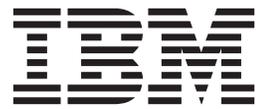


IBM Tivoli OMEGAMON XE for DB2 Performance Expert on
z/OS
IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on
z/OS
Version 5.2.0

Report Command Reference



IBM Tivoli OMEGAMON XE for DB2 Performance Expert on
z/OS
IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on
z/OS
Version 5.2.0

Report Command Reference



Note

Before using this information and the product it supports, read the information in "Notices" on page 173.

First edition, October 2013

This edition applies to the following releases and to all subsequent releases and modifications until otherwise indicated in new editions:

- IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS, version 5, release 2, modification 0 (5655-W37)
- IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS, version 5, release 2, modification 0 (5655-W38)

This edition replaces SH12-6964-00.

© **Copyright IBM Corporation 1985, 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication	v	EXCEPTION and NOEXCEPTION subcommand options	30
Who should read this publication	v	FROM/TO subcommand options	30
Conventions used in the OMEGAMON documentation	v	INCLUDE and EXCLUDE subcommand options	33
Terminology used	vii	INTERVAL subcommand option	37
How to read syntax diagrams	viii	LAYOUT subcommand option	38
Where to find information	ix	LEVEL subcommand option	45
Service updates and support information	x	LIMIT subcommand option	45
Accessibility features	xi	ORDER subcommand option	46
How to send your comments	xii	SCOPE subcommand option	48
		SORTBY subcommand option	48
		SPREADSHEETDD subcommand option	48
		SUMMARIZEBY subcommand option	49
		TOP subcommand option	49
		TYPE subcommand option	50
		WORKLOAD subcommand option	50
What's new	xiii		
Chapter 1. General information about OMEGAMON XE for DB2 PE report sets	1	Chapter 7. OMEGAMON XE for DB2 PE commands	51
Chapter 2. OMEGAMON XE for DB2 PE identifiers	3	Overview of the commands	52
Chapter 3. The OMEGAMON XE for DB2 PE command stream	11	ACCOUNTING command	53
Chapter 4. Output from OMEGAMON XE for DB2 PE reports	13	ACCOUNTING command with REPORT subcommand	55
Chapter 5. DD statements	15	ACCOUNTING command with REDUCE subcommand	62
ACMEMnn statements	15	ACCOUNTING command with SAVE subcommand	64
DPMLOG statement	15	ACCOUNTING command with RESTORE subcommand	66
DPMOUTDD statement	16	ACCOUNTING command with TRACE subcommand	67
DPMPARMS statement	17	ACCOUNTING command with FILE subcommand	71
EXCPTDD statement	17	AUDIT command	73
EXFILDD1 statement	18	AUDIT command with REPORT subcommand	74
EXPRODDA statement	18	AUDIT command with REDUCE subcommand	77
EXTRCDD1 statement	18	AUDIT command with TRACE subcommand	78
INPUTDD statement	19	AUDIT command with FILE subcommand	80
JOBSUMDD statement	19	EXEC command	82
JSSRSDD statement	20	EXPLAIN command	82
SYSIN statement	20	EXPLAIN PLAN specifications	87
SYSOUT statement	21	EXPLAIN PACKAGE specifications	88
SYSRMDDD statement	21	EXPLAIN QMFQUERY specifications	89
ccFILDD1 statements	21	EXPLAIN SQLSTMT specifications	91
ccRPTDD statements	22	EXPLAIN QUERYNO specifications	91
ccRSTDD statements	22	IOACTIVITY command	93
ccSAVDD statements	23	IOACTIVITY command with REPORT subcommand	94
ccTRCDDx statements	23	IOACTIVITY command with REDUCE subcommand	97
ccWORK statements	24	LOCKING command	98
		LOCKING command with REPORT subcommand	100
		LOCKING command with REDUCE subcommand	102
Chapter 6. OMEGAMON XE for DB2 PE subcommand options	25		
BOUNDARY subcommand option	26		
DATEFORMAT subcommand option	27		
DDNAME subcommand option	29		

LOCKING command with TRACE subcommand	103
LOCKING command with FILE subcommand	105
RECTRACE command	105
RECTRACE command with TRACE	
subcommand	107
RECTRACE command with FILE subcommand	108
SQLACTIVITY command	109
SQLACTIVITY command with REPORT	
subcommand	111
SQLACTIVITY command with TRACE	
subcommand	115
STATISTICS command	119
STATISTICS command with REPORT	
subcommand	121
STATISTICS command with REDUCE	
subcommand	125
STATISTICS command with SAVE subcommand	128
STATISTICS command with RESTORE	
subcommand	130
STATISTICS command with TRACE	
subcommand	131
STATISTICS command with FILE subcommand	133
SYSPARMS command	135
SYSPARMS command with TRACE	
subcommand	137
SYSPARMS command with FILE subcommand	137
UTILITY command	138
UTILITY command with REPORT subcommand	139

UTILITY command with REDUCE subcommand	141
UTILITY command with TRACE subcommand	142

Chapter 8. Auxiliary commands . . . 145

CASE command	145
FIELD command	147
GLOBAL command	149
GLOBAL command with EXPLAIN option	154
GLOBAL command with INCLUDE and	
EXCLUDE subcommand options	155
GROUP command	156
Grouping records	159
LIST command	164

Chapter 9. Troubleshooting commands . . . 167

DUMP command	167
TAPECOPY command	169

Notices . . . 173

Trademarks	175
----------------------	-----

Bibliography . . . 177

Index . . . 179

About this publication

This publication describes the report commands and command syntax of IBM® Tivoli® OMEGAMON® XE for DB2® Performance Expert on z/OS® and IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS.

It is closely related to the *IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS*; *IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS: Reporting User's Guide*, which describes how to choose and create reports and the *IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS*; *IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS: Report Reference*, which shows and explains the reports.

For information about collecting report data, including commands and related information for the Workstation Online Monitor, the ISPF Online Monitor, the DB2 START TRACE command, and the Collect Report Data batch program, see the *IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS*; *IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS: Reporting User's Guide*.

The technical changes for this edition are summarized under “What's new” on page xiii. Specific changes since the previous edition of this publication are indicated by a vertical bar (|) to the left of a change.

Always check the DB2 Tools Product Page and Tivoli Documentation Central for the most current version of this publication.

The product often provides context-related online help information that can be invoked from menus, panels, and windows through the F1 PF key or the Help button. Online help information is not necessarily repeated in this information, especially if it is very detailed information that is of interest only when you actively work with a function. You are encouraged to use F1 or Help to see the entire available information.

Who should read this publication

This publication is written for anyone who uses OMEGAMON XE for DB2 PE reports and traces to monitor DB2 activity for:

- Determining DB2 system performance and efficiency
- Tuning DB2
- Identifying bottlenecks
- Measuring performance and resource cost of an application
- Evaluating the effects of an application on other applications and the system

Conventions used in the OMEGAMON documentation

This information uses several conventions for special terms and actions, and operating system-dependent commands and paths.

Panels and figures

The panels and figures in this document are representations. Actual product panels might differ.

Symbols

The following symbols might appear in command syntax:

Symbol	Usage
	<p>The or symbol is used to denote a choice. You can use the argument on the left or the argument on the right. For example:</p> <pre>YES NO</pre> <p>In this example, you can specify YES or NO.</p>
()	<p>Denotes optional arguments. Arguments that are not enclosed in square brackets are required. For example:</p> <pre>APPLDEST DEST (ALTDEST)</pre> <p>In this example, DEST is a required argument and ALTDEST is optional.</p>
{ }	<p>Some documents use braces to denote mandatory arguments, or to group arguments for clarity. For example:</p> <pre>COMPARE {workload} - REPORT={SUMMARY HISTOGRAM}</pre> <p>In this example, the workload variable is mandatory. The REPORT keyword must be specified with a value of SUMMARY or HISTOGRAM.</p>
_	<p>Default values are underscored. For example:</p> <pre>COPY infile outfile - [COMPRESS={YES NO}]</pre> <p>In this example, the COMPRESS keyword is optional. If specified, the only valid values are YES or NO. If omitted, the default is YES.</p>

Notation conventions

The following conventions are used when referring to high-level qualifiers:

hilev A high-level qualifier. The high-level qualifier is the first prefix or set of prefixes in the data set name. Site-specific high-level qualifiers are shown in italics.

For example:

- *thilev* refers to the high-level qualifier for your target data set.
- *rhilev* refers to the high-level qualifier for your runtime data set.
For members in target libraries, the high-level qualifier is *thilev* rather than *rhilev*.
- *shilev* refers to the SMP/E library high-level qualifier.

Typeface conventions

This information uses the following typeface conventions:

Bold

- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Note:**)
- Keywords and parameters in text

Italic

- Words defined in text
- Emphasis of words (for example: Use the word *that* to introduce a restrictive clause.)
- New terms in text (except in a definition list)

Monospaced

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

Significant elements

Recommendation

Provides guidance when more than one option is available.

Related reading

Refers you to other publications that contain relevant information.

Requirement

Identifies a condition that must be met to ensure that the product is functional.

Restriction

Identifies a restriction or limitation with this product or an associated procedure.

Terminology used

IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS can be considered as a functional subset of IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS. Therefore the abbreviation OMEGAMON XE for DB2 PE or DB2 PE is used for both products. If a distinction is required, OMEGAMON XE for DB2 PM or DB2 PM is used explicitly.

The following table shows the products that are described in this publication and the short names with which they are referred to throughout this publication:

Table 1. Product names and their short names

Product name	Short name
IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS	OMEGAMON XE for DB2 PE or DB2 PE
IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS	OMEGAMON XE for DB2 PM or DB2 PM
IBM DB2 Buffer Pool Analyzer for z/OS or a particular subsystem	Buffer Pool Analyzer
IBM DB2 database for z/OS	DB2

- Performance Expert Client and Workstation Online Monitor designate the client component of DB2 PE.

The client component of DB2 PE also designates the end user interface of Performance Expert for Multiplatforms, Performance Expert for Workgroups, and DB2 PE.

- OMEGAMON Collector designates the server component of DB2 PE.

How to read syntax diagrams

The rules in this section apply to the syntax diagrams that are used in this publication.

Arrow symbols

Read the syntax diagrams from left to right, from top to bottom, following the path of the line.

- ▶▶— Two right arrows followed by a line indicate the beginning of a statement.
- ▶ One right arrow at the end of a line indicates that the statement syntax is continued on the next line.
- ▶— One right arrow followed by a line indicates that a statement is continued from the previous line.
- ▶◀ A line followed by a right arrow and a left arrow indicates the end of a statement.

Conventions

- SQL commands appear in uppercase.
- Variables appear in italics (for example, *column-name*). They represent user-defined parameters or suboptions.
- When entering commands, separate parameters and keywords by at least one blank if there is no intervening punctuation.
- Enter punctuation marks (slashes, commas, periods, parentheses, quotation marks, equal signs) and numbers exactly as given.
- Footnotes are shown by a number in parentheses, for example, (1).

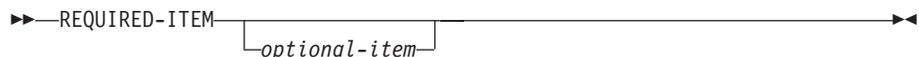
Required items

Required items appear on the horizontal line (the main path).



Optional items

Optional items appear below the main path.



If an optional item appears above the main path, that item has no effect on the execution of the statement and is used only for readability.



Multiple required or optional items

If you can choose from two or more items, they appear vertically in a stack. If you *must* choose one of the items, one item of the stack appears on the stack main path.

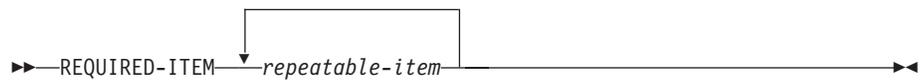


If choosing one of the items is optional, the entire stack appears below the main path.



Repeatable items

An arrow returning to the left above the main line indicates that an item can be repeated.



If the repeat arrow contains a comma, you must separate repeated items with a comma.



If the repeat arrow contains a number in parenthesis, the number represents the maximum number of times that the item can be repeated.



A repeat arrow above a stack indicates that you can specify more than one of the choices in the stack.

Default keywords

IBM-supplied default keywords appear above the main path, and the remaining choices are shown below the main path. In the parameter list following the syntax diagram, the default choices are underlined.



Where to find information

You can access the documentation in several ways.

The documentation for this product is provided in PDF and in HTML format at the following websites:

- Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS information center
- Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS information center

Accessing publications online

IBM posts publications for this and all other Tivoli products, as they become available and whenever they are updated, to the Tivoli software information center website. You can access the Tivoli software information center by going to the Tivoli Documentation Central website and clicking **O** under **Tivoli Documentation A-Z** to access all of the IBM Tivoli OMEGAMON product manuals.

Note: If you print PDF documents on other than letter-sized paper, set the option in the **File > Print** window that allows Adobe Reader to print letter-sized pages on your local paper.

The IBM Software Support website provides the latest information about known product limitations and workarounds in the form of technotes for your product. You can view this information at the Support home website.

Ordering publications

You can order many IBM publications such as product manuals or IBM Redbooks® online at the IBM Publications Center website.

You can also order by telephone by calling one of the following numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968

In other countries, contact your software account representative to order Tivoli publications.

Accessing terminology online

The IBM Terminology website consolidates the terminology from IBM product libraries in one convenient location.

Service updates and support information

You can access support information for IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS and IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS on the Support home website, or you can use the IBM Support Assistant.

Support home

On the Support home website, you can find service updates and support information including software fix packs, PTFs, Frequently Asked Questions (FAQs), technical notes, troubleshooting information, and downloads.

IBM Support assistant

The IBM Support Assistant (ISA) is a free tool that provides access to several IBM support resources in a single location. You can use the ISA tool to quickly access support-related information and serviceability tools for problem determination.

To use ISA, complete the following steps:

1. Download ISA from the IBM Software Support website.
2. Start the ISA tool.
ISA runs as a web application in the default system-configured web browser.
3. Select the Updater tab.
4. Select the New Products and Tools tab.
The plug-in features are categorized by product family.
5. Select **Tivoli > IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS and IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS**.
6. Check the feature(s) to be installed and click **Install**.
7. Restart ISA.

To learn more about how to use ISA, click the Help link in the IBM Support Assistant window.

Accessibility features

Accessibility features help people with a physical disability, such as restricted mobility or limited vision, or with other special needs, to use software products successfully. This information center is developed to comply with the accessibility requirements of software products according to Section 508 of the Rehabilitation Act of the United States.

The accessibility features in this information center enable users to do the following tasks:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. In this information center, all information is provided in HTML format. Consult the product documentation of the assistive technology for details on using assistive technologies with HTML-based information.
- Operate specific or equivalent features using only the keyboard.
- Magnify what is displayed on the screen.

In addition, all images are provided with alternative text so that users with vision impairments can understand the contents of the images.

Navigating the interface by using the keyboard

Standard shortcut and accelerator keys are used by the product and are documented by the operating system. Refer to the documentation provided by your operating system for more information.

Magnifying what is displayed on the screen

You can enlarge information in the product windows using facilities provided by the operating systems on which the product is run. For example, in a Microsoft Windows environment, you can lower the resolution of the screen to enlarge the font sizes of the text on the screen. Refer to the documentation provided by your operating system for more information.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information.

If you have any comments about this information or any other documentation, you can do one of the following actions:

- Complete and submit the Reader Comment Form .
- Send your comments by e-mail to swsdid@de.ibm.com.

Include the documentation name, the part number, the version number, and, if applicable, the specific location of the text you are commenting on (for example, a page number or table number).

What's new

This topic summarizes the significant improvements or enhancements for the product and refers you to the relevant topics for more information.

SH12-6992-00 — October 2013

This edition replaces *IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS*; *IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS: Report Command Reference*, SH12-6964-00.

- The FILE DATATYPE(ACCEL) option for the ACCOUNTING command has been added.
- Usage information about the new Spreadsheet Input Data Generator utility has been added in “SPREADSHEETDD subcommand option” on page 48.
- Changes have been made to “ACCOUNTING command with REPORT subcommand” on page 55.
- The DISTRIBUTE command and all references to it have been removed.
- The SPREADSHEETDD subcommand has been removed from the STATISTICS command.
- The new subcommand option CONVERT has been added to the SAVE subcommand of the ACCOUNTING and STATISTIC commands.
- Modified the “STATISTICS command with REDUCE subcommand” on page 125, “STATISTICS command with REPORT subcommand” on page 121, “STATISTICS command with FILE subcommand” on page 133, and “STATISTICS command with TRACE subcommand” on page 131.

Chapter 1. General information about OMEGAMON XE for DB2 PE report sets

This section describes the subjects that are common to most of OMEGAMON XE for DB2 PE reports or traces.

You can find:

- The OMEGAMON XE for DB2 PE identifiers that are used in reports and traces.
- A description of the input for reports and traces.
- A JCL sample that shows how to produce an OMEGAMON XE for DB2 PE report.
- A list of subcommand options for OMEGAMON XE for DB2 PE.
- The conventions that are used in reports and traces.

For details refer to the following information:

- Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3
- Chapter 3, "The OMEGAMON XE for DB2 PE command stream," on page 11
- Chapter 4, "Output from OMEGAMON XE for DB2 PE reports," on page 13

Chapter 2. OMEGAMON XE for DB2 PE identifiers

This section provides a description of the OMEGAMON XE for DB2 PE identifiers that are used in reports and traces. These identifiers describe the object that OMEGAMON XE for DB2 PE is reporting on.

Usage

DB2 trace records contain identifiers that OMEGAMON XE for DB2 PE uses to:

- Group data
- Order reports
- Identify trace records
- Include or exclude specific data

Identifiers

ACE (Agent control element address)

The absolute hexadecimal address of the DB2 agent control element for the thread. Each work request in DB2 is represented by an agent. When a work request identifies itself to DB2, an agent control element address (ACE) is used to track the agent. The agent can be:

- An allied agent representing a work request that originated in allied address spaces
- A system agent representing a work request internal to DB2

You can use this address to select records for a particular thread. Note that an ACE address can be reused after a thread terminates.

BPID (Buffer pool ID)

The buffer pool ID.

CLASS (DB2 trace class)

DB2 groups records of a similar nature into classes. When running a DB2 performance trace, you can limit the type of information that is traced by selecting one or more trace classes. OMEGAMON XE for DB2 PE can be used to select records by DB2 performance class.

CONNECT (Connection ID)

The connection identifier of the correlation header, which is the ID of the address space that interfaces with DB2. You can, for example, specify the CICS® or IMS™ ID.

CONNTYPE (Connection type)

The type of connection for a thread. You can, for example, specify that you want to include only records that have a connection type of TSO or CICS. Possible values for thread connection types are:

TSO

TSO foreground and background

DB2CALL

DB2 CALL attach

CICS

CICS attach

DLI-BTCH

DL/I batch

IMS-BMP
IMS nontransaction-oriented BMP

IMS-MPP
IMS attach MPP

IMS-CNTL
IMS control region

IMS-TBMP
IMS transaction-oriented BMP

DB2 PRIV
DB2 private protocol

DRDA[®]
DRDA protocol

UTILITY
Utility attach

RRS
Recoverable Resource Manager Services attach

Correlation identifier

A 12-character value identifying the DB2 task in conjunction with the connection ID.

OMEGAMON XE for DB2 PE uses the correlation ID to derive two identifiers: the correlation name and the correlation number:

CORRNAME (Correlation name)

An identifier assigned to a task. This field is a subset of the correlation ID. Its meaning varies with the connection type.

CORRNMBR (Correlation number)

An identifier assigned to a task. This field is a subset of the correlation ID. Its meaning varies with the connection type.

The location of the correlation name and correlation number within the 12-character value depends on the type of connection that the task executes in.

In distributed processing, when the application requester is a DB2 system, the value assigned to the correlation ID at the application server is the same as the value assigned to the application requester. If the application requester is not a DB2 system, the value assigned to the correlation ID at the application server is the name of the job, task, or process the requester is servicing.

For more information about correlation ID translation, see *Monitoring Performance from ISPF*.

DATABASE (Database name)

The name of the DB2 database.

Note that INCLUDE and EXCLUDE processing uses the character names of databases, while DB2 instrumentation records contain the decimal IDs used by DB2. OMEGAMON XE for DB2 PE translates the decimal ID to the character name.

DATASET (Data set name)

The 8-character name of the archive log, active log, or bootstrap data set (BSDS).

ENDUSER (End user ID)

The user ID of the workstation end user. This can be different from the AUTHID used to connect to DB2.

FIELD (Comparison with data in a record field)

This option is used in conjunction with the FIELD command. By using the FIELD command, you can define a value and comparison operator for a data field in a specific IFCID type. You can include or exclude records based on the result of the comparison. For more information, see “FIELD command” on page 147.

GROUP (Group name)

The name of the data-sharing group.

IFCID (Instrumentation Facility Component Identifier)

A decimal identifier that represents a significant DB2 event and appears in the trace records produced by DB2.

INSTANCE (Instance number)

This hexadecimal number is mainly for distributed activity and is part of the LUWID. It can be used to match the activity performed by DBATs, DBAT-distributed threads, and allied-distributed threads. The instance number is allocated at thread creation.

LOCATION (Location name)

The name of a DB2 system. The location name is unique among DB2 systems that can communicate with each other.

If an input data set contains data from several subsystems, at least one of the following identifiers has to be different if OMEGAMON XE for DB2 PE is to distinguish between different subsystems: group name, location name, member name, or subsystem ID.

MAINPACK (Main package)

This identifier can be used to distinguish between plans according to the packages they contain. The representative package is either the first or last package or DBRM in a plan.

This identifier is useful when the name of a plan does not provide satisfactory identification.

The MAINPACK definition is stored in the DPMPARMS member MAINPACK. You can get access to the member by using the MAINPACK Definition Editor panel of the IRF. A different MAINPACK definition can be specified for each unique combination of requester location, connection ID, and plan name.

The default value for MAINPACK is the package ID of the first executed package for any requester location, connection ID, and plan name. If there is no package data available, MAINPACK returns to the default plan name.

MEMBER (Member name)

The name of the data-sharing group member.

OBJECT (Object type)

The type of DB2 resource. Valid values are:

BUFFER

Buffer pool

COLLECT

Collection

DATABASE

Database

DISTTYPE

Distribution type

FUNCTION

Function

PACKAGE

Package

SCHEMA

Schema

APPLPLAN

Application plan

LOBTS

LOB table space

STOGROUP

Storage group

TAB/VIEW

Table or view

USERAUTH

For system privileges like SYSADM or SYSOPR

ORIGAUTH (Original authorization ID)

The original value of the primary authorization ID at the time of connection to DB2 and before it can be changed by any authorization exits.

In distributed processing and if the requester is a DB2 system, the value of the original authorization ID at the application server is the same as the value assigned to the application requester. If the application requester is not a DB2 system, the value of the original authorization ID at the application server is the user ID used during the initial connection with the application server.

PACKAGE (Package information)

This identifier is used to identify a package or DBRM. It is displayed in traces and reports in the headings of corresponding package data blocks.

PAGESET (Page set name)

The name of the page set.

Note that INCLUDE and EXCLUDE processing uses the character names of page sets, while DB2 instrumentation records contain the decimal IDs used by DB2. OMEGAMON XE for DB2 PE translates the decimal ID to the character name.

PARTNBR (Partition number)

PARTNBR identifies a partition of a table space by its number.

Table spaces can be partitioned or non-partitioned. The value for this identifier is zero if the table is not partitioned. This identifier is only valid for the ORDER command used with IOACTIVITY. You must specify PAGESET before you can specify PARTNBR.

PLANNAME (Plan name)

The plan name from the correlation header. A plan is a control structure produced during the bind process and used by DB2 to process SQL statements encountered during statement execution.

To receive better identification and granulation of bind and utility traces and reports, it is necessary to filter and order data. The plan name on a bind event (constant DSNBIND) is replaced by the program name, and the plan name on a utility event (constant DSNUTIL) is replaced by the utility name.

PRIMAUTH (Primary authorization ID) or AUTHID (Authorization ID)

The two terms are interchangeable. This is the primary authorization ID you entered at signon or identify.

REQLOC (Requester location)

For distributed processing, this is the location requesting the work. If the requester location is not a DB2 subsystem, or is not recognized by DB2, the logical unit name from the DRDA LUWID is printed instead of the requester location name. The logical unit name is enclosed in less than (<) and greater than (>) symbols.

For TCP/IP connections, the requester location can contain the dotted-decimal IP address.

Note: The IP address must be specified in a specific format. For example, REQLOC('::10.4.1.237').

RESOURCETYPE (Resource type)

The type of lock resource. You can specify one of the following values:

DATAPAGE

Data page locking

DATABASE

Locking of the DBD

PAGESET

Page set locking

DATASET

Locking of partitioned data sets

SKCT

Skeleton cursor table locking

INDEX

Index page locking

TABLE

Table locking

SKPT

Skeleton package table locking

COLLECT

Collection ID locking

DRAIN

All types of drain locking

ROW

Data row lock

OTHER

All unlisted resource types

RMID (Resource manager identifier)

The decimal identifier of a DB2 resource manager. You can use RMID to select a particular type of activity. For example, RMID 6 selects records associated with storage management.

SQLCODE

The SQL return code. This identifier is only valid with the SQLACTIVITY TRACE or SQLACTIVITY FILE commands. It can be used to include statements that completed, for example, with a specific error.

See "INCLUDE and EXCLUDE subcommand options" on page 33 for syntax details, and "SQLACTIVITY command with TRACE subcommand" on page 115 for examples using this filter.

SUBSYSTEMID (Subsystem ID)

The ID of the DB2 subsystem that generated the data.

THREADTYPE (Thread type category)

The type of thread that you want included in a report. You can specify one or more of the following values:

ALLIED

Threads without distributed activity. An allied thread does not involve distributed activity, that is, it is not initiated by a remote location and does not request data from another location.

ALLIED_DIST

Threads that request work from remote locations. An allied-distributed thread is not initiated by a remote location, but it requests data from one or more server locations.

This category covers thread type ALLDDIST. It can be reported in Accounting Report if ordered by THREADTYPE.

DBAT

Threads performed on behalf of remote locations.

In the Accounting report, if data is ordered by THREADTYPE, this category covers following thread types:

- DBAT - Indicates accumulated data of threads that are initiated, created, and performing work on behalf of a remote (requester) location.
- DBATDP - Indicates accumulated data of DBAT duplicate threads.
- DBATDIST - Indicates accumulated data of DBAT distributed threads that are initiated by a requester location and executed by the server location that in turn requests data from another server location.
- DBATDICP - Indicates accumulated data of DBAT distributed and copy threads.
- DBATDIDP - Indicates accumulated data of DBAT distributed and duplicate threads.

TRANSACT (End user transaction name)

The transaction or application that the client is running.

TYPE (Event type)

Specifies which event types are to be included in, or excluded from, the lock detail trace. Valid event type values are:

IRLMREQ

Lock, unlock, change, query, and notify requests

CLAIMREQ

Claim acquire, claim change, and claim release

DRAINREQ

Drain request and drain release

PLOCKREQ

Page set or partition and page P-lock requests

IRLMSUSP

The beginning of lock, unlock, change, query, and notify suspensions

DRAINSUSP

The beginning of drain suspensions

LATCHSUSP

The beginning of page latch suspensions

IRLMRES

The end (resumption) of lock, unlock, change, query, and notify suspensions

DRAINRES

The end (resumption) of drain suspensions

LATCHRES

The end (resumption) of page latch suspensions

TIMEOUT

Timeouts

DEADLOCK

Deadlocks

LOCKSUMMARY

Lock summary events

LOCKAVOID

Successful lock avoidance events

The default is *all* event types.

Note: TYPE can also be used with the REDUCE and FILE subcommands of LOCKING, with the following limitations:

- Valid types for REDUCE are: IRLMRES, DRAINRES, and LATCHRES.
- Valid types for FILE are: IRLMREQ, CLAIMREQ, DRAINREQ, and LOCKAVOID.

If values other than those listed are used with INCLUDE, REDUCE, or FILE, an empty report or file is produced.

If values not relevant to REDUCE or FILE are used with EXCLUDE, the event type is not filtered.

WSNAME (End user workstation name)

The user's workstation name.

Chapter 3. The OMEGAMON XE for DB2 PE command stream

A JCL command stream consists of DD statements that determine the data sets used with reports and traces, and the commands, subcommands and subcommand options that determine the type of report or trace to be produced.

Usage

You can create OMEGAMON XE for DB2 PE reports and traces as follows:

- Interactively, by choosing the *Create and execute reporting commands* option from the main menu. This invokes the Interactive Report Facility (IRF), which lets you interactively specify the type of report and further options and then composes and submits a corresponding batch report command stream. The IRF is described in the *Reporting User's Guide*.
- Manually, by using the ISPF editor and creating JCL command streams to generate reports and traces. This method is described herein, together with the OMEGAMON XE for DB2 PE commands, subcommands, options and keywords that can be used in commands streams.

Usage notes

- Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See the individual DD statement descriptions for more information.
- The final EXEC statement is required. If you do not include the EXEC statement in your JCL, no report is produced. All statements following the EXEC statement are ignored.
- The syntax of your JCL is checked and written to the DPMLOG data set together with any information, warning or error messages raised.

Example

Figure 1 on page 12 shows a sample of the JCL required to produce OMEGAMON XE for DB2 PE reports and traces. In this sample, several place holders (in lower case, italic characters) are used. The place holders *command*, *subcommand* and *option* determine the type of report or trace to be produced. The actual OMEGAMON XE for DB2 PE commands are described in the relevant sections of this information. Other place holders (*cc* and *x*) are described in context in Chapter 5, "DD statements," on page 15.

```

//          PEMAIN EXEC PGM=FPECMAIN
// * FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=FPE.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=FPE.FPELIB.DPMIN,DISP=SHR
//DPMLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//EXCPTDD DD DSN=FPE.EXCEPT.THRESH,DISP=OLD
//EXTRCDD1 DD SYSOUT=*
//EXFILDD1 DD DSN=FPE.EXCEPT.LOGFILE,DISP=OLD
//SYSPRMDD DD SYSOUT=*
//DPMOUTDD DD DSN=FPE.FPELIB.DPMOUT.DATA,DISP=OLD
//JSSRSDD DD DSN=FPE.FPELIB.JSSRS.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
// * FOLLOWING ARE REPORT SET DDNAMES
//ccWORK DD DSN=FPE.FPELIB.op.WORKDD,DISP=OLD
//ccRPTDD DD SYSOUT=*
//ccTRCDD1 DD SYSOUT=*
//ccSAVDD DD DSN=FPE.FPELIB.opSAV.DATA,DISP=OLD
//ccRSTDD DD DSN=FPE.FPELIB.opRST.DATA,DISP=SHR
//ccFILDDx DD DSN=FPE.FPELIB.op.FILE,DISP=OLD
// * FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *
command
  subcommand
    option
    option
  subcommand
    option
command
  subcommand
    option
    option
  subcommand
    option
EXEC

```

Figure 1. Sample JCL

Chapter 4. Output from OMEGAMON XE for DB2 PE reports

This section helps you understand output values, date formats, and time formats.

Usage

The output from OMEGAMON XE for DB2 PE is one, or more, data sets containing the reports you requested, and a set of log file data sets. These log files are described in the *Report Reference*. The reports and traces produced are described and explained in detail in the relevant sections of this information.

Large and missing values in reports

Values printed on reports can be either total values or average values.

If there is insufficient space to print a value on a report or trace, a rounded value is printed followed by one of the following letters to indicate the magnitude:

K	kilo — 10^3
M	mega — 10^6
G	giga — 10^9
T	tera — 10^{12}
P	peta — 10^{15}
E	exa — 10^{18}

The letter is printed directly after the number, without blank spaces. There can, however, be decimal places, as follows:

- Valid conversions of 12 345 include 12K, 12.35K, and 12.3K.
- Valid conversions of 1 234 567 include 1M, 1.2346M, and 1235K.

If a counter value or specific information in reports, in windows, or on panels is not shown, the following notation is used to indicate the reason:

- N/A** Not applicable is shown if DB2 never produces a counter value in a specific context. Examples are:
- A counter is not available in one DB2 version.
 - Counters are mutually exclusive.
- N/C** Not calculated is shown for a derived field where the value cannot be calculated or is useless. Examples are:
- A divide by zero (percentages, ratios).
 - Suppression of negative elapsed time values.
 - Required counter values for calculation marked as N/A or N/P.
 - Insufficient data or small counter values to allow significant statements (meaningless or misleading averages).
- N/P** Not present is shown for a field where DB2 can present values, but does not in this instance. Examples are:
- When counter values are not generated because of operational conditions (a trace class is not active).
 - An application does not provide a value because it is optional.

Default date format

By default, OMEGAMON XE for DB2 PE shows dates as *mm/dd/yy*. You can change this format to suit your own national language, or corporate convention by using DATEFORMAT. For more information, see “DATEFORMAT subcommand option” on page 27.

This example shows the default date format as printed in the header of an Accounting short report.

```
LOCATION: STM4D61Y          OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V5R1M1)          PAGE: 1-1
GROUP: N/P                ACCOUNTING REPORT - SHORT          REQUESTED FROM: 11/06/12 23:45
MEMBER: N/P                ORDER: ENDUSER-WSNAME-TRANSACT          TO: 11/06/12 23:50
SUBSYSTEM: Y61Y            SCOPE: MEMBER          INTERVAL FROM: 11/06/12 23:48:01.86
DB2 VERSION: 10                                TO: 11/06/12 23:53:34.20
```

Figure 2. Accounting report - short

Elapsed time formats

Time values are presented in one of the following formats:

dd hh:mm:ss.fffffff

where:

dd represents days

hh represents hours

mm represents minutes

ss represents seconds

fffffff represents the fractions of a second up to 8 decimal places.

For example, a time value of 1:30:25.10 represents 1 hour, 30 minutes, and 25.1 seconds.

Some of the reports that use this format might not report days (*dd*) or hours (*hh*).

sssssss.fffffff

where:

sssssss

represents seconds

fffffff represents the fractions of a second up to 8 decimal places.

The actual number of decimal places varies from one field to another.

Some time fields can be rounded. If there is insufficient space to print a time value, the time is rounded by removing decimal places as required. For elapsed times, a rounded value is printed.

Chapter 5. DD statements

Here you find a description of the DD statements used for OMEGAMON XE for DB2 PE.

This section describes the DD statements and the data sets shown in Figure 1 on page 12.

