES or NO indicating that the cluster will be processed for reading
OUTPUT	Set to YES or NO indicating that the cluster will be processed for writing
SEQUENTIAL	Set to YES or NO indicating sequential access to the cluster
RANDOM	Set to YES or NO indicating random access to the cluster
SKIP SEQUENTIAL	Set to YES or NO indicating skip sequential access to the cluster
RLS	Set to YES or NO indicating whether record level sharing is to be used for this cluster
SEQUENTIAL INSERT STRATEGY	Set to YES or NO indicating whether to use sequential insert strategy
DEFER WRITES	Set to YES or NO indicating whether direct PUTS will not or will cause an immediate write to disk of the affected buffer
LSR	Set to YES or NO indicating whether LSR processing will be used for this cluster
GSR	Set to YES or NO indicating whether GSR processing will be used for this cluster
LSR POOL ID	The LSR buffer pool number
EXCLUSIVE CONTROL WAIT	Set to YES or NO indicating whether an LSR request which encounters an exclusive control conflict should be deferred or if a return code should be set and control returned to the requestor.
31-BIT BUFFERS	Set to YES or NO indicating whether to place buffers above or below the 16-megabyte line
31-BIT CONTROL BLOCKS	Set to YES or NO indicating whether to place control blocks above or below the 16-megabyte line
USER BUFFERING	Set to YES or NO indicating whether the user or VSAM will control buffer processing
FIXED BUFFERS	Set to YES or NO indicating whether VSAM control blocks and buffers should be fixed in real storage
STRINGS	The number of strings allocated to this ACB
BASE STRINGS	The number of strings to allocate to processing of the base cluster associated with this cluster if this cluster is a path defined over an alternate index
AIX	Set to YES or NO indicating whether the alternate index or the base cluster of a path is to be processed
DSN NAME SHARING	Set to YES or NO indicating whether subtask control block sharing is to be based on like data set names
EXIT LIST	Set to YES or NO indicating whether an exit list exists for this ACB
EODAD	Set to YES or NO indicating the presence of an EODAD exit for this ACB
JRNAD	Set to YES or NO indicating the presence of a JRNAD exit for this ACB

Table A-8 VSAM Data Set Information Report Field Descriptions (Part 2 of 2)

Field	Description
UPDAD	Set to YES or NO indicating the presence of a UPDAD exit for this ACB
LERAD	Set to YES or NO indicating the presence of a LERAD exit for this ACB
SYNAD	Set to YES or NO indicating the presence of a SYNAD exit for this ACB
ACBOFLGS	Displays a hexadecimal representation of the ACBOFLGS field of the ACB
ACBMACR	Displays a hexadecimal representation of the ACBMACR1-4 fields of the ACB
ACBINFLG	Displays a hexadecimal representation of the ACBINFLG field of the ACB
ACBERFLG	Displays a hexadecimal representation of the ACBERFLG field of the ACB

VSAM Catalog Information Report

The VSAM Catalog Information Report is produced when the Catalog Information Report data set is being processed. The final group of data items in Figure A-2 on page A-12 reflect the Catalog Information Report information. Table A-9 shows the fields that are generated from the catalog or control blocks for the cluster.

Table A-9 VSAM Catalog Information Report Field Descriptions (Part 1 of 2)

Field	Description
ORGANIZATION	The type of the VSAM cluster, either KSDS, RRDS, ESDS or LDS
MAX LRECL	The maximum length record possible for this cluster
INDEX VOLSER	The initial or only volume serial number of the index component of the cluster. If additional volumes exist, a plus sign (+) will follow the volume serial number.
INDEX DEVICE TYPE	The model number or compatible model number of the device as identified by the INDEX VOLSER field
INDEX VOLUME COUNT	The number of volumes allocated to the index component of the cluster. This value represents the volume count after close processing.
INDEX CI SIZE	The control interval size of the index component of the cluster
INDEX CI SPLITS	The number of control interval splits for the index component of the cluster
INDEX CA SPLITS	The number of control area splits for the index component of the cluster
INDEX HI-ALLOC RBA	The highest allocated relative byte address for the index component of the cluster
INDEX HI-USED RBA	The highest used relative byte address for the index component of the cluster
INDEX EXTENTS	The total number of extents which exist for the index component of the cluster
INDEX READS	The number of reads issued against the index component of the cluster
INDEX INSERTS	The number of inserts issued against the index component of the cluster
INDEX UPDATES	The number of updates issued against the index component of the cluster
INDEX DELETES	The number of deletes issued against the index component of the cluster

Table A-9 VSAM Catalog Information Report Field Descriptions (Part 2 of 2)

Field	Description
INDEX RECORDS	The number of records within the index component of the cluster.
INDEX LEVELS	The number of index levels within the cluster
DATA VOLSER	The initial or only volume serial number of the data component of the cluster. If additional volumes exist, a plus sign (+) will follow the volume serial number.
DATA DEVICE TYPE	The model number or compatible model number of the device as identified by the DATA VOLSER field
DATA VOLUME COUNT	The number of volumes allocated to the data component of the cluster. This value represents the volume count after close processing.
DATA CI SIZE	The control interval size of the data component of the cluster
DATA CI SPLITS	The number of control interval splits for the data component of the cluster
DATA CA SPLITS	The number of control area splits for the data component of the cluster.
DATA HI-ALLOC RBA	The highest allocated relative byte address for the data component of the cluster
DATA HI-USED RBA	The highest used relative byte address for the data component of the cluster
DATA EXTENTS	The total number of extents which exist for the data component of the cluster
DATA READS	The number of reads issued against the data component of the cluster
DATA INSERTS	The number of inserts issued against the data component of the cluster
DATA UPDATES	The number of updates issued against the data component of the cluster
DATA DELETES	The number of deletes issued against the data component of the cluster
DATA RECORDS	The number of records within the data component of the cluster
SHAREOPTIONS	The share options defined for the cluster. The first value is the cross-region share option while the second value is the cross-system share option.

Appendix B Navigating the User Interface

This appendix describes the MAINVIEW[®] Batch Optimizer user interface. |
This appendix discusses the following topics:

Overview	B-2
ISPF Conventions	B-2
ISPF Profile	B-2
Panel Layout	B-3
Accessing the User Interface	B-8
Modifying User Interface Processing	B-10
Dialog Options and Settings	B-10
Modifying ISPF Settings	B-11
Modifying Key Lists	B-12
Using the User Interface	B-13
Navigation	B-14
Pop-up Panels	B-15
Accessing the Action Bar	B-16
Point-and-Shoot	B-17
Creating and Editing Control Data Set Members	B-18
Object List Displays	B-20
Online Help	B-24
Panel Help	B-24
Field Help	B-24
Message Help	B-24
Reference Phrase Help	B-25
Other User Interface Tasks	B-25
Product Version and Maintenance Level Information	B-25
Environment Version and Release Information	B-26
BMC Software Product Authorization User Interface	B-27

Overview

This chapter describes the basic operation and facilities of the MAINVIEW® Batch Optimizer user interface (dialog). The interface is an ISPF-based interactive dialog that provides access to the MAINVIEW Batch Optimizer control data set. This is the data set that is allocated to the BSLPLEX data definition (DD) statement of the MAINVIEW Batch Optimizer subsystem started task. Although you can use the ISPF Edit function to view or modify the control data set members, the MAINVIEW Batch Optimizer user interface affords a protected environment for this process. The user interface verifies that any control data set members you create or modify are syntactically correct and do not provide conflicting instructions to the optimization components.

Warning! Comment lines added to the control data set member through ISPF Edit may be lost when editing these members through the customer interface.

ISPF Conventions

This section provides general information about the user interface and the panels that are provided with the dialog. All panels, pull-down menus, and pop-up windows in this interface conform to IBM CUA standards.

ISPF Profile

When you access the ISPF interface for the first time, the MAINVIEW Batch Optimizer dialog creates member BSSBPROF in your ISPF profile data set. This member contains information about your dialog sessions and some of the input field values from your previous session.

For example, the dialog maintains the name of the control data set that was most recently accessed. If you modified MAINVIEW Batch Optimizer dialog keylists, BSSBPROF also contains a copy of member BSSBKEYS from ISPTLIB with your modifications. Each time you access the MAINVIEW Batch Optimizer user interface, it restores values in the data entry and choice entry fields from information saved in the BSSBPROF member. If you delete member BSSBPROF from your profile data set, the MAINVIEW Batch Optimizer user interface always creates a new one for you. The new member always contains MAINVIEW Batch Optimizer default values.

Panel Layout

Figure B-1 shows a MAINVIEW Batch Optimizer Objects List panel.

Figure B-1 MAINVIEW Batch Optimizer Objects List Panel

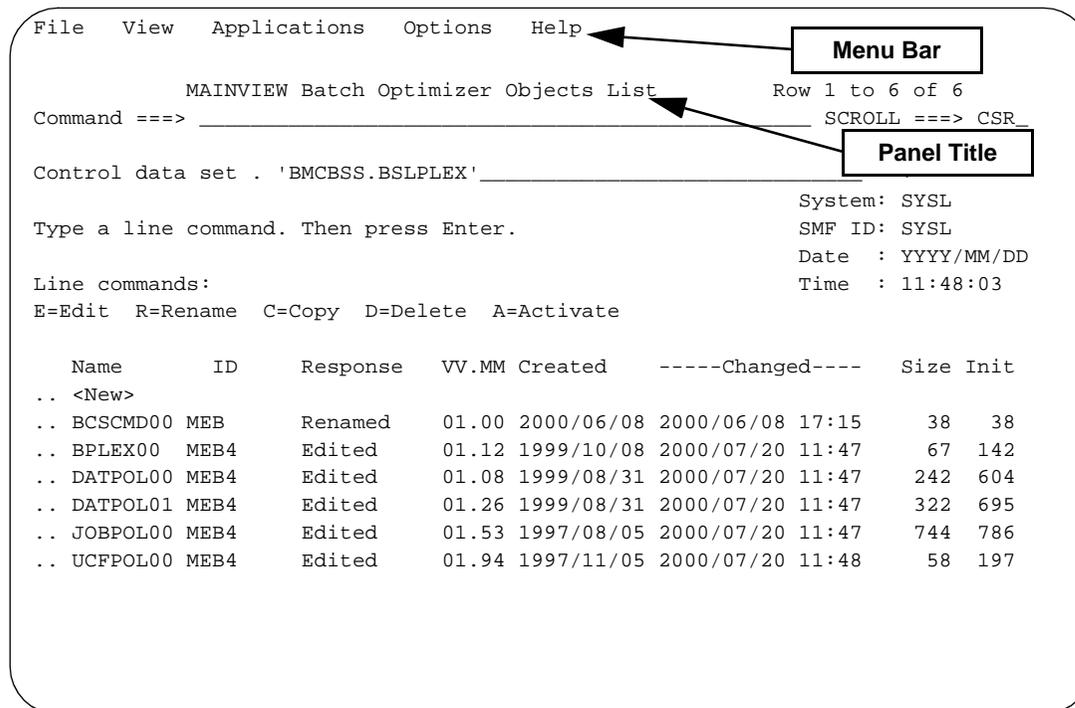


Figure B-2 shows a MAINVIEW Batch Optimizer data entry panel.

Figure B-2 Data Policy Definition Panel

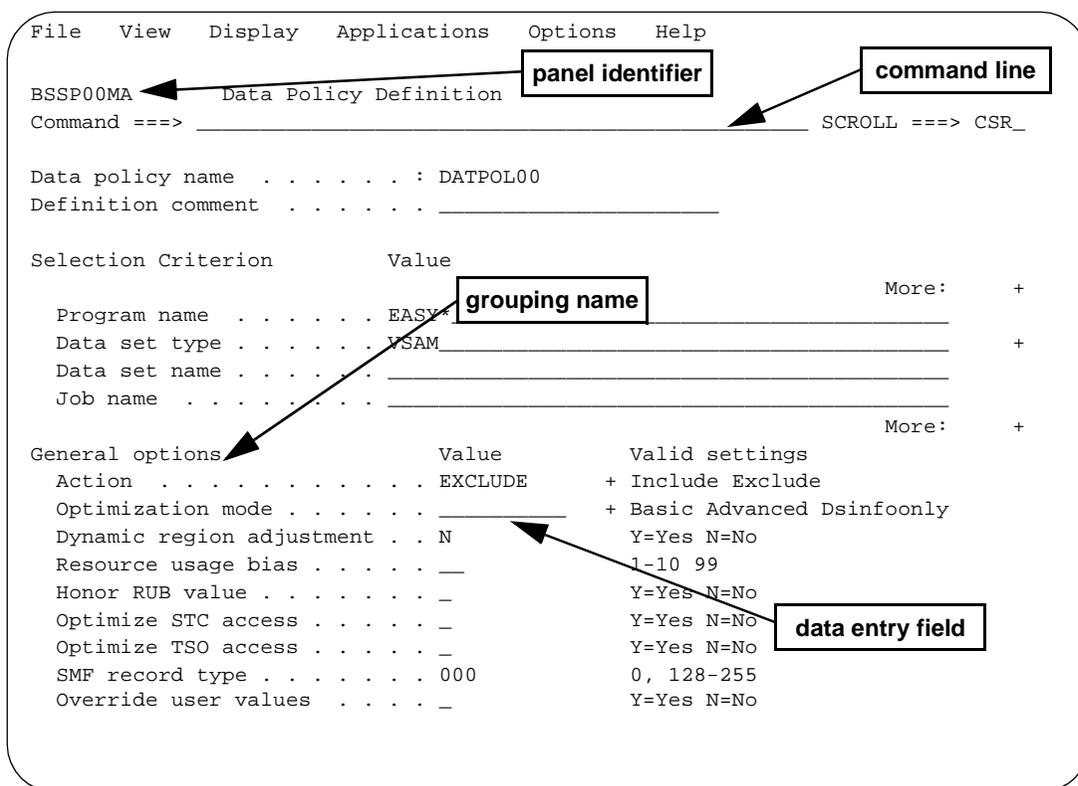


Figure B-3 shows a MAINVIEW Batch Optimizer pop-up panel.

Figure B-3 Data Policy Definition Selection Criterion Pop-up Panel

```

----- Data Policy Definition Selection Criterion -----
Command ==> _____ SCROLL ==> CSR_

Data policy name . . . . . : DATPOL00
Definition comment . . . . . _____

Selection Criterion      Value
                                                                    More:  +
Data set name . . . . . _____
Data definition (DD) name _____
Program name . . . . . IEBDG____
Step name . . . . . _____
Procedure step name . . . _____
Data set type . . . . . _____ +
SMS Data class . . . . . _____
SMS Storage class . . . . . _____
SMS Management class . . . _____
Job name . . . . . _____
User ID . . . . . MEB*____
Group . . . . . _____

```

Action Bar

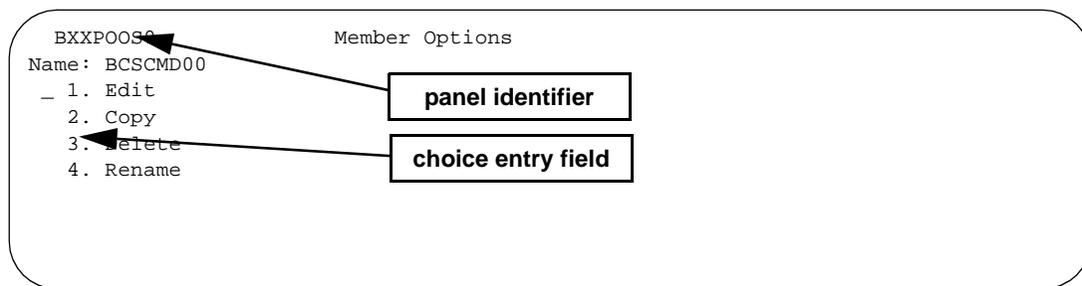
Each panel (except for pop-up panels) has an action bar at the top of the panel. The selections on the action bar enable you to access pull-down menus for navigation or processing purposes. Action bar selections differ according to panel.

Panel Title

Each panel has a unique name that is displayed at the top of the panel.

Panel Identifier

Each panel has a unique panel identifier. The ISPF interface does not display the panel identifier when you first access the interface. The same panel samples in this manual display the panel identifiers for reference purposes. To display the identifier in the upper left corner of the panel, type **PANELID** on the command line and press **Enter**. To remove the panel identifier, type **PANELID** again and press **Enter**. Figure B-4 shows a panel that displays the panel name.

Figure B-4 Member Options Choice Entry Panel

Command Line

You can issue TSO and ISPF commands from any panel with a command line. To issue a command, type the command on the command line, and press **Enter**. For example, to obtain information about a data set, type **TSO LISTD 'ABC.DATA.SET'** on the command line, and press **Enter**. You also can use the command line to issue a command that is equivalent to an F key. The command equivalent to **F1** is **HELP**.

To access help, press **F1** or type **HELP** on the command line and press **Enter**. See “F keys” on page B-7 for a list of all F keys and their command equivalents. You can position the command line at the top or bottom of the panel. For instructions on repositioning the command line, see “Dialog Options and Settings” on page B-10.

Choice Entry Field

Numbered choice entry panels have a single or double character field to the left of the first option in the list. To select an option, type your selection in the choice entry field and press **Enter**. If the point-and-shoot feature is enabled, you can position the cursor to the line corresponding to your choice and press **Enter**.

Grouping Name

The panel grouping name is displayed for organizational and usability purposes.

Data Entry Field

When you can enter data for an option, a data entry field is displayed to the right of the option. This field may be blank or it may contain a default value. To enter or change a value, type a value and clear remaining characters.

Scrolling Indicator

This area of a panel indicates that more information exists outside the visible panel area and shows you the direction to scroll to see that information. If **More** is followed by a minus symbol, press **F7** (UP). If you see **More** followed by a plus symbol, press **F8** (DOWN).

Prompt Indicator

The dialog uses a plus (+) symbol to designate input fields that support the Prompt command. When the dialog appends this indicator to an input field, you can use the prompt facility to obtain a list of valid options for the field. Obtain the list by positioning your cursor in the field and pressing **F4** (Prompt) or by positioning your cursor beneath the plus symbol and pressing **Enter** (if you have enabled the point-and-shoot facility).

F keys

The MAINVIEW Batch Optimizer ISPF dialog uses F keys for navigation and processing purposes. To display the F keys, type **FKA** on the command line, and press **Enter**. The F keys appear at the bottom of each panel. MAINVIEW Batch Optimizer consistently uses the following F key settings but may activate other F keys from time to time, for special functions. The dialog stores these settings in the keylist table. See “Modifying Key Lists” on page B-12 for information about changing the F key settings. You can press the associated F key or type the equivalent command on the command line.

Note: The labels shown with the F keys on panels are not always the equivalent commands.

F1 (HELP)	Display field-level, panel-level, message-level help, or reference phrase help.
F2 (SPLIT)	Split the screen at the cursor position.
F3 (EXIT)	Terminate the process and provide a Confirm Exit pop-up window to save or delete any changes.
F4 (PROMPT)	For input fields followed by a plus (+) symbol, display a list of valid options for the input field.
F7 (UP)	On a panel with a scroll bar, scroll backward in the area—toward the beginning of the entry. Use this key if you see More: - .
F8 (DOWN)	On a panel with a scroll bar, scroll forward in the area—toward the end of the entry. Use this key if you see More: + .

F9 (SWAP) Toggle between ISPF screens.

F12 (CANCEL) Discard any changes made on the panel, and return to the preceding panel.

Messages As you navigate through the MAINVIEW Batch Optimizer panels, the dialogs may display messages that provide information about an action that has occurred. The message could be an informational message to explain that a process has completed successfully or an error message to explain that you have entered invalid data. The dialog displays these messages in the form of short messages or pop-up messages. When you receive a short message, you can press **F1 (HELP)** to view a more descriptive message (a long message). Pop-up messages usually appear at the bottom of the panel, covering your F key selections. To see the F key lines, use the ISPF Window service to move the pop-up or press **F12 (CANCEL)** to remove the message.

Accessing the User Interface

You can access the MAINVIEW Batch Optimizer user interface only if you have installed MAINVIEW Batch Optimizer successfully and have performed the operation verification procedure (see the *MAINVIEW Batch Optimizer Release Notes*).

To access the MAINVIEW Batch Optimizer user interface, complete the following steps:

Step 1 *Required.* Use the REXX Exec or menu selection established during installation to access the MAINVIEW Batch Optimizer customer interface. You can find a sample REXX Exec designed to access the dialog in members BSSBISPI and BSSBISPF in data set BMC.BSS.INSTALL. These members should have been modified as part of installation. Figure B-5 shows the MAINVIEW Batch Optimizer Logon panel.

Figure B-5 MAINVIEW Batch Optimizer Logon panel

```

M M AAA III N N V V III EEEEE W W ®
MM MM A A I NN N V V I E W W
M M M AAAAA I N N N V V I EEEE W W W
M M A A I N NN V V I E WW WW
M M A A III N N V III EEEEE W W

BBBB AAA TTTT CCCC H H
B B A A T C H H
BBBB AAAAA T C HHHH
B B A A T C H H
BBBB A A T CCCC H H

OOO PPPP TTTT III M M III ZZZZ EEEEE RRRR
O O P P T I MM MM I Z E R R
O O PPPP T I M M M I Z EEEE RRRR
O O P T I M M I Z E R R
OOO P T III M M III ZZZZ EEEEE R R

Press Enter to continue.

Contains Licensed Materials - Property of BMC Software, Inc.
© COPYRIGHT BMC Software, Inc. 1987-2002.
All Rights Reserved.
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contracts with BMC Software, Inc.

```

The logo panel contains trademark and copyright information about BMC Software and MAINVIEW Batch Optimizer. The logo panel appears each time you access the dialog unless you disable this facility. See “Dialog Options and Settings” on page B-10 for instructions about disabling the logo panel.

Step 2 *Required.* Press **Enter** to continue. The dialog displays the MAINVIEW Batch Optimizer Objects List. You also can select **F3** (EXIT/END) or **F12** (CANCEL) to leave the ISPF interface.