The values for RECFM, LRECL, and BLKSIZE shown for some data sets are the values that OMEGAMON XE for DB2 PE generates at run time.

Note: The generated value for BLKSIZE is not mandatory, but it is recommended. Do not override the values for RECFM and LRECL.

ACMEM nn statements

The ACMEM nn statements are applicable to Accounting report set.

Usage

If you use the TRACE subcommand, OMEGAMON XE for DB2 PE temporarily uses an ACMEM nn work data set for each DB2 member that occurs in the input records. For example, if the input data set contains, or data sets contain, data from three different members, the ddnames ACMEM01, ACMEM02 and ACMEM03 are used. The work data sets are normally created on the MVS-defined work volumes, and on completion of the task.

Usage notes

- Only include ACMEM nn in your JCL if you want to control the placement or size of the data set. If your input has a large amount of data for a specific member, you might get a B37 abend on the ACMEM nn work data set. In that event, specify ACMEM nn .
- OMEGAMON XE for DB2 PE allocates 68 MB for a work data set by default. Define it as a temporary data set. As a guide, if the number of accounting-related input records for a specific member that satisfy the GLOBAL and ACCOUNTING command criteria exceeds 40 000, specify ACMEM nn .
- Do not specify DUMMY or DISP=MOD for this data set.

RECFM:

VBS

LRECL:

32 756

BLKSIZE:

6 233

DPMLOG statement

The DPMLOG statement is applicable to all report sets.

Usage

OMEGAMON XE for DB2 PE command processor messages are written to DPMLOG.

If DPMLLOG is omitted, it is dynamically allocated to the SYSOUT message class of the job.

Values

RECFM:
FBA
LRECL:
133
BLKSIZE:
6 251

DPMOUTDD statement

The DPMOUTDD statement is applicable to all report sets.

Usage

If you:

- *Do not specify DPMOUTDD:*
All DPMOUT-related processing is completed in storage. Only the records that relate to the report set commands in the same step are processed. No data is externalized.
- *Specify DPMOUTDD:*
All records that satisfy GLOBAL FROM TO and INCLUDE or EXCLUDE selection criteria are reformatted to OMEGAMON XE for DB2 PE trace format, sorted by time sequence, location and group name, and written to the data set specified by DPMOUTDD.

Usage notes

- Include DPMOUTDD in your JCL only if you want to retain a copy of the input data as filtered by any GLOBAL options.
- Do not specify DUMMY or DISP=MOD for DPMOUTDD.
- You can specify a permanent or temporary data set for DPMOUTDD.
- Only specify the temporary DPMOUT data set if the number of filtered input records is likely to exceed 45 000. By default, OMEGAMON XE for DB2 PE allocates up to 68 MB for a work data set.
- The size of the DPMOUT data set depends on the number of input records and the GLOBAL filters. The input records are the IFCID records included in the DB2-related SMF record types 100, 101, and 102. One input record occupies approximately 1.5 KB of the space in the DPMOUT data set.
- When the input data set is composed mostly of DB2-related records, a good size for DPMOUT is 1.6 times the size of the input data set.

Values

RECFM:
VBS
LRECL:
32 756
BLKSIZE:
6 233

DPMPARMS statement

The DPMPARMS statement is applicable to all report sets.

Usage

The DPMPARMS data set is used to store changes that you have made to standard OMEGAMON XE for DB2 PE settings. For example, if you tailor your own report layout, it is stored in the DPMPARMS data set. Specify DPMPARMS if you want to use this layout. The modified OMEGAMON XE for DB2 PE settings stored in DPMPARMS are:

- Time zone processing
- Exception messages
- MAINPACK definition
- Correlation translation
- UTR layouts

Usage notes

Do not specify DUMMY for this data set.

Values

DPMPARMS must be a partitioned data set. Use the following attributes if you want to preallocate a new DPMPARMS data set. You should increase the number of directory blocks if you intend to tailor many report layouts.

RECFM:
 FB
LRECL:
 80
BLKSIZE:
 6 160

EXCPTDD statement

The EXCPTDD statement is applicable to Accounting and Statistics report sets.

Usage

The exception threshold data set EXCPTDD contains the user-defined exception thresholds.

This ddname is required if you want to produce an exception log or if you specified the EXCEPTION subcommand option with the TRACE, REPORT, or FILE subcommand.

Values

RECFM:
 VB
LRECL:
 255
BLKSIZE:
 6 233

EXFILDD1 statement

The EXFILDD1 statement is applicable to Accounting and Statistics report sets.

Usage

The data for the exception log file data set is written to EXFILDD1.

This DD statement is required if you want to produce an exception log file data set.

Usage notes

The exception threshold data set, as defined in EXCPTDD, is also required to produce an exception log file data set.

Values

RECFM:
 VB
LRECL:
 512
BLKSIZE:
 4 096

EXPRODDA statement

The EXPRODDA statement is applicable to Accounting and Statistics report sets.

Usage

The EXPRODDA data set defines the exception profile.

If EXPRODDA is omitted, it is dynamically allocated by OMEGAMON XE for DB2 PE and deleted after the job completes. If your input has a large amount of data for a specific member, you might experience a B37 abend on the work data set. If you want to control the placement or size of this data set, define a temporary data set.

Usage notes

Do not specify DUMMY or DISP=MOD for this data set.

Values

RECFM:
 VBS
LRECL:
 32 756
BLKSIZE:
 6 233

EXTRCDD1 statement

The EXTRCDD1 statement is applicable to Accounting and Statistics report sets.

Usage

The data for the exception log is written to EXTRCDD1 data set. This DD statement is required if you want to produce an exception log. For more information, see the *Report Reference*.

Usage notes

The exception threshold data set, as defined in EXCPTDD, is also required to produce an exception log.

Values

RECFM:

FBA

LRECL:

133

BLKSIZE:

6 251

INPUTDD statement

The INPUTDD statement is applicable to all report sets.

Usage

The INPUTDD data set stores changes that you have made to standard OMEGAMON XE for DB2 PE settings. It lists the input data sets containing the DB2 performance data created by the DB2 trace facility. You can process several input data sets. These data sets are concatenated in the JCL to create one logical data set. The input data sets can be in SMF, GTF, or DPMOUT format or data sets generated by the Collect Report Data function of the Online Monitor. The normal rules for concatenating data sets apply. If DFSORT is used, see the *DFSORT Application Programming Guide* for rules governing the concatenation of data sets.

Note: INPUTDD is not required if you use RESTORE REPORT for Accounting or Statistics.

Values

The default ddname for the input data set is INPUTDD. You can specify another ddname by using the INPUTDD option of the GLOBAL command. If you specify another ddname, make sure that your JCL includes a valid DD statement for the new name. For more information, see "GLOBAL command" on page 149.

JOBSUMDD statement

The JOBSUMDD statement is applicable to all report sets.

Usage

The job summary log and the IFCID frequency distribution log are written to JOBSUMDD.

This ddname is not required unless you want these logs. For more information, see the *Report Reference*.

Values
RECFM:
 FBA
LRECL:
 133
BLKSIZE:
 6 251

JSSRSDD statement

The JSSRSDD statement applies to Accounting and Statistics report sets.

Usage

Job summary data is written to JSSRSDD when a SAVE subcommand is processed, and is restored from JSSRSDD when a RESTORE subcommand is processed.

If you are restoring data, the data set defined by JSSRSDD and the data set defined by *xxRSTDD* should match, that is, be produced by the same Save operation. For more information, see the *Report Reference*.

Usage notes

- JSSRSDD is optional.
- If you omit JSSRSDD, information about the previous processing of saved data is not restored, and information about current processing is not saved.
- The VSAM data set defined by JSSRSDD must exist before you run OMEGAMON XE for DB2 PE. You can do one of the following:
 - Specify an existing data set from a previous OMEGAMON XE for DB2 PE run (when restoring data). If an existing data set is used and the SAVE subcommand is specified, the new job summary data is added to the previous content.
 - Specify a new data set allocated by using the IDCAMS DEFINE CLUSTER function.
- Do not specify DUMMY for JSSRSDD.

For information about the allocation of the save data set, see *Monitoring Performance from ISPF*.

SYSIN statement

The SYSIN statement is applicable to all report sets.

Usage

SYSIN contains the commands of each OMEGAMON XE for DB2 PE report set that are input to OMEGAMON XE for DB2 PE. This DD statement is required.

Values
RECFM:
 FB
LRECL:
 80
BLKSIZE:
 6 160

SYSOUT statement

The SYSOUT statement is applicable to all report sets.

Usage

Messages from DFSORT are written to the ddname SYSOUT.

If SYSOUT is omitted, it is dynamically allocated to the SYSOUT message class of the job.

Values

RECFM:

FBA

LRECL:

133

BLKSIZE:

6 251

SYSRMDD statement

The SYSRMDD statement is applicable to all report sets.

Usage

The System Parameters report is written to SYSRMDD.

Specify this ddname if you want a System Parameters report.

Usage notes

- This ddname is optional.
- The default ddname for the System Parameters report is SYSRMDD.
- You can specify another ddname by using the SYSRMDD option of the OMEGAMON XE for DB2 PE GLOBAL command.
- If you specify another ddname, make sure that your JCL contains a valid DD statement for the new ddname. For information about the GLOBAL command, see “GLOBAL command” on page 149.

Values

RECFM:

FBA

LRECL:

133

BLKSIZE:

6 251

ccFILDD1 statements

The ccFILDD1 statements are applicable to Accounting (*cc=AC*), Audit (*cc=AU*), Locking (*cc=LO*), Record Trace (*cc=RT*), Statistics (*cc=ST*) and System Parameters (*cc=SY*) report sets.

Usage

By default, the output from the FILE subcommand is written to ccFILDD1.

Usage notes

- You can specify a different ddname by using the DDNAME option of the FILE subcommand.
- If you specify a different ddname, your JCL must contain a valid DD statement for the specified ddname.
- If you do not specify a different ddname, your JCL must contain a valid DD statement for the default ddname.

Values

RECFM:
VB
LRECL:
9 072
BLKSIZE:
9 076

ccRPTDD statements

The *ccRPTDD* statements are applicable to Accounting (*cc=AC*), Audit (*cc=AU*), I/O Activity (*cc=IO*), Locking (*cc=LO*), Record Trace (*cc=RT*), SQL Activity (*cc=SQ*), Statistics (*cc=ST*), and Utility Activity (*cc=UT*) report sets.

Usage

ccRPTDD is the default output ddname for the REPORT subcommand.

The reports are written to *ccRPTDD* in the sequence corresponding to the REPORT subcommands. If *ccRPTDD* is omitted, it is dynamically allocated to the SYSOUT message class of the job.

Usage notes

You can specify a different ddname by using the DDNAME option of each REPORT subcommand. If you specify a different ddname, your JCL must contain a valid DD statement for the specified ddname.

Values

RECFM:
FBA
LRECL:
133
BLKSIZE:
6 251

ccRSTDD statements

The *ccRSTDD* statements are applicable to Accounting (*cc=AC*) and Statistics (*cc=ST*) report sets.

Usage

Data processed by the RESTORE subcommand is read from *ccRSTDD* by default. A valid DD statement is required if your job stream contains a RESTORE subcommand.

Usage notes

- You can specify a different ddname by using the DDNAME option of the RESTORE subcommand.
- If your job stream contains a RESTORE subcommand that uses the DDNAME option, your JCL must contain a valid DD statement for the specified ddname. If your job stream contains a RESTORE subcommand that does not use the DDNAME option, your JCL must contain a valid DD statement for the default ddname.
- Do not specify DUMMY for ccRSTDD.

ccSAVDD statements

The ccSAVDD statements are applicable to Accounting (cc=AC) and Statistics (cc=ST) report sets.

Usage

Reduced data processed by the SAVE subcommand is, by default, written as a VSAM data set to ccSAVDD. Batch Accounting and Statistics offer a CONVERT option of the SAVE subcommand that requires to specify a ddname that is assigned to a sequential data set. A valid DD statement is required if your job stream contains a SAVE subcommand.

Usage notes

- You can specify another ddname by using the DDNAME option of the SAVE subcommand. If you specify a different ddname, your JCL must contain a valid DD statement for the specified ddname. If you do not specify a different ddname, your JCL must contain a valid DD statement for the default ddname.
- The VSAM data set (that is the default in Batch Accounting and Statistics) defined by ccSAVDD must exist before you run OMEGAMON XE for DB2 PE. Either specify an existing data set from a previous OMEGAMON XE for DB2 PE run (when restoring data), or specify a new data set allocated by using the IDCAMS DEFINE CLUSTER function. Note that the existing contents of the data set are lost unless the DDNAME options of both SAVE and RESTORE subcommands specify the same ddname or data set.
- Do not specify DUMMY for ccSAVDD.

For information on the allocation of the save data set or sequential data set, see *Report Reference*.

ccTRCDDx statements

The ccTRCDDx statements are applicable to Accounting (cc=AC), Audit (cc=AU), I/O Activity (cc=IO), Locking (cc=LO), Record Trace (cc=RT), SQL Activity (cc=SQ), Statistics (cc=ST), System Parameters (cc=SY) and Utility Activity (cc=UT) report sets.

Usage

ccTRCDD is the output ddname for the TRACE subcommand.

If ccTRCDDx is omitted, it is dynamically allocated to the SYSOUT message class of the job.

Usage notes

- Up to five traces can be run in one job step. The default ddnames for the traces are *ccTRCDD1* to *ccTRCDD5*, consecutively.
- You can specify a different ddname by using the DDNAME option in the corresponding TRACE subcommand. If you specify a different ddname, your JCL must contain a valid DD statement for the specified ddname.

Values

RECFM:

FBA

LRECL:

133

BLKSIZE:

6 251

ccWORK statements

The *ccWORK* statements are applicable to Accounting (*cc=AC*), Audit (*cc=AUD*), I/O Activity (*cc=IO*), Locking (*cc=LO*), Record Trace (*cc=RT*), SQL Activity (*cc=SQL*), Statistics (*cc=ST*), and Utility Activity (*cc=UT*) report sets.

Usage

ccWORK controls the placement or size of the data set.

If you reduce data, OMEGAMON XE for DB2 PE uses a temporary REDUCE work data set to provide virtual storage constraint relief. This is normally created on the MVS-defined work volumes, and deleted by OMEGAMON XE for DB2 PE. Include *ccWORK* in your JCL when you want to control the placement or size of the data set or if you get a B37 abend. This can happen if you are trying to reduce a large amount of data with a large number of different OMEGAMON XE for DB2 PE identifiers and a short reduction interval.

Usage notes

- If you specify *ccWORK*, OMEGAMON XE for DB2 PE allocates up to 68 MB for a work data set by default. Define it as a temporary data set.
- As a guide, if the number of accounting-related input records that satisfy the GLOBAL and report set command criteria exceeds 45 000 and the reduction ratio is very low, specify *ccWORK*.
- Do not specify DUMMY or DISP=MOD for this data set.

Values

RECFM:

VBS

LRECL:

32 756

BLKSIZE:

6 233

Chapter 6. OMEGAMON XE for DB2 PE subcommand options

This section introduces you to subcommand options.

Usage

The following sections provide a comprehensive description of all subcommand options to avoid lengthy recurrences.

Usage notes

The general structure of commands is:

```
command  
  subcommand  
    option(keyword)
```

- A single command can be followed by one or more subcommands.
- A subcommand can be followed by one or more options.
- An option might have further keywords.

The commands are described in Chapter 7, “OMEGAMON XE for DB2 PE commands,” on page 51 and Chapter 8, “Auxiliary commands,” on page 145, together with the associated subcommands and subcommand options.

The following subcommand options can be used to specify how times are reported, and the intervals and time frames of your reports (time functions):

- “DATEFORMAT subcommand option” on page 27
- “FROM/TO subcommand options” on page 30
- “INTERVAL subcommand option” on page 37
- “BOUNDARY subcommand option” on page 26

Supplemental information concerning these subcommand options is provided in “Notes on calculating intervals” on page 26.

The following subcommand options can be used to control the amount of data that is reported and to control the way data is consolidated in reports (filter functions):

- “ORDER subcommand option” on page 46
- “INCLUDE and EXCLUDE subcommand options” on page 33
- “EXCEPTION and NOEXCEPTION subcommand options” on page 30
- “DDNAME subcommand option” on page 29
- “TOP subcommand option” on page 49
- “LAYOUT subcommand option” on page 38
- “SCOPE subcommand option” on page 48
- “TYPE subcommand option” on page 50
- “LEVEL subcommand option” on page 45
- “SUMMARIZEBY subcommand option” on page 49
- “WORKLOAD subcommand option” on page 50
- “SORTBY subcommand option” on page 48
- “LIMIT subcommand option” on page 45

BOUNDARY subcommand option

The BOUNDARY subcommand option controls the alignment of the intervals.

Usage

The BOUNDARY subcommand option is used to control the alignment of the intervals used to summarize records in the reduction process.

Usage notes

- Boundary is ignored for INTERVAL (0). The interval starts at the timestamp of the first record that satisfies FROM.
- If you use the RESTORE and REDUCE subcommands in the same job stream, the INTERVAL and BOUNDARY options specified in REDUCE should match the INTERVAL and BOUNDARY options that were used to reduce the data being restored. If these values are different, the interval and boundary from the restored data is used.

Rules

- The range is from 0 to 60, and indicates minutes past the hour.
- A boundary of 0 specifies that intervals are aligned with the number of minutes in the FROM time.
- A boundary of 60 specifies that intervals are aligned with hour boundaries.
- If no boundary is specified, the default is the boundary specified in the GLOBAL command.
- If no boundary is specified in global, the default is 60.

Notes on calculating intervals

The start time of the first interval processed by REDUCE is influenced by BOUNDARY, INTERVAL, and FROM.

OMEGAMON XE for DB2 PE attempts to reduce all data that falls between FROM and TO dates and times. The first interval processed starts at a time aligned with BOUNDARY, at or before the FROM time. If an interval cannot be aligned with the FROM time, the first properly aligned interval starting before the FROM time is used.

Although there is no restriction on the INTERVAL and BOUNDARY combination, your specification should comply with the following recommendations:

- For intervals of less than 60 (excluding 0), there should be a whole number of intervals in an hour. Choose one of the following values:
 - 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, or 30.
- For intervals of 60 or greater, there should be a whole number of intervals in a day. Choose one of the following values:
 - 60, 120, 180, 240, 360, 480, 720, or 1 440.
- For intervals of one day (1 440) or greater, INTERVAL should be a multiple of 1 440.
- Select your interval and boundary so that the first interval starts at the FROM time.

Example of interval calculation, aligning to the start of the hour

BOUNDARY (60) aligns the start time of intervals at the start of an hour, so the first interval starts at the FROM time (08:00). Subsequent intervals start every 30 minutes (08:30, 09:00, and 09:30 each day).

```
⋮  
REDUCE  
  FROM      (,08:00)  
  TO        (,10:00)  
  INTERVAL  (30)  
  BOUNDARY  (60)  
⋮
```

Example of interval calculation, covering a day

The following defaults are applied:

- For FROM, all dates and a time of 00:00:00.00.
- For TO, all dates and a time of 23:59:59.99.

BOUNDARY (60) aligns the start time of intervals at the start of an hour, so the first interval starts at the FROM time (00:00). Subsequent intervals cover 1 440 minutes or one day; an interval starts at 00:00 each day.

```
⋮  
REDUCE  
  INTERVAL (1440)  
  BOUNDARY (60)  
⋮
```

Example of interval calculation, starting every hour

BOUNDARY (60) aligns the start time of intervals at the start of an hour, so the first interval starts at the hour of the FROM time (08:00). Subsequent intervals start every hour (09:00, 10:00, and 11:00).

```
⋮  
REDUCE  
  FROM      (,08:30)  
  TO        (,12:00)  
  INTERVAL  (60)  
  BOUNDARY  (60)  
REPORT  
  FROM      (,08:30)  
  TO        (,12:00)  
⋮
```

DATEFORMAT subcommand option

The DATEFORMAT subcommand option changes the format of dates.

Usage

The DATEFORMAT subcommand option is used to change the format of dates specified in the FROM and TO subcommand options, and displayed on OMEGAMON XE for DB2 PE reports, traces, and logs.

Format

The DATEFORMAT parameter must be 8 characters long and contain:

dd Day
mm Month
yy Year

Format rules

- You can specify the day, month, and year in any order.
- A single character delimiter is also required in the third and fifth positions.
- You can delimit the day, month, and year with either a slash (/), dash (-), period (.), or any combination of these delimiters.
- The DATEFORMAT parameter can be abbreviated to DF.
- If you use FROM or TO, you must specify the dates in exactly the same format you have defined in the DATEFORMAT parameter. For example, the following formats are all valid DATEFORMAT definitions:
 - *yy-mm-dd*
 - *mm/dd-yy*
 - *dd.mm/yy*
- If you do not specify the DATEFORMAT parameter in your JCL, dates are displayed in the OMEGAMON XE for DB2 PE default format, *mm/dd/yy*.

Sample Accounting JCL with DATEFORMAT subcommand option specified

For example, you might want to display the year before the month and day, separated by a delimiter, on a report. In this case, you would specify *yy/mm/dd* in the DATEFORMAT parameter.

In the JCL example shown below, DATEFORMAT is used to specify the date format as *yy-mm-dd*. The relevant information is highlighted.

```
//PERFORMANCE EXPERT JOB (INSTALLATION DEPENDENCIES)
//*
//* =====*
//*          DB2 PERFORMANCE EXPERT JCL          *
//* =====*
//*
//          PEMAIN EXEC PGM=FPECMMAIN, PARM='DATEFORMAT=YY-MM-DD'
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//INPUTDD DD DSN=FPE.FPELIB.DPMIN,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//ACRPTDD DD SYSOUT=*
//SYSIN DD *
:
:
ACCOUNTING
REPORT
FROM (11-06-12,09:45)
TO (11-06-12,09:55)
:
:
EXEC
```

Accounting Report - Short

The JCL example shown above produces a report similar to the one shown here. The relevant information is highlighted.

LOCATION: MYLOC	OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V5R1M1)	PAGE: 1-1
GROUP: N/P	ACCOUNTING REPORT - SHORT	REQUESTED FROM: 11-06-12 09:47:00.00
MEMBER: N/P		TO: 11-06-12 09:51:00.00
SUBSYSTEM: MYSS	ORDER: ENDUSER-WSNAME-TRANSACTION	INTERVAL FROM: 11-06-12 09:48:43.55
DB2 VERSION: 10	SCOPE: MEMBER	TO: 11-06-12 09:50:52.23

ENDUSER	#OCCURS	#COMMIT	INSERTS	OPENS	PREPARE	CLASS2	EL.TIME	BUF.UPDT	LOCK	SUS
WSNAME	#DISTR	SELECTS	UPDATES	CLOSES	CLASS1	EL.TIME	CLASS2	CPUTIME	SYN.READ	#LOCKOUT
TRANSACTION	#ROLLBK	FETCHES	MERGES	DELETES	CLASS1	CPUTIME	GETPAGES	TOT.PREF		
MYNAME	85210	85439	0.50	0.17		0.17		0.002034	1.84	0.01
MYWS	85210	0.00	0.17	0.00		0.002333		0.000264	0.00	0
MYTX1	0	0.17	0.00	0.17		0.000308		3.01	0.00	

PROGRAM NAME	TYPE	#OCCURS	#ALLOCS	SQLSTMT	CL7	ELAP.TIME	CL7 CPU TIME	CL8 SUSP.TIME	CL8 SUSP
ROLLUP	PACKAGE	85210	85210	N/P		0.002034	0.000264	0.001630	1.20

ANOTHER name	1250	1249	0.00	1.96		1.98		0.082652	61.71	0.06
MYWS	1250	0.00	0.00	0.26		0.164474		0.006964	3.22	0
MYtx2	2	1.92	0.00	0.00		0.007150		271.27	10.31	

PROGRAM NAME	TYPE	#OCCURS	#ALLOCS	SQLSTMT	CL7	ELAP.TIME	CL7 CPU TIME	CL8 SUSP.TIME	CL8 SUSP
ROLLUP	PACKAGE	1250	1250	N/P		0.082652	0.006964	0.065359	8.84

ANOTHER name	108170	0	1.00	0.00		0.00		0.022644	8.98	0.11
MYWS	108170	0.00	0.00	0.00		0.031945		0.000567	0.00	0
MyTx3	108206	0.00	0.00	0.00		0.000643		12.95	0.00	

PROGRAM NAME	TYPE	#OCCURS	#ALLOCS	SQLSTMT	CL7	ELAP.TIME	CL7 CPU TIME	CL8 SUSP.TIME	CL8 SUSP
ROLLUP	PACKAGE	108170	108170	N/P		0.022644	0.000567	0.020832	4.08

*** SUB-TOTAL ***										
ANOTHER name	109420	1249	0.99	0.02		0.02		0.023330	9.58	0.11
MYWS	109420	0.00	0.00	0.00		0.033459		0.000640	0.04	0
	108208	0.02	0.00	0.00		0.000717		15.90	0.12	

PROGRAM NAME	TYPE	#OCCURS	#ALLOCS	SQLSTMT	CL7	ELAP.TIME	CL7 CPU TIME	CL8 SUSP.TIME	CL8 SUSP
ALL PROG	PACKAGE	109420	109420	N/P		0.023330	0.000640	0.021341	4.14

*** GRAND TOTAL ***										
	194630	86688	0.78	0.08		0.09		0.014006	6.19	0.06
	194630	0.00	0.07	0.00		0.019832		0.000475	0.02	0
	108208	0.08	0.00	0.07		0.000538		10.26	0.07	

PROGRAM NAME	TYPE	#OCCURS	#ALLOCS	SQLSTMT	CL7	ELAP.TIME	CL7 CPU TIME	CL8 SUSP.TIME	CL8 SUSP
ALL PROG	PACKAGE	194630	194630	N/P		0.014006	0.000475	0.012711	2.85

DDNAME subcommand option

The DDNAME subcommand option specifies the data set where the report is written.

Usage

Use DDNAME to specify the data set where the report is written. You can specify any valid ddname including the default, provided that your JCL contains a DD statement for it. If a DD statement is omitted, it will be dynamically allocated to the SYSOUT message class of the job.

Default ddnames

The default ddnames are shown in the following table. *x* represents a number equal to or greater than one and equal to or less than the maximum number of subcommand invocations allowed.

Table 2. Default ddnames for DDNAME subcommand option

	Report Set								
	Accounting	Audit	I/O Activity	Locking	Record Trace	SQL Activity	Statistics	System Parameters	Utility Activity
FILE	ACFILDDx	AUFILDDx		LOFILDDx	RTFILDDx	SQFILDDx	STFILDDx	SYFILDDx	
REDUCE	ACWORK	AUDWORK	IOWORK	LOWORK		SQLWORK	STWORK		UTWORK
REPORT	ACRPTDD	AURPTDD	IORPTDD	LORPTDD		SQRPTDD	STRPTDD		UTRPTDD
RESTORE	ACRSTDD						STRSTDD		
SAVE	ACSAVDD						STSAVDD		
TRACE	ACTRCDDx	AUTRCDDx		LOTRCDDx	RTTRCDDx	SQTRCDDx	STRCDDx	SYTRCDDx	UTTRCDDx

EXCEPTION and NOEXCEPTION subcommand options

The EXCEPTION subcommand option reports only the exception status. The NOEXCEPTION subcommand option produces a standard report.

Usage

Specify EXCEPTION if you want to report only those entries on reports with at least one field in exception status. Specify NOEXCEPTION to produce a standard report. NOEXCEPTION is the default.

The thresholds for exception fields are defined in the exception threshold data set. For more information about the exception threshold data set, see the *Report Reference*.

Usage notes

If you use this subcommand option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see Chapter 5, "DD statements," on page 15.

FROM/TO subcommand options

The FROM/TO subcommand options specify the range of record timestamps.

Usage

These subcommand options are used to specify the range of record timestamps that OMEGAMON XE for DB2 PE processes.

These options are used with the GLOBAL command and the FILE, REDUCE, REPORT, and TRACE subcommands.

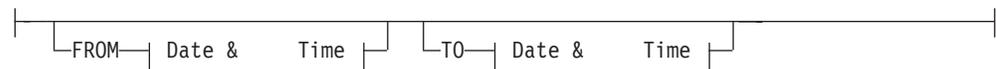
Usage notes

- If used with GLOBAL, all records outside those dates and times are discarded by OMEGAMON XE for DB2 PE before the remaining records are processed by other commands in the same JCL command stream. For more information about the GLOBAL command, see The GLOBAL command.

- You can specify a time adjustment for a DB2 location by using the TIMEZONE option of the GLOBAL command. The time adjustment is applied to record timestamps before other time-related processing is applied.
- You have several choices to specify a time frame. However, you can specify only one time frame with a subcommand.
 - FROM specifies the starting date and time. Records are processed beginning with the first record having a timestamp greater than, or equal to, the FROM date and time. FROM accepts date and time specifications, in any combination, as *absolute* values. Without delimiting TO option, all further records in the input data set are processed.
 - TO specifies the finishing date and time. Records are processed ending with the last record having a timestamp less than the TO date and time. TO accepts date and time specifications, in any combination, as *absolute* values. Without delimiting FROM option, processing begins with the first record in the input data set.
 - FROM and TO specify the starting date and time and the finishing date and time.

Syntax

FROM/TO block:



Date & Time:



Parameters

date

The date in the form *mm/dd/yy*, where *mm* is the month, *dd* is the day, and *yy* is the year. For example, 25 February 2008 is entered as 02/25/08.

OMEGAMON XE for DB2 PE assumes that the year value (*yy*) lies between the system-defined date minus 95 years and the system-defined date plus four years.

You can change the way dates are specified by using DATEFORMAT. For more information, see “DATEFORMAT subcommand option” on page 27.

time

The time in the form *hh:mm:ss.th*, where *hh* is the hour in 24-hour format, *mm* is the minute, *ss* is the second, and *th* is tenths and hundredths of a second. Trailing zeros can be omitted.

Rules

The *date* and *time* variables can be used in various combinations and can also inherit values specified with the GLOBAL command. The table below describes the assumed defaults if either of the options or values is not explicitly specified.

Table 3. Option defaults for date and time variablesw

Subcommand option	Default
FROM(<i>date,time</i>) TO(<i>,time</i>)	If the TO date is not specified, the FROM date is assumed.
FROM(<i>,time</i>) TO(<i>date,time</i>)	If the FROM date is not specified, the TO date is assumed.
FROM(<i>date</i>) TO(<i>date</i>)	If the FROM time is not specified, 00:00:00.00 is assumed. If the TO time is not specified, 23:59:59.99 is assumed.
FROM(<i>,time</i>) TO(<i>,time</i>)	If dates are not specified, all records that comply with the times are processed; the date of records are ignored. REQUESTED ALL DATES is printed on reports and traces in place of REQUESTED FROM.
FROM(<i>,time</i>)	If only the FROM time is specified, 23:59:59.99 is used for the TO time. All records with a timestamp between the FROM time and 23:59:59.99 are processed; the dates of records are ignored. REQUESTED ALL DATES is printed on reports and traces in place of REQUESTED FROM.
TO(<i>,time</i>)	If only the TO time is specified, 00:00:00.00 is used for the FROM time. All records with a timestamp between 00:00:00.00 and the TO time are processed; the dates of records are ignored. REQUESTED ALL DATES is printed on reports and traces in place of REQUESTED FROM.
FROM(<i>date,time</i>) or FROM(<i>date</i>)	If the TO date and time is not specified: <ul style="list-style-type: none"> For all subcommands, the TO values specified in GLOBAL are applied. If no TO values are specified in GLOBAL, processing begins with the first record with a timestamp after the FROM time (if specified) or 00:00:00.00 (if not specified), and ends with the last available record. NOT SPECIFIED is printed on reports and traces in place of REQUESTED TO.
TO(<i>date,time</i>) or TO(<i>date</i>)	If the FROM date and time is not specified: <ul style="list-style-type: none"> For all subcommands, the FROM values specified in GLOBAL are applied. If no FROM values are specified in GLOBAL, processing begins with the first available record, and ends with the last record with a timestamp before the TO time (if specified) or 23:59:59.99 (if not specified) on the TO date. NOT SPECIFIED is printed on reports and traces in place of REQUESTED FROM.

Example using FROM/TO: Specifying a fixed time frame

Processing starts with the first record with a timestamp after 10:00 on 25 February 2008, and ends with the last record with a timestamp before 10:10 on the same day.

```
⋮  
FROM(02/25/08,10:00),TO(,10:10)  
⋮
```

Example using FROM/TO: Specifying a fixed time interval per day

All records with a timestamp at or after 10:00 and before 10:10 are processed, regardless of the date.

```
⋮  
FROM(,10:00)  
TO (,10:10)  
⋮
```

Example using FROM/TO: Specifying a particular start time

Processing starts with the first record with a timestamp after 13:00 on 25 February 2008, and ends with the last record in the input data set.

```
⋮  
FROM(02/25/08,13:00)  
⋮
```

INCLUDE and EXCLUDE subcommand options

The INCLUDE and EXCLUDE subcommand options define a list of identifier values that are included or excluded.

Usage

The INCLUDE and EXCLUDE subcommand options are used to include or exclude data associated with specific OMEGAMON XE for DB2 PE identifiers. If you omit this subcommand option, all records are included.

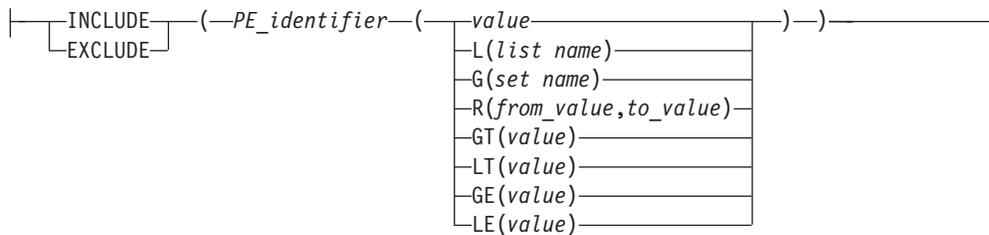
Use INCLUDE to define a list of identifier values that are included in processing. Use EXCLUDE to define a list of identifier values that are excluded from processing.

Usage notes

- You can specify list names, set names, values, or range values for each identifier. If you include some identifier values and omit EXCLUDE, only those records described in INCLUDE are processed. If you exclude some identifier values and omit INCLUDE, all records are processed except those described in EXCLUDE.
- Use GLOBAL INCLUDE, or EXCLUDE carefully with IFCID. Because many IFCID records are paired events, BEGIN and END records are necessary to get a valid report or trace. All IFCIDs that share a common END record must be used together with INCLUDE or EXCLUDE.

Syntax

INCLUDE/EXCLUDE block:



Options

The following options are available with INCLUDE and EXCLUDE:

PE_identifier

Specifies the OMEGAMON XE for DB2 PE identifiers that you want to include in, or exclude from, your output. Identifiers valid for each report are listed in this section. Definitions of identifiers are described in Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

If you do not specify an identifier when using the REDUCE, REPORT, TRACE, or FILE subcommand, OMEGAMON XE for DB2 PE uses the identifiers specified for GLOBAL INCLUDE or EXCLUDE.

You can use an OMEGAMON XE for DB2 PE identifier in either INCLUDE or EXCLUDE, but not with both in the same command or subcommand.

value

A value for the specified OMEGAMON XE for DB2 PE identifier, or an asterisk (*) indicating all values.

Identifier values must consist of the following characters: A — Z, #, \$, @, >, <, or 0 — 9. If the value you want to include contains a character that is not in this list, use an asterisk in its place.

- L** The name of a list containing values for the specified OMEGAMON XE for DB2 PE identifier. The list name must be defined by a LIST command in the same job step.
- G** The name of a set of values for the selected OMEGAMON XE for DB2 PE identifier. The set name must be defined by a GROUP command in the same job step.

If a set name is specified in a value block, only that set is processed. If you want all identifier values that are not contained in any set specification to be reduced and reported individually, enter an asterisk (*) in the INCLUDE or EXCLUDE options following the last set name. For example:

```

:
INCLUDE (
PRIMAUTH (G(AUTHGRP1)
          G(AUTHGRP2)
          *)
:

```

This indicates that the authorization IDs contained in the sets AUTHGRP1 and AUTHGRP2 are processed as sets, and all other authorization IDs are processed individually.

- R** Denotes a range of values beginning with *from_value* and ending with *to_value*. The *from_value* must be less than *to_value*. The generic form can be used only

in the last character position in range values. For example, R(AUTH1*,AUTH2*) is acceptable, but R(AUTH*1,AUTH*2) is not.

GT Denotes values greater than the given value.

This is only valid for SQLCODE.

LT Denotes values less than the given value.

This is only valid for SQLCODE.

GE Denotes values equal to, or greater than the given value.

This is only valid for SQLCODE.

LE Denotes values equal to, or less than the given value.

This is only valid for SQLCODE.

Rules

The following table shows which OMEGAMON XE for DB2 PE identifiers can be used with the different commands and subcommands.

Table 4. OMEGAMON XE for DB2 PE identifiers used with INCLUDE and EXCLUDE subcommand option.

OMEGAMON XE for DB2 PE identifier	ACCT	AUD	I/O	LOCK	RT	SQL	STAT	UTIL
ACE (Agent control element address)					t,f	r,t,p		
CLASS (DB2 trace class)					t,f			
CONNECT (Connection ID)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
CONNTYPE (Connection type)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
CORRNAME (Correlation name)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
CORRNMBR (Correlation number)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
DATABASE (Database name)			r,p	r,t,p,f				
DATASET (Data set name)			r,p					
ENDUSER (End user ID)	r,t,p,f	r,t,p,f		r,t,p,f	t,f	r,t,p,f		
FIELD (Comparison with data in a record field)					t,f			
GROUP (Group name)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p	r,t,p,f	r,t,p
IFCID (Instrumentation Facility Component Identifier)					t,f	r,t,p		
INSTANCE (Instance number)	r,t,f	r,t,p,f		t,f	t,f	r,t,p		r,t,p
LOCATION (Location name)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p	r,t,p,f	r,t,p
MAINPACK (Main package)	r,t,p,f							
MEMBER (Member name)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p	r,t,p,f	r,t,p
ORIGAUTH (Original authorization ID)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
PACKAGE (Package information) or PROGRAM (Program information)	r,t,p,f							
PAGESET (Page set name)			r,p	r,t,p,f				
PLANNAME (Plan name)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p

Table 4. OMEGAMON XE for DB2 PE identifiers used with INCLUDE and EXCLUDE subcommand option (continued).

OMEGAMON XE for DB2 PE identifier	ACCT	AUD	I/O	LOCK	RT	SQL	STAT	UTIL
PRMAUTH (Primary authorization ID) or AUTHID (Authorization ID)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
REQLOC (Requester location)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
RESOURCETYPE (Resource type)				r,t,p,f				
RMID (Resource manager identifier)					t,f			
SQLCODE					t,f	t		
SUBSYSTEMID (Subsystem ID)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p	r,t,p,f	r,t,p
THREADTYPE (Thread type)	r,t,p,f	r,t,p,f	r,p	r,t,p,f	t,f	r,t,p		r,t,p
TRANSACT (End user transaction name)	r,t,p,f	t,p,f	r,t,p,f	t,f	r,t,p,f			
TYPE (Event type)				r,t,f				
WSNAME (End user workstation name)	r,t,p,f	t,p,f		r,t,p,f	t,f	r,t,p,f		

Note:

1. The use of PLANNAME as a filter for Accounting reports in a DDF environment can cause unexpected results. See Special considerations for DDF trace data.
2. The commands are abbreviated as follows: ACCT (ACCOUNTING), AUD (AUDIT), I/O (IOACTIVITY), LOCK (LOCKING), RT (RECTRACE), SQL (SQLACTIVITY), STAT (STATISTICS), and UTIL (UTILITY).
3. Subcommands are abbreviated as r (REDUCE), t (TRACE), p (REPORT), and f (FILE), whereas "all" applies to all subcommands.

Considerations when using the DATABASE, PACKAGE, PROGRAM, PAGESET, or RESOURCETYPE identifier

If you use the INCLUDE or EXCLUDE subcommand option with one of these OMEGAMON XE for DB2 PE identifiers, note that the scope of these identifiers affects only the relevant parts of trace records, but not the entire trace records. This means that *all* records are processed, but only the parts of the records specified by one of the identifiers are included in or excluded from processing. Consequently, remaining record information (specifically included or not excluded) from *all* trace records is shown in the resulting reports.

Example: INCLUDE/EXCLUDE with PACKAGE

As an example, take Accounting trace records that might contain information about packages P1, P2, and P3:

- If you specify INCLUDE PACKAGE(P1), information about packages other than P1 is filtered from each trace record. All trace records (possibly containing information about package P1) are considered for report processing. If you specify INCLUDE PACKAGE(Px), where Px is an identifier that is not included in any trace record, the resulting report shows no package information, but all remaining information.
- If you specify EXCLUDE PACKAGE(P1), information about package P1 is filtered from each trace record. All trace records (possibly containing information about packages other than P1) are considered for report processing. If you

specify EXCLUDE PACKAGE(Px), where Px is an identifier that is not included in any trace record, the resulting report shows package information, if they exist in the trace records.

With these identifiers, specified parts of the trace records are included, respectively excluded. The trace records themselves, with the remaining parts, pass the filter. Consequently, this information is shown in the resulting reports.

Example: INCLUDE and EXCLUDE: Specifying plan and primary authorization ID

The following example produces an Accounting long report. The only plan reported is QMF311. All primary authorization IDs are reported except those on the list called PRODUSER.

```
⋮  
LIST (PRMAUTH(PRODUSER(USER1,USER2)))  
ACCOUNTING  
  REDUCE  
  REPORT  
    LAYOUT (LONG)  
    INCLUDE (PLAN(QMF311))  
    EXCLUDE (PRMAUTH(L(PRODUSER)))  
⋮
```

INTERVAL subcommand option

The INTERVAL subcommand option defines the interval of time when data is summarized.

Usage

The INTERVAL subcommand option is used to define the interval of time when data is summarized.

INTERVAL is used with the GLOBAL command and the REDUCE subcommand.

Usage notes

- If INTERVAL is not specified, the interval specified in the GLOBAL command is used. If no interval is specified in GLOBAL, the default is 0 and all records are summarized and treated as one record.
- INTERVAL processing uses a lot of system resource. Always use the largest interval that meets your reporting requirements. If interval processing is not required, the default INTERVAL (0) is recommended for optimum processing.

Rules

- The interval range is from 0 to 99 999 and is specified in minutes.
- An interval of 15 specifies that entries are calculated within each 15-minute period.
- An interval of 0 specifies that data is summarized over the entire reduction period.

Example: Specifying an interval of two hours

The following example specifies an interval of two hours:

⋮
REDUCE INTERVAL (120)
⋮

LAYOUT subcommand option

The LAYOUT subcommand option specifies the name of a report layout.

Usage

You can select the layout of Accounting and Statistics model reports and traces.

Usage notes

- If none of them suit the requirements at your site, you can tailor your own layouts using the User-Tailored Reporting feature (UTR). For details refer to the *Reporting User's Guide*.
- Migration layouts help you identify OMEGAMON XE for DB2 PE data that was previously shown in the reports of the OMEGAMON Historical Reporter.

Options

You can specify one of the supplied layouts or one that you have previously tailored:

SHORT

Provides a short layout. SHORT is used by default.

Note: If a report does not contain the details that you are looking for (such as thread-related data), run the job again using the LAYOUT subcommand option LONG. You can also use other functions such as Record Trace and SQL Activity to find detailed DB2 trace data.

LONG

Provides detailed thread-related data.

ACCEL

Accounting: Provides detailed thread-related Accelerator activity data.

Statistics: Provides detailed system-related Accelerator activity data.

migration layouts

Supplied migration layouts to help you identify OMEGAMON XE for DB2 PE data that was previously shown in the reports of the OMEGAMON Historical Reporter. The migration layouts include:

- "Options for OMEGAMON Historical Reporter Accounting reports" on page 39
- "Options for OMEGAMON Historical Reporter Statistics reports" on page 42

user-defined layouts

You can customize your own report layouts by specifying which blocks of data and which fields within the blocks are included, and their relative order. To adapt the reports according to your requirements, you use user-tailored reporting (UTR). With UTR, you can control the volume, contents, and layout of Statistics traces and reports.

For information about tailoring report layouts, see the *Reporting User's Guide*.

Options for OMEGAMON Historical Reporter Accounting reports

To find information that was previously contained in the Accounting reports of the OMEGAMON Historical Reporter use the LAYOUT subcommand options shown in the following table.

Table 5. LAYOUT subcommand options for Accounting reports of the OMEGAMON Historical Reporter

LAYOUT subcommand option	OMEGAMON XE for DB2 PE blocks which contain information previously provided in OMEGAMON Historical Reporter
DDF	<p>The Distributed Activity - Requester and Distributed Activity - Server blocks contain information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DISTRIBUTED_DATA_FACILITY_ACTIVITY • DETAIL
LOCK	<p>The Locking block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • LOCK_ACTIVITY • DETAIL • SUMMARY <p>The Drain/Claim block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • LOCK_ACTIVITY • DETAIL <p>The Data Sharing Locking block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • GLOBAL_LOCK_ACTIVITY
PACK	<p>The following blocks:</p> <ul style="list-style-type: none"> • Package Identification • Times - Class 7 - Package Times • Package Suspensions • Global Locking L-Locks (Package) • Global Locking P-Locks (Package) • SQL Activity (Package) • Buffer Pool Activity (Package) • Locking Activity (Package) <p>contain information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DETAIL • PACKAGE_DETAIL
POOL	<p>The Buffer Pool Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • BUFFER_POOL_ACTIVITY • DETAIL • SUMMARY <p>The Group Buffer Pool Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • GROUP_BUFFER_POOL_ACTIVITY

Table 5. LAYOUT subcommand options for Accounting reports of the OMEGAMON Historical Reporter (continued)

LAYOUT subcommand option	OMEGAMON XE for DB2 PE blocks which contain information previously provided in OMEGAMON Historical Reporter
PTA	<p>The Query Parallelism block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • Query Parallelism (SQL_ACTIVITY) • DETAIL <p>The Times - Class 1 - Application Time block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • DETAIL • TIME_SUMMARY <p>The Times - Class 2 - DB2 Time block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • DETAIL • TIME_SUMMARY <p>The Times - Class 3 - Suspensions block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • SUMMARY • TIME_SUMMARY <p>The Highlights block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Miscellaneous Statements (SQL_ACTIVITY), only traces • TERMINATION_SUMMARY
RID	<p>The RID List block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • RID Pool Access (SQL_ACTIVITY) • DETAIL
RLF	<p>The Resource Limit Facility block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Limit Origin (RESOURCE_LIMIT_FACILITY_SUMMARY) • DETAIL

Table 5. LAYOUT subcommand options for Accounting reports of the OMEGAMON Historical Reporter (continued)

LAYOUT subcommand option	OMEGAMON XE for DB2 PE blocks which contain information previously provided in OMEGAMON Historical Reporter
STP	<p>The Times - Class 1 - Application Time block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • DETAIL • TIME_SUMMARY <p>The Times - Class 2 - DB2 Time block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • DETAIL • TIME_SUMMARY <p>The Times - Class 3 - Suspensions block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • SUMMARY • TIME_SUMMARY <p>The Stored Procedures block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Stored Procedures (SQL_ACTIVITY) • DETAIL
SQL	<p>The SQL DCL block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DCL Statements (SQL_ACTIVITY) • DETAIL <p>The SQL DDL block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DDL Statements (SQL_ACTIVITY) • Miscellaneous Statements (SQL_ACTIVITY) • DETAIL <p>The SQL DML block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DML Statements (SQL_ACTIVITY) • DETAIL
TERM	<p>The Application Termination block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • TERMINATION_SUMMARY

Table 5. LAYOUT subcommand options for Accounting reports of the OMEGAMON Historical Reporter (continued)

LAYOUT subcommand option	OMEGAMON XE for DB2 PE blocks which contain information previously provided in OMEGAMON Historical Reporter
TIME	<p>The Times - Class 1 - Application Time block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • DETAIL • TIME_SUMMARY <p>The Times - Class 2 - DB2 Time block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • PARALLEL_TASKS • DETAIL • TIME_SUMMARY <p>The Times - Class 3 - Suspensions block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • SUMMARY • TIME_SUMMARY <p>The Highlights block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Miscellaneous Statements (SQL_ACTIVITY), only traces • TERMINATION_SUMMARY <p>The Global Locking L-Locks (Plan) block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • SUMMARY • TIME_SUMMARY <p>The Global Locking P-Locks (Plan) block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • SUMMARY • TIME_SUMMARY

Options for OMEGAMON Historical Reporter Statistics reports

To find information that was previously contained in the Statistics reports of the OMEGAMON Historical Reporter use the LAYOUT subcommand options shown in the following table.

Table 6. LAYOUT subcommand options for Statistics reports of the OMEGAMON Historical Reporter

LAYOUT subcommand option	OMEGAMON XE for DB2 PE PE blocks which contain information previously provided in OMEGAMON Historical Reporter
BIND	<p>The Plan/Package Processing block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DETAIL • Miscellaneous Statements (SQL_ACTIVITY) • Plan (BIND_ACTIVITY) • Package (BIND_ACTIVITY) <p>The Authorization Management block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Auth Chk (BIND_ACTIVITY)
CMD	<p>The DB2 commands block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • COMMAND_ACTIVITY • DETAIL
DDF	<p>The DRDA Remote Locations block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • 2-Phase Commits (DISTRIBUTED_DATA_FACILITY_ACTIVITY) • DETAIL <p>The Global DDF Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DETAIL
LOCK	<p>The Locking Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • LOCK_ACTIVITY • DETAIL • SUMMARY <p>The Data Sharing Locking block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • GLOBAL_LOCK_ACTIVITY
LOG	<p>The Log Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • LOCK_ACTIVITY • DETAIL
OPCL	<p>The Open/Close Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • OPEN_CLOSE_ACTIVITY • DETAIL

Table 6. LAYOUT subcommand options for Statistics reports of the OMEGAMON Historical Reporter (continued)

LAYOUT subcommand option	OMEGAMON XE for DB2 PE PE blocks which contain information previously provided in OMEGAMON Historical Reporter
POOL	<p>The Buffer Pool General block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • General (BUFFER_POOL_ACTIVITY) • Query Parallelism (BUFFER_POOL_ACTIVITY) • DETAIL • SUMMARY <p>The Buffer Pool Read block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Read Operations (BUFFER_POOL_ACTIVITY) • DETAIL • SUMMARY <p>The Buffer Pool Sort/Merge block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Sort/Merge (BUFFER_POOL_ACTIVITY) • DETAIL • SUMMARY <p>The Buffer Pool Write block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Write Operations (BUFFER_POOL_ACTIVITY) • DETAIL • SUMMARY <p>The Group Buffer Pool Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • GROUP_BUFFER_POOL_ACTIVITY <p>The EDM Pool Activity block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • EDM_POOL_ACTIVITY • DETAIL
QP	<p>The Query Parallelism block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Query Parallelism (SQL_ACTIVITY) • DETAIL
STP	<p>The Stored Procedures block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • Stored Procedures (SQL_ACTIVITY) • DETAIL

Table 6. LAYOUT subcommand options for Statistics reports of the OMEGAMON Historical Reporter (continued)

LAYOUT subcommand option	OMEGAMON XE for DB2 PE PE blocks which contain information previously provided in OMEGAMON Historical Reporter
SQL	<p>The SQL DCL block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DCL Statements (SQL_ACTIVITY) • DETAIL • SUMMARY <p>The SQL DDL block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DDL Statements (SQL_ACTIVITY) • DETAIL • SUMMARY • Miscellaneous Statements (SQL_ACTIVITY) <p>The SQL DML block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • DML Statements (SQL_ACTIVITY) • DETAIL • SUMMARY
SUBS	<p>The Subsystem Services block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • SUBSYSTEM_ACTIVITY • DETAIL
TIME	<p>The CPU Times block contains information previously provided in OMEGAMON Historical Reporter:</p> <ul style="list-style-type: none"> • CPU_TIME_SUMMARY • DETAIL

LEVEL subcommand option

The LEVEL subcommand option is described in command and subcommand options.

Usage

The meaning and usage of LEVEL varies and is described in context with the commands and subcommands using this option.

LIMIT subcommand option

The LIMIT subcommand option sets the number of threads processed by SQL Activity Trace.

Usage

The LIMIT subcommand option sets the number of threads processed by SQL Activity Trace.

ORDER subcommand option

The ORDER subcommand option defines which OMEGAMON XE for DB2 PE identifiers are used to aggregate records.

Usage

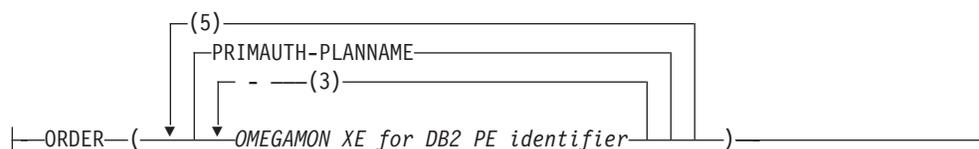
The ORDER subcommand option specifies which OMEGAMON XE for DB2 PE identifiers are used to aggregate records and, unless the TOP(ONLY) subcommand option is specified for Accounting, identifies the presentation sequence of the report entries.

Usage notes

- You can specify one entry of ORDER for each REPORT subcommand.
- You can order by one, two, or three identifiers separated by a dash, and you can specify up to five sets of the identifiers separated by at least one blank, a comma, or a new line.
- For group-scope Locking reports, the default order is DATABASE-PAGESET. For all other reports, the default for ORDER is PRMAUTH-PLANNAME.
- MEMBER is automatically added as the second, third, or fourth identifier in group-scope reports.
- In a distributed environment you should order your reports by REQLOC or CONNTYPE. If REQLOC or CONNTYPE are *not* used in the ORDER subcommand option of REPORT, the accounting portion of *all* threads, including DBATs, where the combination of OMEGAMON XE for DB2 PE identifiers is the same are reported as one entry.

Syntax

ORDER Block:



Rules

The table below shows the OMEGAMON XE for DB2 PE identifiers and criteria that can be used with the ORDER subcommand option.

Table 7. OMEGAMON XE for DB2 PE identifiers used with ORDER subcommand option

OMEGAMON XE for DB2 PE Identifier	Accounting	Audit	I/O Activity	Locking	SQL Activity	Statistics	Utility Activity
ACTNAME (Activity)	●						
BPID (Buffer pool ID)			●				
CLASS (DB2 trace class)	●	●	●	●	●		●
CONNECT (Connection ID)	●	●	●	●	●		●
CONNTYPE (Connection type)	●	●	●	●	●		●
CORRNAME (Correlation name)	●	●	●	●	●		●
CORRNMBR (Correlation number)	●	●	●	●	●		●

Table 7. OMEGAMON XE for DB2 PE identifiers used with ORDER subcommand option (continued)

OMEGAMON XE for DB2 PE Identifier	Accounting	Audit	I/O Activity	Locking	SQL Activity	Statistics	Utility Activity
DATABASE (Database name)			●	●			
ENDUSER (End user ID)	●	●		●	●		
INSTANCE (Instance number)							●
INTERVAL (Interval)	●		●			●	
MAINPACK (Main package)	●						
OBJECT (Object type)		●					
ORIGAUTH (Original authorization ID)	●	●	●	●	●		●
PACKAGE (Package information) or PROGRAM (Program information)	●						
PAGESET (Page set name)			●	●			
PLANNAME (Plan name)	●	●	●	●	●		●
PARTNBR (Partition number)			●				
PRIMAUTH (Primary authorization ID) or AUTHID (Authorization ID)	●		●	●	●		●
REQLOC (Requester location)	●	●	●	●	●		●
THREADTYPE (Thread type)	●						
TRANSACT (End user transaction name)	●	●		●	●		
WSNAME (End user workstation name)	●	●		●	●		

Note:

- For Utility Activity, the ORDER option can be used for both reports and traces.
- Instance number is applicable to Utility trace only.
- For Accounting, the ORDER by THREADTYPE can result in reports for thread types as follows:
 - ALLIED - Indicates accumulated data of allied threads only. These are threads that were not involved in any distributed activity. This type belongs to thread type category ALLIED.
 - ALLDDIST - Indicates accumulated data of threads initiated by DB2 and requested data from one or more server locations. This type belongs to thread type category ALLIED_DIST.
 - DBAT - Indicates accumulated data of threads that are initiated, created, and performing work on behalf of a remote (requester) location. This type belongs to thread type category DBAT.
 - DBATDP - Indicates accumulated data of DBAT duplicate threads. This type belongs to thread type category DBAT.
 - DBATDIST - Indicates accumulated data of DBAT distributed threads that are initiated by a requester location and executed by the server location that in turn requests data from another server location. This type belongs to thread type category DBAT.
 - DBATDICP - Indicates accumulated data of DBAT distributed and copy threads. This type belongs to thread type category DBAT.
 - DBATDIDP - Indicates accumulated data of DBAT distributed and duplicate threads. This type belongs to thread type category DBAT.

Example of the ORDER subcommand option

The following example specifies that two reports are produced:

- Both reports use the SHORT layout by default.
- The first report is ordered by primary authorization ID within plan name within connection ID.
- The second report is ordered by plan name within primary authorization ID.
- Both reports are written to the data set with the default ddname ACRPTDD.

```
⋮  
ACCOUNTING  
REPORT  
ORDER (CONNECTION-PLANNAME-PRMAUTH,PRMAUTH-PLANNAME)  
⋮
```

SCOPE subcommand option

The SCOPE subcommand option specifies the scope of the report in a data-sharing environment.

Usage

Use SCOPE to specifies the scope of the report in a data-sharing environment.

Options

MEMBER

In member-scope reports, a data-sharing group's instrumentation data is presented member by member. The events are reported in the specified ORDER sequence within the DB2 subsystem (member) where they occurred. Member-scope reports are used for DB2 subsystems that are not involved in data sharing.

MEMBER is used by default.

GROUP

In group-scope reports, instrumentation data belonging to individual members is merged and presented for the entire group. The events are reported in the specified ORDER sequence within the DB2 data-sharing group, regardless of which member of the group actually generated the events.

SORTBY subcommand option

The SORTBY subcommand option is described in command and subcommand options.

Usage

The meaning and usage of SORTBY varies and is described in context with the commands and subcommands using this option.

SPREADSHEETDD subcommand option

| The SPREADSHEETDD subcommand is used to write locking related DB2 data to
| a specified data set.

Usage

| The SPREADSHEETDD subcommand option can be used to write locking related
| DB2 data to a specified data set in a format that can be imported by common
| spreadsheet programs. This subcommand option is only applicable to the REPORT
| subcommand of the LOCKING command.

| Instead of this subcommand and for statistics and accounting only, you can use the
| Spreadsheet Input Data Generator utility to create CSV (comma-separated value)
| data for spreadsheet programs as described in *Reporting User's Guide*.

SUMMARIZEBY subcommand option

The SUMMARIZEBY subcommand option is applicable to the SQL Activity report set only.

Usage

Use SUMMARIZEBY to specify the SQL events to be summarized by the REPORT subcommand.

TOP subcommand option

The TOP subcommand option lists all applications or users that have required most use of specific DB2 resources.

Usage

To identify report entries with a high value in certain fields, you can produce an Accounting report or trace with TOP lists. TOP lists are index-like reports at the end of a report or trace pointing out the most interesting entries in the report or trace. You generate such lists using the TOP subcommand option.

Usage notes

- TOP is only used in the Accounting report set.
- Entries with 0 or undetermined values are not shown. Moreover, if your input data contains only 0 or undetermined values for the TOP fields requested, a TOP list is not generated.

Keywords

ONLY

Filters a report or trace such that only the main resource consumers are shown, ordered by descending resource value.

Note: The TOP(ONLY) subcommand option changes the presentation sequence of Accounting reports and traces. Report and trace entries are ordered according to the TOP resource instead of the timestamp or ORDER subcommand options. The summarization in reports is not affected by the changed sequence caused by TOP filtering.

n ONLY

Filters a report or trace such that only the top *n* entries for the default TOP field are shown, where *n* is a number.

TOTAL

The resource values shown are average values. To get total values, that is, the main resource consumers calculated by taking into account how often they consumed resources, you specify the TOP subcommand option with the TOTAL keyword.

Note: TOTAL only applies to reports. If you specify it with a trace, it is ignored.

- * Generates Accounting reports and traces showing TOP lists for all the fields available for use with the TOP subcommand option.

n request type

Produces a trace with a TOP list for the defined number *n* of the specified *request types*, such as Getpage requests.

Examples

For more information and examples refer to the *Reporting User's Guide*.

TYPE subcommand option

The TYPE subcommand option specifies the activity types reported.

Usage

TYPE is used in the Audit report set to control the type of data reported and in the Utility Activity report set to specify the activity types reported.

WORKLOAD subcommand option

The WORKLOAD subcommand option is described in command and subcommand options.

The meaning and usage of WORKLOAD varies and is described in context with the commands and subcommands using this option.

Chapter 7. OMEGAMON XE for DB2 PE commands

This section describes the OMEGAMON XE for DB2 PE report commands, their subcommands, and subcommand options.

The description of the subcommand options is intentionally brief, to avoid recurrences. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions.

The following topics provide additional information:

Report set commands:

- “ACCOUNTING command” on page 53
- “AUDIT command” on page 73
- “EXEC command” on page 82
- “EXPLAIN command” on page 82
- “IOACTIVITY command” on page 93
- “LOCKING command” on page 98
- “RECTRACE command” on page 105
- “SQLACTIVITY command” on page 109
- “STATISTICS command” on page 119
- “SYSPARMS command” on page 135
- “UTILITY command” on page 138

Auxiliary commands:

- “CASE command” on page 145
- “FIELD command” on page 147
- “GLOBAL command” on page 149
- “GROUP command” on page 156
- “LIST command” on page 164

Troubleshooting commands:

- “DUMP command” on page 167
- “TAPECOPY command” on page 169

Overview of the commands

The table below identifies the commands and their related subcommands and subcommand options. There are three types of commands: report set commands, auxiliary commands and troubleshooting commands.

Table 8. Commands, subcommands and subcommand options.

Command types	Commands	Subcommands	Subcommand options
Report set command	ACCOUNTING	REPORT	TO, FROM, DDNAME, ORDER, INCLUDE, EXCLUDE, EXCEPTION, NOEXCEPTION, LAYOUT, TOP, SCOPE
		REDUCE	TO, FROM, INTERVAL, BOUNDARY, INCLUDE, EXCLUDE
		SAVE	DDNAME, DATATYPE
		RESTORE	DDNAME
		TRACE	TO, FROM, DDNAME, INCLUDE, EXCLUDE, EXCEPTION, NOEXCEPTION, LAYOUT, TOP
		FILE	TO, FROM, DDNAME, INCLUDE, EXCLUDE, EXCEPTION, NOEXCEPTION
Report set command	AUDIT	REPORT	TO, FROM, DDNAME, TYPE, LEVEL, ORDER, INCLUDE, EXCLUDE, SCOPE
		REDUCE	TO, FROM, INCLUDE, EXCLUDE, SCOPE
		TRACE	TO, FROM, DDNAME, TYPE, INCLUDE, EXCLUDE, SCOPE
		FILE	TO, FROM; DDNAME, TYPE, INCLUDE, EXCLUDE
Report set command	EXEC	N/A	N/A
Report set command	EXPLAIN	ACCTYPE	MATCHING, NONMATCH, TABSCAN
		DBRM	N/A
		DEGREE	ANY, DSJ, ESJ
		FIRST/LAST	N/A
		FORCE	NO, YES
		FORMAT	NO, YES
		GEN	N/A
		HOSTVAR	NO, YES
		INDEX	ALL, NO, YES
		LEVEL	N/A
Report set command	IOACTIVITY	REPORT	TO, FROM, DDNAME, LEVEL, ORDER, INCLUDE, EXCLUDE
		REDUCE	TO, FROM, INTERVAL, BOUNDARY, INCLUDE, EXCLUDE
Report set command	LOCKING	REPORT	TO, FROM, DDNAME, LEVEL, ORDER, INCLUDE, EXCLUDE, SCOPE, SPREADSHEETDD
		REDUCE	TO, FROM, INTERVAL, BOUNDARY, INCLUDE, EXCLUDE
		TRACE	TO, FROM, DDNAME, LEVEL, INCLUDE, EXCLUDE, SCOPE
		FILE	TO, FROM, DDNAME, INCLUDE, EXCLUDE
Report set command	RETRACE	TRACE	TO, FROM, DDNAME, LEVEL, INCLUDE, EXCLUDE, SORTBY
		FILE	TO, FROM, DDNAME, INCLUDE, EXCLUDE
Report set command	SQLACTIVITY	REPORT	TO, FROM, DDNAME, ORDER, INCLUDE, EXCLUDE, SORTBY, SUMMARIZEBY, WORKLOAD
		TRACE	TO, FROM, DDNAME, INCLUDE, EXCLUDE, LIMIT, SORTBY, SUMMARIZEBY, WORKLOAD, SQLTEXT
Report set command	STATISTICS	REPORT	TO, FROM, DDNAME, ORDER, INCLUDE, EXCLUDE, EXCEPTION, NOEXCEPTION, LAYOUT, SCOPE, DSETSTAT
		REDUCE	TO, FROM, INTERVAL, BOUNDARY, INCLUDE, EXCLUDE
		SAVE	DDNAME, DSETSTAT
		RESTORE	DDNAME
		TRACE	TO, FROM, DDNAME, INCLUDE, EXCLUDE, EXCEPTION, NOEXCEPTION, LAYOUT, DSETSTAT
		FILE	TO, FROM, DDNAME, INCLUDE, EXCLUDE, EXCEPTION, NOEXCEPTION, DSETSTAT

Table 8. Commands, subcommands and subcommand options (continued).

Command types	Commands	Subcommands	Subcommand options
Report set command	SYSPARMS	TRACE	DDNAME
		FILE	DDNAME
Report set command	UTILITY	REPORT	TO, FROM, DDNAME, TYPE, ORDER, INCLUDE, EXCLUDE
		REDUCE	TO, FROM, INTERVAL, BOUNDARY, INCLUDE, EXCLUDE
		TRACE	TO, FROM, DDNAME, TYPE, ORDER, INCLUDE, EXCLUDE, WORKLOAD
Auxiliary commands	CASE	N/A	SENSITIVE, ANY
	FIELD	N/A	N/A
	GLOBAL	N/A	FROM, TO, INTERVAL, PRESORTED, SPANINC, BOUNDARY, INPUTDD, PAGESIZE, TIMEZONE, EXPLAIN, INCLUDE, EXCLUDE
	GROUP	N/A	L, G, R
	LIST	N/A	L, G, R
Troubleshooting commands	DUMP	N/A	EXEC, OFFSET, CODES, SKIP, STOPAFT, MAXDUMP
	TAPECOPY	N/A	EXEC, OFFSET, CODES, SKIP, STOPAFT, NEWCOPY

ACCOUNTING command

This section provides an overview of the ACCOUNTING command.

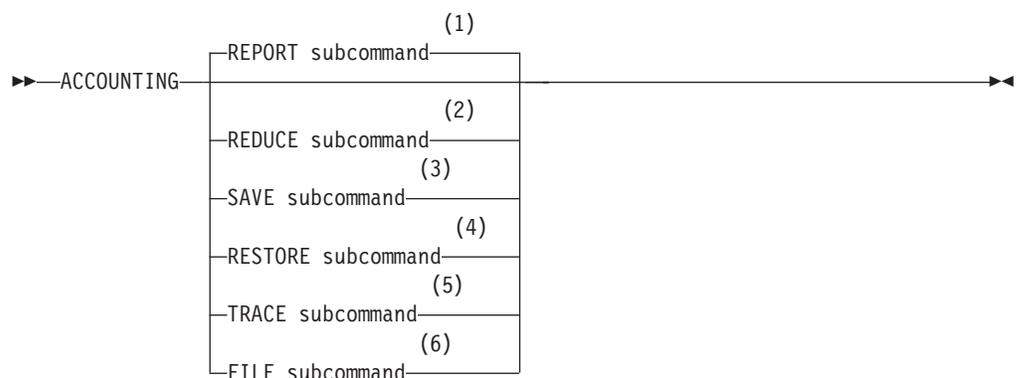
Usage

Use the ACCOUNTING command to generate Accounting reports, traces, and data sets. You can also use it to reduce, save, and restore data.

Usage notes

- This command can be used once in a job step.
- However, it can be used in the same job step with commands of the other report sets.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax



Notes:

- 1 You can specify REPORT up to 35 times.
- 2 You can specify REDUCE only once. If specified, you must also specify REPORT or SAVE at least once.