Modifying User Interface Processing

You can modify certain aspects of the dialog's behavior by using dialog specific options, ISPF options, or program function key (PFK) changes via Keylist modifications. Some of the modifications available in the dialog include enabling/disabling the point-and-shoot facility and cursor positioning.

Dialog Options and Settings

You can access the dialog options by selecting the Dialog pull-down choice of the Options action bar item or by entering the command **DLGOPTNS** on the command line. Figure B-6 shows the Dialog Options panel.

Figure B-6 Dialog Options Panel

```

Command ==> _____

General Options
Command line placement . . ASIS__  (Asis or bottom)
Display logo panel . . . . Y        (Y=Yes or N=No)
Point-n-shoot support . . ON_     (On or Off)
Maintain data set list . . Y        (Y=Yes or N=No)

Policy Edit Options
Display global panel . . . . Y        (Y=Yes or N=No)
Display insert panel . . . . Y        (Y=Yes or N=No)

Debug Options
General . . . . . N                (Y=Yes or N=No)
Initialization . . . . . N          (Y=Yes or N=No)
Take a dump on abends. . . . . N     (Y=Yes or N=No)
    
```

General Options

General Options allows you to control dialog functions such as command line location or displaying the logo panel on entry. You can also choose to disable the point-and-shoot facility. You may want to disable the point-and-shoot facility if you are not using a terminal emulator with mouse support. Since ISPF treats point-and-shoot fields as “input” fields, enabling this options increases the amount of tabbing required to position the cursor.

If you specify a "Y" on the "Maintain data set list" option, the dialog will maintain a list of each control data set you reference. You can access this list by using the Prompt facility for the Control data set input field on the MAINVIEW Batch Optimizer Objects List panel. The dialog maintains this list in an ISPF table named BSSBDSNT that is stored in the data set allocated to your ISPTABL file. If the ISPTABL file is unavailable to your TSO session, the dialog issues a warning message and disables this facility.

Policy Edit Options

The Policy Edit Options allow you to control the amount of assistance the dialog will offer when creating new control data set members. When creating new members, there are logical steps to perform. By enabling the Display global panel, when you elect to create a new control data set member, the dialog will first display the "global specifications" panel for the member. You can modify or accept the displayed values and press **Enter** or enter **F12** (CANCEL) to proceed. By enabling the Display insert panel, when you elect to create a new control data set member, the dialog will first display the "statement insert" panel for the member. If you enable both options, the dialog first displays the "global specifications" panel followed by the "statement insert" panel. If you elect to disable both of these options, when you create a new control data set member, the dialog will place you in an empty table list for the member. This option is useful for those unfamiliar with the requirements of the various control data set members.

Note: These options have no effect on processing the BCSS Commands member.

Debug Options

Leave these options set to "N" unless instructed otherwise by BMC Software Product Support.

Modifying ISPF Settings

You can access the ISPF setting by selecting the ISPF pull-down choice of the Options action bar item or by entering the command **SETTINGS** on the command line. Figure B-7 shows the ISPF Settings panel.

Figure B-7 ISPF Settings Panel

```

Log/List  Function keys  Colors  Environ  Workstation  Identifier  Help

SETTINGS                                ISPF Settings
Command ==>

Options                                  Print Graphics
  Enter "/" to select option              Family printer type 2
    Command line at bottom                Device name . . . .
    Panel display CUA mode                Aspect ratio . . . 0
  / Long message in pop-up
  / Tab to action bar choices
  / Tab to point-and-shoot fields          General
  / Restore TEST/TRACE options            Input field pad . . N
  / Session Manager mode                  Command delimiter . ;
  / Jump from leader dots
  / Edit PRINTDS Command
  / Always show split line
  / Enable EURO sign

Terminal Characteristics
Screen format  1  1. Data    2. Std    3. Max    4. Part

Terminal Type  3      1. 3277    2. 3277A  3. 3278    4. 3278A
                5. 3290A    6. 3278T  7. 3278CF  8. 3277KN
                9. 3278KN   10. 3278AR 11. 3278CY 12. 3278HN
                13. 3278HO  14. 3278IS 15. 3278L2 16. BE163
                17. BE190   18. 3278TH
    
```

This ISPF panel lets you alter your ISPF session defaults. The MAINVIEW Batch Optimizer user interface offers direct access to ISPF settings so that you can enable or disable the “Tab to action bar choices” or “Tab to point-and-shoot fields” features. You can use the "Command line at bottom" option or the "Long message in pop-up" option to change the command line or message placement, respectively.

Note: Changes made on the ISPF Settings dialog apply to your ISPF session and not only your MAINVIEW Batch Optimizer dialog session.

Modifying Key Lists

You can access the key lists by selecting the Keylists pull-down choice of the Options action bar item or by entering the command **KEYLIST** on the command line. Figure B-8 shows the Keylist Utility panel.

Figure B-8 Keylist Utility Panel

```

Keylist Utility
  File  View

SETTINGS          Keylist Utility for BSSB          Row 1 to 14 of 24
Command ==>      Scroll ==> PAGE

Actions:  N=New  E=Edit  V=View  D=Delete  /=None

  Keylist  Type
  BSSBKBK  SHARED
  BSSBKCU  SHARED
  BSSBKDF  SHARED
  BSSBKGU  SHARED
  BSSBKHC  SHARED
  BSSBKHH  SHARED
  BSSBKHI  SHARED
  BSSBKHK  SHARED
  BSSBKHT  SHARED
  BSSBKHX  SHARED
  BSSBKMA  SHARED
  BSSBKMB  SHARED
  BSSBKMC  SHARED
  BSSBKMD  SHARED

```

Warning! BMC Software recommends that you *do not* modify the keylists. When you modify the keylists, MAINVIEW Batch Optimizer copies the content of member BSSBKEYS from your ISPTLIB library and places this information, with your changes, into member BSSBPROF of your ISPF profile data set. From the time you modify the keylists until you delete member BSSBPROF, you will access your personal copy of the keylists. If MAINVIEW Batch Optimizer changes any keylists as part of ongoing maintenance, you will not have access to these changes, which could result in unpredictable results while using the MAINVIEW Batch Optimizer user interface.

Using the User Interface

This section provides general instructions on how to use the user interface. It discusses navigation of the MAINVIEW Batch Optimizer user interface using common ISPF commands and facilities, mechanisms for handling the display of pop-up messages and data entry panels, and different shortcuts provided by the dialog as well as ISPF. See the *IBM ISPF Users Guide* for a detailed discussion of ISPF commands and facilities.

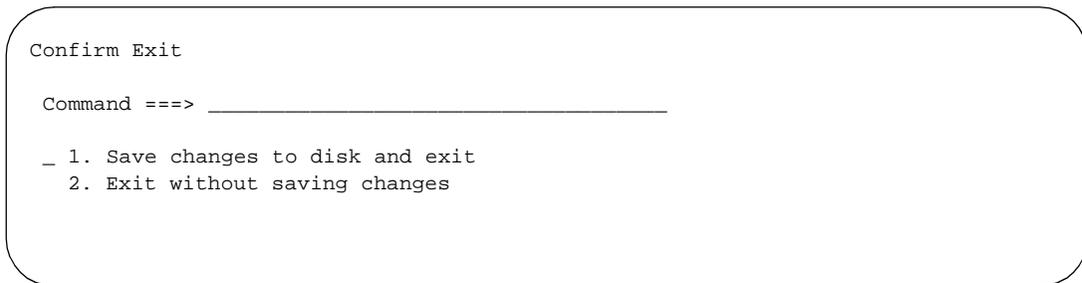
Navigation

Navigate user interface panels with **F3** (EXIT or END), **F12** (CANCEL), or **Enter**.

Exit or End

EXIT or END returns you to the next higher level application or exits the dialog. For example, if you select a Data Policy for editing and issue the EXIT command, the dialog returns to the MAINVIEW Batch Optimizer Objects List panel. If you changed the policy member, the dialog prompts you for confirmation by displaying the pop-up panel before exiting. This pop-up panel is shown in Figure B-9.

Figure B-9 **Exit Confirmation Panel**



Cancel

CANCEL returns you to the previously displayed panel. If the prior panel was part of a higher level application, then you will exit the current application. If you enter CANCEL after making changes to a data entry panel, the dialog prompts you for confirmation. Figure B-10 shows a Cancel Confirmation panel.

Figure B-10 Cancel Confirmation Panel

```
Confirm Cancel
Command ==> _____

_ 1. Retain policy changes and cancel
   2. Cancel without retaining changes

Attention: Changes will only be retained internally.
           Use Exit or Save to save a policy to disk.
```

Enter

Use the **Enter** key to move forward. Either commit data entry panel changes, invoke a command, select a line item from a list (or table) of items, or select one of several numbered choices on a choice entry panel with the **Enter** key.

Pop-up Panels

The MAINVIEW Batch Optimizer user interface makes extensive use of pop-up panels for data entry and messages. The dialog places pop-up error messages near the field in error. It places pop-up prompt lists near the input field. It is best to view the dialog on a screen capable of displaying 43 lines of data (a 3270 Model 4 terminal or terminal emulator). When using a smaller screen size, pop-up panels may obscure underlying data necessary to correct erroneous input. ISPF offers two methods for addressing this issue:

- ISPF Window command
- ISPF Resize command

ISPF Window Command

The ISPF Window command allows you to move an active pop-up to another location within the screen space. Use this command to expose underlying fields obscured by the pop-up. There are two methods of invoking this command. One method is to type **WINDOW** at the command line (or use a function key), move the cursor to the position where you want the pop-up to appear, and press **Enter**. Alternatively, place the cursor anywhere on the active pop-up frame and press **Enter**, (ISPF will issue the message **WINDOW MOVE PENDING**). Move the cursor to the position where you want the pop-up to appear, and press **Enter** again.

For more details about this command, see the *IBM ISPF Users Guide* or the ISPF interactive help facility.

ISPF Resize Command

The MAINVIEW Batch Optimizer user interface limits the line depth of a pop-up screen to 20 lines. Although the actual size of a panel may exceed this, you can only view 20 lines at a time. For pop-up screens that exceed the maximum number of lines, ISPF activates the scrolling facility (indicated by More: +) which allows you to scroll through the entire panel content. To view the entire pop-up, use the **RESIZE** command. The **RESIZE** command causes the pop-up to be displayed in the entire screen space. To restore the pop-up to its original size, re-enter the **RESIZE** command at the command line.

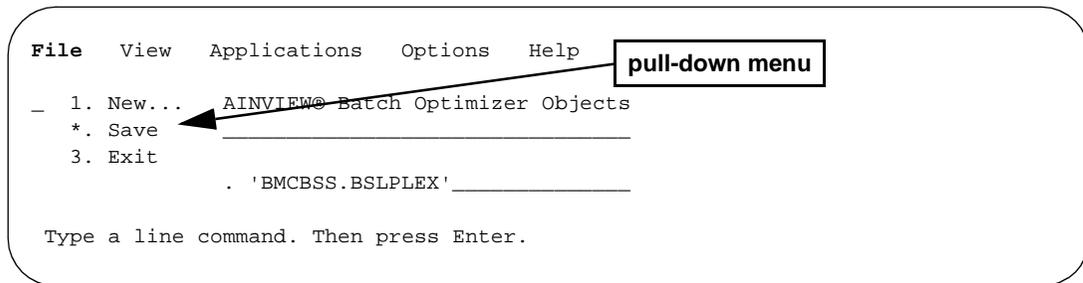
For more details about this command, see the *IBM ISPF Users Guide* or the ISPF interactive help facility.

Accessing the Action Bar

The Action bar offers access to the MAINVIEW Batch Optimizer user interface commands. The action bar is most accessible if you use a 3270 emulator that has mouse support for host sessions. However, if you do not have this support, the ISPF Actions command affords direct access to the Action bar options. To access an Action bar option using the ISPF Actions command, type **ACTIONS** at the command line followed by the "underlined" character of the Action bar option. For example, to access the File action bar option, type **ACTIONS F**.

Figure B-11 shows the result of this command, a File Pull-Down Menu panel. On panels with an Action bar, the dialog sets **F10** to Actions. This allows you to enter the "underlined" character at the command line and enter **F10** to invoke the Actions command.

Figure B-11 File Pull-Down Menu



For more details about this command, see the *IBM ISPF Users Guide* or the ISPF interactive help facility.

Point-and-Shoot

The MAINVIEW Batch Optimizer user interface makes extensive use of the ISPF Point-and-Shoot facility. When enabled, you can access this facility by positioning your cursor beneath a Point-and-Shoot defined field and pressing **Enter**. If your terminal emulator offers mouse support for host sessions, access this facility by moving your mouse pointer to the Point-and-Shoot field and double-click the left mouse button.

What the dialog displays in response to a point-and-shoot request depends upon the field. If you point-and-click a control data set member name on the MAINVIEW Batch Optimizer Objects List panel, the dialog displays Figure B-4 on page B-6. If you point-and-click a prompt indicator, the dialog displays a list of valid options for the input field. For example, if while editing a Data Policy definition, you click the prompt indicator adjacent to the Optimization method input field, the panel shown in Figure B-12 on page B-18 is displayed.

You also can access the point-and-shoot facility within the prompt list to select your choice. When enabled, the point-and-shoot facility allows you to navigate through the dialog with only the **Tab** and **Enter** keys. With a 3270 emulator enabled for host mouse support, you can navigate without using the keyboard until you have to perform a data entry task.

Figure B-12 Data Policy Method Options Pop-up List Panel

```

File  View  Display  Applications  Options  Help
----- Data Policy Method Options-----
BSSP00MA      BSSP00PP
Command ==> _____ Command ==> _____ SCROLL ==> PAGE

Data policy name      Method      Description
Definition comment   . ADVANCED  Use advanced buffering techniques
                    . BASIC     Modify buffer specifications only
Selection Criterio   . DSINFOONLY Gather statistical data only
                    *****  ***** End of Data *****

Program name .
Data set type .
Data set name .
Job name . . .

General options
Action . . . .
Optimization mode . . . . . _____ + Basic Advanced Dsinfoonly
Dynamic region adjustment . . N      Y=Yes N=No
Resource usage bias . . . . . _____ 1-10 99
Honor RUB value . . . . . _____ Y=Yes N=No
Optimize STC access . . . . . _____ Y=Yes N=No
Optimize TSO access . . . . . _____ Y=Yes N=No
SMF record type . . . . . 000          0, 128-255

```

Creating and Editing Control Data Set Members

You use the MAINVIEW Batch Optimizer user interface to create and edit objects. These objects may be control data set members or definition statements within a control data set member. In either case, the dialog provides multiple methods for creating and editing objects.

New Command

The dialog provides a number of ways to create new objects. If you wish to create a new control data set member object, use the NEW command. The dialog provides several methods for accessing this command. You can access this command by using the NEW pull-down choice of the File action bar option, or you can type the command directly at the command line. In cases where the dialog displays a list of objects, to create a new object you can select the <NEW> line with the E line command or position your cursor beneath this text and press **Enter**. When the dialog displays a list of objects, it displays this *special* line at the top of the list.

Note: The <NEW> line typically does not scroll with the data in the list so that it always appears at the top of the list. The exception to this rule is certain prompt list pop-up panels.

Insert Line Command

To create new objects within a control data set member, you can use the **NEW** command or you can use the **I** line command to insert new objects at specific positions. When you use the **New** command to create new objects, the dialog always inserts them at the top of the list. The insert line command gives you the flexibility of choosing where a new object will appear in the list. The insert facility is available when editing policy control data set members where the position of a statement in the list is important to how the product processes the policy. The insert facility is not available for object lists where position in the list is unimportant, i.e., sorted lists such as the BatchPlex Definition MVS images list or the MAINVIEW Batch Optimizer Objects List.

To use the insert line command, place an **I** on the line after which you want to place the new object and press **Enter**. The dialog will then display the appropriate data entry panel. After pressing **Enter** to submit the request, the dialog places the new object after the selected line and positions the cursor to the newly inserted line.

Copy Processing – Named Objects

The dialog allows you to copy control data set members and BatchPlex definition MVS images when creating new members. These two lists are examples of sorted lists that are insensitive to position. To copy a member or MVS image, place a **C** in front of the object you wish to copy and press **Enter**. The dialog displays a data entry panel where you can specify the name of the new object. After completing the data entry process and pressing **Enter**, the dialog inserts a new object in the correct sort order (alphabetic for control data set members and MVS images).

Rename Processing – Named Objects

The dialog allows you to rename objects that have names; such as control data set member names or MVS images. To rename a member or MVS image, type **R** in front of the object you wish to rename and press **Enter**. The dialog will display a data entry panel where you can specify the new name of the object. After completing the data entry process and pressing **Enter**, the dialog renames the object and inserts it in the correct sort order (alphabetic for control data set members and MVS images).

Copy and Paste Processing – Positional Objects

Another mechanism for creating new objects is through the Copy and Paste line commands, **C** and **P** respectively. When displaying a list where both line commands are available, you can use the copy line command to create a copy of an object by placing a **C** in front of the object you wish to copy. The dialog copies this object to a clipboard. Use the paste line command by placing a **P** on the line after which you want to place the contents of the clipboard and press **Enter**. The dialog copies the contents of the clipboard after the selected line and positions the cursor to the newly pasted line.

Cut and Paste Processing – Positional Objects

When available, the Cut and Paste line commands, **X** and **P** respectively, perform in the same manner as the COPY and PASTE processing with one exception. When you cut an object from the list using the X line command, the dialog removes the object from the list and places it on the clipboard. The object remains in the clipboard until you paste it at the desired location.

Warning! If you save the control data set member without pasting the object back into the list, you effectively delete the object from the list.

Object List Displays

The data displayed in an object list depends on the object type. For the MAINVIEW Batch Optimizer Objects List, the data displayed is a member list (see Figure B-1 on page B-3). For the BatchPlex Definition list, the data displayed is a list of MVS images defined to the BatchPlex (see Figure B-13 on page B-21). For a Job Policy list, the data displayed is a list of selection criteria and actions to perform (see Figure B-14 on page B-21). Although each list differs in content, certain concepts remain consistent for all lists. Enter line commands to the left of the line that you wish to select. The dialog displays "last activity" information regarding each line in a Response column. The panel displays the valid line commands for the list.

Figure B-13 BatchPlex Definition Panel

```

File  View  Applications  Options  Help
-----
BSSP00IU                               BatchPlex Definition                Row 1 to 4 of 4
Command ==> _____ SCROLL ==> CSR_
-

BatchPlex Information                               System: SYSM
Name . . . . . : BPLEX00                               SMF ID: SYSM
Comment . . . . . Mainview BatchPlex_____           Date  : 2002/02/26
XCF group name . . . . . BMCXCF_                       Time   : 12:42:55
Data policy name . . . . . DATPOL00 +
Job policy name . . . . . JOBPOL00 + UCF member name . UCFPOL00 +
Pipes initialization parms. PIPPRM00 + Pipes subsystem ID . MBOP
Pipes policy list name . . PIPLST00 +

Line commands:
E=Edit C=Copy D=Delete R=Rename

          ---Subsystems--- Number of
MVS Image BMCP MBOS XJS Initiators Definition Comment      Response
.. <New>
.. SYSA      BMCP MBOS MBOX      40      Production batch
.. SYSD      BMCP MBOS MBOX      40      Production DB2 Image
.. SYSI      BMCP MBOS MBOX      40      Production IMS image
.. SYSP      BMCP MEMD MBOX      40      Test image
***** Bottom of data *****

```

Figure B-14 Job Optimizer Policy Selection Criteria and Action List

```

File  View  Display  Applications  Options  Help
-----
                               Job Optimization Policy                Row 1 to 6 of 42
Command ==> _____ SCROLL ==> PAGE

Job Policy Information                               System: SYSP
Name . . : JOBPOL00                               SMF ID: SYSP
Comment . Removed the equal sign                   Date  : 2000/07/23
Priority . 9998                                     Time   : 19:54:46

Type an action code. Then press Enter.

Edit line commands:                               Statement line commands:
E=All fields                                       C=Copy D=Delete I=Insert
ES=Selection criterion only                       X=Cut P=Paste

          Selection Criteria                               Actions          Response
.. <New>
.. JBN(EQ,PGLSMS08) UID(EQ,PGL*) A01(EQ,5413)      + Readjtl
.. JBN(EQ,PGLMEL04)                               Readjtl
.. JBN(EQ,PGLSCR*)                               Readjtl
.. JBN(EQ,PGLPOPO3) UID(EQ,PGL*)                   Readjtl
.. UID(EQ,PGL*) JBN(EQ,PGLR*)                       Readjtl
.. JBN(EQ,PGLADD*)                               Readjtl

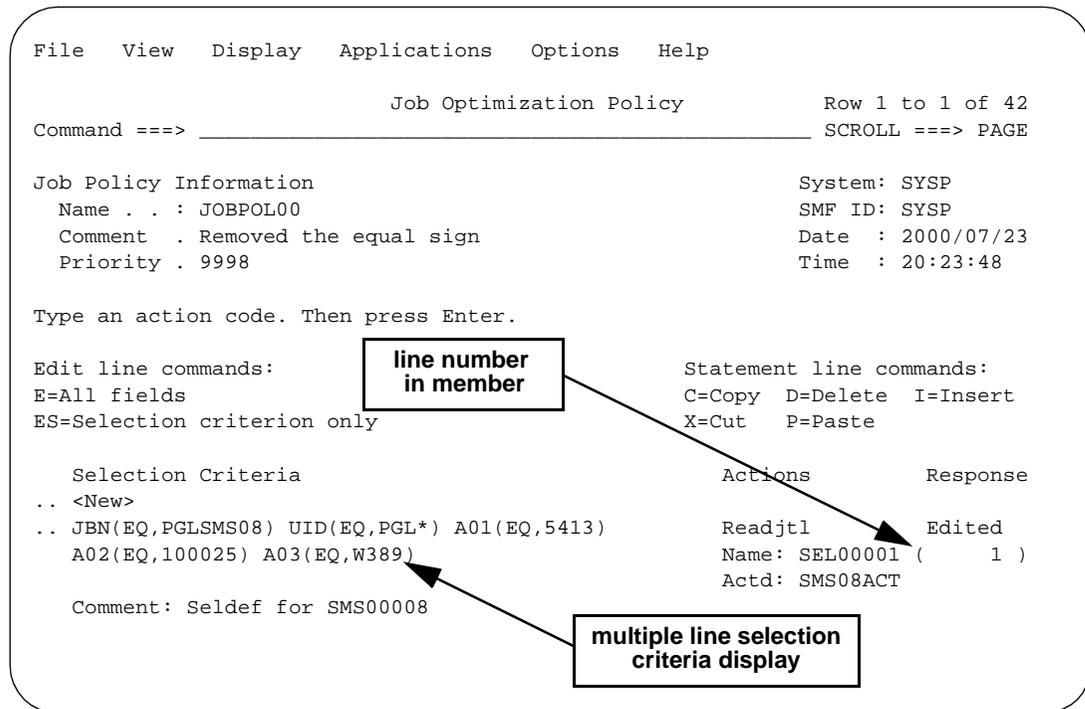
```

Prompt for Valid Line Command

When the dialog displays a list, the key portion of the list (typically highlighted) is defined as a point-and-shoot field. Instead of entering a line command, you can position your cursor beneath the key portion of the line and press **Enter**. The dialog will display a numbered list of valid line commands for the selected line (see Figure B-4 on page B-6). You can select a number that corresponds to the line command that you wish to invoke or position the cursor beneath the line command you want and press **Enter**. The dialog invokes the chosen line command.

Expand and Collapse

With some object lists, the dialog offers the ability to increase (Expand) or decrease (Collapse) the amount of data displayed for each item in the list. Typically, a list contains one line per item in the list. However, most list objects contain a lot more information than can be displayed in one line. For example, the Job Optimization Policy list contains a column that describes selection criteria. However, some selection criteria are too extensive to display in the allotted space. The dialog places a plus (+) symbol next to selection criteria meeting this condition (see Figure B-14 on page B-21). In such cases, you can *expand* the display by selecting the Expand Rows pull-down choice of the Display action bar option. After selecting this option, the list now displays the complete selection criteria for each object as well as diagnostic information such as the object position within the control data set member (see Figure B-15 on page B-23). Use the Collapse Rows pull-down choice of the Display action bar option to reverse the effects of the expand request.

Figure B-15 Job Optimization Policy Multi-line Criteria Display Panel

Row Update Processing

When editing some policy members, you might notice an Update action bar option. When present, this option allows you to enable *power-editing* fields in the list objects. To turn this option on, select the Row update ON pull-down choice of the Update action bar option. Once enabled, you can change fields appearing within the list display portion of the panel as input fields by over-typing them with new data. Once you press **Enter**, the dialog updates the object. When you activate this option, the dialog displays a message indicating that the line update mode is now active.

Row update processing is currently limited to certain policy members and certain fields within those policy members.

Online Help

The MAINVIEW Batch Optimizer user interface provides explanations of all panels, fields, and messages in the ISPF interface and dialogs. Panel-level help displays a pop-up window that explains the purpose of the panel, provides information about using the panel, and provides a description of all fields on the panel. Field-level help displays a pop-up window that describes the purpose of the field and possible values that you can enter in the field. Message-level help displays a pop-up window that describes an error message, the reason for the error message, and possible responses to the error. Reference phrase help displays a pop-up window that describes the highlighted word or phrase. These highlighted phrases typically appear on help panels. Help is also available for the action bar.

You can activate the help feature by, pressing **F1 (HELP)**, selecting the Help option from the Help pull-down menu, or typing **HELP** on the command line and pressing **Enter**. The type of help displayed when you activate the help feature depends on the cursor's position.

Panel Help

The dialog displays the panel-level help when you activate the help feature with the cursor on the command line or on any part of the panel other than the action bar or a choice entry or data entry field.

Field Help

The dialog displays the field-level help when you activate the help feature with the cursor on a choice entry field, a data entry field, the heading of a data entry field, or some output-only fields.

Message Help

The dialog displays the message-level help when you activate the help feature after the dialog displays a "long message" pop-up on the screen. The message help clarifies the error causing the message, describes the reason for the error message, and suggests actions to resolve the error.

Reference Phrase Help

The dialog displays reference phrase help when you position the cursor on a highlighted field within a help panel and activate the help feature. This facility offers further explanations of phrases or words mentioned in an already displayed panel, field, or message help panel.

Other User Interface Tasks

Besides editing the control data set members, the dialog provides access to other facilities of MAINVIEW Batch Optimizer. When contacting BMC Software Product Support, you will need to obtain the product version information. You can obtain this as well as module maintenance and environmental version information through the dialog. Alternatively, you can install BMC Software product authorization codes through the dialog.

Product Version and Maintenance Level Information

You can access the product version and maintenance level information by using the MN line command on the MAINVIEW Batch Optimizer Objects List panel to select the active BatchPlex member (BPLEXnn). Figure B-16 on page B-26 shows the Components Maintenance Information panel.

Figure B-16 Components Maintenance Information Panel

```

File  View  Applications  Options  Help

BSSP00ZX                               Component Maintenance Information  Row 1 to 23 of 52
Command ==>> _____ SCROLL ==>> PAGE

Control data set: 'BMCBSS.BSLPLEX'
Subsystem ID . : BCSS                               System: SYSP
                                                    SMF ID: SYSP
                                                    Date  : 2000/07/23
                                                    Time  : 13:00:35

Component Information
Data Optimizer           Job Optimizer
Version . . . : 1.1.00   Version . . : 1.1.00
Tape created  : 2000/05/19 Tape created : 2000/05/19
Authorization : Licensed  Authorization: Licensed
NonVSAM status: Enabled   Status . . . : Enabled
VSAM status . : Enabled

Load module maintenance information
---Assembly---
Name      FMID      Date      Time      Last PTF   Address
BSLMACTR  BMCBSL3  06/14/00  07.12  BMCBSL3   158101D0
BSLMATCH  BMCBSL3  07/15/00  20.46  BMCBSL3   1586EA68
BSLMATCX  BMCBSL3  06/14/00  07.13  BMCBSL3   158DF010
BSLMBAGT  BMCBSL3  06/14/00  07.13  BMCBSL3   157F2800
BSLMBLTR  BMCBSL3  07/19/00  11.12  BMCBSL3   15840CA0
BSLMBLTU  BMCBSL3  06/14/00  07.14  BMCBSL3   1586E1C8
BSLMCELL  BMCBSL3  06/14/00  07.15  BMCBSL3   12953408
BSLMDAEP  BMCBSL3  07/20/00  11.38  BMCBSL3   157F9F90

```

BMC Software Product Support may ask you to refer to this information when attempting to diagnose a reported problem.

Environment Version and Release Information

You can access the version and release information by selecting the Version information pull-down choice of the Options action bar item or by entering the command **VERSION** on the command line. Figure B-17 shows the Version/Release Information panel.

Figure B-17 Version/Release Information Panel

```

Version/Release Information

Command ==> _____

IBM Products          Version

MVS . . . . . : SP6.0.8
OS/390 . . . . . : 02.08.00
MVS sysplex name . . . : BMCPLEX0
RACF . . . . . : 2.60.8
DFHSM . . . . . : 1.05.0
ISPF . . . . . : 4.8
TSO/E . . . . . : 2.06.0
DFSMS . . . . . : 1.5.0
SMF system identifier : SYSM
SMS subsystem name . . : SMS
ISPF Y2K Support . . . : Y

```

It is useful to have the information displayed by this panel on hand when contacting BMC Software Product Support.

BMC Software Product Authorization User Interface

You can access the product authorization facility by selecting the BMC Software Security Facility pull-down choice of the Applications action bar item or entering the command **SECURITY** on the command line. Figure B-18 shows the Product Authorization Primary Menu panel.

Figure B-18 Product Authorization Primary Menu Panel

```
MAINVIEW Batch Optimizer Product Authorization Primary Menu

Select an option. Type additional information if applicable. Then press
Enter.

Options

_ 1. Process password (Requires product load library and password)
  2. Display product authorization (Requires product load library only)
  3. Display current processor information
  4. Help about...
  5. Exit
Additional information

Product load library . . . 'BMC.BSS.LOAD'

Authorization password . . ____ ____ ____ ____

COMMAND ==> _____
```

In this facility, you can add new licensing information, display existing licensing information, or display information about the current processor. It is on this panel that you will include any licensing passwords provided to you by BMC Software.

Appendix C User Control Exit

This appendix describes how to implement the User Control exit to intercept data sets for I/O optimization processing. This appendix discusses the following topics:

Overview	C-2
User Control Exit Decisions	C-4
Status Change Call-Back Capability	C-4
Testing the User Control Exit	C-5
User Exit Statistics Report	C-6

Overview

In gathering information about a data set, Data Optimizer checks for the User Control exit before checking the data policy. When an OPEN request occurs, Data Optimizer must determine which data sets should be processed by the Data Optimizer Performance component. Typically, these decisions can be made by using the data policy entries combined with the data set's attributes and the type of processing being performed by the accessing program. However, there are times when the user may wish to implement selection algorithms that are too complex for the data policy definition. An example would be if you wanted to include jobs that have jobnames whose fifth character position is a numeric between two and five. Since the source of decision-related information is virtually unlimited, you can create a User Control exit to assist in tailoring Data Optimizer processing to the installation requirements of your site.

Note: Member DAPUEXIT in BMC.DAP.JCL contains a sample User Control exit that factors SMS information into the decision-making process.

The User Control exit executes as an extension of the OPEN SVC instruction used to connect the data set to the application program. The User Control exit must be capable of being loaded into the extended CSA and must be link edited as APF-authorized, RMODE(ANY), AMODE(31) and reentrant.

BMC Software Product Support recommends that you place the User Control exit in the same library as Data Optimizer. Since the User Control exit receives control in the supervisor state and storage protection key of the issuer of the OPEN SVC, you must use extreme care during the design, coding, and testing phases to ensure that you do not compromise the integrity of the application or the system. If you are uncertain about any aspects of a User Control exit that you are writing, contact BMC Software Product Support.

Data Optimizer is shipped with the User Control exit disabled. You must explicitly enable the User Control exit using the DAP USEREXIT ENABLE operator command either from an MVS console or the BCSS commands data set (processed during BCSS subsystem address space initialization). For details, see "Activating a Data Policy" on page 6-18.

The default name of the User Control exit is DAPUEXIT; however, you can assign any name to the User Control exit. If you use a name other than DAPUEXIT, specify the new name immediately after the ENABLE parameter (separated by at least one space); for example, DAP USEREXIT ENABLE *exitname*.

Only one User Control exit per Data Optimizer image may be active at any time. The User Control exit will not be shared between multiple copies of Data Optimizer executing on the same MVS image. You can disable processing with the User Control exit by using the DAP USEREXIT DISABLE command. The DAP USEREXIT,STATUS command provides information about the status of the current user exit environment. For more information, see “Activating a Data Policy” on page 6-18.

The User Control exit is called once during the execution of the DAP USEREXIT ENABLE command to perform global exit initialization and once during the execution of the DAP USEREXIT DISABLE command to perform global exit termination. The initialization call may be used to initialize global (ECSA) storage areas with information that is required for the decision-making process. The termination call should be used to clean up any global resources that are allocated during the initialization and decision-making calls. During these initialization and termination calls, the exit routine will be executing in the supervisor state and storage protection key of the BCSS address space that is hosting the Data Optimizer product image.

Note: This protect key may be different than the protect keys used during the decision-making calls. Since switching execution protection keys can be confusing, BMC Software recommends that any global (ECSA) storage acquired during the initialization call be allocated from storage subpools that are not *fetch-protected* (such as 241).

The installation SAMPLIB data set contains several DAPUE-prefixed members containing sample exit routine code or macros and other assembler-time information that are required to generate the sample User Control exit. The SAMPLIB data set should be allocated as part of SYSLIB concatenation during the assembly process for the User Control exit (sample JCL is also provided). Whenever possible, your User Control exit(s) should make use of the facilities and services that are contained in these members.

Due to the nature of the environment in which the User Control exit is called, there are several general restrictions on the types of MVS services that can be used:

- Do not issue DYNALLOC requests.
- Do not issue data management requests (such as BLDL, FIND, LOAD, OPEN, OPEN TYPE=J, CLOSE, CLOSE TYPE=T, EOVS and FEOVS).
- Do not use ENQs and RESERVEs that might result in address space hangs.
- Do not issue task management requests (such as ATTACH, ATTACHX, LINK, and LINKX).

- Do not use ABEND.

The User Control exit should not attempt to modify itself or any control block structure that is not defined as part of the API explicitly documented with the DAPDEPL macro. Failure to adhere to these guidelines may result in hangs, secondary ABENDs, or system-integrity exposures.

User Control Exit Decisions

The parameter list provided to the User Control exit is documented in the Data Optimizer exit parameter list (DAPDEPL) macro, which contains a request-specific section for each call *event* that will be presented to the User Control exit. All calls, except for the initialization and termination events, expect that a return code will be placed into the DEPL, indicating the decision made by the User Control exit. The default return code 0 indicates that no decision could be provided and that the User Control exit caller should make an intercept decision by using data policy information or DD override statements. Return code 0 causes Data Optimizer to function as if the User Control exit was disabled.

Note: An exit routine that always returns either a positive or a negative intercept decision may result in data policy information not being used.

Status Change Call-Back Capability

Sometimes when the User Control exit may need to modify its operating environment or disable itself altogether. A special call-back facility has been provided to enable these changes during any User Control exit call. For example, if the User Control exit determines that it should decrease the scope of exit routine serialization from the system-level default (allowing one active caller from a single address space) to a step-level (allowing one active caller from a number of address spaces), or to disable serialization altogether (allowing a number of active callers from a number of address spaces), the appropriate bits in the DEPLSCFLAGS field would be set and the routine whose address is contained in DEPLSCENTRY would be called to perform the change. When control is returned from the call-back facility, the status change would be in effect.

Note: The serialization change in the preceding example might typically be performed during the *exit initialization* call if the User Control exit can support concurrent callers from multiple address spaces. The sample exit routine contains this type of logic.

Testing the User Control Exit

Before you implement the User Control exit, BMC Software Product Support recommends that you test it to ensure that it is operating correctly. A driver program that performs this function, DAPUEDRV, is provided. The program executes the DAPUEXIT User Control exit, which is located in the STEPLIB or your MVS linklist.

The program produces a report on the linkage editor attributes of the User Control exit. The program reads the data sets that are specified and indicates whether the data sets would be intercepted, based on exit processing.

Execute the driver program each time you change the User Control exit, and test as many data sets as necessary to ensure that the User Control exit is making the proper intercept decision. Executing the driver program will help to avoid problems when you implement the User Control exit.

Sample JCL to execute the program and produce the report is contained in member DAPUCJCL of data set BMC.BSS.SAMPLIB. When you modify the JCL, indicate the data set names for the data sets you want to test. A sample of the User Exit Statistics Report is shown in Figure C-1 on page C-6.

User Exit Statistics Report

The program DAPUEDRV produces a report indicating the types of processing that occur if the User Control exit is implemented. The report indicates whether the User Control exit is reentrant and whether it is in a 31-bit addressing mode. Both items are requirements for the User Control exit. Figure C-1 shows the report that is produced.

Figure C-1 User Exit Statistics Report

```

month dd, cyy          D A T A   O P T I M I Z E R

TIME hh:mm:ss DATE cyyymmdd   BMC SOFTWARE, INC
                               USER EXIT STATISTICS
                               DAPUEXIT FOUND TO BE REENTRANT AND AMODE31

      DSNAME                                USEREXIT DECISION CODE

DAP.COMPNONE.DATA                NOINTERCEPT
ABC.IPCSDDIR                     NOINTERCEPT
DAP.COMPOPEN.USEREXIT.DATA      INTERCEPT
SYS1.MAN1                        NOINTERCEPT
ABC2.IPCSDDIR                   NOINTERCEPT
DAP.COMPLOAD                    INTERCEPT
SYS1.MAN2                       NOINTERCEPT
ABC3.IPCSDDIR                   NOINTERCEPT
ABC.TEST                        NODECISION
DAP.COMPNONE.USEREXIT.DATA      NOINTERCEPT
SYS1                             NOINTERCEPT
ABC4.IPCSDDIR                   NOINTERCEPT
ABC5.IPCSDDIR                   NODECISION

DATA OPTIMIZER (C) COPYRIGHT 2000.      BMC SOFTWARE, INC.
DATA OPTIMIZER IS A TRADEMARK OF BMC SOFTWARE, INC.

```

The following fields are included in this report:

Table C-1 User Exit Statistics Report Field Descriptions

Field	Description
DSNAME	The data set names specified in the JCL.
USEREXIT DECISION CODE	For each data set name, a user exit decision code is displayed. The following decision codes can be indicated for the data sets: INTERCEPTThe User Control exit will intercept processing for data sets with this name. NOINTERCEPTThe User Control exit will not intercept processing for data sets with this name. NODECISIONThe User Control exit will not make a decision regarding this data set; the data set will be passed to the subsystem for an intercept decision.

Appendix D Special Considerations

This appendix discusses items that may require special consideration in your environment. This appendix also describes certain exclusions that occur during Data Optimizer processing. This appendix discusses the following topics:

Overview	D-2
Buffer Values Coded in the JCL	D-2
Data Integrity with Non-VSAM Advanced Mode Processing	D-3
EXCP and EXCPVR Processing Exclusions	D-4
Excluded Programs	D-6
Excluded DD Statements	D-9
Excluded Data Sets	D-10

Overview

When implementing Data Optimizer, there are several considerations which deserve special attention. Some have to do with the way Data Optimizer behaves when it intercepts an open for optimization and others to do with Data Optimizer's exclusion of known entities that don't optimize well for one reason or another. This section describes these special considerations and exclusions in detail.

Buffer Values Coded in the JCL

For most data centers, the quest for throughput improvements is an ever-present battle. Many have gone through extensive evaluations to determine optimum buffer values and to hard-code them on the DD statements of their production job streams. With Data Optimizer, many of these values are no longer needed. However, removing these buffer values can be a significant task and may not be desirable if Data Optimizer is ever disabled.

By default, Data Optimizer will only alter buffer values when the value it would select is greater than the value specified in the JCL. By specifying the `OVERRIDEUSER` option, Data Optimizer can be instructed to use its selected value regardless of any JCL specifications. This option can be specified for a specific definition or at the global level. It can be particularly valuable for providing virtual storage constraint relief for non-VSAM processing, since Data Optimizer can provide performance benefits for non-VSAM data sets without increasing below-the-line buffers. In fact, this option used in conjunction with the `BUFNO` and `NCP` options specifying low values such as 2, can be used to provide even more virtual storage constraint relief than removing the JCL specifications and allowing the defaults. For more information, see "Override User Values Option" on page 5-10.

Data Integrity with Non-VSAM Advanced Mode Processing

The Advanced mode of Data Optimizer non-VSAM processing uses a variety of techniques to greatly reduce the amount of I/O wait time for the application. One way this reduction is accomplished is by posting the I/O as complete when the data is successfully moved to a data buffer, not necessarily written to the device. This movement allows the application to immediately continue processing without waiting for I/O completion. This situation is similar to the DEFERED-WRITE feature that is provided by VSAM.

The only potential for difficulty occurs when data is being written to (and read from) a data set simultaneously in some synchronized manner. If a function writes data to DASD and subsequently informs a reading function that the updated data has been written, the reading function could read an old copy of the data from the device because the current copy is still in a buffer in the writer's memory.

Data Optimizer tries to prevent this problem in a variety of ways. When the writer and reader of a common data set reside in the same address space as different tasks, Data Optimizer flushes all writes immediately to DASD. This action ensures that the reader will always get the most current copy of any updated records.

When the writer and reader of a common data set are in different address spaces—possibly on different MVS images—Data Optimizer presumes that the MVS disposition processing will ensure that any process which updates a data set has exclusive use (DISP=OLD) of the data set for the duration of the updates. This action prevents any readers from accessing any records in the data set until all records have been written to DASD by the writing process.

If an updating process is allowed to run concurrently (DISP=SHR) with other processes accessing a common data set, Data Optimizer considers this a data integrity exposure and abends the process with a U2210-7709 abend. This action ensures that no reader can access records in a data set while another process may have updated copies buffered.

Warning! This data integrity protection method works only when Data Optimizer is active on all systems that could potentially access any given file concurrently.

When Data Optimizer is not on a system in which a file is processed, Data Optimizer has no way of knowing what type of access is being performed. This condition is analogous to running GRS on some, but not all members of a sysplex.

Fortunately, processes that rely on synchronization while concurrently reading and writing shared files are not common. The most common of these processes are vendor products that use a shared control file or database such as Computer Associates CA-View (SAR) and CA-Deliver (RMO) products. Some in-house written applications may also use this methodology for synchronization. As long as Data Optimizer is active on all systems in a sysplex, it will not allow any data corruption to occur.

If Data Optimizer is not active on all systems, such as during an initial implementation, you can take several approaches to ensure data integrity. The safest approach is to create a data policy entry to set the non-VSAM optimization mode to BASIC. This mode prevents Data Optimizer from posting any write as complete until the data is physically written to DASD.

If you are aware of what applications or vendor products rely on synchronization, you can create data policy entries to restrict only those processes to BASIC mode or you can exclude them from optimization processing. All other processes can continue to use ADVANCED mode as before. When Data Optimizer is active on all systems, this concern is eliminated and all non-VSAM data sets that are eligible for ADVANCED mode can be processed with it.

When Data Optimizer is active on all systems, it detects any potential data integrity exposures and issues a U2210-7709 abend. If you encounter one of these abends, the easiest solution is to exclude all processes that access the data set from Data Optimizer processing. If the process involves a single program, use a data policy entry to exclude the program. If the process involves any program accessing a given file or group of files, use a data policy entry to exclude the data set or group of data sets (through a *dsname* prefix). If the process causing the U2210-7709 abend can read and write to the data set(s) in question without integrity concerns, a data policy entry that includes the SHAREDWRITE option can be added, which tells Data Optimizer to bypass the data integrity check and perform optimization as normal.