- 3 You can specify SAVE only once.
- 4 You can specify RESTORE only once.
- 5 You can specify TRACE up to five times.
- 6 You can specify FILE up to five times.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting Accounting functions

This is a sample of the JCL required to produce Accounting reports and traces. See Chapter 5, “DD statements,” on page 15 for descriptions of the DD statements.

The OMEGAMON XE for DB2 PE command language shown in this example may not be appropriate in all circumstances. You must modify it to meet your requirements. Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See the individual DD statement descriptions for more information.

Note:

1. There is an advantage in omitting DPMOUTDD from your JCL. For more information, see “DPMOUTDD statement” on page 16.
2. If you omit the EXEC statement, a report is not produced, the syntax of the OMEGAMON XE for DB2 PE command stream is checked and written together with any information, warning, or error messages generated to the DBPMLLOG data set. All statements following the EXEC are ignored.

```

| //          PEMAIN EXEC PGM=FPECMAIN
| // * FOLLOWING ARE SYSTEM DDNAMES
| //STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
| //DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
| //INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
| //DPMLLOG DD SYSOUT=*
| //SYSOUT DD SYSOUT=*
| //JOBSUMDD DD SYSOUT=*
| //EXCPTDD DD DSN=MYID.EXCEPT.THRESH,DISP=OLD
| //EXTRCDD1 DD SYSOUT=*
| //EXFILDD1 DD DSN=MYID.EXCEPT.LOGFILE,DISP=OLD
| //DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
| //JSSRSDD DD DSN=MYID.FPELIB.JSSRS.DATA,DISP=OLD
| //SYSUDUMP DD DUMMY
| // * FOLLOWING ARE REPORT SET DDNAMES
| //ACRPTDD DD SYSOUT=*
| //ACTRCDD1 DD SYSOUT=*
| //ACSAVDD DD DSN=MYID.FPELIB.ACSAV.DATA,DISP=OLD
| //ACRSTDD DD DSN=MYID.FPELIB.ACRST.DATA,DISP=SHR
| //ACFILDD1 DD DSN=MYID.FPELIB.ACC.FILE,DISP=OLD
| //ACWORK DD DSN=MYID.FPELIB.ACC.WORKDD,DISP=OLD
| // * FOLLOWING IS THE COMMAND STREAM
| //SYSIN DD *
| ACCOUNTING
| REDUCE
| RESTORE
| TRACE
| FILE
| REPORT
| SAVE
| EXEC

```

Figure 3. Sample JCL for requesting Accounting functions

ACCOUNTING command with REPORT subcommand

This section describes the ACCOUNTING command with the REPORT subcommand.

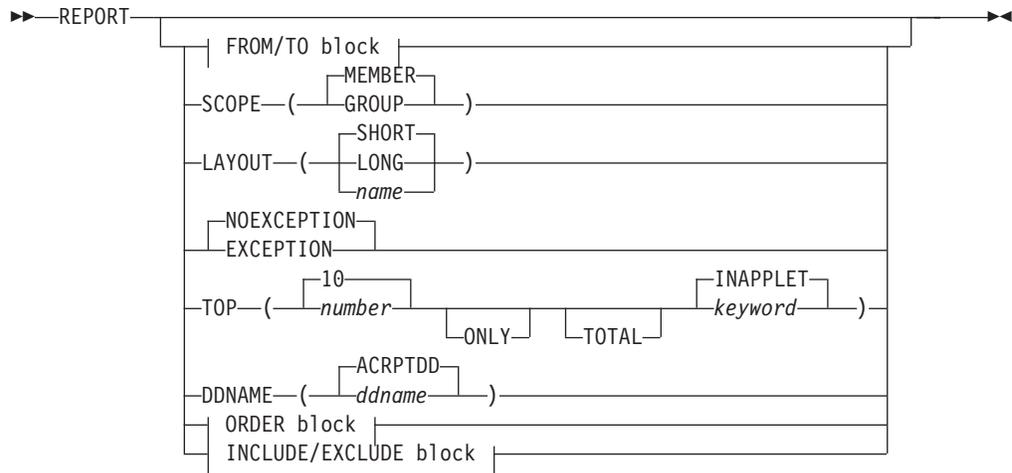
Usage

Use the REPORT subcommand to generate reports from reduced records.

Usage notes

- Up to 35 REPORT subcommands can be specified within each ACCOUNTING command.

Syntax of the REPORT subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reporting process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

You need to specify a REDUCE INTERVAL() for the **FROM/TO** subcommand under ACCOUNTING REPORT. Otherwise, you will see the following message:

```
NO DATA TO REPORT - NO ADEQUATE INPUT DATA OR TOO RESTRICTIVE FILTERS.
```

The reason for this message is that all accounting input data will be reduced by default to only one reporting interval. If you request several ACCOUNTING REPORTs with different FROM/TO subcommands, the INTERVAL() value that you chose should be the smallest of all FROM/TO intervals specified, as in the following example:

```
ACCOUNTING
  REDUCE INTERVAL(15)
  REPORT DDNAME(ACCLONG1) /*ACCOUNTING REPORT*/
        FROM(,10:00:00)
        TO(,10:59:59)
  ...
  REPORT DDNAME(ACCLONG2) /*ACCOUNTING REPORT*/
        FROM(,11:00:00)
        TO(,11:29:59)
  ...
  REPORT DDNAME(ACCLONG3) /*ACCOUNTING REPORT*/
        FROM(,13:00:00)
        TO(,13:14:59)
```

In order to ensure that no extra reporting interval is taken into account, specify either FROM or TO in a way that it does not overlap an adjacent interval boundary. In the following example, the report will show only one reporting interval every day FROM(,13:00:00) TO(,13:14:59):

```

ACCOUNTING
      REDUCE  INTERVAL(15)
      REPORT  DDNAME(ACCLONG3) /*ACCOUNTING REPORT*/
              FROM(,13:00:00)
              TO(,13:14:59)

```

As a comparison, the following report will report two adjacent reporting intervals every day - one from 13:00 to 13:15 and one from 13:15 to 13:30:

```

ACCOUNTING
      REDUCE  INTERVAL(15)
      REPORT  DDNAME(ACCLONG3) /*ACCOUNTING REPORT*/
              FROM(,13:00)
              TO(,13:15)

```

For more details, see “BOUNDARY subcommand option” on page 26 and “INTERVAL subcommand option” on page 37.

SCOPE

Specifies the scope of the report in a data sharing environment.

MEMBER

In member-scope reports, a data sharing group's instrumentation data is presented member by member. The events are reported in the specified ORDER sequence within the DB2 subsystem (member) where they occurred. Member-scope reports are used for DB2 subsystems that are not involved in data sharing.

GROUP

In group-scope reports, instrumentation data belonging to individual members is merged and presented for the entire group. The events are reported in the specified ORDER sequence within the DB2 data sharing group, regardless of which member of the group actually generated the events.

LAYOUT

Specifies the name of a report layout. You can specify one of the supplied layouts or one that you have previously tailored:

SHORT

This is the default.

LONG

This option provides detailed thread-related data. You can also use other functions such as Record Trace and SQL Activity to find detailed DB2 trace data.

Historical Reporter migration layouts

Use the Historical Reporter migration layouts to help you identify OMEGAMON XE for DB2 PE data that was previously shown in the reports of the OMEGAMON Historical Reporter. The migration layouts include:

- “Options for OMEGAMON Historical Reporter Accounting reports” on page 39
- “Options for OMEGAMON Historical Reporter Statistics reports” on page 42

User-defined layouts

You can customize your own report layouts by specifying which blocks of data and which fields within the blocks are included, and their relative order. To adapt the reports according to your requirements, you use user-tailored reporting (UTR). With UTR, you can control the volume, contents, and layout of Statistics traces and reports.

For information about tailoring report layouts, see the *Reporting User's Guide*.

EXCEPTION**NOEXCEPTION**

Specify EXCEPTION if you want to show those records with at least one field in exception status. Otherwise, a standard report is produced.

If you use this option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see Figure 3 on page 55.

TOP

To identify report entries with a high value in certain fields, you can produce an Accounting report with TOP lists. TOP lists indicate which entries on the report have the highest value in the field you have specified by using the TOP subcommand option.

For more information about TOP processing, see Chapter 6, "TOP subcommand option" on page 49 and the *Reporting User's Guide*. You can specify:

number

By default, the TOP list contains the top ten entries, but you can change the number to anything from one to fifty.

- * Specifies any combination of the fields. You can generate reports showing TOP lists for all the fields available for use with the TOP subcommand option.

ONLY

ONLY indicates the use of TOP as a filter. When TOP is requested as a filter, the index is not shown. For example, the report produced by the command following shows only the top 3 entries for the default TOP field, elapsed time in application (INAPPLET).

```
ACCOUNTING
REPORT
TOP (3 ONLY)
```

TOTAL

Produces reports that show total values instead of averages.

keyword

TOP lists for package fields report the maximum values of these fields in individual packages within a report entry. TOP lists for buffer pool fields contain the totals for all buffer pools.

The keywords that specify these fields are as follows:

INAPPLET

The class 1 elapsed time (in an application). This value is an average.

If no TOP subcommand option is specified, the default is class 1 elapsed time (INAPPLET).

INDB2ET

The class 2 elapsed time accumulated in DB2. This value is an average.

OUTDB2ET

The elapsed time outside DB2. This value is an average.

INAPPLPT

The class 1 CPU time in an application. This value is an average.

INDB2PT

The class 2 CPU time in DB2. This value is an average.

OUTDB2PT

The CPU time outside DB2. This value is an average.

INAPPLWT

The class 1 waiting time in an application. This value is an average.

INDB2WT

The class 2 waiting time in DB2. This value is an average.

OUTDB2WT

The waiting time outside DB2. This value is an average.

TOTSUSTM

The waiting time for all class 3 suspensions. This value is an average.

DMLSTAT

The total number of SQL DML statements executed. This value is an average.

DCLSTAT

The total number of DCL statements executed. This value is an average.

DDLSTAT

The total number of DDL statements executed.

UPDPERCM

The sum of SQL UPDATE, SQL INSERT, and SQL DELETE statements executed.

CMUPERUPD

The ratio of the sum of commits and rollbacks to the sum of SQL UPDATE, SQL INSERT, and SQL DELETE statements.

TOTSUSP

The number of suspensions. This value is an average.

GETPAGES

The number of Getpage requests. This value is an average.

NOTACCT

The time not accounted in DB2. You use this time to determine whether there is a large percentage of time that has not been captured within the DB2 Accounting record. This value is an average.

BUFUPDTS

The number of buffer updates. This value is an average.

SYNCREAD

The number of synchronous read I/O operations. This value is an average.

TOTPREF

The number of all types of prefetch requests. This value is an average.

PINDBET

The total elapsed time for executing the package or DBRM. This value is an average.

PINDBPT

The CPU time spent by the package or DBRM (class 7). This value is an average.

PTSUSTME

The waiting time for the package or DBRM caused by a class 8 suspension. This value is an average.

PNOTACCT

The total unaccounted time in DB2 caused by the execution of the package or DBRM. This value is an average.

DDNAME

Specifies the data set where the report is written. The default ddname for report is ACRPTDD.

ORDER

Specifies which OMEGAMON XE for DB2 PE identifiers are used to aggregate Accounting records.

For details, see "ORDER subcommand option" on page 46 and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

Additionally, you can use the REDUCE INTERVAL to order data on Accounting reports and Statistics reports. For example, if you want to report data at daily interval, specify INTERVAL (1440) in the REDUCE subcommand and ORDER(INTERVAL) in the REPORT subcommand.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other allowed identifiers with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Special considerations for DDF trace data:

For single DB2 systems the PLANNAME identifier can be used to filter data and order reports in a meaningful manner.

When DB2 is acting as a DDF server, the PLANNAME identifier has a constant value of DISTSERV for requesters using DRDA. For OM XE for DB2 PE Accounting data, DISTSERV is replaced with the first eight characters of the name of the client application so that filtering on PLANNAME, is meaningful.

Because this replacement is done by the batch accounting component, you need to consider the impact when using the INCLUDE and EXCLUDE subcommand options with other commands. This can cause unexpected results, for example when GLOBAL INCLUDE is used.

The following example shows how PLANNAME can be used in Accounting reports with DDF data to include data from a PLANNAME with the value CAPPNAME:

```
GLOBAL
  INCLUDE (PLANNAME(DISTSERV))      // to include all DDF data
ACCOUNTING
  REDUCE
    INCLUDE (PLANNAME(CAPPNAME))
  REPORT
    INCLUDE (PLANNAME(CAPPNAME))
:
```

When GLOBAL INCLUDE is omitted, a report is produced, containing replaced names.

When REDUCE INCLUDE is omitted, no report is produced. This is because REDUCE is called implicitly by REPORT by using the GLOBAL filter. As accounting works with the replaced values, DISTSERV is not found in the input records.

When INCLUDE is omitted from REPORT, there is no data to report. The GLOBAL filter is used as default, explicitly including DISTSERV. As accounting works with the replaced values, DISTSERV is not found in the input records.

If the use of PLANNAME as a filter causes unexpected results across reports, use REQLOC, CONNTYPE and THREADTYPE, which are interpreted in exactly the same way in all OM XE for DB2 PE reports.

Example using REPORT with LAYOUT, INCLUDE, FROM, and TO options

The following example specifies:

- An Accounting long report

- Data is included that is only associated with the location in the range of LOCN01 to LOCN05
- Using records between and including the FROM and TO times

```

:
:
REPORT
  LAYOUT (LONG)
  INCLUDE (LOCATION(R(LOCN01 LOCN05)))
  FROM (03/18/99,10:00:00.00)
  TO (03/19/99,12:00:00.00)
:
:

```

Example using ACCOUNTING REPORT, specifying two order sets

This example specifies that two reports are produced:

- Both reports use the SHORT layout by default.
- The first report is ordered by primary authorization ID within plan name within connection ID.
- The second report is ordered by plan name within primary authorization ID.
- For both reports, a TOP list is produced identifying the top three report entries for INDB2ET (elapsed time spent in DB2).
- Both reports include data for the following primary authorization IDs:
 - UID0001
 - UID0005
 - UID0009.
- As no ddname was specified, both reports are written to the data set with the default ddname ACRPTDD.

```

:
:
ACCOUNTING
REPORT
  TOP (3 INDB2ET)
  ORDER (CONNECTION-PLANNAME-PRMAUTH,PRMAUTH-PLANNAME)
  INCLUDE (PRMAUTH(UID0001 UID0005 UID0009))
:
:

```

Example using ACCOUNTING with multiple REPORT subcommands

This example demonstrates how you can produce Accounting reports to show all accounting data for different requesting locations.

Because the Statistics report set does not offer a breakdown of the activity of DRDA protocol requests, an Accounting report showing the DDF statistics can be useful.

The first report shows fields summarized by the requesting location.

The information is shown for every requesting location, including the reporting location, regardless of the method of access.

The second report shows accounting data summarized by requesting location for all locations (except the reporting location) using DRDA protocol.

```

:
:
ACCOUNTING
REPORT
  LAYOUT (LONG)
  ORDER (REQLOC)

```

```

REPORT
  LAYOUT (LONG)
  ORDER (REQLOC)
  INCLUDE (CONNTYPE(DRDA))
:

```

ACCOUNTING command with REDUCE subcommand

This section describes the ACCOUNTING command with the REDUCE subcommand.

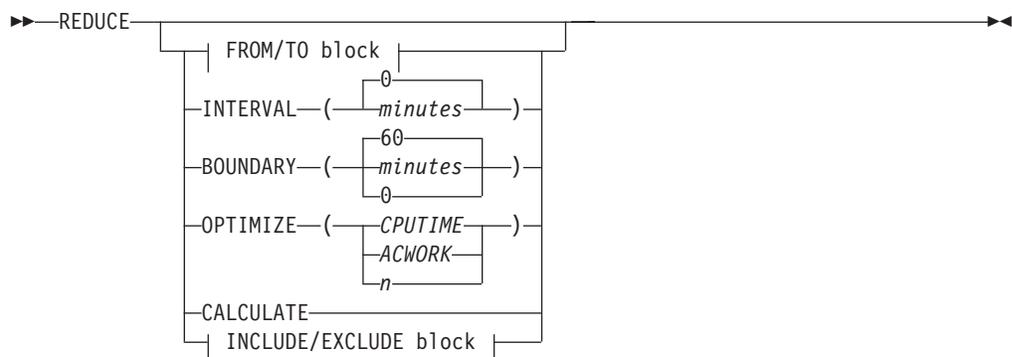
Usage

Use the REDUCE subcommand to reduce the volume of data that is input to the REPORT and SAVE subcommands.

Usage notes

- REDUCE consolidates records with certain common characteristics into one record.
- REDUCE can be used once in an ACCOUNTING command.
- REDUCE is invoked automatically when you use REPORT or SAVE. You must specify REDUCE, however, if you want to:
 - Specify an interval that can be used to order data on Accounting reports.
 - Produce several reports with different time spans.
 - Define the interval and input filters for SAVE. SAVE does not have its own FROM and TO, and INCLUDE or EXCLUDE filters. Instead, it uses REDUCE FROM and TO, and INCLUDE or EXCLUDE filters.
- In most cases, the reduced records contain totals of the values from the individual records. In some cases, the reduced records contain maximum values.

Syntax of the REDUCE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reduction process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

INTERVAL

Defines the interval of time that accounting data is summarized for. Note the remarks about performance impact in “INTERVAL subcommand option” on page 37.

BOUNDARY

Controls the alignment of the intervals used to summarize records in the reduction process.

OPTIMIZE

Activates a balancing process that affects the relation of the CPUTIME and ACWORK consumptions according to the specified option.

CPUTIME

Optimizes the processing of CPU time.

ACWORK

Optimizes the processing for consumption of ACWORK space.

n A balancing value between the optimums of CPUTIME and ACWORK. The value can be between 0 and 100.

For details, see the *IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS*; *IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS: Reporting User's Guide*.

CALCULATE

Indicates that a calculation of consumed or required ACWORK has to be done during the reduction of data for an Accounting REPORT or SAVE function.

For details, see the *IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS*; *IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS: Reporting User's Guide*.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Some care is needed when using PLANNAME as a filter in a DDF environment, see Special considerations for DDF trace data.

Example using REDUCE

This example requests the following:

- Reduce the input data on daily intervals.
- Report the Accounting data by intervals.
- Write the report to the data set defined by the default ddname ACRPTDD.

⋮
ACCOUNTING
 REDUCE
 INTERVAL(1440)
 REPORT

ORDER (INTERVAL)
⋮

ACCOUNTING command with SAVE subcommand

This section describes the ACCOUNTING command with the SAVE subcommand.

Usage

Use the SAVE subcommand (without CONVERT option) to produce a VSAM data set containing reduced records. After the data has been saved, you can:

- Convert the save files to sequential data sets by using the save-file utility and load it to DB2 for subsequent use.
- Restore and combine it with newly reduced data to produce long-term reports.
- Restore it and use it in later reporting.

Use the SAVE subcommand with CONVERT option to produce a sequential data set containing reduced data in records. After the data has been processed, you can:

- Load it to DB2 for subsequent use.

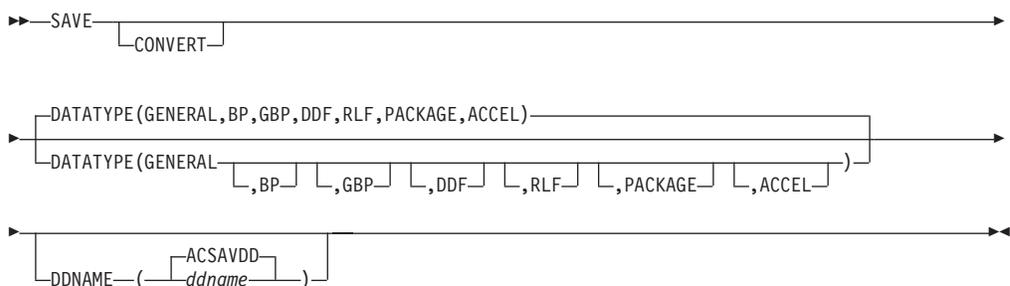
Note: In contrast to the other usage, the conversion by the save-file utility is accomplished directly. VSAM data is not being produced, thus it also not possible to restore it.

You can also use the converted SAVE data sets to generate CSV (comma-separated value) input-data. This CSV data can then be transferred to workstations and imported into spreadsheets to improve DB2 performance analysis using graphical representations or pivot tables. For more information refer to *Reporting User's Guide*.

Usage notes

- VSAM data sets cannot be concatenated.

Syntax of the SAVE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

CONVERT

Specifies that converted reduced data is written directly to a sequential data set.

It is recommended to specify this option if you want to create high amount of loadable reduced data in a sequential data set. It avoids a temporary VSAM data set to be used as a SAVE data set. This options causes OMEGAMON XE for DB2 PE to write converted reduced data directly to a sequential data set in a single step. The resulting output is loadable to the PDB tables. The user may experience performance improvements compared to the default path with a separate SAVE step and the subsequent convert of saved data by the save-file utility.

Note: Not every big trace input results in big reduced data and small trace input in small reduced data. This option becomes effective when the reduction results in high amount of reduced data. For example, it depends on the amount of different criteria that results due to your filtering. It is possible that you experience more performance improvement with a small input trace that results in many different criteria than with a large trace input with only a few different criteria.

This option can generally be used as it is neutral to the performance when processing only a few reduced records, but becomes effective when processing a high amount of reduced data. However, do not use this option if you want to RESTORE and REPORT saved data.

DATATYPE

Specifies which data types are to be written to the VSAM or sequential data set with each record. By default, if DATATYPE is not specified, or if DATATYPE() is specified (without keywords), all data types are written to the reduced data set.

This subcommand option lets you improve the system performance while the VSAM or sequential data set is generated by explicitly specifying only the required data types. If performance is not critical, you can safely ignore this subcommand option. Use one or more of the following keywords to specify the required data types.

Note: Note that a later restore of partially saved data does not reconstruct the unsaved data types.

GENERAL

General performance data. Always required for a save data set. Note that the GENERAL keyword needs to be specified explicitly.

BP Buffer pool performance data

GBP

Group buffer pool performance data

DDF

Distributed Data Facility (DDF) data

RLF

Resource Limit Facility (RLF) data

PACKAGE

Package execution data

ACCEL

Accelerator data

DDNAME

Specifies the ddname where the save data is written. The default ddname is ACSAVDD.

Without CONVERT option:

The VSAM data set defined by the default ddname must already exist when you run OMEGAMON XE for DB2 PE. Either specify an existing data set from a previous OMEGAMON XE for DB2 PE run (when restoring data), or specify a new data set allocated by using the IDCAMS DEFINE CLUSTER function.

Note: If ddname is assigned to a non-VSAM file, you receive an error message and the job terminates.

With CONVERT option:

The ddname needs to be assigned to a physical sequential data set. This data set can be used for a subsequent load to Accounting SAVE tables.

Note: If ddname is assigned to a nonsequential data set, you receive an error message and the job terminates.

Example using SAVE

This example requests the following:

- Save aggregated Accounting data to the VSAM save data set with ddname ACSAVDD1.

```
⋮  
ACCOUNTING  
  SAVE  
    DDNAME(ACSAVDD1)  
⋮
```

This example requests the following:

- Convert reduced Accounting data and write it to the sequential data set with ddname ACSAVDD2.

```
⋮  
ACCOUNTING  
  SAVE  
    CONVERT  
      DDNAME(ACSAVDD2)  
⋮
```

ACCOUNTING command with RESTORE subcommand

This section describes the ACCOUNTING command with the RESTORE subcommand.

Usage

Use the RESTORE subcommand to reload previously saved data (in a VSAM data set) for additional processing. After the data is restored, you can produce reports from the restored data alone, or from the restored data combined with newly reduced data.

Usage notes

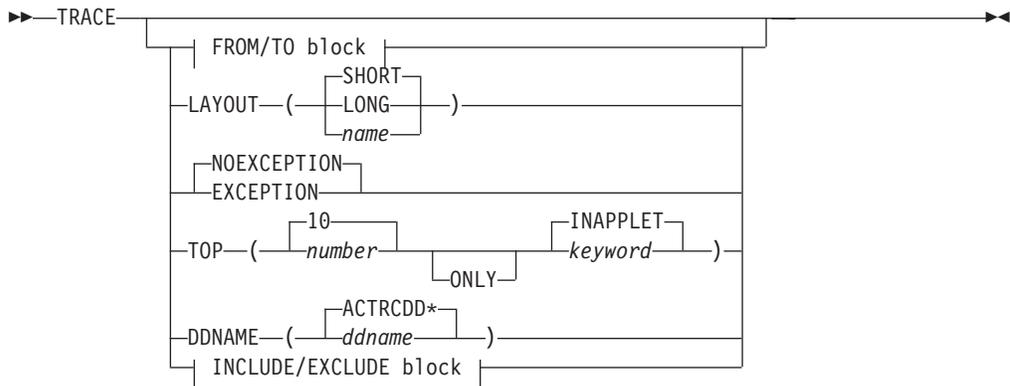
- Saved data can be restored as often as required.

Syntax of the RESTORE subcommand

Usage notes

- An entry in an Accounting trace is referred to as a logical Accounting record as it can consist of several physical Accounting records (IFCIDs 3 and 239).
- Up to five traces can be requested in a job step.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the trace by date and time.

For details, see “FROM/TO subcommand options” on page 30.

LAYOUT

Specifies the name of a report layout. You can specify one of the supplied layouts or one that you have previously tailored:

SHORT

This is the default.

LONG

This option provides detailed thread-related data. You can also use other functions such as Record Trace and SQL Activity to find detailed DB2 trace data.

Historical Reporter migration layouts

Use the Historical Reporter migration layouts to help you identify OMEGAMON XE for DB2 PE data that was previously shown in the reports of the OMEGAMON Historical Reporter. The migration layouts include:

- “Options for OMEGAMON Historical Reporter Accounting reports” on page 39
- “Options for OMEGAMON Historical Reporter Statistics reports” on page 42

User-defined layouts

You can customize your own report layouts by specifying which blocks of data and which fields within the blocks are included, and their relative order. To adapt the reports according to your requirements, you use user-tailored reporting (UTR). With UTR, you can control the volume, contents, and layout of Statistics traces and reports.

For information about tailoring report layouts, see the *Reporting User's Guide*.

EXCEPTION
NOEXCEPTION

Specify EXCEPTION if you want to show those accounting trace entries with at least one field in exception status. Otherwise, a standard report is produced.

If you use this option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see Figure 3 on page 55.

TOP

To identify trace entries with a high value in certain fields, you can produce an accounting trace with TOP lists. TOP lists indicate which entries on the trace have the highest value in the field you have specified by using the TOP subcommand option.

For more information about TOP processing, see Chapter 6, "TOP subcommand option" on page 49 and the *Reporting User's Guide*. You can specify:

number

By default, the TOP list contains the top ten entries, but you can change the number to anything from one to fifty.

ONLY

ONLY indicates the use of TOP as a filter. When TOP is requested as a filter, the index is not shown. For example, the trace produced by the command following only shows the top 3 entries for the default TOP field, elapsed time in application (INAPPLET). The keyword ONLY is valid wherever a resource name can be used.

```
ACCOUNTING
TRACE
TOP (3 ONLY)
```

- * Specifies any combination of the fields. You can generate reports showing TOP lists for all the fields available for use with the TOP subcommand option.

keyword

TOP lists for package fields report the maximum values of these fields in individual packages within a trace entry. TOP lists for buffer pool fields contain the totals for all the buffer pools.

The following list shows the keywords that these fields are specified by:

INAPPLET

The class 1 elapsed time (in an application). This value is an average.

If no TOP subcommand option is specified, the default is class 1 elapsed time (INAPPLET).

INDB2ET

The class 2 elapsed time accumulated in DB2. This value is an average.

OUTDB2ET

The elapsed time outside DB2. This value is an average.

INAPPLPT

The class 1 CPU time in an application. This value is an average.

INDB2PT

The class 2 CPU time in DB2. This value is an average.

OUTDB2PT

The CPU time outside DB2. This value is an average.

INAPPLWT

The class 1 waiting time in an application. This value is an average.

INDB2WT

The class 2 waiting time in DB2. This value is an average.

OUTDB2WT

The waiting time outside DB2. This value is an average.

TOTSUSTM

The waiting time for all class 3 suspensions. This value is an average.

DMLSTAT

The total number of SQL DML statements executed. This value is an average.

DCLSTAT

The total number of DCL statements executed. This value is an average.

DDLSTAT

The total number of DDL statements executed.

UPDPERCM

The sum of SQL UPDATE, SQL INSERT, and SQL DELETE statements executed.

CMUPERUPD

The ratio of the sum of commits and rollbacks to the sum of SQL UPDATE, SQL INSERT, and SQL DELETE statements.

TOTSUSP

The number of suspensions. This value is an average.

GETPAGES

The number of Getpage requests. This value is an average.

NOTACCT

The time not accounted in DB2. You use this time to determine whether there is a large percentage of time that has not been captured within the DB2 Accounting record. This value is an average.

BUFUPDTS

The number of buffer updates. This value is an average.

SYNCREAD

The number of synchronous read I/O operations. This value is an average.

TOTPREF

The number of all types of prefetch requests. This value is an average.

PINDBET

The total elapsed time for executing the package or DBRM. This value is an average.

PINDBPT

The CPU time spent by the package or DBRM (class 7). This value is an average.

PTSUSTME

The waiting time for the package or DBRM caused by a class 8 suspension. This value is an average.

PNOTACCT

The total unaccounted time in DB2 caused by the execution of the package or DBRM. This value is an average.

DDNAME

Specifies the data set where the trace is written. The default ddname for the first trace is ACTRCDD1. The default ddnames for the second to fifth traces are ACTRCDD2 through ACTRCDD5.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Some care is needed when using PLANNAME as a filter in a DDF environment, see Special considerations for DDF trace data.

Example using TRACE without subcommand options

The following example specifies:

- An accounting short trace
- Exception data is not presented
- Written to ACTRCDD1

```

:
:
TRACE
:

```

ACCOUNTING command with FILE subcommand

This section describes the ACCOUNTING command with the FILE subcommand.

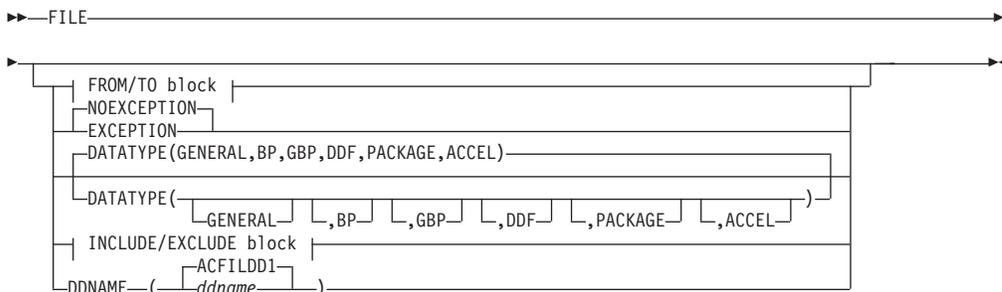
Usage

Use the FILE subcommand to format unreduced DB2 data and store it in sequential data sets suitable for use by the DB2 load utility.

Usage notes

- In the case of CP parallelism, the logical Accounting record (aggregation of all the activity within the thread) is stored in the data set.
- The records can be placed in DB2 tables, and you can produce reports by using a reporting facility such as Query Management Facility (QMF™).
- You can also use the FILE data sets to generate CSV (comma-separated value) input-data. This CSV data can then be transferred to workstations and imported into spreadsheets to improve DB2 performance analysis using graphical representations or pivot tables. For more information refer to *Reporting User's Guide*.

Syntax of the FILE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25

for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the data set by date and time.

For details, see “FROM/TO subcommand options” on page 30.

If you are reducing data, the times specified in REDUCE affect the data available for filing.

EXCEPTION

NOEXCEPTION

Specify EXCEPTION if you want to include only those file entries with values outside the user-specified limits. Otherwise, all records are included.

If you use this option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see Figure 3 on page 55.

DATATYPE

Specifies which data types are to be written to the FILE data set specified in the DDNAME subcommand option. By default, if DATATYPE is not specified, or if DATATYPE() is specified (without keywords), all data types are written.

This subcommand option lets you improve the system performance while the data set is generated by explicitly specifying only the required data types. If performance is not critical, you can safely ignore this subcommand option. Use one or more of the following keywords to specify the required data types.

GENERAL

General performance data

BP Buffer pool performance data

GBP

Group buffer pool performance data

DDF

Distributed Data Facility (DDF) data

PACKAGE

Package execution data

ACCEL

Accelerator data

DDNAME

Specifies the ddname where the file data set is written. The default ddname is ACFILDD1 for the first file instance, ACFILDD2 for the second file instance, and so on.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using FILE with EXCEPTION and INCLUDE options

This example specifies a file that includes only exception records for User10, User11, and User12. The file is generated on ddname ACFILDD1 by default.

```
⋮  
FILE
```

```

EXCEPTION
INCLUDE (PRIMAUTH(USER10 USER11 USER12))
:

```

AUDIT command

This section provides an overview of the AUDIT command.

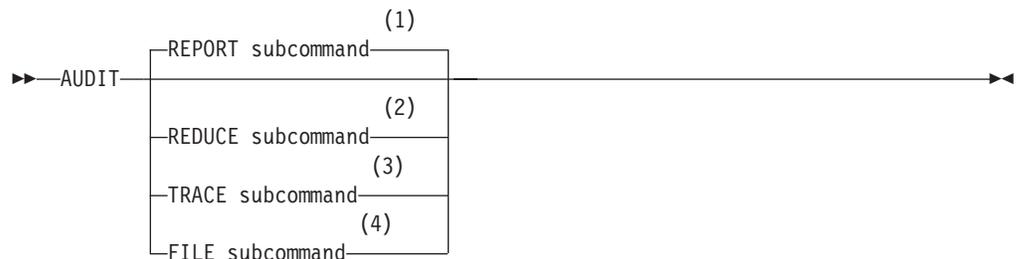
Usage

Use the AUDIT command to reduce data and generate Audit reports, traces, and file data sets. The subcommands are described in detail, together with their various options, in the following sections.