EXCP and EXCPVR Processing Exclusions

Data Optimizer provides I/O performance processing to applications that use EXCP and EXCPVR processing and that process significant amounts of data through general-purpose channel programs.

EXCP and EXCPVR applications that manipulate control files, or perform unusual or non-standard channel program sequences, might not receive the expected performance improvements. Additionally, EXCP/EXCPVR applications that perform control and diagnostic-type operations, such as write record zero, reserve, reset allegiance, and diagnostic control are not supported and should be excluded from Data Optimizer processing.

Table D-1 shows a list of CCW op-codes that Data Optimizer does not support.

Table D-1 Excluded CCW Op-codes

Op-Code	Description
X'02'	read IPL
X'09'	diagnostic write home address
X'0A'	diagnostic read home address
X'0F'	space count
X'13'	recalibrate
X'14'	unconditional reserve
X'15'	write record zero
X'19'	write home address
X'1F'	set file mask
X'27'	perform subsystem function
X'34'	sense path group ID
X'3B'	set high performance storage limits
X'3E'	read subsystem data
X'44'	reset allegiance
X'4E'	read message ID
X'54'	sense subsystem status
X'5B'	suspend multipath reconnection
X'63'	define extent
X'74'	sense subsystem counts
X'87'	set subsystem mode
X'94'	device release
X'A4'	read and reset buffered log
X'AF'	set path group ID
X'B4'	device reserve
X'C4'	diagnostic sense/read
X'F3'	diagnostic control

If an application uses standard access methods that you would like to optimize, but also performs EXCP/EXCPVR processing on other files and uses one or more of these CCWs, you can disable performance for the EXCP/EXCPVR accesses only. To do so, add the following DD statement to the step:

```
//DAP@NXCP DD DUMMY
```

EXCP/EXCPVR processing is bypassed, but all standard access method processing continues to be optimized.

Excluded Programs

Table D-2 shows the programs that Data Optimizer automatically excludes from non-VSAM I/O performance processing.

Table D-2 Excluded Programs - Non-VSAM Processing (Part 1 of 3)

Program	Description
ACK	BMC Software utilities that start with these characters
ADR	IBM DFDSS DASD backup and restore utility or any load module whose name starts with these characters
ADSMI002	Sterling Software DMS DASD space management utility
ADUU	BMC Software utilities that start with these characters
AESATTCH	BMC Software ARC utility
AESCTRL	BMC Software ARC utility
AESIRRC0	BMC Software ARC utility
AFRMAIN	BMC Software DB2 utility
AMASPZAP	IBM Super ZAP utility
AMU	BMC Software utilities which start with these characters
APLEND	Mainstar ASAP Disaster Recovery Utility
ARC	IBM DF/HSM DASD space management utility or any load module whose name starts with these characters
ARU	BMC Software DB2 utilities that start with these characters
ASF	IBM SmartBatch components or any load module whose name starts with these characters
ASUU	BMC Software utilities that start with these characters
BCSM	BMC Software Consolidated Subsystem (BCSS) load modules or any load module whose name starts with these characters
CAPU	BMC Software Change Accumulation utilities which start with these characters

Table D-2 Excluded Programs - Non-VSAM Processing (Part 2 of 3)

Program	Description
CAxxA001	Computer Associates Sort product Note: xx is typically replaced with the product release number.
CTT	BMC Software Control-T utility or any load module whose name starts with these characters
DBUTLTY	Computer Associates Datacom Product
DBMUFPR	Computer Associates Datacom product
DFHSI	CICS
DFS	IBM IMS and its utilities which start with these characters
DSIATTMT	Tivoli NetView program product
DSNUTIL	IBM DB2 utilities that start with these characters
EASY	Computer Associates CA-Easytrieve utilities that start with these characters
EDM	BMC Software utilities that start with these characters
EYU9XECS	CICS SYSPLEX/SM
EZT	Computer Associates Easytrieve utilities that start with these characters
FABAD	Fast Path Analyzer
FABCU	Fast Path Unload/Reload
FDR	Innovation DASD backup and restore utility or any load module whose name starts with these characters
SYNC	Syncsort Sort product
ICKDSF	IBM ICKDSF DASD utility
ICPU	BMC Software Image Copy Plus utilities that start with these characters
ICE	IBM DFSORT or any load module whose name starts with these characters
IDMS	Computer Associates IDMS utilities that start with these characters
IEFIIC	IBM MVS initiator
IERRCO00	alternate name for Sort utilities
IGHRCO00	alternate name for Sort utilities
LDPU	BMC Software Load Plus utilities that start with these characters
NDVR	Computer Associates Endeavor utilities that start with these characters
\$PDSFA	SEA PDSFAST utility or any load module whose name starts with these characters
PIPESORT	SyncSort Sort for Pipes utility

Table D-2 Excluded Programs - Non-VSAM Processing (Part 3 of 3)

Program	Description
PLSORT0	Phase Linear Sort product
PDSEEDIT	PDSMAN PDSEEDIT
PDSM	PDSMAN
PFPMMAIN	BMC Fast Path Utilities
PTL	Computer Associates/Platinum DB2 utilities that start with these characters
RCU	BMC Software DB2 utilities which start with these characters
RSPSMFTK	Mainstar ASAP Disaster Recovery Utility
RVPU	BMC Software Recovery Plus utilities which start with these characters
SCRIPT	IBM DCF/SCRIPT mainframe word processing utility
SORT	alternate name for Sort utilities
TSSAF	Fast Path Fast Recovery
TSSR	Fast Path Recovery Control
UPPU	BMC Software utilities which start with these characters
UTL	Computer Associates/Platinum DB2 utilities which start with these characters

Table D-3 lists the programs that Data Optimizer automatically excludes from VSAM I/O performance processing.

Table D-3 Excluded Programs - VSAM Processing (Part 1 of 2)

Program	Description
AESATTCH	BMC Software ARC utility
AESCTRL	BMC Software ARC utility
AESIRRC0	BMC Software ARC utility
AFRMAIN	BMC Software DB2 utility
ARC	IBM's DF/HSM DASD space management utilities which start with these characters
DFHSI	CICS
IDMS	Computer Associates IDMS utilities which start with these characters
DBIUBC00	BMC Software Data Base Integrity utility
DFS	IBM IMS and its utilities which start with these characters
DSIATTMT	Tivoli NetView program product
DSNUTIL	IBM DB2 utilities which start with these characters
EYU9XECS	CICS Sysplex/SM

Table D-3 Excluded Programs - VSAM Processing (Part 2 of 2)

Program	Description
FABAD	Fast Path Analyzer
FABCU	Fast Path Unload/Reload
IEFIIC	IBM MVS initiator
LDPU	BMC Software Load Plus utilities which start with these characters
NDVR	Computer Associates Endeavor utilities which start with these characters
PFPMAIN	BMC Fast Path Utilities
RCU	BMC Software DB2 utilities which start with these characters
TSSAP	Fast Path Fast Recovery
TSSR	Fast Path Recovery Control

Excluded DD Statements

Table D-4 lists the DD statements that Data Optimizer automatically excludes from both VSAM and non-VSAM performance processing.

Table D-4 Excluded DD Statements (Part 1 of 2)

DD Statement	Description
ABNLTERM	ABEND-AID log data set
ASFDSTAT	IBM SmartBatch formatted statistical report
DAP@DSI	Data Optimizer data set information report
DAPMESGS	Data Optimizer messages listing
DAPNSTAT	Data Optimizer statistical report
DAPRPTS	Data Optimizer reports listing
DAPSTATS	Data Optimizer statistical report
DAPVSTAT	Data Optimizer statistical report
DPOSTATS	Data Optimizer statistical report
DPVSTATS	Data Optimizer statistical report
JOBCAT	job-level catalog DD statement
JOBLIB	job-level load library DD statement
PDX	BMC PDX utility library
STEPCAT	step-level catalog DD statement
STEPLIB	step-level load library DD statement

Table D-4 Excluded DD Statements (Part 2 of 2)

DD Statement	Description
SYSABEND	formatted system dump data set
SYSCHK	checkpoint data set
SYSCKEOV	checkpoint at end-of-volume facility data set
SYSEXEC	REXX procedure library
SYSLBC	TSO broadcast data set
SYSMDUMP	machine-readable system dump data set
SYSPROC	CLIST procedure library
SYSUADS	TSO userid data set
SYSUDUMP	formatted system dump data set
SYSUEXEC	REXX procedure library
SYSPROC	CLIST procedure library
VVHC\$IOS	HIPER-CACHE statistics listing
UCC7CMD5	CA-7 batch terminal interface commands
VVHC\$M\$\$	HIPER-CACHE message listing
VVHC\$MSG	HIPER-CACHE message listing

Excluded Data Sets

Data Optimizer automatically excludes the following data set names from VSAM and non-VSAM performance processing:

- NULLFILE
- any data set whose name begins with “SYS1.VVDS.V”
- any data set whose name has “SARDBASE.” or “RMODBASE.” anywhere in it

Appendix E SMF Recording

This appendix describes the SMF recording function of Data Optimizer. This appendix discusses the following topics:

Overview	E-1
SMF Record Header	E-2
Self-Defining Section	E-3
Product Section	E-5
Subsystem Identification Section	E-5
Job Information Section	E-6
Subsystem Processing Options Section	E-7
General Data Set Information Section	E-11
DCB Information Section	E-11
ACB Information Section	E-13
Non-VSAM I/O Statistics Section	E-14
VSAM I/O Statistics Section	E-15
VSAM LSR Statistics Section	E-15
VSAM Catalog Information Section	E-16

Overview

Data Optimizer will produce SMF records at CLOSE time for data sets that it has intercepted for performance processing. You can control the record type of the SMF record by using the SMF Record Type option.

Data Optimizer's SMF recording feature provides a wide range of informational and performance-related data that is presented in a number of different self-defining sections within the record. The following types of information are produced:

- job characteristics
- data policy options in effect

- ACB/DCB information
- logical I/O counts
- physical I/O counts
- catalog information

In addition to defining the SMF record type to be used, you must also activate SMF recording for the selected record type. You can activate SMF recording by initializing the SMFPRMxx member of PARMLIB in one of the following two ways:

- Specify SUBSYS (mbos, ... TYPE(..xxx..) ..) in SMFPRMxx. *mbos* is the name of the MBO subsystem and *xxx* is the SMF record type.
- Omit the SUBSYS parameter and include the SMF record type in the SYS parameter.

Once the SMF data has been recorded, you can report on it using the provided SMF Report Utility or one that you create. For more information on the provided SMF Report Utility, see Appendix G, “Batch Utilities.”

SMF Record Header

The SMF record header appears at the beginning of each record. SMF header field descriptions are shown in Table E-1.

Table E-1 SMF Header Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFRCDDO	180	structure	Record header and self-defining sections
0 0	SMFDOHDR	28	structure	Record header
0 0	SMFDOLEN	2	binary	Record length
2 2	SMFDOSEG	2	binary	Segment descriptor
4 4	SMFDOFLG	1	binary	System indicator BitDefinition 0Subsystem identification follows system identification 1Subtypes used 2Reserved 3–6MVS Version indicators (Set by SMF) 7Reserved. Note: For MVS/ESA SP Version 4, bits 3, 4, 5, and 6 are on.
5 5	SMFDORTY	1	binary	Record type

Table E-1 SMF Header Field Descriptions

Offset	Name	Length	Format	Description
6 6	SMFDOTME	4	binary	Local time, in hundredths of a second, record was moved to SMF buffer.
10 A	SMFDODTE	4	packed	Date that the record was moved to the SMF buffer form is 0cyyddF F is the sign
14 E	SMFDOSID	4	character	System identification (taken from SID parameter)
18 12	SMFDOWID	4	character	Subsystem identification for the address space STCStarted Task
22 16	SMFDOSTP	2	binary	Record subtype 1Data Optimizer Performance Statistics Note: This field defines which record mapping is being used.
24 18	SMFDOSDL	4	binary	Length of self-defining section

Self-Defining Section

The self-defining section consists of a set of triplet fields that you use to locate the various sections within the record.

The first field of each triplet is the offset from the beginning of the record (field SMFDOxxO) to the start of the section. The second field (SMFDOxxL) contains the length of the section, while the third field (SMFDOxxN) contains the number of the particular type of sections that contiguously exist. A nonzero value in the number field means that the section exists in the record; a zero value indicates that the section does not exist in the record.

The self-defining section field descriptions are shown in Table E-2.

Table E-2 Self-Defining Section Field Descriptions (Part 1 of 2)

Offset	Name	Length	Format	Description
0 0	SMFDOSDS	152	structure	Self-defining section
0 0	SMFDOPRO	4	binary	Offset of product section
4 4	SMFDOPRL	2	binary	Length of product section
6 6	SMFDOPRN	2	binary	Number of product sections - '1'
8 8	SMFDOSIO	4	binary	Offset of subsystem identification section
12 C	SMFDOSIL	2	binary	Length of subsystem identification section
14 E	SMFDOSIN	2	binary	Number of subsystem identification sections - '1'

Table E-2 Self-Defining Section Field Descriptions (Part 2 of 2)

Offset	Name	Length	Format	Description
16 10		64	binary	Reserved
80 50	SMFDOJIO	4	binary	Offset of job information section
84 54	SMFDOJIL	2	binary	Length of job information section
86 56	SMFDOJIN	2	binary	Number of job information sections - '1'
88 58	SMFDOSOO	4	binary	Offset of subsystem processing options section
92 5C	SMFDOSOL	2	binary	Length of subsystem processing options section
94 5E	SMFDOSON	2	binary	Number of subsystem processing options sections - '1'
96 60	SMFDOGDO	4	binary	Offset of general data set information section
100 64	SMFDOGDL	2	binary	Length of general data set information section
102 66	SMFDOGDN	2	binary	Number of general data set information sections - '1'
104 68	SMFDODIO	4	binary	Offset of DCB information section
108 6C	SMFDODIL	2	binary	Length of DCB information section
110 6E	SMFDODIN	2	binary	Number of DCB information sections 0-1
112 70	SMFDOAIO	4	binary	Offset of ACB information section
116 74	SMFDOAIL	2	binary	Length of ACB information section
118 76	SMFDOAIN	2	binary	Number of ACB information sections 0-1
120 78	SMFDONIO	4	binary	Offset of non-VSAM I/O statistics section
124 7C	SMFDONIL	2	binary	Length of non-VSAM I/O statistics section
126 7E	SMFDONIN	2	binary	Number of non-VSAM I/O statistics sections 0-1
128 80	SMFDOVIO	4	binary	Offset of VSAM I/O statistics section
132 84	SMFDOVIL	2	binary	Length of VSAM I/O statistics section
134 86	SMFDOVIN	2	binary	Number of VSAM I/O statistics sections 0-1
136 88	SMFDOVLO	4	binary	Offset of VSAM LSR statistics section
140 8C	SMFDOVLL	2	binary	Length of VSAM LSR statistics section
142 8E	SMFDOVLN	2	binary	Number of VSAM LSR statistics sections 0-1
144 90	SMFDOVCO	4	binary	Offset of VSAM Catalog statistics section
148 94	SMFDOVCL	2	binary	Length of VSAM Catalog statistics section
150 96	SMFDOVCN	2	binary	Number of VSAM Catalog statistics sections 0-1

Product Section

The Product section appears in all records and provides general information about Data Optimizer and your system. Product section field descriptions are shown in Table E-3.

Table E-3 Product Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDOPR	68	structure	Product section
0 0	SMFDOPNM	10	char	Product name - SmartBatch
10 A	SMFDORVN	2	binary	Record version number Data Optimizer
12 C	SMFDOOSL	8	char	MVS system level (for example: 'SP 4.2.0')
20 14	SMFDOSYN	8	char	MVS system name (SYSNAME from IEASYSxx)
28 1C	SMFDOINT	8	binary	Interval timestamp (STCK format of GMT) – non-interval records
36 24	SMFDOINO	8	binary	Local time offset for SMFDOINT field (STCK format) – non-interval records
44 2C	SMFDOIST	4	binary	Interval start time (local, in .01 seconds since midnight) – non-interval records
48 30	SMFDOISD	4	packed	Interval start date (in the form 0cyydddF, F is the sign)—non-interval records
52 34	SMFDOIET	4	binary	Interval end time (local, in .01 seconds since midnight) – non-interval records Note this time may differ from the value in SMFDOINT.
56 38	SMFDOIED	4	packed	Interval end date (in the form 0cyydddF, F is the sign) – non-interval records
60 3C	SMFDOSYP	8	char	Sysplex name

Subsystem Identification Section

The Subsystem Identification section appears in all records and provides information about the MAINVIEW Batch Optimizer subsystem and the Data Optimizer component. Subsystem Identification section field descriptions are shown in Table E-4.

Table E-4 Subsystem Identification Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDOSI	15	structure	Subsystem identification section
0 0	SMFDOSNM	4	char	Subsystem name - from START command
4 4	SMFDOSCC	8	char	Command control character, blank if none specified
12 C	SMFDOSST	1	binary	Subsystem status - 'Active'
13 D	SMFDOSTR	1	binary	Type of trace active BitDefinition 0Error (default) 1Function 2Flow
14 E	SMFDXST	1	binary	Cross-system mode BitDefinition 0Local Mode 1Cross-System Mode

Job Information Section

The Job Information section provides information about the job and step associated with the data set being processed by Data Optimizer. The Job Information section appears in all records. Job information section field descriptions are shown in Table E-5.

Table E-5 Job Information Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDOJI	68	structure	Job Information Section
0 0	SMFDOJBN	8	char	Job Name
8 8	SMFDOSTN	8	char	Step Name
16 10	SMFDOPSN	8	char	PROC Step Name
24 18	SMFDOSNO	2	binary	Step Number
26 1A	SMFDOJST	4	binary	JES Job Reader start time (in .01 seconds)
30 1E	SMFDOJSD	4	packed	JES Job Reader start date (0cyydddF)
34 22	SMFDOJNM	8	char	JES Job ID
42 2A	SMFDOJCL	1	char	JES Job Class
43 2B	SMFDOUSR	8	char	Userid
51 33	SMFDOGRP	8	char	Security System Group

Table E-5 Job Information Section Field Descriptions

Offset	Name	Length	Format	Description
59 3B	SMFDOJF1	1	binary	Address Space Type Flag
60 3C	SMFDOPGM	8	char	Program Name

Subsystem Processing Options Section

The Subsystem Processing Options section provides information about the Data Optimizer processing options in effect when the data set was processed. The Subsystem Processing Options section appears in all records. Subsystem Processing Options section field descriptions are shown in Table E-6.

Table E-6 Subsystem Processing Options Section Field Descriptions (Part 1 of 4)

Offset	Name	Length	Format	Description
0 0	SMFDOSO	188	structure	Subsystem processing options section
0 0	SMFDOSO1	1	binary	Processing options flag 1 BitDefinition 0dynamic region adjustment option active 1STC support 2TSO support 3MSGLEVEL0 4MSGLEVEL1 5MSGLEVEL2 6MSGLEVEL3 7Unused
1 1	SMFDOSO2	1	binary	Processing options flag 2 BitDefinition 0non-VSAM actual I/O counts 1shared write 2statistics report 3optimize tape data sets 4optimize PO data sets 5optimize EXCP access 6optimize DA data sets 7unused

Table E-6 Subsystem Processing Options Section Field Descriptions (Part 2 of 4)

Offset	Name	Length	Format	Description
2 2	SMFDOSO3 SMFDOP30 SMFDOP31 SMFDOP32 SMFDOP33 SMFDOP34 SMFDOP35 SMFDOP36 SMFDOP37	1	binary	Processing options flag 3 BitDefinition 0VSAM deferred write 1VSAM SHAREOPTIONS(3,3) 2disable LSR 3prefer LSR 4require LSR 531-bit control blocks 6FORCE LSR 7use hiperspace LSR buffers
3 3	SMFDOSO4 SMFDOP40 SMFDOP41 SMFDOP42	1	binary	Processing options flag 4 BitDefinition 0unused 1Basic mode 2Advanced mode
3 3	SMFDOSO4 SMFDOP40 SMFDOP41 SMFDOP42 SMFDOP43 SMFDOP44 SMFDOP45 SMFDOP46 SMFDOP47	1	binary	Processing options flag 4 BitDefinition 0unused 1Basic mode 2Advanced mode 3Data set info report only 4JCL requested interception 5user exit requested interception 6Data Policy requested interception 7Batch LSR was disabled
4 4	SMFDOSO5 SMFDOP50 SMFDOP51 SMFDOP52 SMFDOP53 SMFDOP54 SMFDOP55 SMFDOP56 SMFDOP57	1	binary	Processing options flag 5 BitDefinition 0Honor RUB 1unused 2unused 3unused 4unused 5unused 6unused 7unused
5 5	SMFDOSO6 SMFDOP60 SMFDOP61 SMFDOP62 SMFDOP63 SMFDOP64 SMFDOP65 SMFDOP66 SMFDOP67	1	binary	Processing options flag 6 BitDefinition 0override user values 1sequential insert strategy 2page-fix LSR IOBs 3multiple strings allowed for LSR 4allow LSR for alternate indexes 5override user's LSR definition 6disable Batch LSR 7unused

Table E-6 Subsystem Processing Options Section Field Descriptions (Part 3 of 4)

Offset	Name	Length	Format	Description
6 6	SMFDOSO7 SMFDOP70 SMFDOP71 SMFDOP72 SMFDOP73 SMFDOP74 SMFDOP75 SMFDOP76 SMFDOP77	1	binary	Processing options flag 7 BitDefinition 0unused 1unused 2Data set info report 3Reports to SYSOUT 4Messages to SYSOUT 5Options Report 6Catalog Info Report 7unused
7 7	SMFDOSO8 SMFDOP80 SMFDOP81 SMFDOP82 SMFDOP83 SMFDOP84 SMFDOP85 SMFDOP86 SMFDOP87	1	binary	Processing options flag 8 BitDefinition 0unused 1unused 2unused 3unused 4unused 5unused 6unused 7unused
8 8	SMFDOPTMS	4	binary	Time range start in minutes since midnight
12 C	SMFDOPTME	4	binary	Time range end in minutes since midnight
16 10	SMFDOBFNO	4	binary	BUFNO
20 14	SMFDONCP	4	binary	NCP
24 18	SMFDOBFNS	4	binary	BUFNO storage amount
28 1C	SMFDONCPS	4	binary	NCP storage amount
32 20	SMFDOMBFNO	4	binary	Maximum BUFNO
36 24	SMFDOMNCP	4	binary	Maximum NCP
40 28	SMFDOMBFNS	4	binary	Maximum BUFNO storage
44 2C	SMFDOMNCPS	4	binary	Maximum NCP storage
48 30	SMFDOS24R	4	binary	24-bit storage reserve
52 34	SMFDOS31R	4	binary	31-bit storage reserve
56 38	SMFDOS24M	4	binary	24-bit storage usage maximum
60 3C	SMFDOS31M	4	binary	31-bit storage usage maximum
64 40	SMFDOPGLM	4	binary	Paging limit
68 44	SMFDOCLM	4	binary	CPU limit
72 48	SMFDOMLDB	4	binary	Maximum LSR data buffers
76 4C	SMFDOMLIB	4	binary	Maximum LSR index buffers
80 50	SMFDOMLDS	4	binary	Maximum LSR data buffer storage
84 54	SMFDOMLIS	4	binary	Maximum LSR index buffer storage

Table E-6 Subsystem Processing Options Section Field Descriptions (Part 4 of 4)

Offset	Name	Length	Format	Description
88 58	SMFDOALDB	4	binary	Number of LSR data buffers
92 5C	SMFDOALIB	4	binary	Number of LSR index buffers
96 60	SMFDOALDS	4	binary	LSR data buffer storage to allocate
100 64	SMFDOALIS	4	binary	LSR index buffer storage to allocate
104 68	SMFDOPLIB	4	binary	Percent of index records for which to provide LSR buffers
108 6C	SMFDOMNDB	4	binary	Maximum NSR data buffers
112 70	SMFDOMNIB	4	binary	Maximum NSR index buffers
116 74	SMFDOMNDS	4	binary	Maximum NSR data buffer storage
120 78	SMFDOMNIS	4	binary	Maximum NSR index buffer storage
124 7C	SMFDOANDB	4	binary	Number of NSR data buffers
128 80	SMFDOANIB	4	binary	Number of NSR index buffers
132 84	SMFDOANDS	4	binary	NSR data buffer storage to allocate
136 88	SMFDOANIS	4	binary	NSR index buffer storage to allocate
140 8C	SMFDOVPCT	4	binary	VSAM optimization percent
144 90	SMFDONPCT	4	binary	Non-VSAM optimization percent
148 94	SMFDOARUB	4	binary	RUB value
152 98	SMFDOLINE	4	binary	Data Policy line number causing interception
156 9C	SMFDOMSGC	1	char	Message report JES output class
157 9D	SMFDOSTAC	1	binary	Reports JES output class
158 9E	SMFDOSTAC	8	binary	Data Policy member name from which options were used
166 A6	SMFDOLSRP	1	char	User-specified LSR pool ID
167 A7		1	char	Unused
168 A8	SMFDOHPCT	4	binary	Hiperspace buffer percent
172 AC	SMFDOHSTG	4	binary	Hiperspace buffer storage
176 B0	SMFDOMAXE	4	binary	Hiperspace cluster maximum expanded storage
180 B4	SMFDOESMX	4	binary	Hiperspace maximum system expanded storage
184 B8	SMFDOESMP	4	binary	Hiperspace maximum system expanded storage percent

General Data Set Information Section

The General Data Set Information section provides general information about the data set being processed by Data Optimizer. The General Data Set Information section appears in all records. General Data Set information section field descriptions are shown in Table E-7.

Table E-7 General Data Set Information Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDOGD	137	structure	General data set information section
0 0	SMFDODSN	44	char	Data set name (entity)
44 2C	SMFDODDN	8	char	DD name
52 34	SMFDOOPN	8	binary	Time the data set was opened in STCK format
60 3C	SMFDOCLS	8	binary	Time the data set was closed in STCK format
68 44	SMFDODTP SMFDOTP0 SMFDOTP1 SMFDOTP2 SMFDOTP3 SMFDOTP4 SMFDOTP5 SMFDOTP6 SMFDOTP7	1	binary	Data set type BitDefinition 0VSAM data set 1Non-VSAM data set 2unused 3unused 4unused 5unused 6unused 7unused
69 45	SMFDMGMT	8	char	SMS management class
77 4D	SMFDSTOR	8	char	SMS storage class
85 55	SMFDDATA	8	char	SMS data class
93 5D	SMFDODSNB	44	char	Data set name (base cluster)

DCB Information Section

The DCB Information section provides information from the DCB for the data set being processed by Data Optimizer. The DCB Information section appears for non-VSAM data sets only. DCB Information section field descriptions are shown in Table E-8.

Table E-8 DCB Information Section Field Descriptions (Part 1 of 3)

Offset	Name	Length	Format	Description
0 0	SMFDODI	25	structure	DCB information section
0 0	SMFDODMC	2	binary	DCBMACR field from the DCB

Table E-8 DCB Information Section Field Descriptions (Part 2 of 3)

Offset	Name	Length	Format	Description
0 0	SMFDODM1	1	binary	DCBMACR1 field from the DCB
1 1	SMFDODM2	1	binary	DCBMACR2 field from the DCB
2 2	SMFDODSO	2	binary	DCBDSORG field from the DCB
2 2	SMFDODS1	1	binary	DCBDSOR1 field from the DCB
3 3	SMFDODS2	1	binary	DCBDSOR2 field from the DCB
4 4	SMFDODLN	2	binary	DCBLRECL field from the DCB
6 6	SMFDODBS	2	binary	DCBBLKSI field from the DCB
8 8	SMFD OVCT	2	binary	Volume count
10 A	SMFDODNC	1	binary	DCBNCP field from the DCB
11 B	SMFDODBF	1	binary	DCBBUFNO field from the DCB
12 C	SMFDORCF	1	binary	DCBRECFM field from the DCB
13 D	SMFDOBFA	1	binary	DCBBFALN field from the DCB
14 E	SMFDOERO	1	binary	DCBEROPT field from the DCB
15 F	SMFDOUTY	1	binary	UCBTBYT4 from the UCB for the volume
16 10	SMFDOOOP	1	binary	Open option
17 11	SMFDOOCD	1	binary	DCBOPTCD field from the DCB
18 12	SMFDODFL SMFDOFL0 SMFDOFL1 SMFDOFL2 SMFDOFL3 SMFDOFL4 SMFDOFL5 SMFDOFL6 SMFDOFL7	1	binary	Flag byte BitDefinition 0DCBE existed 1SYNAD exit existed 2EODAD exit existed 3DCB open exit existed 4DCB abend exit existed 5DCB JFCB exit existed 6DCB JFCBE exit existed 7DCB EOVS exit existed
19 13	SMFD OVOL	6	char	Initial volume VOLSER
20 14	SMFDODF2	1	binary	Flag byte <u>BitDefinition</u> 0only BUFNO/NCP modified 1unused 2unused 3unused 4unused 5unused 6unused 7unused
21 15	SMFDOONC	1	binary	Original NCP value
22 16	SMFDOOBF	1	binary	Original BUFNO value

Table E-8 DCB Information Section Field Descriptions (Part 3 of 3)

Offset	Name	Length	Format	Description
25 19	SMFDOD2	1	binary	Flag byte BitDefinition 0only BUFNO/NCP modified 1unused 2unused 3unused 4unused 5unused 6unused 7unused
26 1A	SMFDOONC	1	binary	original NCP value
27 1B	SMFDOOBF	1	binary	original BUFNO value
28 1C	SMFDORNC	1	binary	recommended NCP value
29 1D	SMFDORBF	1	binary	recommended BUFNO value

ACB Information Section

The ACB Information section provides information from the ACB for the data set being processed by Data Optimizer. The ACB Information section appears only for VSAM data sets. ACB Information section field descriptions are shown in Table E-9.

Table E-9 ACB Information Section Field Descriptions (Part 1 of 2)

Offset	Name	Length	Format	Description
0 0	SMFDOAI	20	structure	ACB information section
0 0	SMFDOABI	2	binary	ACBBUFNI field from the ACB
2 2	SMFDOABD	2	binary	ACBBUFND field from the ACB
4 4	SMFDOABP	4	binary	ACBBUFSP field from the ACB
8 8	SMFDOABS	1	binary	ACBBSTNO field from the ACB
9 9	SMFDOAST	1	binary	ACBSTRNO field from the ACB
10 A	SMFDOAMC	2	binary	ACBMACRF field from the ACB
10 A	SMFDOAM1	1	binary	ACBMACR1 field from the ACB
11 B	SMFDOAM2	1	binary	ACBMACR2 field from the ACB
12 C	SMFDOAM3	1	binary	ACBMACR3 field from the ACB
13 D	SMFDOAOR	2	binary	ACBDSORG field from the ACB
13 D	SMFDOAO1	1	binary	ACBDSOR1 field from the ACB
14 E	SMFDOAO2	1	binary	ACBDSOR2 field from the ACB

Table E-9 ACB Information Section Field Descriptions (Part 2 of 2)

Offset	Name	Length	Format	Description
15 F	SMFDOAM4	1	binary	ACBMACR4 field from the ACB
16 10	SMFDOAOF	1	binary	ACBOFLGS field from the ACB
17 11	SMFDOAIF	2	binary	ACBINFLG field from the ACB
17 11	SMFDOAI1	1	binary	ACBINFL1 field from the ACB
18 12	SMFDOAI2	1	binary	ACBINFL2 field from the ACB
19 13	SMFDOSH P	1	binary	ACBSHRP field from the ACB
20 14	SMFDOAFL SMFDOAF0 SMFDOAF1 SMFDOAF2 SMFDOAF3 SMFDOAF4 SMFDOAF5 SMFDOAF6 SMFDOAF7	1	binary	Flag byte BitDefinition 0Converted ACB to LSR 1EODAD exit existed 2JRNAD exit existed 3UPAD exit existed 4LERAD exit existed 5SYNAD exit existed 6OPENAD exit existed 7ACB originally used LSR
21 15	SMFDOOBI	2	binary	original BUFNI value
23 17	SMFDOOBF	2	binary	original BUFND value
25 19	SMFDOAER	1	binary	ACBERFLG field from the ACB

Non-VSAM I/O Statistics Section

The Non-VSAM I/O Statistics section provides statistical I/O information for non-VSAM data sets being processed by Data Optimizer. The Non-VSAM I/O Statistics section appears only for non-VSAM data sets. Non-VSAM I/O Statistics section field descriptions are shown in Table E-10.

Table E-10 SMF Record Non-VSAM I/O Statistics Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDONI	12	structure	Non-VSAM I/O statistics information section
0 0	SMFDOLGB	4	binary	Logical blocks processed
4 4	SMFDOLEX	4	binary	Logical EXCPs intercepted
8 8	SMFDOPEX	4	binary	Physical EXCPs performed by Data Optimizer

VSAM I/O Statistics Section

The VSAM I/O Statistics section provides statistical I/O information for the VSAM data set being processed by Data Optimizer. The VSAM I/O Statistics section appears only for VSAM data sets. The VSAM I/O Statistics Information section field descriptions are shown in Table E-11.

Table E-11 SMF Record VSAM I/O Statistics Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDOVI	56	structure	VSAM I/O statistics information section
0 0	SMFDORN	4	binary	Number of random non-update reads
4 4	SMFDORUR	4	binary	Number of random update reads
8 8	SMFDOSNR	4	binary	Number of sequential non-update reads
12 C	SMFDOSUR	4	binary	Number of sequential update reads
16 10	SMFDORIN	4	binary	Number of random inserts
20 14	SMFDOSIS	4	binary	Number of sequential inserts
24 18	SMFDORUP	4	binary	Number of random updates
28 1C	SMFDOSUP	4	binary	Number of sequential updates
32 20	SMFDOSKP	4	binary	Number of skip sequential requests
36 24	SMFDOPNT	4	binary	Number of POINT requests
40 28	SMFDOEND	4	binary	Number of ENDREQ requests
44 2C	SMFDODEL	4	binary	Number of ERASE requests
48 30	SMFDOEXI	4	binary	Number of index EXCPs
52 34	SMFDOEXD	4	binary	Number of data EXCPs

VSAM LSR Statistics Section

The VSAM LSR Statistics section provides statistical LSR buffer pool information for the VSAM data set being processed by Data Optimizer. The VSAM LSR Statistics section appears only for VSAM data sets which have been processed using Local Shared Resources. The VSAM LSR Statistics Information section field descriptions are shown in Table E-12.

Table E-12 VSAM LSR Statistics Section Field Descriptions

Offset	Name	Length	Format	Description
0 0	SMFDOVL	40	structure	VSAM LSR statistics information section
0 0	SMFDOBFL	4	binary	Number of data buffer hits

Table E-12 VSAM LSR Statistics Section Field Descriptions

Offset	Name	Length	Format	Description
4 4	SMFDOBFR	4	binary	Number of data buffer reads
8 8	SMFDONNW	4	binary	Number of non-user writes
12 C	SMFDONUW	4	binary	Number of user writes
16 10	SMFDOMXS	4	binary	Maximum number of concurrent strings
20 14	SMFDOIFL	4	binary	Number of index buffer hits
24 18	SMFDOIFR	4	binary	Number of index buffer reads
28 1C	SMFDOHSB	4	binary	Number of hiperspace buffers
32 20	SMFDOHSR	4	binary	Number of hiperspace reads
36 24	SMFDOHSW	4	binary	Number of hiperspace writes

VSAM Catalog Information Section

The VSAM Catalog Information section provides catalog information for the VSAM data set being processed by Data Optimizer. The VSAM Catalog Information section appears only for VSAM data sets. The VSAM Catalog Information section field descriptions are shown in Table E-13.

Table E-13 VSAM Catalog Information Section Field Descriptions (Part 1 of 2)

Offset	Name	Length	Format	Description
0 0	SMFDOVC	152	structure	VSAM catalog information section
0 0	SMFDOIVL	6	char	First index VOLSER
6 6	SMFDODVL	6	char	First data VOLSER
12 C	SMFDOIVC	4	binary	Index volume count
16 10	SMFDODVC	4	binary	Data volume count
20 14	SMFDOICI	4	binary	Index CI splits
24 18	SMFDODCI	4	binary	Data CI splits
28 1C	SMFDOICA	4	binary	Index CA splits
32 20	SMFDODCA	4	binary	Data CA splits
36 24	SMFDOIAR	8	binary	Index highest allocated RBA
44 2C	SMFDODAR	8	binary	Data highest allocated RBA
52 34	SMFDOIUR	8	binary	Index high used RBA
60 3C	SMFDODUR	8	binary	Data high used RBA
68 44	SMFDOIXT	4	binary	No. of index extents
72 48	SMFDODXT	4	binary	No. of data extents

Table E-13 VSAM Catalog Information Section Field Descriptions (Part 2 of 2)

Offset	Name	Length	Format	Description
76 4C	SMFDOIRD	4	binary	No. of index reads
80 50	SMFDODRD	4	binary	No. of data reads
84 54	SMFDOIAD	4	binary	No. of index inserts
88 58	SMFDODAD	4	binary	No. of data inserts
92 5C	SMFDOIUP	4	binary	No. of index updates
96 60	SMFDODUP	4	binary	No. of data updates
100 64	SMFDOIDL	4	binary	No. of index deletes
104 68	SMFDODDL	4	binary	No. of data deletes
108 6C	SMFDOIRC	4	binary	No. of index records
112 70	SMFDODRC	4	binary	No. of data records
116 74	SMFDOICZ	4	binary	Index CI size (base cluster)
120 78	SMFDODCZ	4	binary	Data CI size (base cluster)
124 7C	SMFDOIXL	4	binary	No. of index levels
128 80	SMFDOMLR	4	binary	Maximum logical record length
132 84	SMFDOIVT	4	char	First index volume device type
136 88	SMFDODVT	4	char	First data volume device type
140 8C	SMFDOSHR	1	char	Inter-region SHAREOPTION
141 8D	SMFDOSHS	1	char	Inter-system SHAREOPTION
142 8E	SMFDOCF0 SMFDOCF1 SMFDOCF2 SMFDOCF3 SMFDOCF4 SMFDOCF5 SMFDOCF6 SMFDOCF7	1	binary	Flag byte BitDefinition 0ESDS organization 1KSDS organization 2RRDS organization 3Linear organization 4unused 5unused 6unused 7unused
143 8F		1	char	unused
144 90	SMFDOICZE	4	binary	Index CI size (entity)
148 94	SMFDODCZE	4	binary	Data CI size (entity)

Appendix F **HIPER-CACHE to Data Optimizer Conversion Utility**

This appendix describes the function of the HIPER-CACHE to Data Optimizer conversion utility and explains how to use it. This appendix discusses the following topics:

Overview	F-1
Defaults, Definitions, and JCL Overrides	F-2
Sample JCL	F-3
Required DD Statements	F-4
Optional DD Statements	F-5
Reports	F-6
Format	F-7
Report Sample	F-9
Conversion Procedure	F-17
HIPER-CACHE to Data Optimizer Conversion Tables	F-18

Overview

The function of the HIPER-CACHE Conversion Utility is to convert user HIPER-CACHE tables to a Data Optimizer Data Policy. The HIPER-CACHE tables are read, parsed, validated, and then used to build the equivalent Data Optimizer Data Policy. A report is generated highlighting the HIPER-CACHE tables that are processed and the resultant data policy.

Defaults, Definitions, and JCL Overrides

HIPER-CACHE applies the processing options by using a multi-level structure defined by the GLOBAL default table, the contents of a set of tables and JCL overrides.

The conversion utility builds Data Policy defaults from the HIPER-CACHE GLOBAL default table and the defines are built from the JOBNAME, PROGRAM, STORCLAS, DDNAME and DSNAME tables.

The conversion program will not attempt to generate defines to emulate the effect of HIPER-CACHE multi-level structure; HIPER-CACHE options are converted on a one-for-one basis.

HIPER-CACHE options that cannot be converted are highlighted on the conversion report. This may indicate that the associated function is not supported by Data Optimizer or that the function is provided through a different mechanism. For more information, see “HIPERCACHE to Data Optimizer Conversion Utility Required DD Statements” on page F-4.

Users should review the conversion program reports and Data Policy to ensure that the correct dataset will be selected for optimization and that the desired Data Optimizer options will be used. To assist your review, the original HIPER-CACHE statement will be converted into a comment and inserted into the Data Policy prior to the associated Definition.

Data Optimizer will run in the equivalent of HIPER-CACHE explicit mode to minimize the possibility of optimizing the wrong data sets.

The following tables will not be converted:

- HIPER-CACHE PROTECTED program table

The PROTECTED program table is unique to HIPER-CACHE and reflects product-specific restrictions. Data Optimizer has an internal PROTECTED program table.

The PROTECTED program table will be scanned for any user additions. User additions will not be converted, but users may want to consider adding them to the Data Optimizer Data Policy.

- TIME-OF-DAY table

The conversion report will highlight the following:

- HIPER-CACHE options that failed validation.

These options will not be converted. During startup, HIPER-CACHE ignores any options with validation errors. The conversion program validation logic is based on HIPER-CACHE code and will trap any errors and generate an error messages.

- HIPER-CACHE options that cannot be converted, because Data Optimizer does not have an equivalent option.
- HIPER-CACHE processing control table entries that can only be partially converted because one or more of the options specified cannot be converted.
- Non-printing comment lines starting with a percent sign will be reported but will not be converted, comment lines starting with a single or double asterisks will be reported and converted.

Reporting and conversion of all comment lines can be suppressed by using the DONOCOM DD dummy statement.

Sample JCL

Figure F-1 on page F-4 shows the sample JCL for the Conversion Utility that is shipped with Data Optimizer. The sample JCL can be found in the DAPHCNVR member of BMC.BSS.SAMPLIB.