Usage notes

- This command can be used once in a job step.
- However, it can be used in the same job step with commands of the other report sets.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax



Notes:

- 1 You can specify REPORT up to 5 times.
- 2 You can specify REDUCE only once. If specified, you must also specify REPORT, TRACE, or FILE at least once.
- 3 You can specify TRACE up to 5 times.
- 4 You can specify FILE up to 7 times.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting Audit functions

The OMEGAMON XE for DB2 PE command language shown in this example is not appropriate in all circumstances. You must modify it to meet your requirements.

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See Chapter 5, “DD statements,” on page 15 for full descriptions of the DD statements.

Note:

1. There is an advantage in omitting DPMOUTDD from your JCL. For more information, see “DPMOUTDD statement” on page 16.
2. The OMEGAMON XE for DB2 PE command stream is only processed if EXEC is included as the last command. Otherwise, OMEGAMON XE for DB2 PE only checks the syntax and writes the command stream together with any information, warning, or error messages generated to the job summary log.

This is a sample of the JCL required to produce Audit reports and traces. A description of the DD statements follows the sample.

```
//          PEMAIN EXEC PGM=FPECMAIN
//* FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DPMPARMS DD DSN=MYID.FPELIB.DPMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
//DPMLLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
//* FOLLOWING ARE REPORT SET DDNAMES
//AURPTDD DD SYSOUT=*
//AUTRCDDx DD SYSOUT=*
//AUFILDDx DD DSN=MYID.FPELIB.AUDIT.FILE,DISP=OLD
//AUDWORK DD DSN=MYID.FPELIB.AUDIT.WORKDD,DISP=OLD
//* FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *

:
AUDIT
  REDUCE
  REPORT
  TRACE
  FILE
:
EXEC
```

Figure 4. Sample JCL for requesting Audit functions

AUDIT command with REPORT subcommand

This section describes the AUDIT command with the REPORT subcommand.

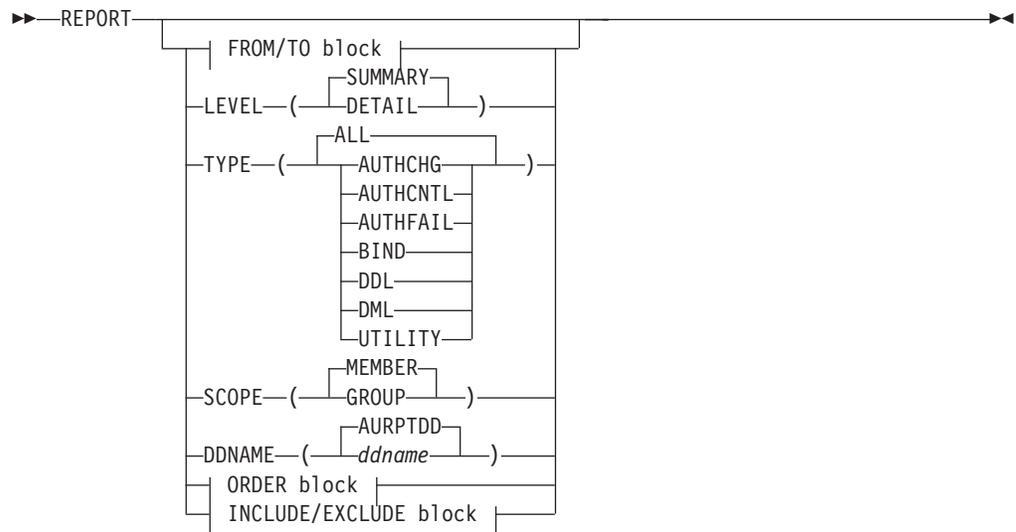
Usage

Use the REPORT subcommand to generate reports from records.

Usage notes

- Up to five REPORT subcommands can be specified within each AUDIT command.

Syntax of the REPORT subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reporting process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

SCOPE

Specifies the scope of the report in a data sharing environment.

MEMBER

GROUP

LEVEL

Specifies the level of the report. You can enter either of the following:

SUMMARY

Produces reports of aggregated audit data.

DETAIL

Produces reports of non-aggregated audit data in timestamp sequence within the requested TYPE and ORDER. Events within reports are sorted by timestamp within OMEGAMON XE for DB2 PE identifiers.

Note: There is a dependency between the two commands AUDIT REPORT LEVEL(SUMMARY) and AUDIT REPORT LEVEL(DETAIL). If you request a SUMMARY report and a DETAIL report for the same audit TYPE, then first specify the DETAIL report and then the SUMMARY report in order to receive meaningful data in the SUMMARY report. If you request a SUMMARY report without a DETAIL report, then there is no dependency and the SUMMARY report delivers valuable results.

TYPE

Specifies the type of audit data to be reported. You can enter one or more of the following:

ALL
All audit categories are reported (the default)

AUTHCHG
Changes to authorization identifiers

AUTHCNTL
GRANTS and REVOKEs of privileges

AUTHFAIL
Authorization failure

BIND
DML statements at bind of auditable DB2 tables

DDL
DDL operations against auditable DB2 tables

DML
Read/write access against auditable DB2 tables

UTILITY
Utility access against auditable DB2 tables

Note:

1. If you select audit types in REPORT that were not specified in the TYPE option of REDUCE, blank reports are generated for those audit types.
2. If TYPE is specified with LEVEL(SUMMARY), a report of aggregated audit data is produced for each type requested.
3. If TYPE is not specified with LEVEL(SUMMARY), one report showing aggregated totals for all audit types is produced.
4. If TYPE is specified with LEVEL(DETAIL), one report showing nonaggregated audit data of all the requested types is produced.
5. If TYPE is not specified with LEVEL(DETAIL), a report showing nonaggregated audit data of all types is produced.

DDNAME
Specifies the data set where the report is written.

ORDER
Specifies the OMEGAMON XE for DB2 PE identifiers and their sequence for sorting the report and, in summary reports, which identifiers are used for aggregation.

For details, see "ORDER subcommand option" on page 46 and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

The default for ORDER varies with the LEVEL and TYPE specified.

The definition of an object depends on the LEVEL and TYPE specified.

INCLUDE/EXCLUDE
Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see "INCLUDE and EXCLUDE subcommand options" on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

Example using REPORT with LEVEL and TYPE option

The following example uses LEVEL and TYPE:

```
⋮  
AUDIT
```

```

REPORT
  LEVEL (DETAIL)
  TYPE  (AUTHFAIL,AUTHCNTL)
:

```

This example produces a detail report for all authorization failures and authorization GRANTs and REVOKEs found in the input data. By default, the reports are in PRMAUTH-PLANNAME sequence and, for summary reports, the output is aggregated by PRMAUTH-PLANNAME. The output is sent to default ddname AURPTDD.

Example using REPORT with ORDER option

In the following example ORDER specifies that three authorization failure detail reports are to be produced.

```

:
REPORT
  LEVEL (DETAIL)
  TYPE  (AUTHFAIL)
  ORDER (PRMAUTH-PLANNAME-REQLOC
         CONNECT-PLANNAME REQLOC-PRMAUTH)
:

```

- The first report is ordered by requesting location within plan name within primary authorization ID.
- The second report is ordered by plan name within connection ID.
- The third report is ordered by primary authorization ID within requesting location.

AUDIT command with REDUCE subcommand

This section describes the AUDIT command with the REDUCE subcommand.

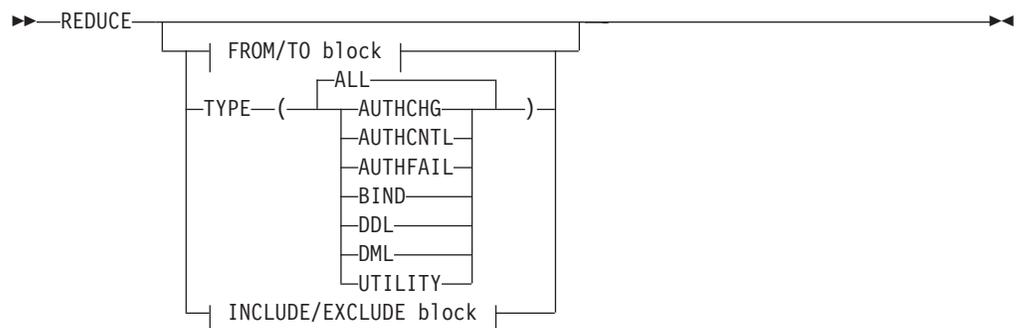
Usage

You use the REDUCE subcommand to reduce the volume of data that is input to subsequent subcommands.

Usage notes

- REDUCE can be used once in an AUDIT command.

Syntax of the REDUCE subcommand



The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25

for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

Subcommand options

FROM/TO

Limits the range of records included in the reduction process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

TYPE

Specifies the type of audit data to be reduced. One or more categories can be selected. Only the data selected on the REDUCE subcommand is available to subsequent REPORT subcommands. You can enter one or more of the following:

ALL

All audit categories are reported (the default)

AUTHCHG

Changes to authorization identifiers

AUTHCNTL

GRANTs and REVOKEs of privileges

AUTHFAIL

Authorization failure

BIND

DML statements at bind of auditable DB2 tables

DDL

DDL operations against auditable DB2 tables

DML

Read/write access against auditable DB2 tables

UTILITY

Utility access against auditable DB2 tables

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using REDUCE

This example requests the following:

- Reduce only input data of type AUTHCHG (Authority Change).
- Include only data of subsystem DSN1.
- Write the report to the data set defined by the default ddname AUDRPTDD.

```
⋮  
AUDIT  
  REDUCE  
    TYPE      (AUTHCHG)  
    INCLUDE   (SUBSYSTEM(DSN1))  
  REPORT  
⋮
```

AUDIT command with TRACE subcommand

This section describes the AUDIT command with the TRACE subcommand.

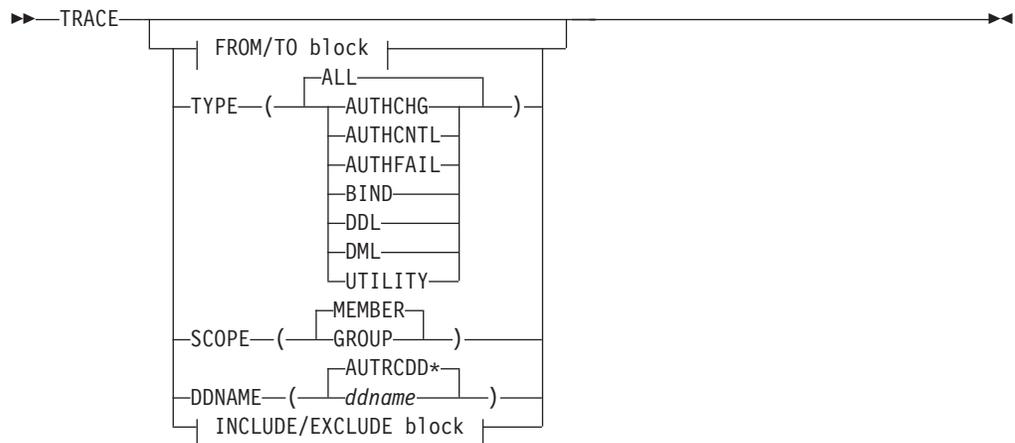
Usage

Use the TRACE subcommand to produce traces with an entry for each DB2 Audit record.

Usage notes

- Up to five traces can be requested in a job step.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the trace by date and time.

For details, see “FROM/TO subcommand options” on page 30.

TYPE

Identifies the type of data traced. You can enter one or more of the following:

ALL

All audit categories are reported (the default)

AUTHCHG

Changes to authorization identifiers

AUTHCNTL

GRANTs and REVOKEs of privileges

AUTHFAIL

Authorization failure

BIND

DML statements at bind of auditable DB2 tables

DDL

DDL operations against auditable DB2 tables

DML

Read/write access against auditable DB2 tables

UTILITY

Utility access against auditable DB2 tables

SCOPE

Specifies the scope of the trace.

MEMBER**GROUP****DDNAME**

Specifies the data set where the trace is written.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using TRACE with TYPE option

This command traces only authorization failures (in the order that they occur). The output goes to the default ddname AUTRCDD1.

```

:
AUDIT
  TRACE
  TYPE (AUTHFAIL)
:

```

AUDIT command with FILE subcommand

This section describes the AUDIT command with the FILE subcommand.

Usage

Use the FILE subcommand to format unreduced DB2 data and store it in sequential data sets suitable for use by the DB2 load utility. The records can be placed in DB2 tables, and you can produce reports by using a reporting facility such as Query Management Facility (QMF).

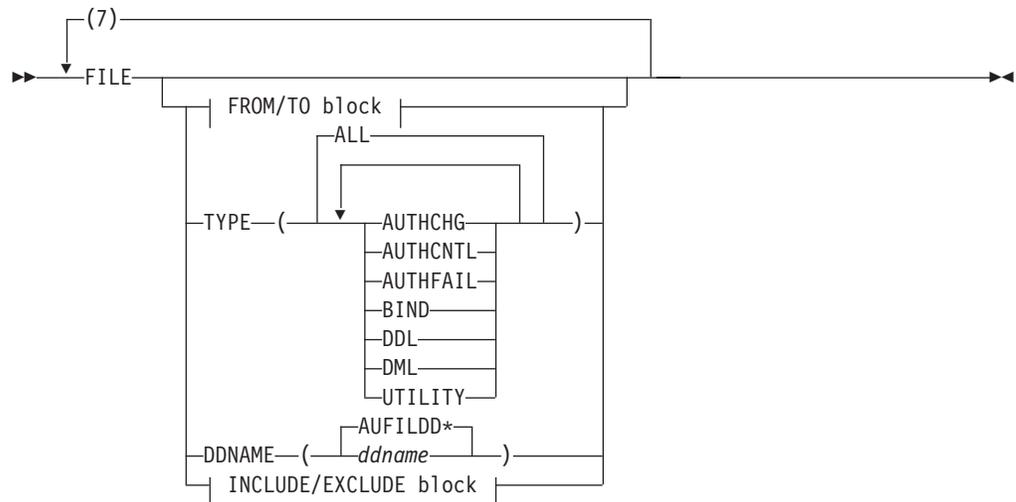
Usage notes

By using the FILE subcommand you can:

- Process the different audit types separately by specifying one audit type for each FILE subcommand.
- Process the different audit types simultaneously by specifying any number of the available audit types in each FILE subcommand.

The FILE subcommand can occur a maximum of seven times in a job step.

Syntax of the FILE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the data set by date and time.

For details, see “FROM/TO subcommand options” on page 30.

TYPE

Selects the audit category. It identifies the type of data included in the data set. You can enter one or more of the following:

ALL

All audit categories are reported (the default)

AUTHCHG

Changes to authorization identifiers

AUTHCNTL

GRANTS and REVOKES of privileges

AUTHFAIL

Authorization failure

BIND

DML statements at bind of auditable DB2 tables

DDL

DDL operations against auditable DB2 tables

DML

Read/write access against auditable DB2 tables

UTILITY

Utility access against auditable DB2 tables

DDNAME

Specifies the ddname where the file data set is written. The default ddname is AUFILDD1 for the first file data set, and AUFILDD2 to AUFILDD7 for the second to seventh file data sets.

You can specify a different ddname by using the DDNAME option in the FILE subcommand. In this case, your JCL must contain a valid DD statement for the ddname you specify. If you do not specify a different ddname, your JCL must contain a valid DD statement for the corresponding default ddname. For

example, if you omit the DDNAME option from the third FILE subcommand in the job stream, your JCL must contain a valid DD statement for AUFILDD3.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using FILE with TYPE option

This command generates a data set in the default ddname AUFILDD1. The data set contains one or more records for each authorization failure: an AUTHFAIL record and any matching records containing the text of the SQL statement that caused the authorization failure.

```
⋮  
FILE  
  TYPE (AUTHFAIL)  
⋮
```

EXEC command

This section describes the EXEC command.

Usage

Typically, the EXEC command is the last statement in the OMEGAMON XE for DB2 PE command stream. This command causes the generation of any previously specified reports. It takes no arguments.

Usage notes

- When the EXEC command is not present, no report is produced. The syntax of the command stream is checked and written out, together with any information, warning, or error messages to the DPMOUTDD data set.
- All statements following the EXEC are ignored.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax

▶▶ EXEC ◀◀

EXPLAIN command

This section provides an overview of the EXPLAIN command.

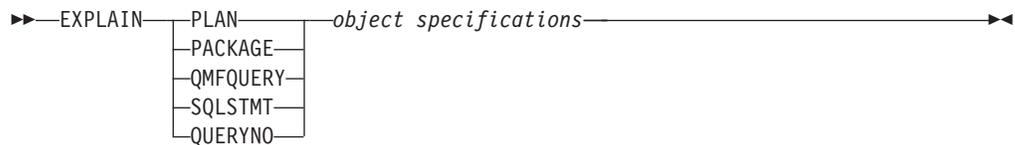
Usage

Use the OMEGAMON XE for DB2 PE EXPLAIN command to produce explain reports.

Usage notes

- You can specify any number of EXPLAIN commands in the OMEGAMON XE for DB2 PE command stream.
- For each EXPLAIN command a separate OMEGAMON XE for DB2 PE explain report is produced.
- Some parameters can include a wildcard (*) provided that it is the last character in the text string.

Syntax



Subcommands and options

Each of the functions is identified with an appropriate keyword, followed by various subcommands and options, which identify the object being explained and control the amount of detail being produced.

The following sections identify the subcommands available for each function. Here you find a description of the EXPLAIN subcommands and options that can be specified with the EXPLAIN and GLOBAL commands:

ACCTYPE

Can be specified for the PLAN and PACKAGE keywords to control the statements that are to be explained within the plan or package based on the chosen access path. The following values are available:

ALL

For each explainable SQL statement in the plan, a report is produced. This is the default value.

MATCHING

Only explainable SQL statements where an access path of matching index scan has been selected, are processed and reported.

NONMATCH

Only explainable SQL statements where an access path of nonmatching index scan has been selected, are processed and reported.

TABSCAN

Only explainable SQL statements where an access path of table space scan has been selected, are processed and reported.

Note: ACCTYPE selection, as with any other OMEGAMON XE for DB2 PE explain selection, applies to individual PLAN_TABLE rows, not to the entire SQL statement. For example, if a particular SQL statement is executed in two steps, the first using matching index scan and the second using nonmatching index scan, and ACCTYPE(MATCHING) is specified, only the first step is reported.

DBRM

Can be specified for the PLAN keyword to control the DBRMs to be explained within the plan. If DBRM is not specified, all DBRMs within the plan are

explained. If only a given DBRM within the plan is to be explained, you must specify the actual DBRM name in DBRM. If all DBRMs with a given name pattern are to be explained, a wildcard (*) can be used.

DEGREE

Can be specified for the QMFQUERY and SQLSTMT keywords to indicate whether the SQL statement or statements are eligible for query parallelism. The following values are available:

1 The SQL statement does not use query parallelism. This is the default.

ANY

The SQL statement is eligible for query parallelism.

DSJ

Disable star join.

ESJ

Enable star join.

FIRST/LAST

Specifies the number of the first and the last statement in the plan or package to be explained, to control the range of SQL statements to be explained within the plan or package.

If FIRST is not specified, a value of 1 is used. If LAST is not specified, a value of 999 999 999 is used. If FIRST is greater than LAST, both parameters are set to the value of FIRST.

FORCE

Can be specified for the PACKAGE keyword to control the explanation of the SQL statements in a package. The following values are available:

NO If more than one package is referenced, the statements are only explained if the total number of SQL statements is less than 300. This is the default value.

YES

All statements in all packages that conform to the specification are explained.

FORMAT

Can be specified for the PLAN and PACKAGE keywords to control the formatting of the SQL statements in the plan or package. The following values are available:

YES

The SQL statement is formatted so that a new line is started for SQL keywords such as SELECT, INTO, FROM, WHERE. Subselects, however, are not indented. This is the default value.

NO The SQL statement is formatted so that a new line is only started for the SQL keywords SELECT and UNION. By using this option, the SQL statement uses minimum page space.

GEN

Controls the number of version generations in a package to be explained, with PLAN and PACKAGE.

Specify a value from 1 to 99 to overwrite the default. For PLAN, if GEN is not specified, the default is 1.

For PACKAGE, if GEN is not specified, the default is 1, when no wildcard is used in the version ID of the package. If a wildcard is present, a value of 99 is used.

HOSTVAR

Can be specified for the PLAN and PACKAGE keywords to control the listing of host variable specifications. The following values are available:

NO No host variable definitions are listed in the report. This is the default value.

YES

A listing of all (maximum 500) host variables used in the SQL statement is produced. This listing includes the definition type, length, and null characteristics.

INDEX

Controls the level of index information unless the index data block is excluded by using the LEVEL option. The following values are available:

YES

If DB2 has selected a matching or nonmatching index scan, OMEGAMON XE for DB2 PE explain shows all index information for the selected index including key column information. If DB2 has selected a table space scan, OMEGAMON XE for DB2 PE explain shows detailed index information for all indexes of the accessed table. This is the default value.

NO No index information is shown in the report.

ALL

Information for all indexes created for the table (including key column information) is shown in the report. The information is listed after the table details.

LEVEL

Controls which of the following blocks of data the OMEGAMON XE for DB2 PE explain report contains:

- “Raw” SQL EXPLAIN data as found in the PLAN_TABLE
- Access path data
- Table and table space data
- Index data
- Key data
- Distribution of the ten mostly used keys
- Plan/package bind data, if applicable
- Host variables data, if applicable
- Summary report

The table below summarizes which values can be specified and which blocks are reported.

Table 9. LEVEL values

LEVEL	Raw SQL Explain	Access Path Data	Table (Space) Data	Index Data	Key Data	Top Ten Key Dist.	Plan/Package Data	Host Variables	Summary Report
SUMMARY									●
SQL		●							●
BASIC	●	●							●
INDEXES	●	●	●	●			●	●	●
DETAIL	●	●	●	●	●		●	●	●
NORAWXPL		●	●	●	●		●	●	●

Table 9. LEVEL values (continued)

LEVEL	Raw SQL Explain	Access Path Data	Table (Space) Data	Index Data	Key Data	Top Ten Key Dist.	Plan/Package Data	Host Variables	Summary Report
KEYDIST	●	●	●	●	●	●	●	●	●
(None)	●	●	●	●	●		●	●	

LOCATION

Can be specified for the PLAN keyword to determine the location of the plan. If LOCATION is not specified, the local location, that is, the location specified in the SSID option, is used.

When a valid location name is specified, OMEGAMON XE for DB2 PE explain connects to the specified location and EXPLAIN PLAN processing continues at that location. To be valid, the location name must appear in the LOCATION column of the SYSIBM.SYSLOCATIONS table, or be the local DB2 subsystem's location name.

OWNER

Can be specified for the GLOBAL and QUERYNO keywords to control the authorization ID of the PLAN_TABLE for the current request. To select the correct PLAN_TABLE, OMEGAMON XE for DB2 PE explain looks for the authorization ID specified in the OWNER option of QUERYNO. If it is not specified, the OWNER value in GLOBAL is used. If OWNER is not specified in GLOBAL, OMEGAMON XE for DB2 PE explain uses the authorization ID of the job submitter.

PACKAGES

Can be specified for the PLAN keyword to control the explanation of statements in packages within a plan. The following values are available:

YES

For each package in the plan, only statements in the most recent version are explained. This is the default value.

NO No statements in the packages are explained.

ALL

All statements in all versions of the packages are explained.

PACKLIMIT

Controls the number of packages to be explained.

A particular plan can consist of more packages than has been specified in PACKLIMIT (default is 100). In this case, a report listing all packages is produced but no SQL statements in these packages are explained. PACKLIMIT does not affect DBRMs.

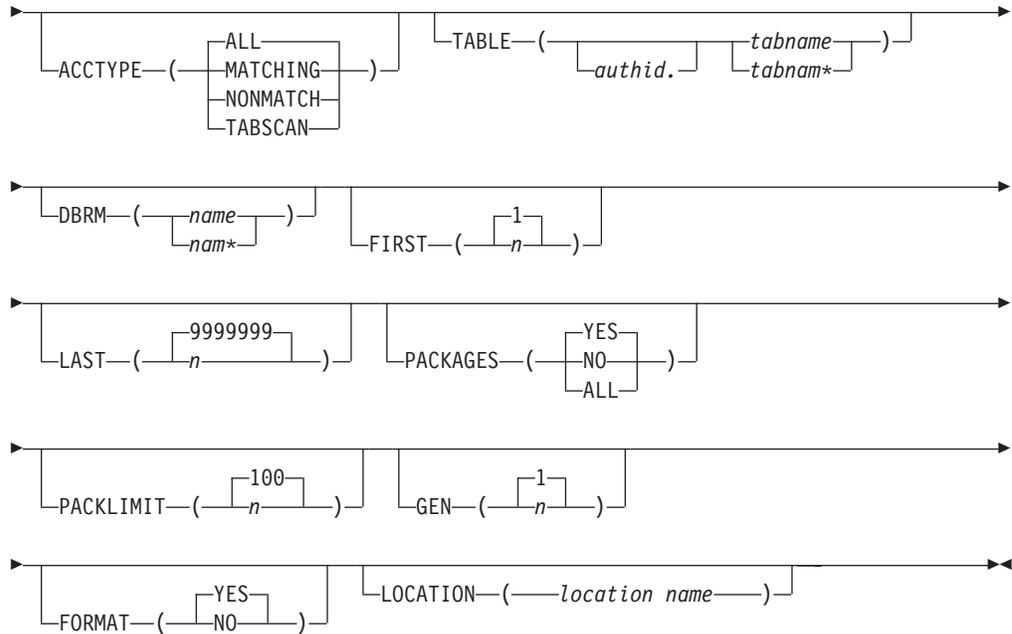
PLANEXPLAIN

One of the steps in OMEGAMON XE for DB2 PE installation is to bind the OMEGAMON XE for DB2 PE explain application. The default name for this plan is KO2EXPL.

If you do not want to use the default name, you can specify the name of the OMEGAMON XE for DB2 PE explain plan at OMEGAMON XE for DB2 PE execution time. This is done by specifying PLANEXPLAIN(xxxxxxxx) in GLOBAL, where xxxxxxxx is the OMEGAMON XE for DB2 PE explain plan name.

SQLID

Defines the current SQL ID. If you specify SQLID(USER), the primary SQL ID



Notes:

- 1 SSID is required unless specified in a preceding GLOBAL command.

Subcommand options

Note: This section only explains specific subcommands. The other subcommands and options are described in “Subcommands and options” on page 83. The following list gives additional or specific descriptions of selected options, where appropriate.

planname

Identifies the plan that statements are to be explained for.

EXPLAIN PACKAGE specifications

This section describes the EXPLAIN PACKAGE specifications.

Usage

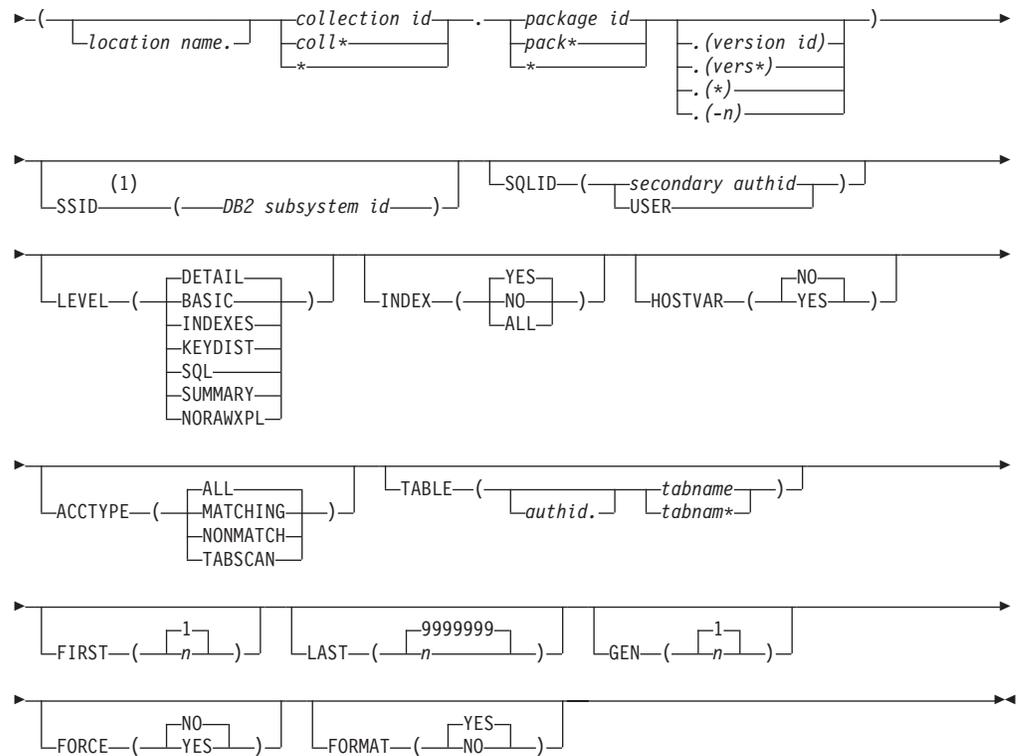
Use this specification to investigate all or selected SQL statements for given packages.

Usage notes

- OMEGAMON XE for DB2 PE EXPLAIN verifies that the specified package exists in the catalog table SYSIBM.SYSPACKAGE.
- If it does, OMEGAMON XE for DB2 PE EXPLAIN further verifies that the package has been bound with the EXPLAIN(YES) option.

Syntax of the EXPLAIN PACKAGE specification





Notes:

- 1 SSID is required unless specified in a preceding GLOBAL command.

Subcommand options

Note: This section only explains specific subcommands. The other subcommands and options are described in “Subcommands and options” on page 83. The following list gives additional or specific descriptions of selected options, where appropriate.

collection id
package id
version id

Identify the package that statements are to be explained for.

The asterisk is allowed as wildcard character for *collection id*, *package id*, and *version id*.

The *version id* string can contain the following special characters: underscore (_), at (@), number (#), dollar (\$), dash (-), and period (.)

Use *-n* (the version generation number preceded by a minus sign) to specify the version id. The version generation number is the line counter of a package list sorted by the precompile date in descending order. The newest package version has the generation number "0" and the oldest package version has the generation number "-n".

EXPLAIN QMFQUERY specifications

This section describes the EXPLAIN QMFQUERY specifications.

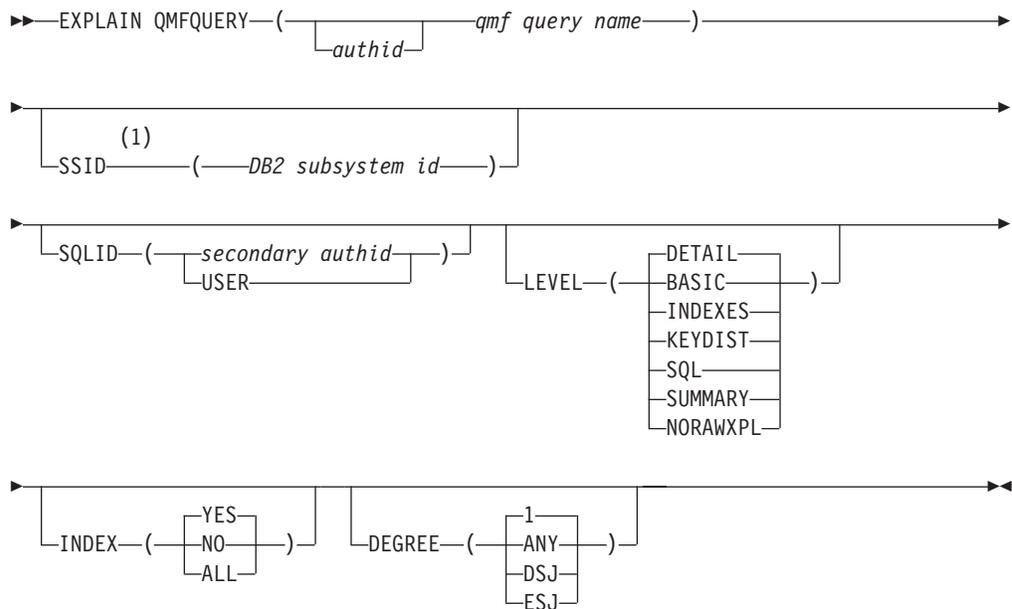
Usage

Use this specification to check if a saved QMF query is written in the SQL language.

Usage notes

- QBE and PROMPTED queries must be converted to SQL before they can be explained.
- Apart from explaining your own saved QMF queries, you can explain a query created by other users provided that the query was saved with SHARE=YES.
- The QMF query can contain parameters, for example PARM1 and &PARM1; These parameters can also substitute column names in the select list. However, the query must not contain literals and other strings with one or more ampersands (&), enclosed between quotes ('). If there is more than one entry in the PLAN_TABLE with the same identifiers (for example, query number, plan name, and program name), a report for the most recent entry is produced.

Syntax of the EXPLAIN QMFQUERY specification



Notes:

- 1 SSID is required unless specified in a preceding GLOBAL command.

Subcommand options

Note: This section only explains specific subcommand options. The other subcommand options are described in “Subcommands and options” on page 83. The following list gives additional or specific descriptions of selected options, where appropriate.

authid

qmf query name

Identify the saved QMF query to be explained.

EXPLAIN SQLSTMT specifications

This section describes the EXPLAIN SQLSTMT specifications.

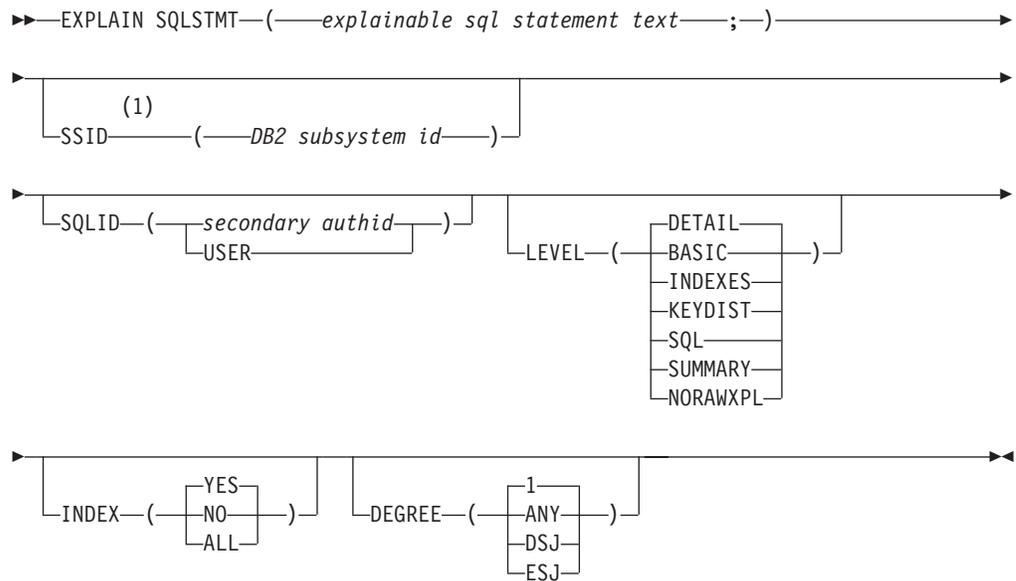
Usage

Use this specification to investigate a particular SQL statement supplied in its text form.

Usage notes

- The supplied SQL statement is explained by SQL EXPLAIN using a query number of 999735911. If this number already exists in the PLAN_TABLE of the job submitter, OMEGAMON XE for DB2 PE explain deletes the rows before processing the SQL statement. After successful execution of SQL EXPLAIN, the newly created rows in the PLAN_TABLE of the job submitter are used to produce the usual OMEGAMON XE for DB2 PE Explain report.

Syntax of the EXPLAIN SQLSTMT specification



Notes:

- 1 SSID is required unless specified in a preceding GLOBAL command.

Subcommand options

Note: This section only explains specific subcommands. The other subcommands and options are described in “Subcommands and options” on page 83. The following list gives additional or specific descriptions of selected options, where appropriate.

explainable SQL statement text
Can span several lines.

EXPLAIN QUERYNO specifications

This section describes the EXPLAIN QUERYNO specifications.

Usage

Use this specification to check if a particular SQL statement that is identified by its query number. You can obtain the query number as follows:

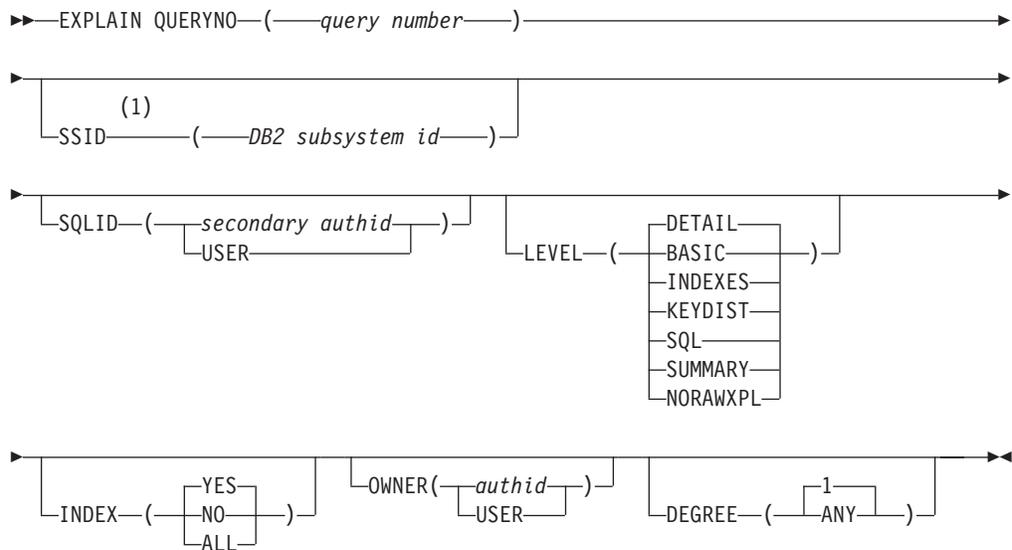
- A dynamic SQL EXPLAIN statement has been executed with a given query number. The SQL EXPLAIN statement might have been executed from either DB2I or QMF. If the query number is not specified in the SQL EXPLAIN statement, DB2 assigns a number. You can obtain the query number directly from the PLAN_TABLE.
- The application has been bound with the EXPLAIN(YES) option on the BIND or REBIND commands. The query number is the statement number that was assigned by the precompiler and placed in the DBRM.

Note: If a statement belonging to a particular plan or package is to be explained, EXPLAIN PLAN or EXPLAIN PACKAGE specifications are better suited than EXPLAIN QUERYNO. By using EXPLAIN PACKAGE or EXPLAIN PLAN options such as FIRST, LAST, or DBRM, the statement can be better identified.

Usage notes

- OMEGAMON XE for DB2 PE EXPLAIN searches for the query number in the job submitter's PLAN_TABLE, unless the OWNER keyword with a different user ID is specified.
- If the job submitter has SELECT authorization to another user's PLAN_TABLE, the job submitter can select EXPLAIN information from this table, by specifying the other user's authorization ID as the OWNER option.
- If the specified query number does not exist in the PLAN_TABLE, a warning message is issued. OMEGAMON XE for DB2 PE resumes processing with the next request.

Syntax of the EXPLAIN QUERYNO specification



Notes:

- 1 SSID is required unless specified in a preceding GLOBAL command.

Subcommand options

Note: This section only explains specific subcommands. The other subcommands and options are described in “Subcommands and options” on page 83. The following list gives additional or specific descriptions of selected options, where appropriate.

query number

Identifies a query number in the PLAN_TABLE being accessed.

IOACTIVITY command

This section provides an overview of the IOACTIVITY command.

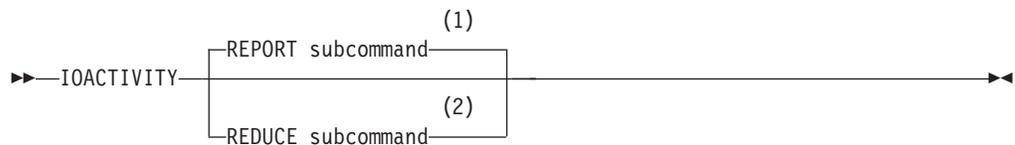
Usage

Use the IOACTIVITY command to reduce data and generate I/O Activity reports.

Usage notes

- You can specify this command only once in a job step.
- However, you can use it in the same job step with commands of the other report sets.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax



Notes:

- 1 You can specify REPORT up to 5 times.
- 2 You can specify REDUCE only once. If specified, you must also specify REPORT at least once.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample command stream for IOACTIVITY

See Chapter 5, “DD statements,” on page 15 for full descriptions of the DD statements contained in this sample.

```

//          EXEC PGM=FPECMAIN
//* FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
//DPMLG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
//* FOLLOWING ARE REPORT SET DDNAMES
//* FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *
IOACTIVITY
  REDUCE
  REPORT
EXEC

```

Figure 5. Sample command stream for IOACTIVITY

IOACTIVITY command with REPORT subcommand

This section describes the IOACTIVITY command with the REPORT subcommand.

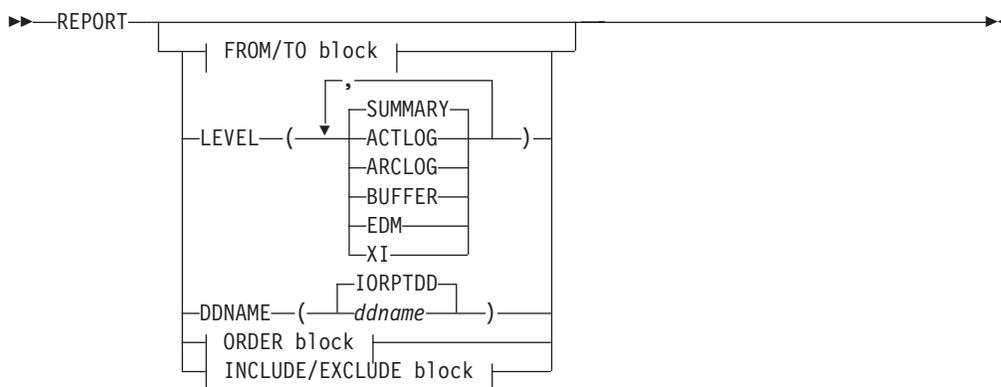
Usage

Use the REPORT subcommand to generate reports from records.

Usage notes

- Up to five REPORT subcommands can be specified within each IOACTIVITY command.

Syntax of the REPORT subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reporting process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

LEVEL

Specifies the level of the report. The following keywords can be used with LEVEL:

SUMMARY

Gives a summary of all I/O activity for the following categories:

- Buffer pool
- EDM pool
- Active log
- Archive log and BSDS
- Cross invalidation

SUMMARY also gives the grand totals of all the I/O statistics for the reporting interval for each category. It provides a quick overview of system-wide I/O activity to help monitor trends and identify potential problem areas.

The following levels generate detail reports. See the ORDER subcommand option for default ordering of detail reports.

ACTLOG

The active log report presents counts and average elapsed time for reads, writes, and other I/O activity.

ARCLOG

The archive log/BSDS report presents the archive log and BSDS read and write requests.

BUFFER

The buffer pool report presents successful and unsuccessful asynchronous and synchronous (prefetch) read requests and write requests.

EDM

The environmental descriptor manager pool report presents cursor table (CT), package table (PT), and database descriptor (DBD) references, loads from a hard disk drive, elapsed time for each load, and average section lengths.

XI The cross-invalidation report presents buffer refresh events caused by aggregated cross invalidation.

The default for LEVEL is SUMMARY.

DDNAME

Specifies the data set where the report is written. You can specify any valid ddname including the default, provided that your JCL contains a DD statement for it. If a DD statement is omitted, it will be dynamically allocated to the SYSOUT message class of the job. The default ddname for report is IORPTDD.

ORDER

Specifies the OMEGAMON XE for DB2 PE identifiers and their sequence for sorting detail reports. Note that ORDER cannot be used with LEVEL(SUMMARY), which is the default. To use ORDER, you must specify LEVEL with one or more of the detail levels.

For details, see “ORDER subcommand option” on page 46 and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

The default for ORDER (if not explicitly specified) depends on the specification of LEVEL:

- LEVEL(ACTLOG) orders by DATASET.
- LEVEL(ARCLOG) orders by DATASET.
- LEVEL(BUFFER) orders by PRIMAUTH-PLANNAME-BPID.
- LEVEL(EDM) orders by PRIMAUTH-PLANNAME.
- LEVEL(XI) orders by PRIMAUTH-PLANNAME.
- If LEVEL specifies multiple detail reports, default ordering is by PRIMAUTH-PLANNAME.

If ORDER is used once, and if LEVEL specifies multiple detail reports, each detail report is ordered by the OMEGAMON XE for DB2 PE identifiers specified with ORDER.

If ORDER is used multiple times, and if LEVEL specifies one or more detail reports, each detail report is repeated for each specified ordering.

Note: The EDM pool report contains the plan name and the database name. Ordering by these identifiers results in repetitious information.

In the following example ORDER specifies that three EDM pool reports are produced.

```

:
:
REPORT
  LEVEL (EDM)
  ORDER (PRIMAUTH-PLANNAME-REQLOC
        CONNECT-PLANNAME INTERVAL-PRIMAUTH-PLANNAME)
:

```

- The first report is ordered by requesting location within plan name within primary authorization ID.
- The second report is ordered by plan name within connection ID.
- The third report is ordered by requesting location within primary authorization ID within interval.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using REPORT with FROM, TO, LEVEL, and INCLUDE options

This example specifies the following:

- A summary report
- Using records with the time and date range of 10:00 a.m. on 22 November 2002 to noon on 23 November 2002
- Data is included that is only associated with the location in the range of LOCN01 to LOCN05

```

REPORT
FROM (11/22/02,10:00:00.00)
TO (11/23/02,12:00:00.00)
LEVEL (SUMMARY)
INCLUDE (LOCATION(R(LOCN01 LOCN05)))

```

Example using REPORT with FROM, TO, LEVEL, ORDER, and EXCLUDE options

This example specifies the following:

- A buffer pool report
- Sorted by plan name within primary authorization ID within correlation name
- Records are used with the time and date range of 10:00 a.m. on 18 March 2002 to noon on 19 March 2002
- Data is excluded that is associated with the following locations:
 - LOCN10
 - LOCN12
 - LOCN15
 - LOCN20

```
REPORT
FROM    (03/18/02,10:00:00.00)
TO      (03/19/02,12:00:00.00)
LEVEL   (BUFFER)
ORDER   (CORRNAME-PRMAUTH-PLANNAME)
EXCLUDE (LOCATION(LOCN10 LOCN12 LOCN15 LOCN20))
```

Example using IOACTIVITY with REDUCE and REPORT subcommands

This example requests the following:

- Reduce only input data with a timestamp within the time range of 10:30 to 11:00 on 14 May 2002.
- Create a summary report (the default) with the default order of PLANNAME within PRMAUTH.
- Write the report to the data set defined by the default ddname IORPTDD.

```
IOACTIVITY
REDUCE
FROM (05/14/02,10:30:00.00)
TO   (,11:00:00.00)
REPORT
```

IOACTIVITY command with REDUCE subcommand

This section describes the IOACTIVITY command with the REDUCE subcommand.

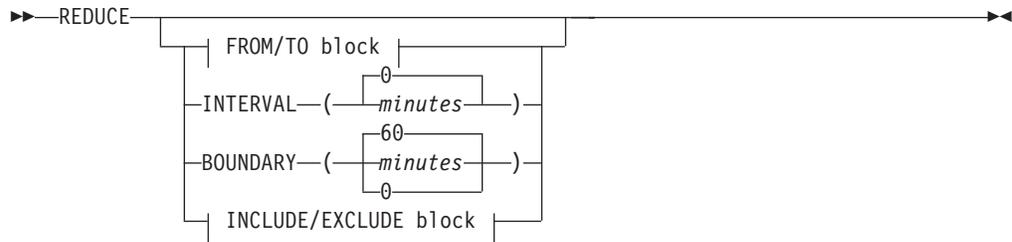
Usage

Use the REDUCE subcommand to reduce the volume of data that is input to subsequent subcommands. REDUCE consolidates records with certain common characteristics into one record.

Usage notes

- You can specify REDUCE only once in an IOACTIVITY command.

Syntax of the REDUCE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reduction process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

INTERVAL

Defines the time interval for consolidating and averaging records. Note the remarks about performance impact in “INTERVAL subcommand option” on page 37.

BOUNDARY

Controls the alignment of the intervals used to summarize records in the reduction process.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using REDUCE with REPORT

This example specifies the following:

- Reduce only input data with a timestamp within the time range of 10:00 to 11:00 on 14 May 2002.
- Create a summary report (the default) with the default order of PLANNAME within PRIMAUTH
- Write the report to the data set defined by the default ddname IORPTDD.

```

IOACTIVITY
  REDUCE
    FROM (05/14/02,10:00:00.00)
    TO (,11:00:00.00)
  REPORT
  
```

LOCKING command

This section provides an overview of the LOCKING command.

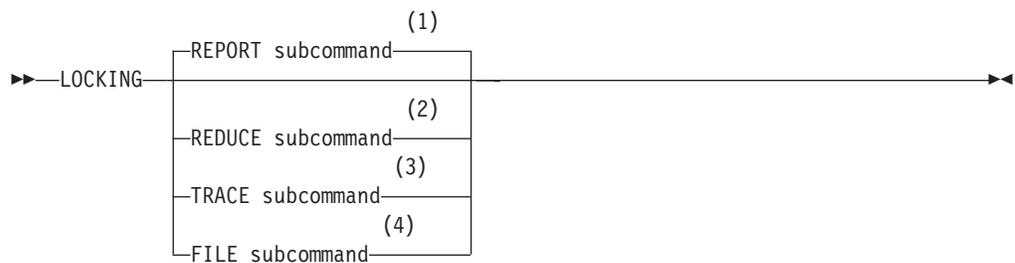
Usage

Use the LOCKING command to generate Locking reports, traces, and file data sets on locking-related DB2 data.

Usage notes

- This command can be used once in a job step.
- However, it can be used in the same job step with commands of the other report sets.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax



Notes:

- 1 You can specify REPORT up to 5 times.
- 2 You can specify REDUCE only once. If specified, you must also specify REPORT at least once.
- 3 You can specify TRACE up to 5 times.
- 4 You can specify FILE only once.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting Locking functions

The following is a sample of the JCL required to produce Locking reports and traces. A description of the DD statements follows the sample.

```

//          PEGMAIN EXEC PGM=FPECMMAIN
/** FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMAIN,DISP=SHR
//DPMLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
/** FOLLOWING ARE DDNAMES
//LOWORK DD DSN=MYID.FPELIB.LO.WORKDD,DISP=OLD
//LORPTDD DD SYSOUT=*
//LOTRCDD1 DD SYSOUT=*
//LOTRCDD2 DD SYSOUT=*
//LOTRCDD3 DD SYSOUT=*
//LOTRCDD4 DD SYSOUT=*
//LOTRCDD5 DD SYSOUT=*
//LOFILDD1 DD DSN=MYID.FPELIB.LOFIL.DATA,DISP=OLD
/** FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *

:
LOCKING
  TRACE
  REPORT
  FILE
:
EXEC

```

Figure 6. Sample JCL for requesting Locking functions

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See Chapter 5, “DD statements,” on page 15 for descriptions and more information.

Note:

1. There is an advantage in omitting DPMOUTDD from your JCL. For more information, see “DPMOUTDD statement” on page 16.
2. If you do not include the EXEC statement in your JCL, no report is produced. The syntax of your JCL is checked and written to the DPMLOG data set together with any information, warning, or error messages raised.

LOCKING command with REPORT subcommand

This section describes the LOCKING command with the REPORT subcommand.

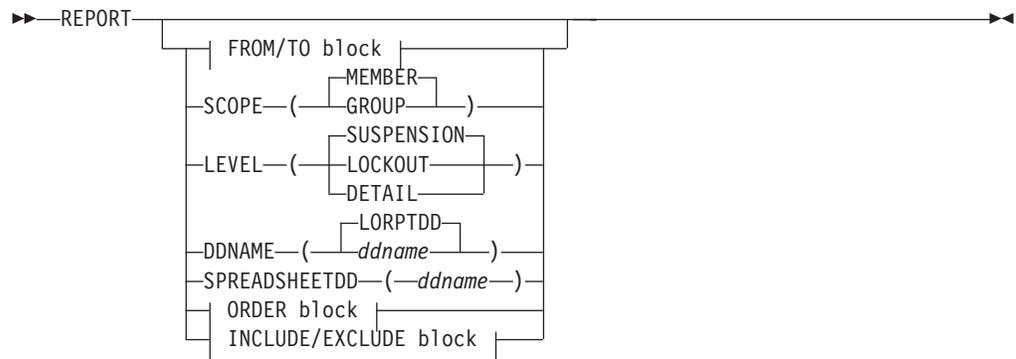
Usage

Use the REPORT subcommand to generate reports from records.

Usage notes

- Up to five REPORT subcommands can be specified within each LOCKING command.

Syntax of the REPORT subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

LEVEL

Specifies the level of the report. Both the **SUSPENSION** and **LOCKOUT** keyword can be used together in the same **REPORT** subcommand.

SUSPENSION

Generates a lock suspension report. This is the default level.

LOCKOUT

Generates a lockout report.

DETAIL

Generates a detail report.

SPREADSHEETDD

Specifies that lock suspension report data should be written to the specified data set. The data is written in a format that can be imported by spreadsheet programs. The data set should be allocated using the following attributes:

RECFM:

VB

LRECL:

1 200 or greater

Use this subcommand option only together with **LEVEL(SUSPENSION)**, which is the default if **LEVEL(LOCKOUT)** is not specified. See also “Example using **REPORT** with **LEVEL**, **ORDER**, and **SPREADSHEETDD** options” on page 102.

This subcommand option accumulates additional locking data. When written to the specified data set, the individual fields are separated by a separator character (a colon) so that workstation-based spreadsheet programs can import the data for further evaluation. When you download the data set to a client, ensure that you specify **ascii** or **text**, not **binary**, as transfer type in your file transfer program. This ensures that the necessary EBCDIC to ASCII conversion is performed.

For more information about using lock suspension report data, see the Report Reference (“Locking Report”).

Example using REPORT with LEVEL, ORDER, and INCLUDE options

This example specifies the following:

- A lock suspension report
- Sorted by plan name
- Data is included that is only associated with the location in the range of LOCN01 to LOCN05

```
REPORT
  LEVEL   (SUSPENSION)
  ORDER   (PLANNAME)
  INCLUDE (LOCATION(R(LOCN01 LOCN05)))
```

Example using REPORT with FROM, TO, LEVEL, ORDER, and EXCLUDE options

This example specifies the following:

- A lockout report
- Sorted by plan name within primary authorization ID within correlation name
- Records are used with the time and date range of 10:00 a.m. on 18 March 2001 to noon on 19 March 2001
- Data is excluded that is associated with the following locations:
 - LOCN10
 - LOCN12
 - LOCN15
 - LOCN20

```
REPORT
  FROM   (03/18/01,10:00:00.00)
  TO     (03/19/01,12:00:00.00)
  LEVEL  (LOCKOUT)
  ORDER  (CORRNAME-PRMAUTH-PLANNAME)
  EXCLUDE (LOCATION(LOCN10 LOCN12 LOCN15 LOCN20))
```

Example using REPORT with LEVEL, ORDER, and SPREADSHEETDD options

This example specifies the following:

- A lock suspension report
- A sort order by PAGESET identifier within DATABASE identifier, which is the default, but any other ordering will do
- The data set TEST1DD where the Locking report is written in a spreadsheet-compatible format

```
REPORT
  LEVEL   (SUSPENSION)
  ORDER   (DATABASE-PAGESET)
  SPREADSHEETDD (TEST1DD)
```

LOCKING command with REDUCE subcommand

This section describes the LOCKING command with the REDUCE subcommand.

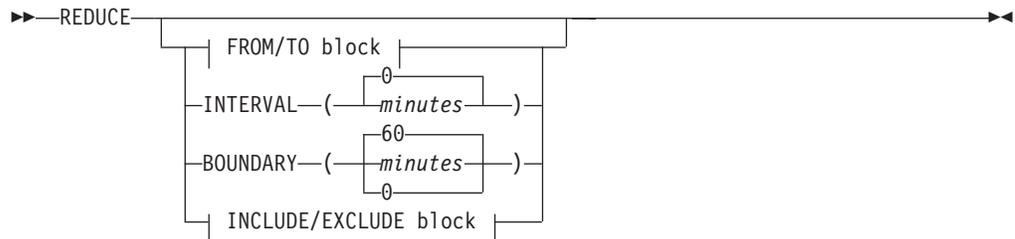
Usage

You use the REDUCE subcommand to limit the volume of data that is input to subsequent subcommands.

Usage notes

- REDUCE consolidates records with certain common characteristics into one record.
- REDUCE can be used once in a LOCKING command.

Syntax of the REDUCE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options.

Example using REDUCE with FROM, TO, and INTERVAL options

This example specifies that data is to be reduced between 10:00 a.m. on 4 February 2001 and noon on 5 February 2001. The records are to be reduced into 60-minute intervals.

```
:
:
REDUCE
  FROM      (02/04/01,10:00:00.00)
  TO        (02/05/01,12:00)
  INTERVAL  (60)
:
```

LOCKING command with TRACE subcommand

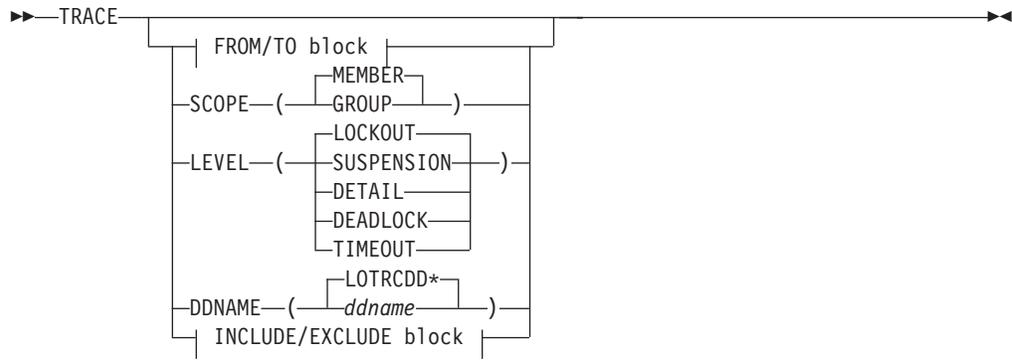
This section describes the LOCKING command with the TRACE subcommand.

Use the TRACE subcommand to produce traces with an entry for every DB2 locking event.

Usage notes

- Up to five traces can be requested in a job step.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

LEVEL

Specifies the type of trace and the amount of detail. One LEVEL option can be specified for each TRACE subcommand:

DETAIL

Generates a lock detail trace.

SUSPENSION

Generates a lock suspension trace.

LOCKOUT

Generates a lockout trace. This is the default.

DEADLOCK

Generates a deadlock trace.

TIMEOUT

Generates a timeout trace.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Note that the identifier TYPE can only be used with LEVEL(DETAIL).

Example using TRACE with FROM and TO options

This example specifies that:

- Three lockout (default) traces are produced for the time intervals specified.
- Each trace covers a unique 10-minute period of time on the same day.
- Each trace is written to the data set defined by the default ddname for the trace.

```

:
TRACE
FROM (03/18/00,10:00:00.00)
TO (03/18/00,10:10:00.00)

```

```
TRACE
  FROM (03/18/00,11:00:00.00)
  TO   (03/18/00,11:10:00.00)
TRACE
  FROM (03/18/00,12:00:00.00)
  TO   (03/18/00,12:10:00.00)
:
```

LOCKING command with FILE subcommand

This section describes the LOCKING command with the FILE subcommand.

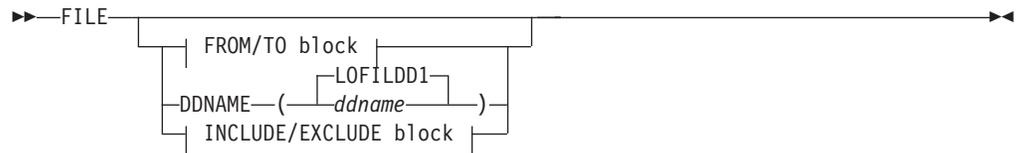
Usage

Use the FILE subcommand to format unreduced DB2 data and store it in sequential data sets suitable for use by the DB2 load utility. The records can be placed in DB2 tables and you can produce reports by using a reporting facility such as Query Management Facility (QMF).

Usage notes

- FILE is used to format lock detail records (from IFCIDs 21, 211, 212, and 223) describing Internal Resource Lock Manager (IRLM) requests, claim and drain requests, and successful lock avoidance events.
- Only one file data set can be generated in a job step.

Syntax of the FILE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options.

Example using FILE with INCLUDE option

This command generates a data set on the default ddname LOFILDD1. The data set contains detailed information about the locking requests made by the primary authorization ID SYSADM.

```

:
:
FILE
  INCLUDE (PRIMAUTH(SYSADM))
:
```

RECTRACE command

This section provides an overview of the RECTRACE command.

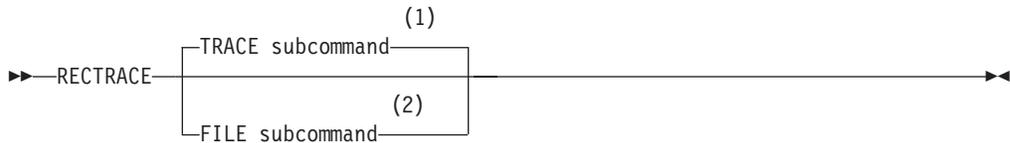
Usage

Use the RECTRACE command to generate Record traces and file data sets.

Usage notes

- The command can be used once in a job step.
- However, it can be used in the same job step with commands of the other report sets.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax of the RECTRACE command



Notes:

- 1 You can specify TRACE up to 5 times.
- 2 FILE can be specified only once.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting Record trace functions

The following is a sample of the JCL required to produce Record traces. For details about the JCL and DD statements used, see Chapter 3, “The OMEGAMON XE for DB2 PE command stream,” on page 11.

```
//          PEGMAIN EXEC PGM=FPECMMAIN
//* FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
//DPMLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
//* FOLLOWING ARE REPORT SET DDNAMES
//RTTRCDD1 DD SYSOUT=*
//RTFILDD1 DD DSN=MYID.FPELIB.RTFIL.DATA,DISP=OLD
//RTWORK DD DSN=MYID.FPELIB.RT.WORKDD,DISP=OLD
//* FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *
RECTRACE
TRACE
FILE
EXEC
```

Figure 7. Sample JCL for requesting Record trace functions

The OMEGAMON XE for DB2 PE command language shown in this section is not appropriate in all circumstances. You must modify it to meet your requirements.

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See the individual DD statement descriptions for more information.

Note: The OMEGAMON XE for DB2 PE command stream is only processed if EXEC is included as the last command. Otherwise, OMEGAMON XE for DB2 PE only checks the syntax.

RECTRACE command with TRACE subcommand

This section describes the RECTRACE command with the TRACE subcommand.

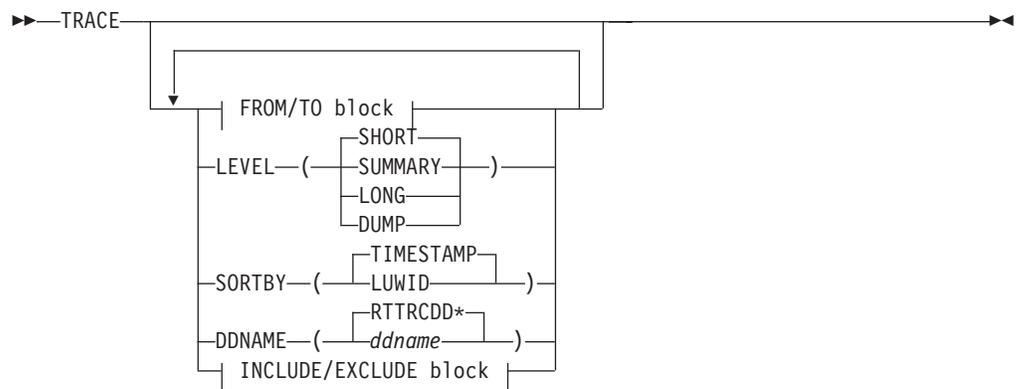
Usage

Use the TRACE subcommand to produce traces with an entry for every IFCID available.

Usage notes

- Up to five traces can be requested in a job step.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the trace by date and time.

For details, see “FROM/TO subcommand options” on page 30.

LEVEL

Selects the amount of detail in the trace. You can specify one of the following keywords:

SUMMARY

Produces a trace that lists all user-selected records in the input data set.

SHORT

Produces a trace that includes all user-selected records that are used in other OMEGAMON XE for DB2 PE reports. This is the default.

LONG

Produces a trace that includes all user-selected records. Serviceability data is also reported.

DUMP

Produces a trace that presents all user-selected records in hexadecimal dump format. The entire record is dumped.

SORTBY

Sorts the events within each location. You can specify either of the following:

LUID

Records are sorted by timestamp within thread within location.

TIMESTAMP

Records are sorted by timestamp within location. **TIMESTAMP** is the default.

One entry of **SORTBY** can be specified for each **TRACE** subcommand.

DDNAME

Specifies the data set where the trace is written.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “**INCLUDE** and **EXCLUDE** subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “**OMEGAMON XE for DB2 PE identifiers**,” on page 3.

Examples

For an example refer to “**Example using RECTRACE with TRACE and FILE subcommands**” on page 109

RECTRACE command with FILE subcommand

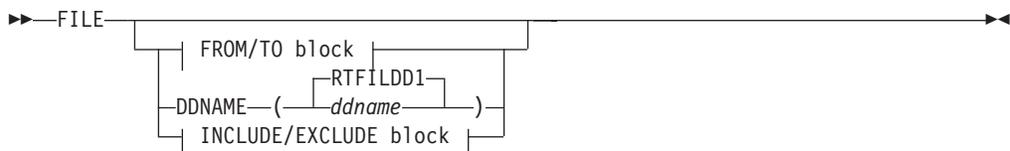
This section describes the **RECTRACE** command with the **FILE** subcommand.

Usage

Use the **FILE** subcommand to format DB2 data and store it in sequential data sets suitable for use by the DB2 load utility. The records can be placed in DB2 tables and you can produce reports by using a reporting facility such as Query Management Facility (QMF).

Usage notes

- **FILE** is used to format Record Trace records from IFCIDs 22, 63, 96, and 125.
- Only one file data set can be generated in a job step.

Syntax of the FILE subcommand

Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the data set by date and time.

For details, see “FROM/TO subcommand options” on page 30.

DDNAME

Specifies the ddname where the file data set is written. The default ddname is RTFILDD1.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using RECTRACE with TRACE and FILE subcommands

This example requests a Record trace that:

- Includes records generated between 08:00 and 08:10 on each day included
- Includes records with PRMAUTH:
 - UID0001
 - UID0002
 - UID0003
- Excludes records with PLANNAME:
 - PLIT2A01
 - PLIT2A02
- Generates a data set on ddname RTFILDD1 by default. The data set contains detailed information about the IFCIDs generated by the primary authorization ID SYSADM.

Note: If the EXEC statement is omitted, no trace is produced. The syntax of your command stream is checked and is written to the DPMLOG data set together with any error, warning, or information messages produced. All statements after the EXEC are ignored.

```
⋮  
RECTRACE  
  TRACE  
    FROM  (,08:00:00)  
    TO    (,08:10:00)  
    INCLUDE (PRMAUTH(UID0001 UID0002 UID0003))  
    EXCLUDE (PLANNAME(PLIT2A01 PLIT2A02))  
  FILE  
    INCLUDE (PRMAUTH(SYSADM))  
  ⋮  
EXEC
```

SQLACTIVITY command

This section provides an overview of the SQLACTIVITY command.

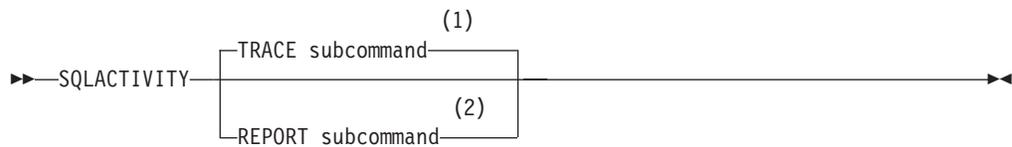
Usage

Use the SQLACTIVITY command to reduce data, and to generate reports and traces.