You must make the following changes to the JCL to run the utility:

- change *?USERID?* to the user ID to be used for this job
- change *?DOLOAD?* to the name of the load library containing the Data Optimizer load modules
- change *?HCTABLES?* to the name of the PDS containing the HIPER-CACHE tables to be converted
- change *?CONTROLDATASET?* to the name of the Control Data Set
- change *?DATAPOLMEM?* to the name of the desired Data Policy member name
- activate optional DD statements as needed

Figure F-1 HIPER-CACHE to Data Optimizer Conversion Utility Sample JCL

```
//DAPHCNRV JOB (ACCT), 'NAME', <== MODIFY
//          MSGCLASS=X, <== CHECK
//          REGION=2048K,
//          NOTIFY=?USERID?, <== MODIFY
//          CLASS=A <== CHECK
//*
//* THIS JOB WILL CONVERT HIPER-CACHE TABLES TO DATA OPTIMIZER
//* DATA POLICY DEFINITIONS
//*
//* PLEASE REFER TO THE DOCUMENTATION FOR DETAILED INSTRUCTIONS
//* ON HOW TO RUN THE CONVERSION UTILITY.
//*
//HCONV EXEC PGM=DAPHCNV
//STEPLIB DD DISP=SHR,DSN=?DOLOAD? <== MODIFY
//VHC$GLBL DD DISP=SHR,DSN=?HCTABLES?(GLOBALS) <== MODIFY
//DOOPTNS DD DISP=SHR,DSN=?CONTROLDATASET?(?DATAPOLMEM?) <== MODIFY
//SYSPRINT DD SYSOUT=*
//*
//* THE FOLLOWING DD STATEMENTS ARE OPTIONAL
//*
//*VHC$JOB DD DISP=SHR,DSN=?HCTABLES?(JOBNAMES) <== CHECK
//*VHC$DDN DD DISP=SHR,DSN=?HCTABLES?(DDNAMES) <== CHECK
//*VHC$DSN DD DISP=SHR,DSN=?HCTABLES?(DSNAMES) <== CHECK
//*VHC$PGM DD DISP=SHR,DSN=?HCTABLES?(PROGRAM) <== CHECK
//*VHC$PRO DD DISP=SHR,DSN=?HCTABLES?(PROTECTD) <== CHECK
//*VHC$STO DD DISP=SHR,DSN=?HCTABLES?(STORCLAS) <== CHECK
//*VHC$TOD DD DISP=SHR,DSN=?HCTABLES?(TIMEODAY) <== CHECK
//*DOTRACE DD DUMMY <== CHECK
//*DONOCOM DD DUMMY <== CHECK
//
```

Required DD Statements

The following DD statements are required:

Table F-1 HIPERCACHE to Data Optimizer Conversion Utility Required DD Statements

DD Statement	Description
VHC\$GLBL	points to a PDS member containing the HIPER-CACHE global default table
SYSPRINT	points to SYSOUT or a sequential dataset that is to contain the HIPER-CACHE to Data Optimizer Conversion report
DOOPTNS	points to a PDS member that is to contain the data policy created from the HIPER-CACHE tables

Optional DD Statements

The following DD statements are optional:

Table F-2 HIPERCACHE to Data Optimizer Conversion Utility Optional DD Statements

DD Statement	Description
VHC\$JOB	points to a PDS member containing the HIPER-CACHE Job table The table will be converted and reported on the Conversion report.
VHC\$DDN	points to a PDS member containing the HIPER-CACHE DDNAME table The table will be converted and reported on the Conversion report.
VHC\$DSN	points to a PDS member containing the HIPER-CACHE DSNNAME table This table will be converted and reported on the Conversion report.
VHC\$PGM	points to a PDS member containing the HIPER-CACHE Program table The table will be converted and reported on the Conversion report.
VHC\$PRO	points to a PDS member containing the HIPER-CACHE Protected Program Table The Protected Program table will not be converted, but the Conversion report will highlight all user additions to the table.
VHC\$STO	points to a PDS member containing the HIPER-CACHE STORCLAS (SMS Storage Class) table The table will be converted and highlighted on the Conversion report.
VHC\$TOD	points to a PDS member containing the HIPER-CACHE TIME OF DAY table The TIME OF DAY table will not be converted. The contents of this table will be reported on the Conversion report.
DOTRACE	turns on the Trace Facility This facility will report how each HIPER-CACHE table entry is converted into Data Optimizer options. Each table will appear on the Conversion Report followed by the associated Data Optimizer options.
DONOCOM	turns on the Comment Suppression Facility This facility will suppress the reporting and conversion of comment lines. Comments on HIPER-CACHE option statements will always be reported and converted.

Table F-3 lists the valid return codes for the conversion utility.

Table F-3 Conversion Utility Return Codes

Return Code	Definition
0	execution completed successfully, no errors or warnings were encountered
4	program identified one or more HIPER-CACHE options that are not supported by Data Optimizer; warning messages were generated.
8	program identified one or more HIPER-CACHE options with validation errors; error messages were generated

Reports

The HIPER-CACHE Conversion report contains the following sections:

- Global Default table
- DDNAME table
- DSNNAME table
- JOBNAME table
- PROGRAM table
- STORCLAS table
- PROTECTED table
- TIME OF DAY table
- Data Policy definitions

The sections of the report contain the following information:

- The Global Default table and Data Policy definitions sections are always created.

The other sections are optional and are only created if the associated HIPER-CACHE tables are selected to be converted.

- HIPER-CACHE tables can be selected to be converted by allocating the tables in the conversion program execution JCL.
- The Global Default table section highlights how HIPER-CACHE global defaults are converted to Data Optimizer data policy defaults.
- The DDNAME, DSNNAME, JOBNAME, PROGRAM and STORCLAS table sections highlight how HIPER-CACHE options are converted to Data Optimizer data policy definitions.

- The PROTECTED table section highlights user additions to the protected program table. HIPER-CACHE and Data Optimizer each have unique protected program tables. The contents of the protected program tables are not converted. Users can convert the user additions by manually adding DEFINES to the Data Policy.
- The TIME-OF-DAY table section is optional and highlights the contents of the HIPER-CACHE time-of-day table. Time-of-day table entries are not converted.
- The Data Policy definitions' section highlights the Data Policy definitions created by the conversion program.

Format

Figure F-2 shows the report detail line for the tables sections. This line consists of a HIPER-CACHE table entry record and a status description.

Figure F-2 Report Detail Line

```

CONVERSION
=====
DATE: 05/02/2000
=====
                                PAGE: 4
                                =====
                                =====
CONVERTED          SUPPD,DYNR          DYNAMIC REG
CONVERTED          SUPPD,NODYNR        DYNAMIC REG
CONVERTED          SUPPD,LSR          LSR SUPPORT
CONVERTED          SUPPD,NOLSR        LSR SUPPORT
CONVERTED          SUPPD,XABUF        VSAM XA BUF

<---status description--><-----Hiper-Cache Table
Entry----->

```

Table F-4 lists the statuses that can be returned via the HIPER-CACHE Conversion report.

Table F-4 Report Status Descriptions

Status	Description
CONVERTED	HIPER-CACHE table entry was successfully converted.
PARTIALLY CONVERTED	HIPER-CACHE table entry was partially converted; one or more of the options specified are not supported by Data Optimizer.
NOT CONVERTED	HIPER-CACHE option could not be converted. This status will appear when Data Optimizer does not have an equivalent option. This may mean that the associated function is not present in Data Optimizer or that the function is provided through a different mechanism. Refer to "HIPER-CACHE Global Process Options Not Supported by Data Optimizer" on page F-22 for specific information.
CHECK CONTINUATION	Status indicates that the associated HIPER-CACHE table entry is part of a set. The last entry in the set will show the status for the set.
USER CHANGE	A user change was found in the HIPER-CACHE Protected Program table
ERROR	An error was found in the HIPER-CACHE table entry. Error messages are generated by the HIPER-CACHE parsing and validation logic and will appear on the next line of the report, see the HIPER-CACHE User's Guide for an explanation of the messages
WARNING	HIPER-CACHE parsing and validation logic, or the parameter conversion logic, found a situation that may need attention An explanation of the warning will appear on the next line
DO TRACE.....>	Identifies a trace entry created by the trace facility The trace facility is activated using the //DOTRACE DD DUMMY statement. This facility reports how each HIPER-CACHE table entry is converted into Data Optimizer Data Policy definitions.

The format of the tables section can be changed by using the trace and suppress comment facilities. As shown in Figure F-3, these facilities are activated by using dummy DD statements in the HIPER-CACHE conversion program Job step.

Figure F-3 Examples of TRACE and Comment Suppression Facilities

```

CONVERSION
=====
DATE: 05/01/2000
          PAGE: 4
=====
          =====
          =====
          CONVERTED          SUPPD , DYNR          DYNAMIC
REGION SIZE ADJUSTMENT
DO TRACE.....>          SUPPD , DYNR          DYNAMIC
REGION SIZE ADJUSTMENT
DO TRACE.....>
DO TRACE.....>          DEFINE
DO TRACE.....>          PGM=SUPPD
DO TRACE.....>          INCLUDE
DO TRACE.....>          ALLFILES
DO TRACE.....>          DRA
DO TRACE.....>

```

Figure F-4 shows a comment block that highlights the program return code and suggests appropriate action.

Figure F-4 Conversion Report Comment Block

```

*****
*****
**
** THE CONVERSION PROGRAM HAS COMPLETED WITH A RETURN CODE OF 8.
**
** A RETURN CODE OF 8 INDICATES THAT ONE OR MORE OF THE HIPER-CACHE OPTIONS
** HAS VALIDATION ERRORS, THE REPORT MAY ALSO HIGHLIGHT HIPER-CACHE OPTIONS
** THAT ARE NOT SUPPORTED BY DATA OPTIMIZER.
**
** PLEASE REVIEW THE REPORT CAREFULLY AND THEN MAKE ANY NECESSARY CHANGES TO
** THE DATA OPTIMIZER OPTIONS.
**
*****
*****

```

Report Sample

Figure F-5 shows a sample Conversion Report.

Figure F-5 Conversion Report Sample (Part 1 of 4)

CONVERSION		HIPER-CACHE TO DATA OPTIMIZER
=====		
DATE: 09/05/2000		GLOBAL DEFAULT
TABLE	PAGE: 1	=====
	=====	=====
CONVERTED	DFR-SUPPORT=YES	DEFERRED
WRITES IN LSR SUPPORT		
CONVERTED	DYNAMIC-REGION=AUTO	DYNAMIC
REGION ADJUSTMENT		
CONVERTED	LSR-SUPPORT=YES	LSR SUPPORT
CONVERTED	MESSAGES=NO	HIPER-CACHE
MESSAGES		
CONVERTED	MESSAGES-ROUTE=WTO	MESSAGES
ROUTING TYPE		
CONVERTED	MESSAGES-CLASS=*	JES CLASS
FOR MESSAGES		
CONVERTED	OPTPCT-VSAM-LSR-INDEX=100	VSAM LSR
INDX COMPONENT OPT. PCT.		
NOT SUPPORTED	PROSMS-INTERFACE=NO	PROSMS
INTERFACE OPTION		
CONVERTED	RESERVED-STORAGE=1024	24-BIT
STORAGE RESERVE (IN KB)		
CONVERTED	SMF-RECTYPE=213	SMF RECORD
NUMBER		
CONVERTED	SMF-STATISTICS=NO	SMF
STATISTICS OPTION		
CONVERTED	TSO-SUPPORT=NO	TSO SESSION
SUPPORT		
CONVERTED	VSAM-OPTIMIZER=EXPLICIT	VSAM
OPTIMIZATION TYPE		
CONVERTED	STATISTICS=YES	HIPER-CACHE
STATISTICS OPTION		
CONVERTED	STATISTICS-ROUTE=WTO	STATISTICS
ROUTING OPTION		
CONVERTED	STATISTICS-CLASS=*	STATISTICS
JES CLASS		
NOT SUPPORTED	ACCESS-RESYNC=(0009,67)	ACCESS
RESYNC		
SUPPORTED	AIX-SUPPORT=YES	ALTERNATE
INDEX SUPPORT		
NOT SUPPORTED	DYNAMIC-ACCESS=AUTO	"ACCESS IS
DYNAMIC" SUPPORT		
SUPPORTED	EXPANDED-STORAGE-MAX=010	MAXIMUM
EXPANDED STORAGE USAGE		
SUPPORTED	HIPERSPACE=NO	HIPERSPACE
SUPPORT		

SUPPORTED	HIPERSPACE-PERCENT=025	% OF
DATASET INTO HIPERSPACE		
NOT SUPPORTED	MESSAGES-TYPE=VERBOSE	MESSAGES
VOLUME INDICATOR		
SUPPORTED	MULTISTRING=NO	MULTISTING
PROTECTION OPTION		
NOT SUPPORTED	OPTPCT-SAM=050	SAM
OPTIMIZATION PERCENTAGE		
NOT SUPPORTED	OPTPCT-SAM-XA=050	XA-SAM
OPTIMIZATION PERCENTAGE		
NOT SUPPORTED	OPTPCT-VSAM-LSR-DATA=100	VSAM LSR
DATA COMPONENT OPT. PCT.		
NOT SUPPORTED	OPTPCT-VSAM-NSR-DATA=100	VSAM NSR
DATA COMPONENT OPT. PCT.		
NOT SUPPORTED	OPTPCT-VSAM-NSR-INDEX=100	VSAM NSR
INDX COMPONENT OPT. PCT.		
SUPPORTED	OVERRIDE-LSR=NO	OVERRIDE
USER LSR SPECIFICATION		
SUPPORTED	PAGE-FIX=NO	PAGE-FIX
BUFFERS OPTION		
CONVERTED	SAM-OPTIMIZER=EXPLICIT	SAM
OPTIMIZATION TYPE		
NOT SUPPORTED	SHARE-OPTIONS-4=NO	SHR(4,X)
DATASET SUPPORT		
SUPPORTED	SIS-SUPPORT=NO	
SEQUENTIAL-INSERT-STRATEGY		
NOT SUPPORTED	STATS-EXTENDED=NO	EXTENDED
STATISTICS OPTION		
NOT SUPPORTED	XA-BUFFERS-VSAM=YES	VSAM XA
BUFFER SUPPORT		
NOT SUPPORTED	XA-REGION-SIZE=128	XA-REGION
SIZE VALUE (IN MEG)		

HIPER-CACHE TO DATA OPTIMIZER

CONVERSION

```

=====
DATE: 09/05/2000
PAGE: 2
=====
=====

```

JOBNAME TABLE

=====

CONVERTED	MYPGM1 , NODFR , VSAM , SAM
CONVERTED	MYPGM2 , DYNR , VSAM , SAM
CONVERTED	MYPGM3 , LSR , VSAM , SAM
CONVERTED	MYPGM4 , MSGCLS=F , VSAM , SAM
CONVERTED	MYPGM5 , SAM
CONVERTED	MYPGM6 , NOSAM

Figure F-5 Conversion Report Sample (Part 2 of 4)

```

                                                    HIPER-CACHE TO DATA OPTIMIZER
CONVERSION

=====
DATE: 09/05/2000                                DATA POLICY
DEFINITIONS                                     PAGE:    3
=====
=====

                                DATAPOLICY
                                DATAPOLICYEND

* > -----*
                                * > DATA POLICY DEFAULTS AND DEFINES
*
* > -----*
```

DEFAULTS

- LSR
- DRA
- STATRPT
- NOOPTIONRPT
- NODSINFORPT
- NOCATRPT
- MSGLEVEL0
- MSGDEST=SYSOUT
- MSGCLASS=*
- RPTDEST=SYSOUT
- RPTCLASS=*
- SMFRECTYPE=000
- STG24RESERVE=1024K
- PCTLSRIBUF=100
- NOTSO
- VDW
- NOAIXLSR
- STGEMAXPCT=010
- HIPERPCT=025
- NOMULTISTRING
- NOLSROVERRIDE
- DISABLEBATCHLSR=NO
- NOHIPERSPACE
- NOPAGEFIX
- NOSIS
- EXCLUDE
- ADVANCED
- RUB=99
- NORUBHONOR

NOSTC
NOSHAREDWRITE
NOAIC
NOOVERRIDEUSER
BUFNO=0
BUFNOSTG=0
MAXBUFNO=0
MAXBUFNOSTG=0
NCP=0
NCPSTG=0
MAXNCP=0
MAXNCPSTG=0
PAGINGLIMIT=0
CPULIMIT=101
STG31RESERVE=0
STG24MAX=0
STG31MAX=0
NOSHROPT33
NOVSAMCB31
NSRIBUF=0
NSRISTG=0
MAXNSRIBUF=0
MAXNSRISTG=0
STGEMAX=0

Figure F-5 Conversion Report Sample (Part 3 of 4)

```

CONVERSION
HUPER-CACHE TO DATA OPTIMIZER

=====
DATE: 09/05/2000
DEFINITIONS
=====
PAGE: 4
DATA POLICY
=====

NSRDBUF=0
NSRDSTG=0
MAXNSRDBUF=0
MAXNSRDSTG=0
LSRIBUF=0
LSRISTG=0
MAXLSRIBUF=0
MAXLSRISTG=0
LSRDBUF=0
LSRDSTG=0
MAXLSRDBUF=0
MAXLSRDSTG=0
NONVSAMOPTPCT=100
VSAMOPTPCT=100

*>MYPGM1 , NODFR , VSAM , SAM

DEFINE
JOBNAME=MYPGM1
INCLUDE
ALLFILES
NOVDW

*>MYPGM2 , DYNR , VSAM , SAM

DEFINE
JOBNAME=MYPGM2
INCLUDE
ALLFILES
DRA

*>MYPGM3 , LSR , VSAM , SAM

DEFINE
JOBNAME=MYPGM3
INCLUDE
ALLFILES
LSR

*>MYPGM4 , MSGCLS=F , VSAM , SAM

```

```
DEFINE
  JOBNAME=MYPGM4
  INCLUDE
  ALLFILES
  MSGCLASS=F

*>MYPGM5 , SAM

DEFINE
  JOBNAME=MYPGM5
```

Figure F-5 Conversion Report Sample (Part 4 of 4)

```

                                                                    HIPER-CACHE TO DATA OPTIMIZER
CONVERSION
=====
DATE: 09/05/2000
DEFINITIONS
=====
                                                                    DATA POLICY
                                                                    =====
                                                                    =====

                                INCLUDE
                                NONVSAM

                                DEFINE
                                JOBNAME=MYPGM5
                                EXCLUDE
                                VSAM

                                *>MYPGM6 ,NOSAM

                                DEFINE
                                JOBNAME=MYPGM6
                                EXCLUDE
                                ALLFILES

*>-----*
                                *> DEFINE(S) TO ENFORCE HIPER-CACHE DEFAULTS
*
*>-----*

                                DEFINE
                                JOBNAME=*
                                EXCLUDE
                                ALLFILES

*****
*****
*****
*****
**
**
** THE CONVERSION PROGRAM HAS COMPLETED WITH A RETURN CODE OF 4 .
**
**
** A RETURN CODE OF 4 INDICATES THAT ONE OR MORE OF THE HIPER-CACHE
OPTIONS **
```

```

** BEING USED ARE NOT SUPPORTED BY DATA OPTIMIZER.
**
**
**
** PLEASE REVIEW THE REPORT CAREFULLY AND THEN MAKE ANY NECESSARY CHANGES
TO **
** THE DATA OPTIMIZER OPTIONS.
**
**
**
*****
*****
*****
*****

```

Conversion Procedure

The following procedure provides step-by-step instructions for converting HIPER-CACHE tables to a Data Optimizer data policy:

- Step 1** Identify the current HIPER-CACHE tables.
- Step 2** Identify the HIPER-CACHE tables that you want to convert:
- The GLOBAL default table entries must be converted.
 - The JOBNAME, PROGRAM, STORCLAS, DDNAME and DSNAME tables are optional and should be converted if they contain HIPER-CACHE options.
 - The PROTECTED table will not be converted, but the conversion program will highlight any user additions or changes.
 - The TIME-OF-DAY table will not be converted.
- Step 3** Review the HIPER-CACHE tables, correct any errors, and remove any redundant or out of date options.
- Step 4** Review the comments in the non-global tables. Add comments where necessary. Comments are copied into the data policy and can help identify the function of the DEFINES after the conversion.
- Step 5** Copy the conversion program JCL from DAPHCNVR member of BMC.BSS.SAMPLIB, review the instructions contained in the JCL, then make the necessary changes. For more information, see “Sample JCL” on page F-3.

Step 6 Run the conversion program.

Step 7 Review the output carefully, and make any necessary changes to the data policy.

Note: Review the conversion program reports and Data Policy to ensure that the correct dataset will be selected for optimization and that the desired options will be used.

- To prevent any potential conflicts with any existing HIPER-CACHE SMF records, SMF records will not be created. SMF recording can be activated by updating the SMFRECTYPE= default to use an assigned SMF record type.
- Started task support will be deactivated. Started task support can be activated by changing the NOSTC default to STC.
- Data-optimizer runs in the equivalent of HIPER-CACHE's EXPLICIT mode. Optimization is controlled by DEFINE EXCLUDE/INCLUDE statements.

Step 8 Convert HIPER-CACHE exit code to run with Data Optimizer.

Step 9 Use the MAINVIEW Batch Optimizer user interface or an operator command to activate the data policy.

HIPER-CACHE to Data Optimizer Conversion Tables

The following tables define how the HIPER-CACHE Global processing options and Processing Control table options are converted.

Table F-5 lists the HIPER-CACHE Global Process options that Data Optimizer supports.

Table F-5 HIPER-CACHE Global Process Options Supported by Data Optimizer (Part 1 of 3)

HIPER-CACHE	Definition	Data Optimizer
AIX-SUPPORT=YES/NO	alternate index support	AIXLSR/NOAIXLSR
DFR-SUPPORT=YES/NO	deferred writes in LSR support	YES = VDW NO = NOVDW
DYNAMIC-REGION=YES/NO/AUTO	dynamic region adjustment	YES = DRA NO = NODRA
EXPANDED-STORAGE-MAX=NNN	maximum expanded storage usage	STGEMAXPCT=

Table F-5 HIPER-CACHE Global Process Options Supported by Data Optimizer (Part 2 of 3)

HIPER-CACHE	Definition	Data Optimizer
FORCE-LSR=YES/NO	override user LSR specification	FORCELSR
HIPERSPACE-PERCENT=NNN	% of dataset into HIPERSPACE	HIPERPCT=
HIPERSPACE=YES/NO	HIPERSPACE support	HIPERSPACE/NOHIPERSPACE
LSR-SUPPORT=YES/NO	LSR support	YES = LSR NO = NOLSR
MESSAGES-CLASS=X	JES class for messages	MSGCLASS=X
MESSAGES-ROUTE=JES/WTO	message route	MSGDEST=SYSOUT
MESSAGES=YES/NO/ALL	HIPER-CACHE messages	NO = MSGLEVEL0 YES = MSGLEVEL2 ALL = MSGLEVEL2
MULTISTRING=YES/NO	multistring protection option	MULTISTRING/NOMULTISTRING
OPTPCT-VSAM-LSR-INDEX=NNN	VSAM LSR index component optimization percentage	PCTLSRIBUF=NNN MAX VALUE IS 100%
OVERRIDE-LSR=YES/NO	override user LSR specification	LSROVERRIDE/NOLSOVERRIDE DISABLEBATCHLSR=YES/NO
PAGE-FIX=YES/NO	Page-fix Buffers option	PAGEFIX/NOPAGEFIX
RESERVED-STORAGE=NNNN	24-bit storage reserve (in KB)	STG24RESERVE=NNNNK
SAM-OPTIMIZER=EXPLICIT/IMPLICIT	SAM optimization type	NOT SUPPORTED BY DATA OPTIMIZER DEFAULTS. THE OPTION IS PROPOGATED INTO THE DATA OPTIMIZER DEFINES
SIS-SUPPORT=YES/NO	sequential-insert-strategy	SIS/NOSIS
SMF-RECTYPE=213	SMF record number	SMFRECTYPE=000
SMF-STATISTICS=YES/NO/ONLY	SMF Statistics option	SMFRECTYPE=000
STATISTICS-CLASS=X	statistics JES class	RPTCLASS=X
STATISTICS-ROUTE=JES/WTO	Statistics Routing option	RPTDEST=SYSOUT
STATISTICS=YES/NO/ALL	HIPER-CACHE Statistics option	YES = STATRPT NOPTIONRPT NODSINFORPT NOCATRPT NO = NOSTATRPT NOPTIONRPT NODSINFORPT NOCATRPT ALL = STATRPT OPTIONRPT DSINFORPT CATRPT

Table F-5 HIPER-CACHE Global Process Options Supported by Data Optimizer (Part 3 of 3)

HIPER-CACHE	Definition	Data Optimizer
TSO-SUPPORT=YES/NO	TSO session support	YES = TSO NO = NOTSO
VSAM-OPTIMIZER=EXPLICIT/IMPLICIT	VSAM optimization type	NOT SUPPORTED AS A DATA OPTIMIZER DEFAULT. THE OPTION IS PROPOGATED INTO THE DATA OPTIMIZER DEFINES

Table F-6 lists the HIPER-CACHE Processing Control Table options that are supported by Data Optimizer.

Table F-6 HIPER-CACHE Processing Control Table Options Supported by Data Optimizer (Part 1 of 2)

HIPER-CACHE	Definition	Data Optimizer
AIX/NOAIX	alternate index support	AIXLSR/NOAIXLSR
DFR/NODFR	deferred writes in LSR support	DFR = VDW NODFR = NOVW
DYNR/NODYNR	dynamic region adjustment	DYNR = DRA NODYNR = NODRA
FLSR/NOFLSR	force data sets into LSR	FORCELSR
HSPC/NOHSPC	HIPERSPACE support	HIPERSPACE/NOHIPERSPACE
HSPC=	percent of dataset into HIPERSPACE	HIPERPCT=
LSR/NOLSR	LSR support	LSR = LSR NOLSR = NOLSR
MSGCLS=X	JES class for messages	MSGCLASS=X
MSGJES/MSGWTO	message route	MSGDEST=SYSOUT
MSGS/NOMSGS	HIPER-CACHE messages	MSG = MSGLEVEL2 NOMSG = MSGLEVEL0
MSTR/NOMSTR	Multistring Protection option	MULTISTRING/NOMULTISTRING
OLSR/NOOLSR	Override User LSR Specification	LSROVERRIDE/NOLSR OVERRIDE DISABLEBATCHLSR=YES/NO
PFIX/NOFIX	Page-fix Buffers option	PAGEFIX/NOPAGEFIX
OVLI=NNN	VSAM LSR index component optimization percentage	PCTLSRIBUF=NNN MAX VALUE IS 100%
RSTG=NNNN	24-bit storage reserve (in KB)	STG24RESERVE=NNNNK
SAM/NOSAM	optimization SAM EXPL/IMPL	GLOBAL AND TABLE OPTIONS FOR SAM AND VSAM ARE EVALUATED TOGETHER, POSSIBLE RESULTS ARE INCLUDE/EXCLUDE FOR VSAM/NONVSAM/ALLFILES
SIS/NOSIS	sequential insert strategy	SIS/NOSIS

Table F-6 HIPER-CACHE Processing Control Table Options Supported by Data Optimizer (Part 2 of 2)

HIPER-CACHE	Definition	Data Optimizer
SMF/NOSMF/SMFONLY	SMF Statistics option	SMFRECTYPE=000
STAT/NOSTAT	HIPER-CACHE Statistics option	STAT= STATRPT OPTIONRPT DSINFORPT CATRPT NOSTAT= NOSTATRPT NOOPTIONRPT NODSINFORPT NOCATRPT
STATCLS=X	statistics JES class	RPTCLASS=X
STATJES/STATWTO	Statistics Routing option	RPTDEST=SYSOUT
VBUF=	obsolete	CONVERTED AS OVLI=NNN
VSAM/NOVSAM	optimize VSAM EXPL/IMPL	GLOBAL AND TABLE OPTIONS FOR SAM AND VSAM ARE EVALUATED TOGETHER, POSSIBLE RESULTS ARE INCLUDE/EXCLUDE FOR VSAM/NONVSAM/ALLFILES

Table F-7 lists the HIPER-CACHE Global Process options that are not converted, but the associated functions are provided through a different mechanism.

Table F-7 HIPER-CACHE Global Process Options Not Converted

HIPER-CACHE Option	Purpose
ACCESS-RESYNC=(NNNN,PP)	optimization technique evaluator
AUTOMATIC-REBLOCKING=YES/NO	SAM data set reblocking
BLOCKS-PER-TRACK=NN	reblocking blocks per track
DYNAMIC-ACCESS=YES/NO/AUTO	"ACCESS IS DYNAMIC" support
MESSAGES-TYPE=VERBOSE/TERSE	volume of messages to produce
OPTPCT-SAM-XA=NNN	XA-SAM optimization percentage
OPTPCT-SAM=NNN	SAM optimization percentage
OPTPCT-VSAM-LSR-DATA=NNN	VSAM LSR data component optimization percentage
OPTPCT-VSAM-NSR-DATA=NNN	VSAM NSR data component optimization percentage
OPTPCT-VSAM-NSR-INDEX=NNN	VSAM NSR index component optimization percentage
PAGING-CONTROL=XXXX	paging control
STATS-EXTENDED=YES/NO	Extended Statistics option
USER-EXIT=NONE/NNNNNNNN	user exit program
XA-BUFFERS-SAM=YES/NO	SAM XA buffer support

Table F-7 HIPER-CACHE Global Process Options Not Converted

HIPER-CACHE Option	Purpose
XA-BUFFERS-VSAM=YES/NO	VSAM XA buffer support
XA-REGION-SIZE=NNN	XA region size specification

Table F-8 lists the HIPER-CACHE Processing Control table options that are not converted, but the associated functions are provided through a different mechanism.

Table F-8 HIPER-CACHE Processing Control Table Options Not Converted

HIPER-CACHE Option	Purpose
BLKSTRK=	reblocking blocks per track
DYNA/NODYNA	ACCESS IS DYNAMIC Support
EMAX=	maximum expanded storage usage
MSGVERB/MSGTERS	volume of messages to be produced
OVLD=	VSAM LSR data component optimization percentage
OVND=	VSAM NSR data component optimization percentage
OVNI=	VSAM NSR index component optimization percentage
REBLK/NOREBLK	SAM dataset reblocking
SBUF=	% optimization for SAM
STEX/NOSTEX	Extended Statistics option
SXBUF=	XA-SAM optimization percentage

Table F-9 lists the HIPER-CACHE Global Process options that are not supported by Data Optimizer.

Table F-9 HIPER-CACHE Global Process Options Not Supported by Data Optimizer (Part 1 of 2)

HIPER-CACHE Option	Purpose
PROSMS-INTERFACE=YES/NO/ONLY	activate interface to PROSMS
SAM-FREEPOOL=YES/NO	freepool issuance for Q/BSAM
SHARE-OPTIONS-4=YES/NO	SHR(4,X) dataset support ¹
SVHC\$OFF=SVHC\$OFF	Extended Statistics option
STATS-EXTENDED=YES/NO	Extended Statistics option
SVHC\$OFF=SVHC\$OFF	DDname to turn SAM optimization off
SVHCON=SVHCON	DDname to turn SAM optimization on
VVHC\$IOS=VVHC\$IOS	DDname to turn statistics on
VVHC\$M\$S=VVHC\$M\$S	DDname to turn statistics and messages on
VVHC\$MESSG=VVHC\$MESSG	DDname to turn messages on

Table F-9 HIPER-CACHE Global Process Options Not Supported by Data Optimizer (Part 2 of 2)

HIPER-CACHE Option	Purpose
VVHC\$OFF=VVHC\$OFF	DDname to turn VSAM optimization off
VVHCON=VVHCON	DDname to turn VSAM optimization on (NO XA)
VVHCONXA=VVHCONXA	DDname to turn VSAM optimization on (W/ XA)

Table F-10 lists the HIPER-CACHE Processing Control table options that are not supported by Data Optimizer.

Table F-10 HIPER-CACHE Processing Control Table Options Not Supported by Data Optimizer

HIPER-CACHE Option	Purpose
DOW=	Day of Week (TIME OF DAY table)
FPOOL/NOFPOOL	freepool issuance for Q/BSAM
FT=	from time (TIME OF DAY TABLE)
IMS/NOIMS	optimize VSAM component of IMS DB
NOPGCTL	Page Control option
RSYNC=	Optimization Technique Evaluator
SHR4/NO SHR4	SHR(4,X) dataset support ¹
TT=	to time (time of day table)
XABUF/NOXABUF	VSAM XA buffer support
XAREGN=	XA region size specification
XSAM/NOXSAM	SAM XA buffer support

Appendix G Batch Utilities

This appendix describes and provides examples of the batch utility programs that are provided with Data Optimizer. The SMF Report Utility is detailed in this appendix. This appendix discusses the following topics:

Overview	G-2
JOB Control Statements and Definitions	G-2
SYSIN Control Card Format	G-4
Keyword Specifications	G-5
SMF Report for VSAM Data Sets	G-6
SMF Report for Non-VSAM Data Sets	G-6

Overview

The SMF report utility that is provided with Data Optimizer can be used to report on the SMF data which is written by Data Optimizer as part of the SFM recording options. Record layouts of the SMF records that are written by Data Optimizer can be found in Appendix E, “SMF Recording,” for use when writing additional reporting facilities for this information. A sample of the JCL to execute the supplied reporting program can be found in the Batch Optimizer SAMPLIB as member DAPUSMFR.

JOB Control Statements and Definitions

Table G-1 shows JOB control statements and definitions.

Table G-1 **JOB Control Statements and Definitions**

Statement	Definition
//USMF1 EXEC PGM=DAPUSMF1	Execute statement
/*STEPLIB DD DISP=SHR, // *DSN=?SHAREDQUAL?.BMC.BSL.LOAD	STEPLIB DD statement that identifies the library containing the DATA OPTIMIZER modules
//SYSPRINT DD SYSOUT=*	output report DD statement RECFM that can be F (fixed) or V (variable), Blocked / Unblocked ANSI, or Machine At least RECFM and LRECL must be specified on non-SYSOUT datasets
//SMFFILE DD DISP=SHR, //DSN=?SHAREDQUAL?.SELECT	DD statement that defines the SMF input dataset (non-VSAM only)
//SYSIN DD * * (CONTROL CARDS) * * FROMDATE=(01/01/1990) * FROMTIME=(00:00:00) * TODATE=(12/31/2099) * TOTIME=(24:00:00) * JOBNAME=* * JOBSTEP=* * PROGRAM=* * RECTYPE=241 * RPTTYPE=(STATS,CAT,DSINFO,OPTIONS) * DSTYPE=(VSAM,NONVSAM) * LC=55 * LIST=(PARMS,NOHDRS,SUM) /*	input control cards to select report options Note: The same control card options can be specified on EXEC PARM= with commas separating Keyword=Value ex, PARM='FROMDATE=(01/12/2001),LIST=PARMS'. The control cards that are shown here have been commented out. These control cards can be used with EXEC PARM. The KEYWORD=VALUE that is specified on the EXEC PARM will override SYSIN control cards.

Example

This example will show how to list parms by using an EXEC PARM= statement as shown in Table G-1 on page G-2.

To list parms, execute the following statement:

```
//USMF1 EXEC PGM=DAPUSMF1,
      PARM' FROMRATE=(01/12/2002),LIST=PARMS'
```

The output will appear as follows:

```

                                DATA OPTIMIZER  SMF REPORT WRITER          PAGE: 1
                                BMC SOFTWARE, INC

PARAMETER(S) SPECIFIED
<IN EFFECT>

FROMDATE=(01/12/2002)
FROMTIME=(00:00:00)
  TODATE=(12/31/2099)
  TOTIME=(24:00:00)
JOBNAME=( *          )
JOBSTEP=( *          )
PROGRAM=( *          )
RECTYPE=(          241)
RPTTYPE=(STATS,CAT,DSINFO,OPTIONS)
DSTYPE=(VSAM, NONVSAM)
  LC=(          55)
  LIST=(PARMS)
```

Note: The results shown here are in addition to the normal data set reports.

Example

This example will show how to obtain summary results by using an EXEC PARM=STATEMENT as shown in Table G-1 on page G-2.

To obtain summary results, execute the following statement:

```
//USMF1 EXEC PGM=DAPUSMF1,
      PARM' FROMRATE=(01/12/2002),LIST=SUM'
```

The output will appear as follows:

```
1                                DATA OPTIMIZER  SMF REPORT WRITER                                PAGE: 4
                                BMC SOFTWARE, INC

DATA OPTIMIZER SMF TOTAL..... 127
DATA OPTIMIZER SMF SELECTED..... 127
                                SMF TOTAL..... 23,278
```

Note: The results shown here are in addition to the normal data set reports.

SYSIN Control Card Format

Various types of control card keywords can be specified for this utility. All keywords are optional. Types of control card keywords may include:

- Keywords must begin in column one.
- Each control card can contain comments by leaving at least one blank following the last keyword specification.
- Each keyword can be specified only one time per execution. If multiples are specified, only the last one will be accepted.
- EXEC PARM keywords override SYSIN control cards.

- Comment cards are specified by placing an asterisk (*) in column one of the card.
- Each keyword can be on a separate control card or separated by commas on the same line.

Keyword Specifications

Table G-2 identifies special characters that allow generic selection to be specified for SYSIN control card keywords JOBNAME, JOBSTEP, and PROGRAM.

Table G-2 Special Characters for SYSIN Control Card Keywords

Special Character	Definition
?	Skip character. The character occupying this position in the name is to be ignored during the comparison process.
*	Generic character. The name that is specified is a generic name, and the asterisk (*) delimits the portion of the generic name to be compared.

Table G-3 identifies control card keywords that can be used in the SYSIN control card format or EXEC PARM.

Table G-3 Control Card Keywords (Part 1 of 2)

Keyword	Definition
FROMDATE= <i>mm/dd/yyyy</i> Default 01/01/1990	This keyword selects a date range of records for reporting. All SMF records that are written by Data Optimizer before the specified date are ignored by the reporting program.
FROMTIME= <i>hh:mm:ss</i> Default 00:00:00	This keyword selects a time range of records for reporting. All SMF records that are written by Data Optimizer are ignored by the reporting program.
TODATE= <i>mm/dd/yyyy</i> Default 1/31/2099	This keyword selects a date range of records for reporting. All SMF records that are written by Data Optimizer after the specified date are ignored by the reporting program.
TOTIME= <i>hh:mm:ss</i> Default 24:00:00	This keyword selects a time range of records for reporting. All SMF records that are written by Data Optimizer after the specified time are ignored by the reporting program.
JOBNAME= <i>nnnnnnnn</i> Default (*)	This keyword selects records that pertain to jobs with specific names or generic name patterns (where <i>nnnnnnnn</i> is a specific job name or the generic job name pattern that contains (?) or (*)). Example: JOBNAME=(A*) JOBNAME=(B??L).
JOBSTEP= <i>nnnnnnnn</i> Default is (*)	This keyword selects records that pertain to jobsteps with specific names or generic name patterns (where <i>nnnnnnnn</i> is a specific jobstep or the generic jobstep pattern that contains (?) or (*)). Example: JOBSTEP=(SMF*) JOBSTEP=(B??L).

Table G-3 Control Card Keywords (Part 2 of 2)

Keyword	Definition
PROGRAM= <i>nnnnnnnn</i> Default (*)	This keyword selects records that pertain to programs with specific names or generic name patterns (where <i>nnnnnnnn</i> is a specific program or the generic program pattern that contains (?) or (*)). Example: PROGRAM=(IEF*) JOBSTEP=(IEF??14).
RECTYPE= <i>nnn</i> Default 241	This keyword identifies the SMF record type of the records written by Data Optimizer. The value specified must correspond to value that is specified by the SMF record type option.
RPTTYPE=(STATS, CAT, DSINFO, OPTIONS, ALL) Default ALL	This keyword selects SMF reports that can be obtained: <ul style="list-style-type: none"> • STATS reports on VSAM / NONVSAM statistics. • CAT provides catalog information on VSAM. • DSINFO provides general data set information. • OPTIONS provides data set options in regard to the data set being processed. • If nothing is specified, a summary report is generated. The ALL keyword specifies STATS, CAT, DSINFO, and OPTIONS.
DSTYPE=(VSAM, NONVSAM, ALL) Default ALL	This keyword selects which dataset type to report on. The ALL keyword specifies VSAM and NONVSAM.
LC=NNNN Default 55	This keyword specifies the line count per page.
LIST=(PARMS, NOHDRS, SUM) No default	This keyword specifies further reporting options: <ul style="list-style-type: none"> • PARMS lists out the keywords in effect for this run. • NOHDRS allows users to obtain reports without page headers. • SUM prints out record totals.

SMF Report for VSAM Data Sets

The VSAM Performance Report is actually a group of reports, each of which you can request or exclude. For more information, see Appendix A, “Data Optimizer Reports.”

SMF Report for Non-VSAM Data Sets

The Non-VSAM Performance Report is actually a group of reports, each of which you can request or exclude. For more information, see Appendix A, “Data Optimizer Reports.”

Index

A

- Above 16 Megabyte Maximum option 5-16
- Above 16 Megabyte Reserve option 5-15
- accessing Action Bar B-16
- Action 2-9
- action option 5-5
- activating a data policy using a command 4-18
- Actual I/O Counts option 5-22
- Adjust Region option 5-8
- Allow Shared Writes option 5-23

B

- Below 16 Megabyte Maximum option 5-15
- Below 16 Megabyte Reserve option 5-14
- BMC Software Primary Subsystem 1-4
- BMCP
 - displaying status 6-6
 - SHUTDOWN BMCP command 6-9
 - SHUTDOWN BMCPPROC command 6-8
 - START BMCPPROC command 6-6
- BMCP SHUTDOWN command 6-8
- BMCP STATUS command 6-7
- BUFNO Storage option 5-24
- BUFNO Value option 5-24

C

- CANCEL BMCPPROC command 6-9
- commands
 - activate 4-18

- activating a data policy 4-18
- BMCP SHUTDOWN 6-8
- BMCP STATUS 6-7
- CANCEL BMCPPROC 6-9
- determining which data policy is active 4-20
- displaying the Status of Data Optimizer 6-15
- enabling and disabling components 6-16
- MBOS CANCEL 6-14
- MBOS DAP DISABLE 6-16
- MBOS DAP ENABLE 6-16
- MBOS DAP HISTKEY DSNPGM 6-21
- MBOS DAP HISTKEY FULL 6-21
- MBOS DAP POLICY STATUS 6-17
- MBOS DAP STATUS 6-15, 6-18
- MBOS DAP USEREXIT DISABLE 6-20
- MBOS DAP USEREXIT ENABLE 6-19, 6-20
- MBOS SHUTDOWN 6-13
- MBOS START 6-10
- MBOS STATUS 6-12
- SHUTDOWN BMCP 6-9
- SHUTDOWN BMCPPROC 6-8
- START BCSSPROC 6-10
- START BMCPPROC 6-6
- status 4-20
- working with User Control Exit 6-17, 6-18
- CPU Utilization Threshold option 5-14

D

- DAPUEXIT
 - see user control exit
- Data 1-9
- Data Optimizer
 - benefits 1-10
 - features 1-9
 - overview 1-5, 2-2
 - performance processing 1-2
- data policies
 - about 4-2
 - activating 4-17
 - definitions 2-2
 - ordering definitions 4-3
 - setting global options 4-4
 - specifying selection criteria 4-2
- Data Policy Definition panel
 - default values 4-10

example 4-8
Data Policy Definition Selection Criterion panel
4-10, 4-11
Data Policy Definitions panel 4-7
data policy global options 4-15
Data Policy Global Options pop-up 4-16
Data Policy Global Options pop-up 4-16
DD override statements 2-3
DD overrides
DAP@MSG0 5-20
DAP@MSG1 5-20
DAP@MSG2 5-20
DAP@MSG3 5-20
DAP@NNPO 5-5
DAP@NPO 5-5
DAP@NVPO 5-6
DD statements
DAP\$VRxx 5-13
DAP@AIC 5-6, 5-23
DAP@BASC 5-6
DAP@CAT 5-19
DAP@DRA 5-8
DAP@MSG0 5-20
DAP@MSG1 5-20
DAP@MSG2 5-20
DAP@MSG3 5-20
DAP@NAIC 5-23
DAP@NDRA 5-8
DAP@NDSI 5-18
DAP@NNPO 5-5
DAP@NOVU 5-10
DAP@NPO 5-5
DAP@NRxx 5-13
DAP@NSHR 5-34, 5-44
DAP@NSMF 5-10
DAP@NSTA 5-17
DAP@NVDW 5-33
DAP@NVPO 5-6
DAP@NXCP D-6
DAP@Nxxx 5-7
DAP@OVU 5-10
DAP@RBH 5-11
DAP@RPTO 5-6
DAP@SHR 5-34, 5-44
DAP@STA 5-17
DAP@VDW 5-33
DAP@VPO 5-5
DAP@Vxxx 5-7

DAPMESGS 5-22
DAPSTATS 5-17
DASP@DSI 5-18
position 5-2
definition
data policy 2-2
determining which data policy is active 4-20
Dynamic Region Adjust option 2-10

E

EXCP D-4

F

formats, syntax diagrams 1-xix

G

general options 2-9
global options
VSAM SHAREOPTIONS(3,3) 5-34

H

HIPER-CACHE Conversion Utility F-1
historical data 2-8
Honor Resource Usage Bias Option 5-11

I

implementation
goal-setting 3-2
planning 3-2
strategies 3-6
Information Report option 5-18
interface
user exit 2-3
ISPF
conventions B-2
profile B-2