Usage notes

- This command can be used once in a job step, it can be used in the same job step with other report set commands.
- For migration purposes you can use the SQLACTIVITY command to find information that was previously contained in the following report subjects for ACCOUNTING report types of the OMEGAMON Historical Reporter:
 - SCAN_ACTIVITY
 - SORT_ACTIVITY
 - SUMMARY
 - DETAIL
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax of the SQLACTIVITY command



Notes:

- 1 You can specify TRACE up to 5 times.
- 2 You can specify REPORT up to 5 times.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting SQL Activity functions

The following figure is a sample of the JCL required to produce SQL Activity reports and traces. A description of the DD statements follows the sample.

```

//          PEMAIN EXEC PGM=FPECMMAIN
// * FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
//DPMLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//SYSPRMD DD SYSOUT=*
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
// * FOLLOWING ARE REPORT SET DDNAMES
//SQTRCDD1 DD SYSOUT=*
//SQRPTDD DD SYSOUT=*
//SQLWORK DD DSN=MYID.FPELIB.SQL.WORKDD,DISP=OLD
// * FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *
:
:
SQLACTIVITY
TRACE
REPORT
:
:
EXEC

```

Figure 8. Sample JCL for requesting SQL Activity functions

The OMEGAMON XE for DB2 PE command language shown in this section is not appropriate in all circumstances. You must modify it to meet your requirements.

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE.

Note:

1. There is an advantage in omitting DPMOUTDD from your JCL. For more information, see “DPMOUTDD statement” on page 16.
2. The OMEGAMON XE for DB2 PE command stream is only processed if EXEC is included as the last command. Otherwise, no report is generated. OMEGAMON XE for DB2 PE checks the syntax of your command stream and writes it, together with any information, warning, or error messages generated.

SQLACTIVITY command with REPORT subcommand

This section describes the SQLACTIVITY command with the REPORT subcommand.

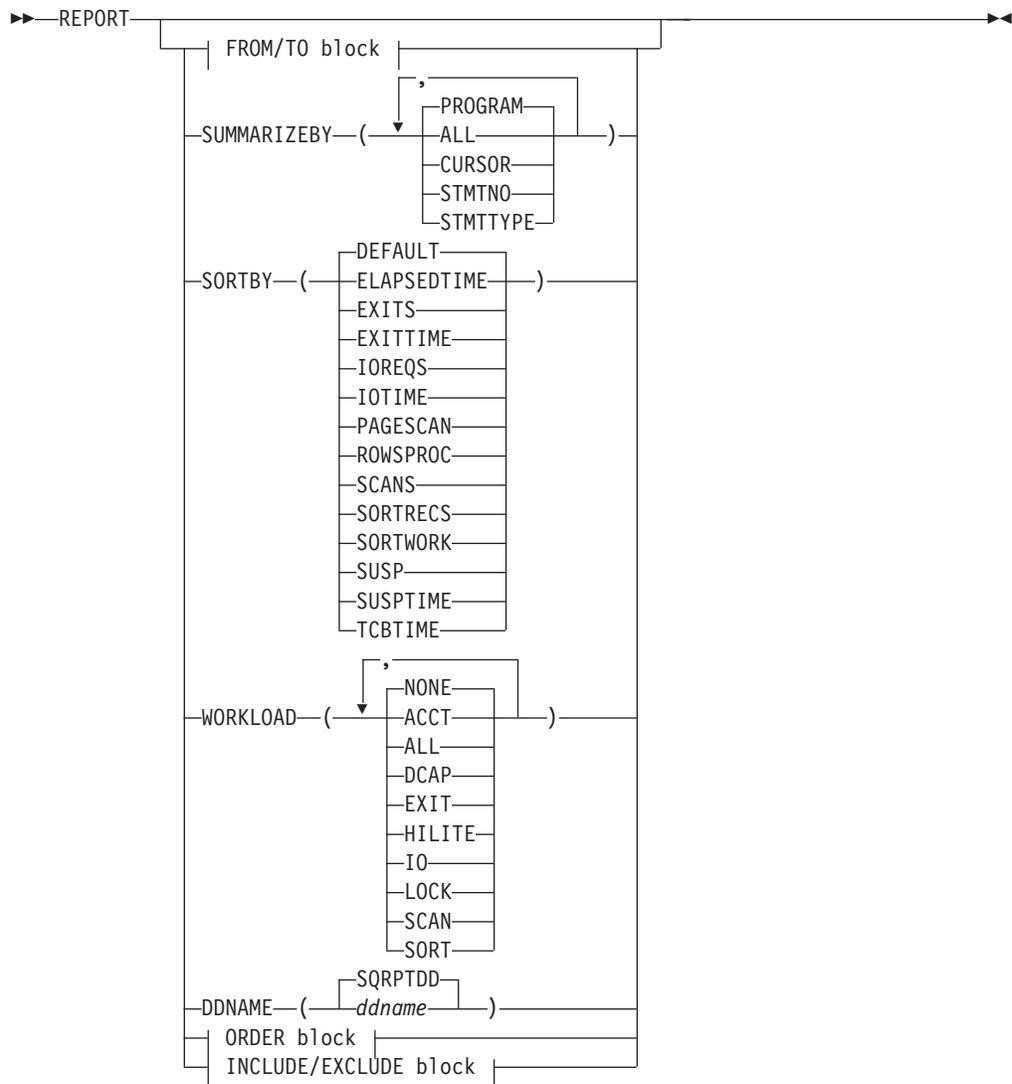
Usage

Use the REPORT subcommand to generate reports from records.

Usage notes

- Up to five REPORT subcommands can be specified within each SQLACTIVITY command.

Syntax of the REPORT subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reporting process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

SUMMARIZEBY

Selects the SQL events to be summarized. You can specify one entry of SUMMARIZEBY for each REPORT subcommand. The following events can be specified:

PROGRAM

This is the default.

CURSOR

STMTNO
Statement number
STMTTYPE
Statement type
ALL
All of the above

For more information about summarization, see the *Reporting User's Guide*.

SORTBY

Sorts the SQL events within each summary level within each thread. You can specify one entry of SORTBY for each REPORT subcommand. One of the following events can be specified:

DEFAULT

The default sort sequence depends on the summary level specified.

ELAPSED TIME

Average elapsed time

EXITS

Number of exits

EXIT TIME

Elapsed time for each exit

I/O REQUESTS

I/O requests

I/O TIME

Elapsed time for each I/O request

PAGESCAN

Pages scanned

ROWS PROCESSED

Rows processed

SCANS

Number of scans

SORT RECORDS

Records sorted

SORTWORK

Workfiles sorted

SUSPENSIONS

Lock suspensions

SUSPENSION TIME

Elapsed time for each lock suspension

TCB TIME

Average TCB time

For more information about sorting and default, see the *Reporting User's Guide*.

WORKLOAD

Displays workload detail for each event. The following detail can be reported:

NONE

No workload activity. This is the default.

ACCT

Accounting

ALL

All workload activity

Including MINIBIND if IFCID 22 is included in the input.

Including UDF activity if IFCID 324 is included in the input.

Note that the more workload detail you request, the more time OMEGAMON XE for DB2 PE requires for processing your request. It is recommended that you do not specify WORKLOAD(ALL) with a large amount of input data unless absolutely necessary.

DCAP

Data capture activity

EXIT

Exit activity

HILITE

Workload highlights

IO I/O activity

LOCK

Lock suspension and page and row locking activity

SCAN

Scan activity, RID list activity, and query parallelism activity

SORT

Sort activity

Note: When IFCID 3 is included in the input, Accounting Trace activity is automatically included as part of the workload detail.

For more information about workload detail, see the *Reporting User's Guide* and the *Report Reference*.

DDNAME

Specifies the data set where the report is written.

ORDER

Specifies the OMEGAMON XE for DB2 PE identifiers and their sequence for sorting the report and, in summary reports, which identifiers are used for aggregation.

For details, see "ORDER subcommand option" on page 46 and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see "INCLUDE and EXCLUDE subcommand options" on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

Example using REPORT with FROM, TO, and EXCLUDE options

This example specifies the following:

- Records are used with the time and date range of 10:00 a.m. on 18 March 2002 to noon on 19 March 2002
- Data is excluded that is associated with the following locations:
 - LOCN10
 - LOCN12
 - LOCN15
 - LOCN20

```
⋮  
REPORT  
FROM (03/18/02,10:00:00.00)  
TO (03/19/02,12:00:00.00)
```

```
EXCLUDE (LOCATION(LOCN10 LOCN12 LOCN15 LOCN20))  
⋮
```

SQLACTIVITY command with TRACE subcommand

This section describes the SQLACTIVITY command with the TRACE subcommand.

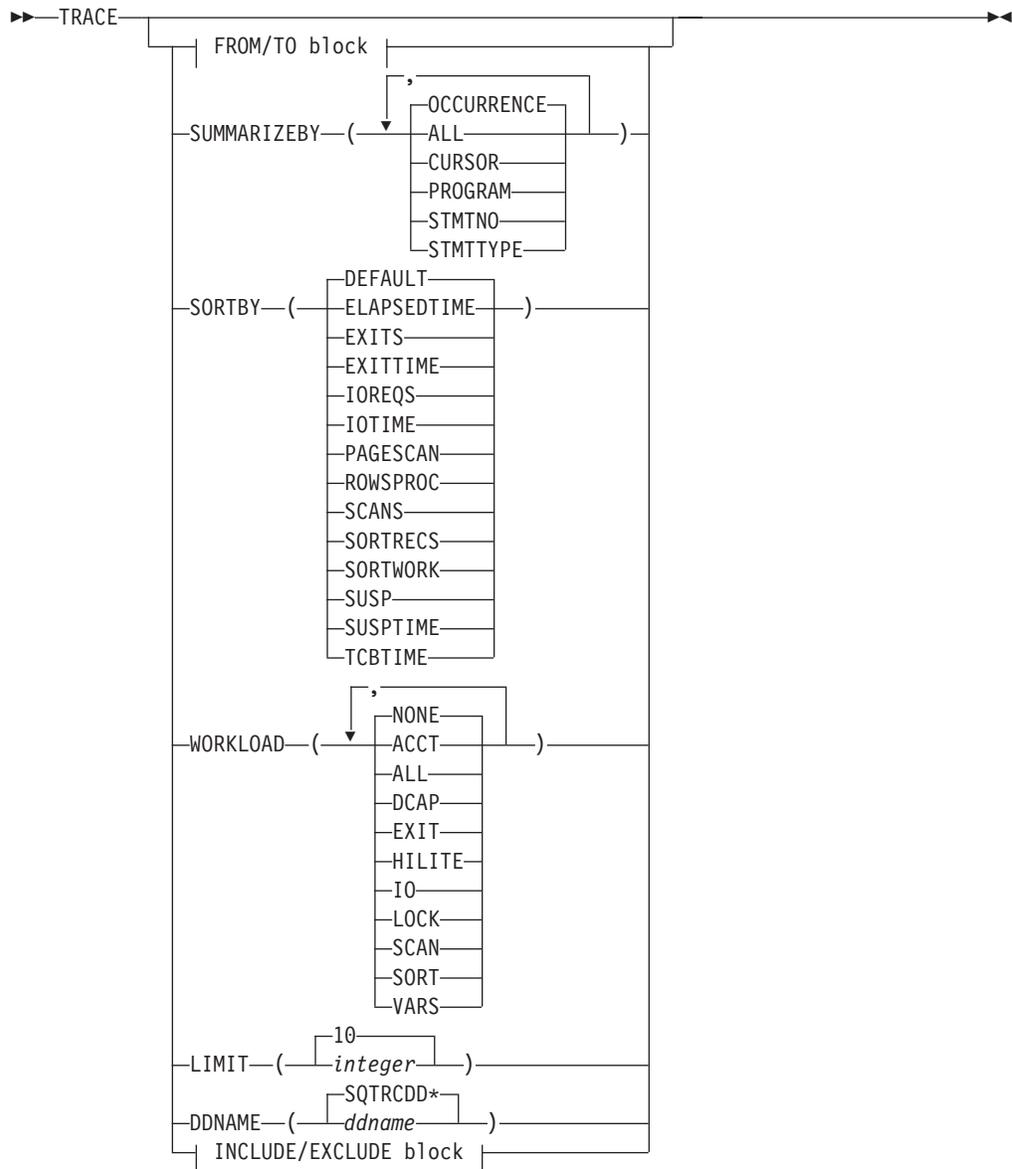
Usage

Use the TRACE subcommand to produce traces with an entry for every DB2 SQL event.

Usage notes

- Up to five traces can be requested in a job step.
- For more information about including or excluding IFCIDs, see the *Reporting User's Guide*.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the trace by date and time.

For details, see “FROM/TO subcommand options” on page 30.

SUMMARIZEBY

Selects the SQL events to be summarized. You can specify one entry of SUMMARIZEBY for each TRACE subcommand. The following events can be specified:

OCCURRENCE

SQL statement occurrence. This is the default.

PROGRAM**CURSORS****STMTNO**

Statement number

STMTTYPE

Statement type

ALL

All of the above

For more information about summarization, see the *Reporting User's Guide*.

SORTBY

Sorts the SQL events within each summary level within each thread. You can specify one entry of SORTBY for each TRACE subcommand. One of the following events can be specified:

DEFAULT

The default sort sequence depends on the summary level specified.

ELAPSEDTIME

Average elapsed time

EXITS

Number of exits

EXITTIME

Elapsed time for each exit

IOREQS

I/O requests

IOTIME

Elapsed time for each I/O request

PAGESCAN

Pages scanned

ROWSPROC

Rows processed

SCANS

Number of scans

SORTRECS

Records sorted

SORTWORK

Workfiles sorted

SUSP

Lock suspensions

SUSPTIME

Elapsed time for each lock suspension

TCBTIME

Average TCB time

For more information about sorting and default, see the *Report Reference*.

WORKLOAD

Displays the workload detail for each event. The following detail can be reported:

ACCT

Accounting

ALL

All workload activity

Including MINIBIND if IFCID 22 is included in the input.

Including UDF activity if IFCID 324 is included in the input.

Note that the more workload detail you request, the more time OMEGAMON XE for DB2 PE requires for processing your request. It is recommended that you do not specify WORKLOAD(ALL) with a large amount of input data unless absolutely necessary.

DCAP

Data capture activity

EXIT

Exit activity

HILITE

Workload highlights

IO I/O activity

LOCK

Lock suspension and page and row locking activity

NONE

No workload activity. This is the default.

SCAN

Scan activity, RID list activity, and query parallelism activity.

SORT

Sort activity

VARs

Show host variables data, if host variables are used by the SQL statements.

If option WORKLOAD(VARS) is specified, option SUMMARIZEDBY(OCCURRENCE) is required (the default).

Note: When IFCID 3 is included in the input, Accounting Trace activity is automatically included as part of the workload detail.

For more information about workload detail, see the *Reporting User's Guide* and the *Report Reference*.

DDNAME

Specifies the data set where the trace is written.

LIMIT

Sets the maximum number of threads processed by TRACE. The range is 1 to 99 999. If, for example, 3 is specified for LIMIT, no more than three threads are reported. A different limit can be set for each of the five possible traces. The default is 10.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see "INCLUDE and EXCLUDE subcommand options" on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

In addition to the common identifiers used with these options, the SQLACTIVITY TRACE command also has an SQLCODE identifier. This identifier can be used to include statements that completed, for example, with a specific error. The following are usage examples of the SQLCODE identifier.

SQLACTIVITY TRACE EXCLUDE(SQLCODE(0))

This excludes any SQL statement that completed without warnings or errors.

SQLACTIVITY TRACE INCLUDE(SQLCODE(R(-251,-203)))

This includes any SQL statements that ended with SQL error codes in the range -251 through -203.

SQLACTIVITY TRACE INCLUDE(SQLCODE(-805,-204,495))

This includes only SQL statements that ended with the specific error codes -805, -204, and 495.

SQLACTIVITY TRACE INCLUDE(SQLCODE(LT(0)))

This includes only SQL statements that ended with an error (-n).

SQLACTIVITY TRACE INCLUDE(SQLCODE(GT(0)))

This includes only SQL statements that ended with a warning.

Example using TRACE with SUMMARIZEBY, SORTBY, and LIMIT options

```
⋮  
TRACE  
  FROM      (,08:00:00.00)  
  TO        (,08:10:00.00)  
  SUMMARIZEBY (STMTNO)  
  SORTBY    (ELAPSED TIME)  
  LIMIT     (20)  
  INCLUDE   (PRIMAUTH(UID0001 UID0003 UID0005)  
            PLANNAME(PLIT2A01,PLIT2A02))  
⋮
```

It is summarized by statement number and sorted by elapsed time.

It reports the ten-minute period from 8:00 a.m. to 8:10 a.m. and includes only data that contains any of the following primary authorization IDs:

- UID0001
- UID0003
- UID0005

Using any of the following plan names:

- PLIT2A01
- PLIT2A02

LIMIT has set the maximum number of threads processed to 20.

The trace is written to the data set defined by the default ddname SQTRCDD1.

STATISTICS command

This section provides an overview of the STATISTICS command.

Usage

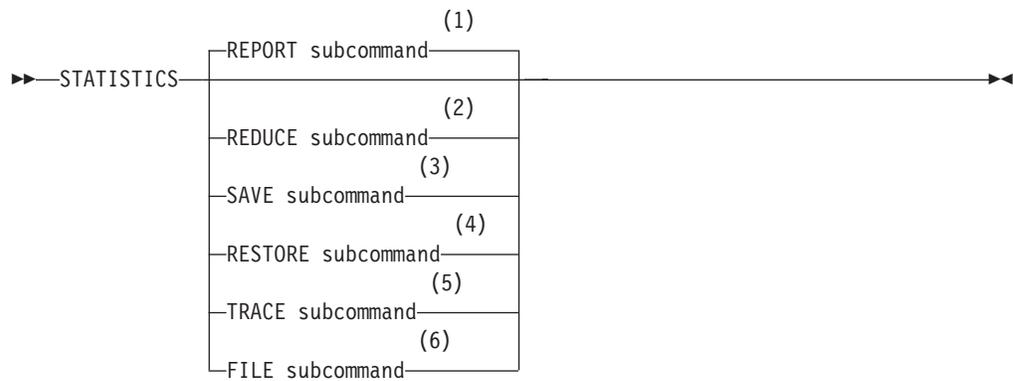
Use the STATISTICS command to generate Statistics reports, traces, and file data sets. You can also use it to reduce, save, and restore data.

Usage notes

- This command can be used once in a job step.
- However, it can be used in the same job step with commands of the other report sets.

- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax of the STATISTICS command



Notes:

- 1 You can specify REPORT up to 5 times.
- 2 You can specify REDUCE only once. If specified, you must also specify REPORT or SAVE at least once.
- 3 You can specify SAVE only once.
- 4 You can specify RESTORE only once.
- 5 You can specify TRACE up to 5 times.
- 6 You can specify FILE only once.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting Statistics functions

Figure 9 on page 121 is a sample of the JCL required to produce statistic reports and traces. See Chapter 5, “DD statements,” on page 15 for descriptions of the DD statements.

```

//          PEMAIN EXEC PGM=FPECMMAIN
//* FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
//DPMLLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//EXCPTDD DD DSN=MYID.EXCEPT.THRESH,DISP=OLD
//EXTRCDD1 DD SYSOUT=*
//EXFILDD1 DD DSN=MYID.EXCEPT.LOGFILE,DISP=OLD
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//JSSRSDD DD DSN=MYID.FPELIB.JSSRS.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
//* FOLLOWING ARE REPORT SET DDNAMES
//STRPTDD DD SYSOUT=*
//STRCDD1 DD SYSOUT=*
//STSAVDD DD DSN=MYID.FPELIB.STSAV.DATA,DISP=OLD
//STRSTDD DD DSN=MYID.FPELIB.STRST.DATA,DISP=SHR
//STFILDD1 DD DSN=MYID.FPELIB.STFIL.DATA,DISP=SHR
//STWORK DD DSN=MYID.FPELIB.STA.WORKDD,DISP=OLD
//* FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *
:
:
STATISTICS
  REDUCE
  TRACE
  REPORT
  RESTORE
  SAVE
  FILE
:
:
EXEC

```

Figure 9. Sample JCL for requesting Statistics functions

The OMEGAMON XE for DB2 PE command language shown in this section is not appropriate in all circumstances. You must modify it to meet your requirements.

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See the individual DD statement descriptions for more information.

Note:

1. There is a performance advantage in omitting DPMOUTDD from your JCL. For more information, see “DPMOUTDD statement” on page 16.
2. The OMEGAMON XE for DB2 PE command stream is only processed if EXEC is included as the last command. Otherwise, OMEGAMON XE for DB2 PE only checks the syntax.

STATISTICS command with REPORT subcommand

This section describes the STATISTICS command with the REPORT subcommand.

Usage

Use the REPORT subcommand to present statistics interval records. For the definition of the interval records see *Report Reference*.

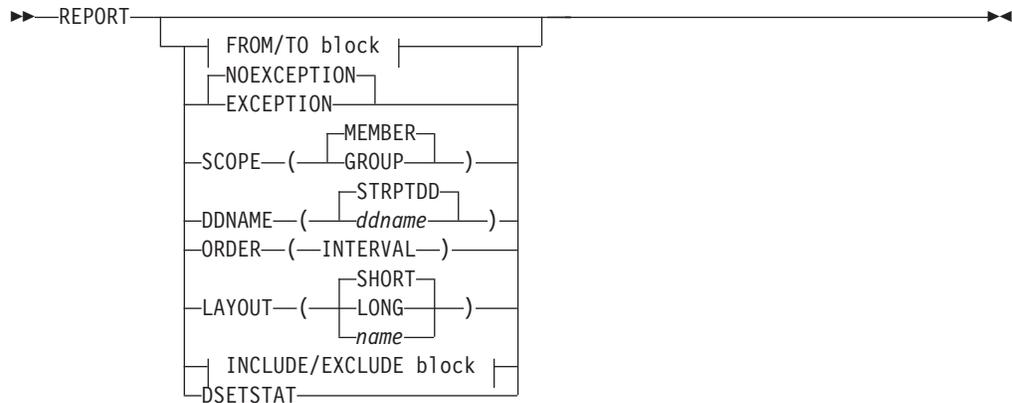
The duration and alignment of the interval records are specified by the REDUCE subcommand.

The interval records can be either presented individually (if the ORDER(INTERVAL) subcommand option is specified) or aggregated into a new interval record for the duration that is specified with the FROM and TO options in the REPORT subcommand.

Usage notes

- Up to five REPORT subcommands can be specified within a STATISTICS command.

Syntax of the REPORT subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of interval records included in the reporting process by date and time.

Combinations of FROM date and time, and TO date and time can be specified. Interval records are included from the first with an ending timestamp greater than or equal to the FROM date and time, to the last with an ending timestamp less than the TO date and time.

If you do not specify a date and time, FROM/TO defaults to the dates and times specified in the GLOBAL command. If FROM/TO is not specified in GLOBAL, all records are included in the report.

The specified FROM/TO dates and times are printed on the report. If FROM/TO is not specified in REPORT or GLOBAL, NOT SPECIFIED is printed on the report. If only the FROM date and time or only the TO date and time has been specified, NOT SPECIFIED is printed for the unspecified value.

Note that GLOBAL FROM and TO filters individual DB2 Statistics records whereas REPORT FROM and TO filters interval records. For example, with a REPORT FROM time of 14:15, an interval record starting at 14:00 and ending at 14:30 would appear on thereport (based on the end interval record time). However, a GLOBAL FROM time of 14:15 would filter the DB2 Statistics records at 14:00 and the interval record (14:00 to 14:30) might not be created (subject to the BOUNDARY specification).

You can specify a time adjustment for a DB2 location using the TIMEZONE option of the GLOBAL command. The time adjustment is applied to the record timestamp before FROM/TO processing.

If you are reducing data, the FROM/TO times specified in REDUCE can affect the data available for reporting.

Refer to “GLOBAL command” on page 149 for more information.

For details on how to use the FROM/TO option, see “FROM/TO subcommand options” on page 30.

LAYOUT

Specifies the name of a report layout. You can specify one of the supplied layouts or one that you have previously tailored:

SHORT

This is the default.

LONG

This option provides detailed thread-related data. You can also use other functions such as Record Trace and SQL Activity to find detailed DB2 trace data.

Historical Reporter migration layouts

Use the Historical Reporter migration layouts to help you identify OMEGAMON XE for DB2 PE data that was previously shown in the reports of the OMEGAMON Historical Reporter. The migration layouts include:

- “Options for OMEGAMON Historical Reporter Accounting reports” on page 39
- “Options for OMEGAMON Historical Reporter Statistics reports” on page 42

User-defined layouts

You can customize your own report layouts by specifying which blocks of data and which fields within the blocks are included, and their relative order. To adapt the reports according to your requirements, you use user-tailored reporting (UTR). With UTR, you can control the volume, contents, and layout of Statistics traces and reports.

For information about tailoring report layouts, see the *Reporting User's Guide*.

SCOPE

Specifies the scope of the report in a data sharing environment.

MEMBER

Member-scope reporting presents DB2 statistics on a per-member basis without aggregating data-sharing-related counters for the entire data-sharing group.

GROUP

Group-scope reporting presents data-sharing-related counters for an entire data-sharing group.

EXCEPTION

NOEXCEPTION

Specify EXCEPTION if you want to report only those interval records on Statistics reports with at least one field in exception status. Otherwise, a standard report is produced.

If you use this subcommand option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see Figure 9 on page 121.

DDNAME

Specifies the data set where the report is written. You can specify any valid ddname including the default, provided that your JCL contains a DD statement for it. If a DD statement is omitted, it will be dynamically allocated to the SYSOUT message class of the job. The default ddname for report is STRPTDD.

ORDER (INTERVAL)

Specifies that your report contains statistics data for each interval record that satisfies input filters.

For details, see “ORDER subcommand option” on page 46 and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

If you want to report the statistics data for the entire period covered by the interval records, do not specify this option.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

DSETSTAT

Include buffer pool data set statistics.

Example using REDUCE, and REPORT with INCLUDE option

This example specifies the following:

- One interval record is derived. It covers the entire period that the DB2 statistics data is available for.
- The report uses the default SHORT layout.
- Data is included that is only associated with the location in the range of LOCN01 to LOCN05

```
REDUCE
REPORT
  INCLUDE (LOCATION(R(LOCN01 LOCN05)))
```

Example using REDUCE with INTERVAL option, and REPORT with several options

This example specifies the following:

- The input DB2 statistics data is distributed over 60-minute intervals aligned with hour boundaries.
- The interval records between 10:00 a.m. on 18 March 1996 and noon on 19 March 1996 are considered for reporting, subject to the exception criteria.
- Each of the qualifying interval records is reported.
- The report uses the LONG layout.

```
REDUCE
  INTERVAL(60)
REPORT
  FROM (09/18/98,10:00:00.00)
  TO (09/19/98,12:00:00.00)
  LAYOUT (LONG)
  EXCEPTION
  ORDER (INTERVAL)
```

STATISTICS command with REDUCE subcommand

This section describes the STATISTICS command with the REDUCE subcommand.

Usage

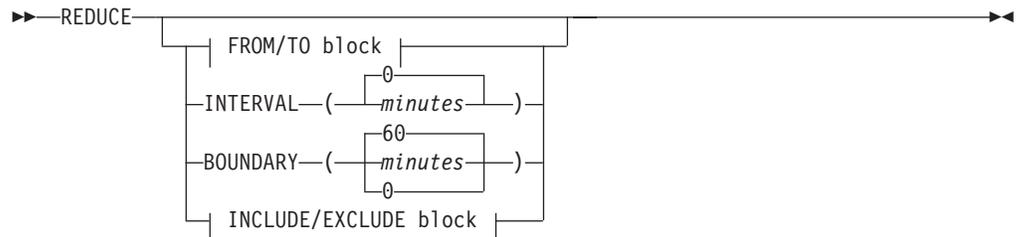
Specify REDUCE to control the following:

- The duration and alignment of statistics interval records. For a definition of the interval record, see *Report Reference*.
- The volume of data to be presented in traces and reports and stored in file and save data sets.

Usage notes

- REDUCE is invoked automatically when you use the REPORT or SAVE subcommand.
- REDUCE can be used once in a STATISTICS command.

Syntax of the REDUCE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, "OMEGAMON XE for DB2 PE subcommand options," on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of DB2 statistics records to process by date and time.

Combinations of FROM date and time and TO date and time can be specified.

Note: DB2 statistics records are included from the last with a timestamp less than the FROM date and time to the last with a timestamp less than the TO date and time. The FROM/TO dates and times specified in REDUCE affect the other functions.

The DB2 statistics records available to REDUCE are limited by the GLOBAL FROM/TO dates and times.

If you do not specify a date and time, FROM/TO defaults to the date and time specified in the GLOBAL command. If dates and times are not specified in the GLOBAL command, all DB2 statistics records are included in the reduction process.

You can specify a time adjustment for a DB2 location using the TIMEZONE option of the GLOBAL command. The time adjustment is applied to the DB2 statistics record timestamp before FROM/TO processing.

Refer to "GLOBAL command" on page 149 for more information.

For details on how to use the FROM/TO option, see “FROM/TO subcommand options” on page 30.

INTERVAL

Defines the duration of statistics interval records. The range is from 0 to 99 999 and is specified in minutes. When the DB2 statistics data is distributed to the interval records, they can be presented in Statistics reports and stored in a save data set.

For example, if INTERVAL(15) is specified, 15-minute intervals are created over the period that the DB2 statistics data is available for and the DB2 data is distributed to these intervals.

INTERVAL(0) specifies that only one interval record is created, starting with the first and ending with the last DB2 Statistics record pair.

If no interval is specified, the interval specified in the GLOBAL command is used. If no interval is specified in GLOBAL, the default is 0.

INTERVAL has an impact on performance. Always use the largest interval that meets your reporting requirements. If interval processing is not required, the default INTERVAL (0) is recommended for optimum performance. For more information about intervals, see “Processing intervals.”

The following example uses INTERVAL to specify an interval of two hours:

```
⋮  
INTERVAL (120)  
⋮
```

BOUNDARY

Controls the alignment of the statistic interval records defined by the INTERVAL option.

For more information on boundaries, see “Processing intervals.”

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Processing intervals

The start time of the first interval processed by REDUCE is influenced by BOUNDARY, INTERVAL, and FROM.

OMEGAMON XE for DB2 PE attempts to reduce all data that falls between FROM and TO dates and times. The first interval processed starts at a time aligned with BOUNDARY, at or before the FROM time. If an interval cannot be aligned with the FROM time, the first properly aligned interval starting before the FROM time is used.

Although there is no restriction on the INTERVAL and BOUNDARY combination, your specification should comply with the following recommendations:

- For intervals of less than 60 (excluding 0), there should be a whole number of intervals in an hour. Choose one of the following values:
 - 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, or 30.

- For intervals of 60 or greater, there should be a whole number of intervals in a day. Choose one of the following values:
 - 60, 120, 180, 240, 360, 480, 720, or 1 440.
- For intervals of one day (1440) or greater, INTERVAL should be a multiple of 1 440.
- Select your interval and boundary so that the first interval starts at the FROM time.

The first interval record starts at a time aligned with BOUNDARY, at or before the FROM time. If an interval record cannot start at the FROM time, it starts before, at the nearest time that satisfies the BOUNDARY specification. Note that an interval record starting before the FROM time only contains data between the FROM time and the end of the interval record. Input data before the FROM time is not processed. An interval record ending after the TO time only contains data between the beginning of the interval record and the TO time. Input data after the TO time is not processed.

Example using REDUCE with FROM, TO, INTERVAL, BOUNDARY, and INCLUDE options

BOUNDARY(60) aligns the start time of interval records at the start of an hour, so the first interval record starts at the FROM time (08:00). Subsequent interval records start every 30 minutes (08:30, 09:00, and 09:30 each day). Only the DB2 statistics for location SYDNEY is processed.

```

:
REDUCE
  FROM      (,08:00)
  TO        (,10:00)
  INTERVAL  (30)
  BOUNDARY  (60)
  INCLUDE   (LOCATION(SYDNEY))
:

```

Example using REDUCE with options, together with SAVE

The following defaults are applied in this example:

- For FROM, all dates and a time of 00:00:00.00.
- For TO, all dates and a time of 23:59:59.99.

BOUNDARY(60) aligns the start time of interval records at the start of an hour, so the first interval record starts at the FROM time (00:00). Subsequent interval records cover 1 440 minutes or one day; an interval starts at 00:00 each day.

The interval records are saved in a data set for subsequent use.

```

:
REDUCE
  INTERVAL (1440)
  BOUNDARY (60)

```

```
SAVE
⋮
```

Example using REDUCE with options, together with REPORT with options

BOUNDARY(60) aligns the start time of interval records at the start of an hour, so the first interval record starts at the first properly aligned time before the FROM time, which is 08:00. Subsequent interval records start every two hours (10:00 and 12:00).

A report presenting the statistics data for each of the 2-hour periods is produced.

```
⋮
REDUCE
  FROM    (,08:15)
  TO      (,13:00)
  INTERVAL (120)
  BOUNDARY (60)
REPORT
  ORDER   (INTERVAL)
⋮
```

STATISTICS command with SAVE subcommand

This section describes the STATISTICS command with the SAVE subcommand.

Usage

Use the SAVE subcommand (without CONVERT option) to produce a VSAM data set containing interval records. After the data has been saved, you can:

- Convert the save files to sequential data sets by using the save-file utility and load them into DB2 for subsequent use.
- Restore and combine it with newly reduced data to produce long-term reports.
- Restore it and use it in later reporting.

Use the SAVE subcommand with CONVERT option to produce a sequential data set containing reduced data in records. After the data has been processed, you can:

- Load it to DB2 for subsequent use.

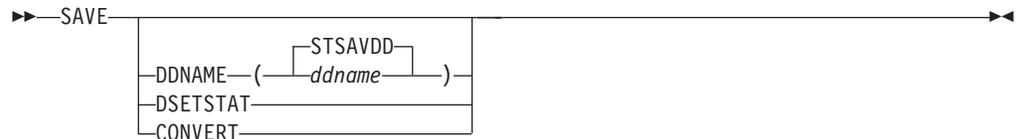
Note: In contrast to the other usage, the conversion by the save-file utility is accomplished directly. VSAM data is not being produced, thus it also not possible to restore it.

You can also use the converted SAVE data sets to generate CSV (comma-separated value) input-data. This CSV data can then be transferred to workstations and imported into spreadsheets to improve DB2 performance analysis using graphical representations or pivot tables. For more information refer to *Reporting User's Guide*.