J

Job Optimizer 1-3
job statement overrides
 DAP@NSMF 5-10
job step overrides
 DAP\$VRxx 5-13
 DAP@AIC 5-6, 5-23
 DAP@BASC 5-6
 DAP@CAT 5-19
 DAP@DRA 5-8
 DAP@DSI 5-18
 DAP@NAIC 5-23
 DAP@NCAT 5-19
 DAP@NDRA 5-8
 DAP@NDSI 5-18
 DAP@NOVU 5-10
 DAP@NRxx 5-13
 DAP@NSHR 5-34, 5-44
 DAP@NSTA 5-17
 DAP@NVDW 5-33
 DAP@Nxxx 5-7
 DAP@OVU 5-10
 DAP@RBH 5-11
 DAP@RPTO 5-6
 DAP@SHR 5-34, 5-44
 DAP@STA 5-17
 DAP@VDW 5-33
 DAP@VPO 5-5
 DAP@Vxxx 5-7
 DAPMESGS 5-22
 DAPSTATS 5-17

L

LSR Data Buffers option 5-41
LSR Data Storage option 5-42
LSR Index Buffers option 5-38
LSR Index Storage option 5-39

M

MAINVIEW Batch Optimizer 1-4
 components 1-2
 user interface 1-4

MAINVIEW Batch Optimizer subsystem
(MBOS)
 DAP ENABLE command 6-16
 DAP HISTKEY DSNPGM command 6-21
 DAP HISTKEY FULL command 6-21
 DAP POLICY STATUS command 6-17
 DAP STATUS command 6-18
 DAP USEREXIT DISABLE command 6-20
 DAP USEREXIT ENABLE command 6-19,
 6-20
Maximum BUFNO Storage option 5-25
Maximum BUFNO Value option 5-25
Maximum Data Storage option 5-43
Maximum Index Buffers option 5-40
Maximum Index Storage option 5-36
Maximum LSR Index Storage option 5-41
Maximum NSR Data Buffers option 5-37
Maximum NSR Data Storage option 5-38
Maximum NSR Index Buffers option 5-35
MBOS
 MBOS STATUS command 6-12
 START BCSSPROC command 6-10
MBOS CANCEL command 6-12, 6-13, 6-14
MBOS DAP DISABLE command 6-16
MBOS DAP ENABLE command 6-16
MBOS DAP HISTKEY DSNPGM command
 6-21
MBOS DAP HISTKEY FULL command 6-21
MBOS DAP POLICY STATUS command 6-17
MBOS DAP STATUS command 6-15, 6-18
MBOS DAP USEREXIT DISABLE command
 6-20
MBOS DAP USEREXIT ENABLE command
 6-19, 6-20
MBOS START command 6-10
Message Destination option 5-21
Message SYSOUT Class option 5-22
messages
 informational 2-13

N

NCP Storage option 5-26
NCP Value option 5-26
non-VSAM optimization 1-7
NSR Data Buffers option 5-36
NSR Data Storage option 5-37

NSR Index Buffers option 5-34
NSR Index Storage option 5-35

O

online documentation 1-xvii
optimization
 components 1-2
 non-VSAM 1-7
 VSAM 1-6
optimization method options 2-9, 5-6
Optimization Percentage option 5-7
Optimize STC Access option 5-8
optimize STC access option 2-10
Optimize TSO Access option 5-9
optimize TSO access option 2-10, 5-9
options
 Above 16 Megabyte Maximum 5-16
 Above 16 Megabyte Reserve 5-15
 action 5-5
 Actual I/O Counts 5-22
 Allow Shared Writes 5-23
 Below 16 Megabyte Maximum 5-15
 Below 16 Megabyte Reserve 5-14
 BUFNO Storage 5-24
 BUFNO Value 5-24
 CPU Utilization Threshold 5-14
 data policy global 4-15
 Description Format 5-1
 Dynamic Region Adjust 2-10
 general 2-9
 Honor Resource Usage Bias 5-11
 Information Report 5-18
 LSR Data Buffers 5-41
 LSR Data Storage 5-42
 LSR Index Buffers 5-38
 LSR Index Storage 5-39
 Maximum BUFNO Storage 5-25
 Maximum BUFNO Value 5-25
 Maximum Data Storage 5-43
 Maximum Index Buffers 5-40
 Maximum Index Storage 5-36
 Maximum LSR Index Storage 5-41
 Maximum NSR Data Buffers 5-37
 Maximum NSR Data Storage 5-38
 Maximum NSR Index Buffers 5-35
 Message Destination 5-21

Message SYSOUT Class 5-22
NCP Storage 5-26
NCP Value 5-26
NSR Data Buffers 5-36
NSR Data Storage option 5-37
NSR Index Buffers 5-34
NSR Index Storage 5-35
optimization method 5-6
Optimization Percentage 5-7
Optimize STC Access 5-8
optimize STC access 2-10
Optimize TSO Access 5-9
optimize TSO access 2-10, 5-9
Override User Values 2-11, 5-10
Page Fault Rate Threshold 5-13
Percent of Index to Buffer 5-40
performance 2-8
report
 non-VSAM A-6, A-16
Report SYSOUT Class 5-18
Resource Usage Bias 5-11
Resource Utilization 5-13
RMODE31 Control Blocks 5-33
SMF Record Type
 SMF Record Type option 5-9
SMF Record Type option 2-10
Statistics Report 5-17
Statistics Report Destination 5-18
VSAM SHAREOPTIONS(3,3) 5-34
override statements
 DD 2-3
Override User Values option 2-11, 5-10

P

Page Fault Rate Threshold option 5-13
panel flow
 Data Policy registration 4-4
panel flow for a Data Policy registration 4-4
panel layout B-3
panels
 Data Policy Definition 4-8
 Data Policy Definition Selection Criterion
 4-10, 4-11
 Data Policy Definitions 4-7
 Data Policy Global Options pop-up 4-16
panel layout B-3

- Policy Activation BatchPlex Name pop-up 4-19
- Policy Activation Images List Pop-up 4-20
- Percent of Index to Buffer option 5-40
- performance options 2-8
- Policy Activation BatchPlex Name pop-up 4-19
- Policy Activation Images List Pop-up 4-20
- pop-ups
 - Data Policy Global Options 4-16
- program registration
 - Honor Resource Usage Bias option 5-11

R

- RACF access requirements ?? to 6-5
- related documentation 1-xvii
- release notes 1-xvii
- Report SYSOUT Class option 5-18
- reporting options 5-16
- reports
 - non-VSAM
 - Options A-6, A-16
 - Statistics A-4, A-15
 - overview A-1
 - User Exit Statistics C-6
- Resource Usage Bias option 5-11
- Resource Utilization options 5-13
- RMODE31 Control Blocks option 5-33

S

- selection criteria
 - order 2-4
 - specifying 4-2
- selection processing 2-2
- SMF Record Type option 2-10
- START BMCPPROC command 6-6
- statistics
 - report
 - non-VSAM A-4, A-15
- Statistics Report Destination option 5-18
- Statistics Report option 5-17
- syntax
 - format for diagrams 1-xix
- system management facilities 5-9
- system resource utilization 2-7

U

- User Control exit
 - definition C-2
 - enabling and disabling 6-20
- user exit interface 2-3
- User Exit Statistics report C-6
- user-specified options 2-7
- utilities
 - candidate 3-5
- utilization
 - system resource 2-7

V

- VSAM
 - optimization 1-6
- VSAM Optimization 1-6
- VSAM Reports A-11
- VSAM SHAREOPTIONS(3,3) option 5-34

W

- wildcard characters 5-4

END USER LICENSE AGREEMENT NOTICE

BY OPENING THE PACKAGE, INSTALLING, PRESSING "AGREE" OR "YES" OR USING THE PRODUCT, THE ENTITY OR INDIVIDUAL ENTERING INTO THIS AGREEMENT AGREES TO BE BOUND BY THE FOLLOWING TERMS. IF YOU DO NOT AGREE WITH ANY OF THESE TERMS, DO NOT INSTALL OR USE THE PRODUCT, PROMPTLY RETURN THE PRODUCT TO BMC OR YOUR BMC RESELLER, AND IF YOU ACQUIRED THE LICENSE WITHIN 30 DAYS OF THE DATE OF YOUR ORDER CONTACT BMC OR YOUR BMC RESELLER FOR A REFUND OF LICENSE FEES PAID. IF YOU REJECT THIS AGREEMENT, YOU WILL NOT ACQUIRE ANY LICENSE TO USE THE PRODUCT.

This Agreement ("**Agreement**") is between the entity or individual entering into this Agreement ("You") and BMC Software Distribution, Inc., a Delaware corporation located at 2101 CityWest Blvd., Houston, Texas, 77042, USA or its affiliated local licensing entity ("BMC"). "You" includes you and your Affiliates. "Affiliate" is defined as an entity which controls, is controlled by or shares common control with a party. IF MORE THAN ONE LICENSE AGREEMENT COULD APPLY TO THE PRODUCT, THE FOLLOWING ORDER OF LICENSE AGREEMENT PRECEDENCE APPLIES: (1) WEB BASED LICENSE AGREEMENT WITH BMC, (2) WRITTEN LICENSE AGREEMENT WITH BMC, (3) SHRINK-WRAP LICENSE AGREEMENT WITH BMC PROVIDED WITH THE PRODUCT, AND (4) THIS ELECTRONIC LICENSE AGREEMENT WITH BMC. In addition to the restrictions imposed under this Agreement, any other usage restrictions contained in the Product installation instructions or release notes shall apply to Your use of the Product.

PRODUCT AND CAPACITY. "**Software**" means the object code version of the computer programs provided, via delivery or electronic transmission, to You. Software includes computer files, enhancements, maintenance modifications, upgrades, updates, bug fixes, and error corrections.

"**Documentation**" means all written or graphical material provided by BMC in any medium, including any technical specifications, relating to the functionality or operation of the Software.

"**Product**" means the Software and Documentation.

"**License Capacity**" means the licensed capacity for the Software with the pricing and other license defining terms, including capacity restrictions, such as tier limit, total allowed users, gigabyte limit, quantity of Software, and/or other capacity limitations regarding the Software. For licenses based on the power of a computer, You agree to use BMC's current computer classification scheme, which is available at <http://www.bmc.com> or can be provided to You upon request.

ACCEPTANCE. The Product is deemed accepted by You, on the date that You received the Product from BMC.

LICENSE. Subject to the terms of this Agreement, as well as Your payment of applicable fees, BMC grants You a non-exclusive, non-transferable, perpetual (unless a term license is provided on an order) license for each copy of the Software, up to the License Capacity, to do the following:

- (a) install the Software on Your owned or leased hardware located at a facility owned or controlled by You in the country where You acquired the license;
- (b) operate the Software solely for processing Your own data in Your business operations; and
- (c) make one copy of the Software for backup and archival purposes only (collectively a "**License**").

If the Software is designed by BMC to permit you to modify such Software, then you agree to only use such modifications or new software programs for Your internal purposes or otherwise consistent with the License. BMC grants You a license to use the Documentation solely for Your internal use in Your operations.

LICENSE UPGRADES. You may expand the scope of the License Capacity only pursuant to a separate agreement with BMC for such expanded usage and Your payment of applicable fees. There is no additional warranty period or free support period for license upgrades.

RESTRICTIONS: You agree to **NOT**:

- (a) disassemble, reverse engineer, decompile or otherwise attempt to derive any Software from executable code;
- (b) distribute or provide the Software to any third party (including without limitation, use in a service bureau, outsourcing environment, or processing the data of third parties, or for rental, lease, or sublicense); or
- (c) provide a third party with the results of any functional evaluation or benchmarking or performance tests, without BMC's prior written approval, unless prohibited by local law.

TRIAL LICENSE. If, as part of the ordering process, the Product is provided on a trial basis, then these terms apply: (i) this license consists solely of a non-exclusive, non-transferable evaluation license to operate the Software for the period of time specified from BMC or, if not specified, a 30 day time period ("**Trial Period**") only for evaluating whether You desire to acquire a capacity-based license to the Product for a fee; and (ii) Your use of the Product is on an AS IS basis without any warranty, and **BMC, ITS AFFILIATES AND RESELLERS, AND LICENSORS DISCLAIM ANY AND ALL WARRANTIES (INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT) AND HAVE NO LIABILITY WHATSOEVER RESULTING FROM THE USE OF THIS PRODUCT UNDER THIS TRIAL LICENSE ("Trial License").** BMC may terminate for its convenience a Trial License upon notice to You. When the Trial Period ends, Your right to use this Product automatically expires. If You want to continue Your use of the Product beyond the Trial Period, contact BMC to acquire a capacity-based license to the Product for a fee.

TERMINATION. This Agreement shall immediately terminate if You breach any of its terms. Upon termination, for any reason, You must uninstall the Software, and either certify the destruction of the Product or return it to BMC.

OWNERSHIP OF THE PRODUCT. BMC or its Affiliates or licensors retain all right, title and interest to and in the BMC Product and all intellectual property, informational, industrial property and proprietary rights therein. BMC neither grants nor otherwise transfers any rights of ownership in the BMC Product to You. Products are protected by applicable copyright, trade secret, and industrial and intellectual property laws. BMC reserves any rights not expressly granted to You herein.

CONFIDENTIAL AND PROPRIETARY INFORMATION. The Products are and contain valuable confidential information of BMC (“**Confidential Information**”). Confidential Information means non-public technical and non-technical information relating to the Products and Support, including, without limitation, trade secret and proprietary information, and the structure and organization of the Software. You may not disclose the Confidential Information to third parties. You agree to use all reasonable efforts to prevent the unauthorized use, copying, publication or dissemination of the Product.

WARRANTY. Except for a Trial License, BMC warrants that the Software will perform in substantial accordance with the Documentation for a period of one year from the date of the order. This warranty shall not apply to any problems caused by software or hardware not supplied by BMC or to any misuse of the Software.

EXCLUSIVE REMEDY. BMC’s entire liability, and Your exclusive remedy, for any defect in the Software during the warranty period or breach of the warranty above shall be limited to the following: BMC shall use reasonable efforts to remedy defects covered by the warranty or replace the defective Software within a reasonable period of time, or if BMC cannot remedy or replace such defective copy of the Software, then BMC shall refund the amount paid by You for the License for that Software. BMC’s obligations in this section are conditioned upon Your providing BMC prompt access to the affected Software and full cooperation in resolving the claim.

DISCLAIMER. EXCEPT FOR THE EXPRESS WARRANTIES ABOVE, THE PRODUCT IS PROVIDED “AS IS.” BMC, ITS AFFILIATES AND LICENSORS SPECIFICALLY DISCLAIM ALL OTHER WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. BMC DOES NOT WARRANT THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR FREE, OR THAT ALL DEFECTS CAN BE CORRECTED.

DISCLAIMER OF DAMAGES. IN NO EVENT IS BMC, ITS AFFILIATES OR LICENSORS LIABLE FOR ANY SPECIAL, INDIRECT, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES RELATING TO OR ARISING OUT OF THIS AGREEMENT, SUPPORT, AND/OR THE PRODUCT (INCLUDING, WITHOUT LIMITATION, LOST PROFITS, LOST COMPUTER USAGE TIME, AND DAMAGE OR LOSS OF USE OF DATA), EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, AND IRRESPECTIVE OF ANY NEGLIGENCE OF BMC OR WHETHER SUCH DAMAGES RESULT FROM A CLAIM ARISING UNDER TORT OR CONTRACT LAW.

LIMITS ON LIABILITY. BMC’S AGGREGATE LIABILITY FOR DAMAGES IS LIMITED TO THE AMOUNT PAID BY YOU FOR THE LICENSE TO THE PRODUCT.

SUPPORT. If Your order includes support for the Software, then BMC agrees to provide support (24 hours a day/7 days a week) (“**Support**”). You will be automatically re-enrolled in Support on an annual basis unless BMC receives notice of termination from You as provided below. There is a free support period during the one year warranty period.

(a) **Support Terms.** BMC agrees to make commercially reasonable efforts to provide the following Support: (i) For malfunctions of supported versions of the Software, BMC provides bug fixes, patches or workarounds in order to cause that copy of the Software to operate in substantial conformity with its then-current operating specifications; and (ii) BMC provides new releases or versions, so long as such new releases or versions are furnished by BMC to all other enrolled Support customers without additional charge. BMC may refuse to provide Support for any versions or releases of the Software other than the most recent version or release of such Software made available by BMC. Either party may terminate Your enrollment in Support upon providing notice to the other at least 30 days prior to the next applicable Support anniversary date. If You re-enroll in Support, BMC may charge You a reinstatement fee of 1.5 times what You would have paid if You were enrolled in Support during that time period.

(b) **Fees.** The annual fee for Support is 20% of the Software’s list price less the applicable discount or a flat capacity based annual fee. BMC may change its prices for the Software and/or Support upon at least 30 days notice prior to Your support anniversary date.

VERIFICATION. If requested by BMC, You agree to deliver to BMC periodic written reports, whether generated manually or electronically, detailing Your use of the Software in accordance with this Agreement, including, without limitation, the License Capacity. BMC may, at its expense, perform an audit, at your facilities, of Your use of the Software to confirm Your compliance with the Agreement. If an audit reveals that You have underpaid fees, You agree to pay such underpaid fees. If the underpaid fees exceed 5% of the fees paid, then You agree to also pay BMC’s reasonable costs of conducting the audit.

EXPORT CONTROLS. You agree not to import, export, re-export, or transfer, directly or indirectly, any part of the Product or any underlying information or technology except in full compliance with all United States, foreign and other applicable laws and regulations.

GOVERNING LAW. This Agreement is governed by the substantive laws in force, without regard to conflict of laws principles: (a) in the State of New York, if you acquired the License in the United States, Puerto Rico, or any country in Central or South America; (b) in the Province of Ontario, if you acquired the License in Canada (subsections (a) and (b) collectively referred to as the “**Americas Region**”); (c) in Singapore, if you acquired the License in Japan, South Korea, Peoples Republic of China, Special Administrative Region of Hong Kong, Republic of China, Philippines, Indonesia, Malaysia, Singapore, India, Australia, New Zealand, or Thailand (collectively, “**Asia Pacific Region**”); or (d) in the Netherlands, if you acquired the License in any other country not described above. The United Nations Convention on Contracts for the International Sale of Goods is specifically disclaimed in its entirety.

ARBITRATION. ANY DISPUTE BETWEEN YOU AND BMC ARISING OUT OF THIS AGREEMENT OR THE BREACH OR ALLEGED BREACH, SHALL BE DETERMINED BY BINDING ARBITRATION CONDUCTED IN ENGLISH. IF THE DISPUTE IS INITIATED IN THE AMERICAS REGION, THE ARBITRATION SHALL BE HELD IN NEW YORK, U.S.A., UNDER THE CURRENT COMMERCIAL OR INTERNATIONAL, AS APPLICABLE, RULES OF THE AMERICAN ARBITRATION ASSOCIATION. IF THE DISPUTE IS INITIATED IN A COUNTRY IN THE ASIA PACIFIC REGION, THE ARBITRATION SHALL BE HELD IN SINGAPORE, SINGAPORE UNDER THE CURRENT UNCITRAL ARBITRATION RULES. IF THE DISPUTE IS INITIATED IN A COUNTRY OUTSIDE OF THE AMERICAS REGION OR ASIA PACIFIC REGION, THE ARBITRATION SHALL BE HELD IN AMSTERDAM, NETHERLANDS UNDER THE CURRENT UNCITRAL ARBITRATION RULES. THE COSTS OF THE ARBITRATION SHALL BE BORNE EQUALLY PENDING THE ARBITRATOR’S AWARD. THE AWARD RENDERED SHALL BE FINAL AND BINDING UPON THE PARTIES AND SHALL NOT BE SUBJECT TO APPEAL TO ANY COURT, AND MAY BE ENFORCED IN ANY COURT OF COMPETENT JURISDICTION. NOTHING IN THIS AGREEMENT SHALL BE DEEMED AS PREVENTING EITHER PARTY FROM SEEKING INJUNCTIVE RELIEF FROM ANY COURT HAVING JURISDICTION OVER THE PARTIES AND THE SUBJECT MATTER OF

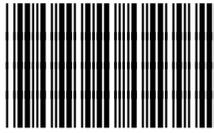
THE DISPUTE AS NECESSARY TO PROTECT EITHER PARTY'S CONFIDENTIAL INFORMATION, OWNERSHIP, OR ANY OTHER PROPRIETARY RIGHTS. ALL ARBITRATION PROCEEDINGS SHALL BE CONDUCTED IN CONFIDENCE, AND THE PARTY PREVAILING IN ARBITRATION SHALL BE ENTITLED TO RECOVER ITS REASONABLE ATTORNEYS' FEES AND NECESSARY COSTS INCURRED RELATED THERETO FROM THE OTHER PARTY.

U.S. GOVERNMENT RESTRICTED RIGHTS. The Software under this Agreement is "commercial computer software" as that term is described in 48 C.F.R. 252.227-7014(a)(1). If acquired by or on behalf of a civilian agency, the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of this Agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the Federal Acquisition Regulations ("**FAR**") and its successors. If acquired by or on behalf of any agency within the Department of Defense ("**DOD**"), the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of this Agreement as specified in 48 C.F.R. 227.7202 of the DOD FAR Supplement and its successors.

MISCELLANEOUS TERMS. You agree to pay BMC all amounts owed no later than 30 days from the date of the applicable invoice, unless otherwise provided on the order for the License to the Products. You will pay, or reimburse BMC, for taxes of any kind, including sales, use, duty, tariffs, customs, withholding, property, value-added (VAT), and other similar federal, state or local taxes (other than taxes based on BMC's net income) imposed in connection with the Product and/or the Support. This Agreement constitutes the entire agreement between You and BMC and supersedes any prior or contemporaneous negotiations or agreements, whether oral, written or displayed electronically, concerning the Product and related subject matter. No modification or waiver of any provision hereof will be effective unless made in a writing signed by both BMC and You. You may not assign or transfer this Agreement or a License to a third party without BMC's prior written consent. Should any provision of this Agreement be invalid or unenforceable, the remainder of the provisions will remain in effect. The parties have agreed that this Agreement and the documents related thereto be drawn up in the English language. Les parties exigent que la présente convention ainsi que les documents qui s'y rattachent soient rédigés en anglais.

SW Click EULA 071102

Notes



28328