Usage notes

- VSAM data sets cannot be concatenated.

Syntax of the SAVE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

DDNAME

Specifies the ddname where the save data is written. The default ddname is STSAVDD. You can specify any valid ddname provided that your JCL contains a valid DD statement for it.

Without CONVERT option:

When using SAVE, your JCL can also include a valid DD statement for JSSRSDD, so that the related job summary information can be saved.

Note: Do not specify DUMMY in the JCL for either STSAVDD or JSSRSDD. When you are not using SAVE, omit these ddnames from your JCL.

The VSAM data set defined by the default ddname must already exist when you run OMEGAMON XE for DB2 PE. See the *Report Reference* for a description of VSAM data sets. Either specify an existing data set from a previous OMEGAMON XE for DB2 PE run (when restoring data), or specify a new data set allocated using the IDCAMS DEFINE CLUSTER function.

Note: If ddname is assigned to a non-VSAM file, you receive an error message and the job terminates.

With CONVERT option:

The ddname needs to be assigned to a physical sequential data set. This data set can be used for a subsequent load to Accounting SAVE tables.

Note: If ddname is assigned to a nonsequential data set, you receive an error message and the job terminates.

DSETSTAT

Include buffer pool data set statistics.

CONVERT

Specifies that converted reduced data is written directly to a sequential data set.

It is recommended to specify this option if you want to create high amount of loadable reduced data in a sequential data set. It avoids a temporary VSAM data set to be used as a SAVE data set. This options causes OMEGAMON XE for DB2 PE to write converted reduced data directly to a sequential data set in a single step. The resulting output is loadable to the PDB tables. The user may experience performance improvements compared to the default path with a separate SAVE step and the subsequent convert of saved data by the save-file utility.

Note: Not every big trace input results in big reduced data and small trace input in small reduced data. This option becomes effective when the reduction results in high amount of reduced data. For example, it depends on the amount of different criteria that results due to your filtering. It is possible that you experience more performance improvement with a small input trace that results in many different criteria than with a large trace input with only a few different criteria.

This option can generally be used as it is neutral to the performance when processing only a few reduced records, but becomes effective when processing a high amount of reduced data. However, do not use this option if you want to RESTORE and REPORT saved data.

Example using REDUCE, and SAVE with DDNAME and CONVERT option

This example specifies that one interval record covering the entire input statistics data (because INTERVAL(0) is issued) is written to the external data set defined by the ddname SAVDSNDD. In this case it has to be a sequential data set.

```

:
REDUCE
SAVE
  DDNAME (SAVDSNDD)
  CONVERT
:

```

STATISTICS command with RESTORE subcommand

This section describes the STATISTICS command with the RESTORE subcommand.

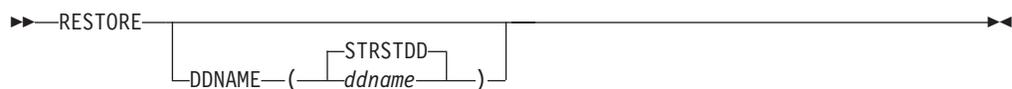
Usage

Use the RESTORE subcommand to reload previously saved data for additional processing.

Usage notes

- After the data is restored, you can produce reports from the restored data alone, or from the restored data combined with newly reduced data.
- Saved data can be restored as often as required.

Syntax of the RESTORE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, "OMEGAMON XE for DB2 PE subcommand options," on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

DDNAME

Specifies the ddname used to reload a previously saved file for additional use.

The ddname specifies the data set being restored. The default ddname is STRSTDD. You can specify any valid ddname, provided your JCL contains a valid DD statement for it.

When using RESTORE, your JCL can also include a valid DD statement for JSSRSDD, so that the related job summary information can also be restored.

Example using RESTORE with DDNAME option

This example specifies that the previously saved, reduced data is read from the data set defined by the ddname RESDSNDD.

Note:

1. Do not specify DUMMY in the JCL for either STRSTDD or JSSRSDD. When you are not using RESTORE, omit these ddnames from your JCL.
2. If you use the RESTORE and REDUCE subcommands in the same job stream, the INTERVAL and BOUNDARY options specified in REDUCE should match the INTERVAL and BOUNDARY options that were used to reduce the data being restored. If these values are different, the interval and boundary from the restored data is used.
3. Data from previous versions of DB2 PM cannot be restored until it has been changed to the current OMEGAMON XE for DB2 PE format by using the migrate function of the save-file utility. For information about migrating data, see the *Report Reference*.

```

:
RESTORE
  DDNAME (RESDSNDD)
:

```

STATISTICS command with TRACE subcommand

This section describes the STATISTICS command with the TRACE subcommand.

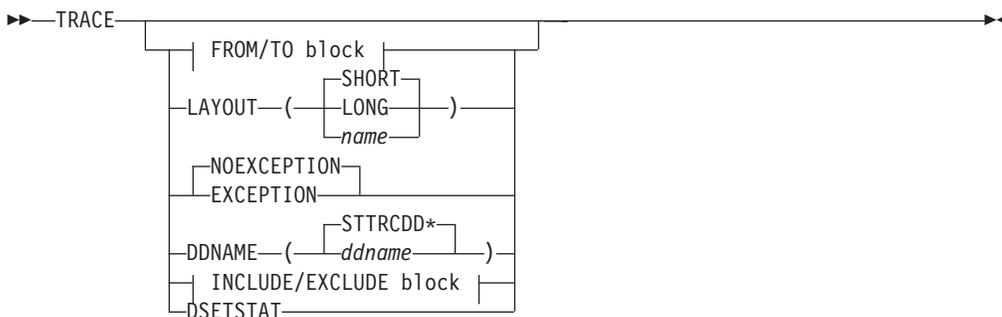
Usage

Use the TRACE subcommand to present statistics delta records.

Usage notes

- Up to five traces can be requested in a job step.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of delta records included in the trace by date and time.

Combinations of FROM date and time, and TO date and time can be specified. Delta records are included from the first with an ending timestamp greater than or equal to the FROM date and time, to the last with an ending timestamp less than the TO date and time.

If you do not specify a date and time, FROM/TO defaults to the dates and times specified in the GLOBAL command. If FROM/TO is not specified in GLOBAL, all records are included in the trace.

The specified FROM/TO dates and times are printed on the trace. If FROM/TO is not specified in TRACE or GLOBAL, NOT SPECIFIED is printed on the trace. If only the FROM date and time or only the TO date and time has been specified, NOT SPECIFIED is printed for the unspecified value.

Note that GLOBAL FROM and TO filters individual DB2 Statistics records whereas TRACE FROM and TO filters delta records. For example, with a TRACE FROM time of 14:15, a delta record starting at 14:00 and ending at 14:30 would appear on the trace (based on the end delta record time). However, a GLOBAL FROM time of 14:15 would filter the DB2 Statistics records at 14:00 and the delta record (14:00 to 14:30) might not be created (subject to the BOUNDARY specification).

You can specify a time adjustment for a DB2 location using the TIMEZONE option of the GLOBAL command. The time adjustment is applied to the record timestamp before FROM/TO processing.

If you are reducing data, the FROM/TO times specified in REDUCE can affect the data available for tracing.

Refer to “GLOBAL command” on page 149 for more information.

For details on how to use the FROM/TO option, see “FROM/TO subcommand options” on page 30.

LAYOUT

Specifies the name of a report layout. You can specify one of the supplied layouts or one that you have previously tailored:

SHORT

This is the default.

LONG

This option provides detailed thread-related data. You can also use other functions such as Record Trace and SQL Activity to find detailed DB2 trace data.

Historical Reporter migration layouts

Use the Historical Reporter migration layouts to help you identify OMEGAMON XE for DB2 PE data that was previously shown in the reports of the OMEGAMON Historical Reporter. The migration layouts include:

- “Options for OMEGAMON Historical Reporter Accounting reports” on page 39

- “Options for OMEGAMON Historical Reporter Statistics reports” on page 42

User-defined layouts

You can customize your own report layouts by specifying which blocks of data and which fields within the blocks are included, and their relative order. To adapt the reports according to your requirements, you use user-tailored reporting (UTR). With UTR, you can control the volume, contents, and layout of Statistics traces and reports.

For information about tailoring report layouts, see the *Reporting User's Guide*.

EXCEPTION

NOEXCEPTION

Specify EXCEPTION if you want to report only those delta records on Statistics reports with at least one field in exception status. Otherwise, a standard report is produced.

If you use this subcommand option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see Figure 3 on page 55.

DDNAME

Specifies the data set where the trace is written. The default ddname for the first trace is STTRCDD1. The default ddnames for the second to fifth traces are STTRCDD2 through STTRCDD5.

You can specify a different ddname by using the DDNAME subcommand option in the TRACE subcommand. In this case, your JCL must contain a valid DD statement for the ddname you specify.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

DSETSTAT

Include buffer pool data set statistics.

Example using TRACE with DSETSTAT, LAYOUT, and DDNAME

This example:

- Includes the buffer pool data set statistics.
- Specifies the default (short) report format.
- Writes the data to the STTRCDD1 dataset.

```

:
:
TRACE
DSETSTAT
LAYOUT(SHORT)
DDNAME(STTRCDD1)
:
:

```

STATISTICS command with FILE subcommand

This section describes the STATISTICS command with the FILE subcommand.

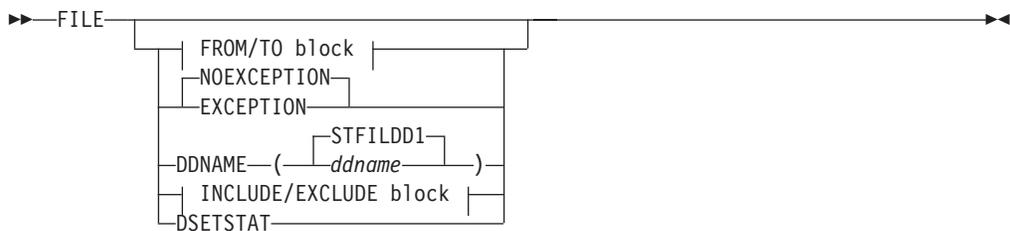
Usage

- Use the FILE subcommand to format unreduced DB2 data and store it in sequential data sets suitable for use by the DB2 load utility. The records can be placed in DB2 tables, and you can produce reports by using a reporting facility such as Query Management Facility™ (QMF).
- You can also use the FILE data sets to generate CSV (comma-separated value) input-data. This CSV data can then be transferred to workstations and imported into spreadsheets to improve DB2 performance analysis using graphical representations or pivot tables. For more information refer to *Reporting User's Guide*.

Usage notes

- The data is stored in delta records in a sequential data set.

Syntax of the FILE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of delta records included in the data set by date and time.

Combinations of FROM date and time, and TO date and time can be specified. Delta records are included from the first with an ending timestamp greater than or equal to the FROM date and time, to the last with an ending timestamp less than the TO date and time.

Any FROM/TO dates and times specified in GLOBAL can restrict the range of delta records included in the file data set. GLOBAL FROM/TO filters the raw DB2 statistics records used to make up the delta records. Therefore, all records with begin or end times outside the GLOBAL dates and times are discarded. If the date and time is not specified in GLOBAL, all records in the input data are available.

You can specify a time adjustment for a DB2 location using the TIMEZONE option of the GLOBAL command. The time adjustment is applied to the record timestamp before FROM/TO processing.

If you are reducing data, the FROM/TO times specified in REDUCE can affect the data available for filing.

Refer to “GLOBAL command” on page 149 for more information.

For details on how to use the FROM/TO option, see “FROM/TO subcommand options” on page 30.

EXCEPTION**NOEXCEPTION**

Specifies EXCEPTION to include only those file entries containing fields with values outside user-specified limits. Otherwise, all records are included.

If you use this subcommand option, your JCL must contain a valid DD definition for the ddname EXCPTDD. For more information about required ddnames, see “EXCPTDD statement” on page 17.

DDNAME

Specifies the ddname where the file data set is written. The default ddname is STFILDD1.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

DSETSTAT

Includes buffer pool data set statistics.

Example using FILE with DSETSTAT, and DDNAME

This example:

- Includes buffer tool data set statistics.
- Writes the data to the STFILDD1 dataset.

```

.
.
FILE
  DSETSTAT
  DDNAME(STFILDD1)
.
.

```

SYSPARMS command

This section provides an overview of the SYSPARMS command.

Usage

Use the SYSPARMS command to generate System Parameter reports. This command replaces the SYSPRMDD card. You can still use the SYSPRMDD in existing jobs, but not concurrently with the SYSPARMS command.

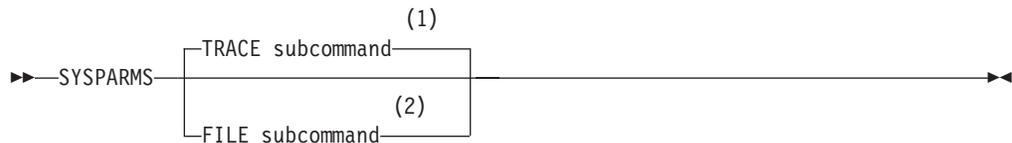
Usage notes

- With the SYSPARMS command, you can use the FILE subcommand to store system parameter information in FILE data sets, which can be loaded into the OMEGAMON XE for DB2 PE Performance Database.
- With the exception of the INCLUDE and EXCLUDE subcommand options with GROUP and MEMBER, the GLOBAL command does not affect the System Parameters reports. For example, if FROM or TO is used to limit the record timestamps, System Parameters reports are generated even when the timestamps of the records are outside these times. When INCLUDE or EXCLUDE is used

with LOCATION, GROUP, MEMBER, or SUBSYSTEMID, no System Parameters reports are generated for those locations, groups, members, or subsystems excluded by this process.

- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax of the SYSPARMS command



Notes:

- 1 You can specify TRACE only once.
- 2 You can specify FILE only once.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting SYSPARMS functions

The following is a sample of the JCL required to produce SYSPARMS reports and traces. See Chapter 5, “DD statements,” on page 15 for descriptions of the DD statements.

```

//SYSPBTCH JOB ,MSGCLASS=X,CLASS=A,NOTIFY=USER1,REGION=0M
//* -----
//* BATCH SYSTEM PARAMETERS JCL.
//* -----
//PEMAIN EXEC PGM=FPECMAIN
//STEPLIB DD DSN=FPE.FPLIB.RKANMOD,DISP=SHR
//* -----
//INPUTDD DD DISP=SHR,DSN=DEA.DPMOUT.V8I106
//SYSOUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//DPMLLOG DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//* FOLLOWING ARE REPORT AND FILE DD NAMES
//SYFILDD1 DD SYSOUT=* -> Used for FILE
//SYTRCDD1 DD SYSOUT=* -> Used for TRACE
//* -----
//SYSIN DD *
SYSPARMS
TRACE
FILE
EXEC
/*
  
```

Figure 10. Sample JCL for requesting SYSPARMS functions

The OMEGAMON XE for DB2 PE command language shown in this section is not appropriate in all circumstances. You must modify it to meet your requirements.

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See the individual DD statement descriptions for more information.

Note:

1. There is a performance advantage in omitting DPMOUTDD from your JCL. For more information, see “DPMOUTDD statement” on page 16.
2. The OMEGAMON XE for DB2 PE command stream is only processed if EXEC is included as the last command. Otherwise, OMEGAMON XE for DB2 PE only checks the syntax.

SYSPARMS command with TRACE subcommand

This section describes the SYSPARMS command with the TRACE subcommand.

Usage

Use the TRACE subcommand to produce traces with an entry for every DB2 system parameter event.

Syntax of the SYSPARMS TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options.

Example using TRACE

The following is an example of the SYSPARMS command with the TRACE subcommand.

```
.  
.  
SYSPARMS  
TRACE  
EXEC
```

SYSPARMS command with FILE subcommand

This section describes the SYSPARMS command with the FILE subcommand.

Usage

To create system parameter data for the Performance Database, use the SYSPARMS command with the FILE subcommand to produce a data set suitable for loading into DB2.

Syntax of the SYSPARMS FILE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options.

Example using FILE

The following is an example of the SYSPARMS command with the FILE subcommand.

```

.
.
SYSPARMS
FILE
EXEC

```

UTILITY command

This section provides an overview of the UTILITY command.

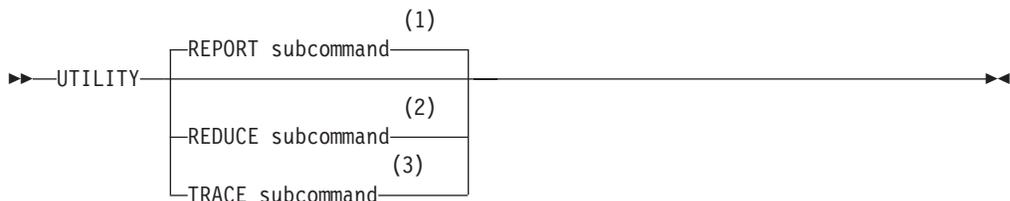
Usage

Use the UTILITY command to generate Utility Activity reports and traces.

Usage notes

- The command can be used once in a job step.
- However, it can be used in the same job step with commands of the other report sets.
- You can filter records with the GLOBAL command first. This can minimize your report output and reduce processing time. See “GLOBAL command” on page 149 for more information.

Syntax of the UTILITY command



Notes:

- 1 You can specify REPORT up to 5 times.
- 2 You can specify REDUCE only once. If specified, you must also specify REPORT at least once.
- 3 You can specify TRACE up to 5 times.

Subcommands

The subcommands are described in detail, together with their various options, in the following sections.

Sample JCL for requesting Utility Activity functions

The following figure is a sample of the JCL required to produce Utility Activity reports and traces. A description of the DD statements can be found in Chapter 5, “DD statements,” on page 15.

```
//          PEGMAIN EXEC PGM=FPECMMAIN
// * FOLLOWING ARE SYSTEM DDNAMES
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//DMPARMS DD DSN=MYID.FPELIB.DMPARMS,DISP=SHR
//INPUTDD DD DSN=MYID.FPELIB.DPMIN,DISP=SHR
//DPMLLOG DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//JOBSUMDD DD SYSOUT=*
//SYSPRMD DD SYSOUT=*
//DPMOUTDD DD DSN=MYID.FPELIB.DPMOUT.DATA,DISP=OLD
//SYSUDUMP DD DUMMY
// * FOLLOWING ARE REPORT SET DDNAMES
//UTRPTDD DD SYSOUT=*
//UTTRCDD1 DD SYSOUT=*
//UTWORK DD DSN=MYID.FPELIB.UT.WORKDD,DISP=OLD
// * FOLLOWING IS THE COMMAND STREAM
//SYSIN DD *
:
UTILITY
  REDUCE
  TRACE
  REPORT
:
EXEC
```

Figure 11. Sample JCL for requesting Utility Activity functions

The OMEGAMON XE for DB2 PE command language shown in this section is not appropriate in all circumstances. You must modify it to meet your requirements.

Most of the DD statements with a SYSOUT destination do not have to be specified because they are dynamically allocated by OMEGAMON XE for DB2 PE. See the individual DD statement descriptions for more information.

Note:

1. There is an advantage in omitting DPMOUTDD from your JCL. For more information, see the “DPMOUTDD statement” on page 16.
2. The OMEGAMON XE for DB2 PE command stream is only processed if EXEC is included as the last command. If you omit the EXEC statement, no report is generated. OMEGAMON XE for DB2 PE checks the syntax of the job stream and writes it, together with any information, warning, or error messages generated to the job summary log.

UTILITY command with REPORT subcommand

This section describes the UTILITY command with the REPORT subcommand.

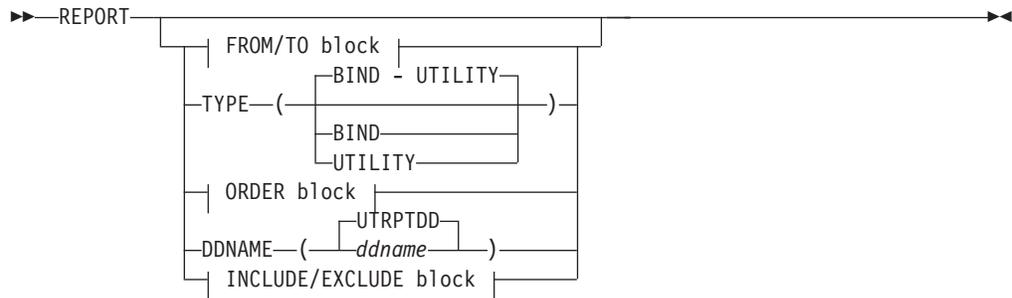
Usage

Use the REPORT subcommand to generate reports from records.

Usage notes

- Up to five REPORT subcommands can be specified within each UTILITY command.

Syntax of the REPORT subcommand



The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the reporting process by date and time.

For details, see “FROM/TO subcommand options” on page 30.

TYPE

Specifies the activity types. You can select one or both of these activity types with each REPORT subcommand:

BIND

Gives the elapsed time for each occurrence of a bind event. This includes information on the number of bound and nonbound events and the distribution of the bind into various bind subevents.

UTILITY

Gives the elapsed time for each occurrence of a utility event. Also provides information on the performance and resource usage of the various utility events.

If the TYPE subcommand option is omitted, both activity types are reported.

ORDER

Specifies the OMEGAMON XE for DB2 PE identifiers and their sequence for sorting the report, and in summary reports, the identifiers are used for aggregation.

For details, see “ORDER subcommand option” on page 46 and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

In the following example ORDER specifies that three Utility reports are to be produced.

```
⋮  
REPORT  
  TYPE (UTILITY)  
  ORDER (PRIMAUTH-PLANNAME-REQLOC  
         CONNECT-PLANNAME REQLOC-PRIMAUTH)  
  ⋮
```

- The first report is ordered by requesting location within plan name within primary authorization ID.
- The second report is ordered by plan name within connection ID.
- The third report is ordered by primary authorization ID within requesting location.

DDNAME

Specifies the data set where the report is written.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see “INCLUDE and EXCLUDE subcommand options” on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, “OMEGAMON XE for DB2 PE identifiers,” on page 3.

Example using REPORT with FROM, TO, ORDER, and EXCLUDE options

This example specifies the following:

- A report
- Sorted by plan name within primary authorization ID within correlation name
- Records are used with the time and date range of 10:00 a.m. on 18 March 1999 to noon on 19 March 1999
- Data is excluded that is associated with the following locations:
 - LOCN10
 - LOCN12
 - LOCN15
 - LOCN20

```
⋮  
UTILITY  
  REPORT  
    FROM (03/18/02,10:00:00.00)  
    TO (03/19/02,12:00:00.00)  
    ORDER (CORRNAME-PRIMAUTH-PLANNAME)  
    EXCLUDE (LOCATION(LOCN10 LOCN12 LOCN15 LOCN20))  
  ⋮
```

UTILITY command with REDUCE subcommand

This section describes the UTILITY command with the REDUCE subcommand.

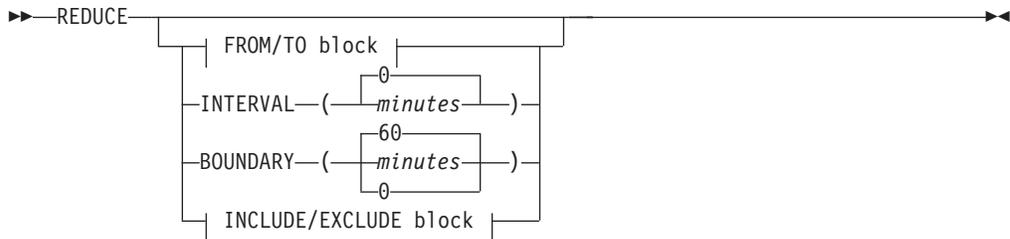
Usage

Use the REDUCE subcommand to reduce the volume of data that is input to subsequent subcommands. REDUCE consolidates records with certain common characteristics into one record.

Usage notes

- REDUCE can be used once in a UTILITY command.

Syntax of the REDUCE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

Example using REDUCE with FROM, TO, and INTERVAL options

This example specifies that data is to be reduced between 10:00 a.m. on 4 February 1999 and noon on 5 February 1999. The records are to be reduced into 60-minute intervals.

```

:
REDUCE
  FROM      (02/04/02,10:00:00.00)
  TO        (02/05/02,12:00)
  INTERVAL  (60)
:

```

UTILITY command with TRACE subcommand

This section describes the UTILITY command with the TRACE subcommand.

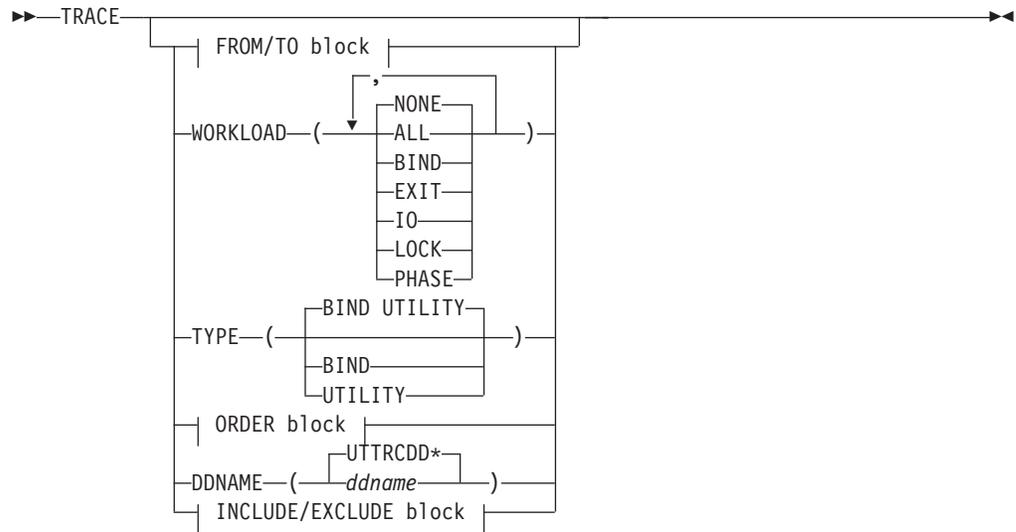
Usage

Use the TRACE subcommand to produce traces with an entry for each DB2 utility or bind event.

Usage notes

- Up to five traces can be requested in a job step.

Syntax of the TRACE subcommand



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

Limits the range of records included in the trace by date and time.

For details, see “FROM/TO subcommand options” on page 30.

WORKLOAD

Displays the workload detail for each event. The following detail can be specified:

NONE

No workload activity. This is the default.

BIND

Bind activity

EXIT

Exit activity

IO

I/O activity

LOCK

Lock suspension and page and row locking activity

PHASE

Utility phases

ALL

All workload activity

TYPE

There are two activity types. You can select one or both of these activity types with each TRACE subcommand. If the TYPE subcommand option is omitted, both activity types are reported.

BIND

Gives the elapsed time for each occurrence of a bind event for each trace entry. This includes information on the number of bound and nonbound events and the distribution of the bind into various bind subevents.

UTILITY

Gives the elapsed time for each occurrence of a utility event for each trace

entry, LISTDEF information, and information on the performance and resource and data set usage of the various utility events.

ORDER

Specifies the OMEGAMON XE for DB2 PE identifiers reported in the trace.

For details, see "ORDER subcommand option" on page 46 and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

Note: Traces are printed in the order that the threads end in and are not sorted by these identifiers.

DDNAME

Specifies the data set where the trace is written.

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

For details, see "INCLUDE and EXCLUDE subcommand options" on page 33, which lists other identifiers allowed with this command and subcommand combination, and Chapter 2, "OMEGAMON XE for DB2 PE identifiers," on page 3.

Example using TRACE with FROM and TO options

This example specifies:

- A trace with no workload.
- Including records from 10:00 to 10:15 on the input data set regardless of the date, because it was not specified.
- Sent to UTTRCDD1 by default.

```
TRACE
FROM (,10:00)
TO (,10:15)
```

Chapter 8. Auxiliary commands

This section provides an overview of the Auxiliary commands.

Usage

This section describes auxiliary processing commands shared by various report sets. You can use these commands to streamline the generation of reports. For each command its subcommand options and keywords and its relationship to the report set command is described. The description of the subcommand options is intentionally brief, to avoid recurrences. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions.

Usage notes

The table below shows which auxiliary commands can affect which OMEGAMON XE for DB2 PE report sets.

Table 10. Auxiliary commands valid with different report sets

Auxiliary Command	Accounting	Audit	I/O Activity	Locking	Record Traces	SQL Activity	Statistics	System Parameters	Utility Activity
CASE	●	●	●	●	●	●	●	●	●
FIELD					●				
GLOBAL	●	●	●	●	●	●	●	●	●
GROUP	●		●	●		●			●
LIST	●	●	●	●	●	●	●	●	●

The following topics provide additional information:

- “CASE command”
- “FIELD command” on page 147
- “GLOBAL command” on page 149
- “GROUP command” on page 156
- “LIST command” on page 164

CASE command

This section provides an overview of the Auxiliary command CASE.

Usage

Use the CASE command to differentiate between lowercase and uppercase values for DB2 PE identifiers. Specify CASE (SENSITIVE) before other commands if you want the commands to differentiate between uppercase and lowercase values. If no CASE command is specified or if you specify CASE (ANY), there is no differentiation between lowercase characters and uppercase characters.

Usage notes

- The CASE command affects all of the commands listed below it.
- You should usually list the CASE command first so that it affects all of the commands in the execution.

Syntax

►► CASE—(—^{SENSITIVE}—_{ANY}—)——►►

Subcommand options

The syntax diagram shows the options that are available with this subcommand. The following list gives descriptions of these options.

SENSITIVE

Differentiates between uppercase and lowercase values.

ANY

Does not differentiate between uppercase and lowercase values.

Examples using the CASE command

In these examples, there are two correlation names: driver and DRIVER. They are differentiated by their case.

In this example, CASE(SENSITIVE) is specified. Data related to the lower case correlation name driver is retrieved:

```
⋮  
CASE(SENSITIVE)  
GLOBAL  
  INCLUDE (IFCID (3,239))  
  INCLUDE (CORRNAME (driver))  
EXEC  
/*
```

In this example, CASE(SENSITIVE) is specified. Data related to the uppercase correlation name DRIVER is retrieved:

```
⋮  
CASE(SENSITIVE)  
GLOBAL  
  INCLUDE (IFCID (3,239))  
  INCLUDE (CORRNAME (DRIVER))  
EXEC  
/*
```

In this example, CASE(ANY) is specified. Data related to the lowercase correlation name driver and the uppercase correlation name DRIVER is retrieved:

```
⋮  
CASE(ANY)  
GLOBAL  
  INCLUDE (IFCID (3,239))  
  INCLUDE (CORRNAME (driver))  
EXEC  
/*
```

In this example, CASE is not specified. Data related to the lowercase correlation name driver and the uppercase correlation name DRIVER is retrieved:

```
⋮  
GLOBAL
```

```

INCLUDE (IFCID (3,239))
INCLUDE (CORRNAME (driver))
EXEC
/*

```

FIELD command

This section provides an overview of the Auxiliary command FIELD.

Usage

Use the FIELD command to define exception conditions. You can then filter individual records meeting these conditions by using the INCLUDE or EXCLUDE subcommand option.

Usage notes

- The FIELD command is used only with the INCLUDE and EXCLUDE subcommand option with RECTRACE. By using FIELD, you define the location of the data in a particular IFCID type, a comparison operator, and a value to compare the data against. The OMEGAMON XE for DB2 PE Record traces produced include or exclude those records that meet the comparison value. For example:

```

:
//SYSIN DD *
FIELD (
    QW0018ID,
    18,
    2,
    0,
    3,
    C,
    EQ,
    INDX
)
RECTRACE
TRACE
    LEVEL (SHORT)
    INCLUDE (IFCID(18))
FIELD (QW0018ID))
:
EXEC

```

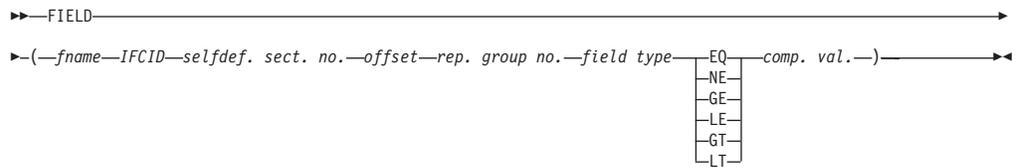
This example reads the scan end record IFCID 18, looking for all records that contain index data (INDX). It produces a short Record trace containing only those scan end records that meet this condition.

- Effective use of the FIELD command requires a detailed knowledge of the DB2 IFCID record formats.
- You can specify the following with FIELD:
 - A value you want used as a comparison value.
 - A location within a particular type of performance input record (identified by IFCID) that you want compared to the comparison value.
 - The type of comparison that you want to make between the previous two values (that is, data in the record equal to the comparison value).
- The decision to include a record in the trace is based on whether the comparison was true or false (see Table 11 on page 148) and whether the field name for the field definition was referenced in an INCLUDE or an EXCLUDE statement.

Table 11. Comparison values

Comparison True?	Field name Location	Decision
Yes	in INCLUDE	Record is used in the trace.
No	in INCLUDE	Record is not used in the trace.
Yes	in EXCLUDE	Record is not used in the trace.
No	in EXCLUDE	Record is used in the trace.

Syntax



Parameters

The syntax diagram shows the parameters that are available with this subcommand. The following list gives additional or specific descriptions of selected parameters, where appropriate.

field name

An 8-character name identifying a particular field.

IFCID

The decimal IFCID number of the performance record.

self defining section number

The decimal number identifying the self-defining data section that points to the data section containing the field to be compared. If the value is 0, the self-defining section identified is the first one that points to the product section. If the value is 1, the self-defining section is the first one that points to a data section.

offset

The decimal offset into the data section of the starting byte of the record field to be compared.

repeating group number

The number of repeating data sections where the comparison is made. If this field contains a value of 0, the comparison is made in all repeating data sections. If the value is, for example, 12, the self-defining section is the twelfth one that points to a data section.

Valid values are 0-99.

field type

A character indicating the type of data to be compared:

- C** Character data
- X** Hexadecimal data
- F** Fullword binary data
- H** Halfword binary data

comparison type

The type of comparison to be made between the field in the performance record and the comparison value in the definition:

- EQ** Equal to
- NE** Not equal to
- GE** Greater than or equal to
- LE** Less than or equal to
- GT** Greater than
- LT** Less than

comparison value

The value to be compared to the defined field in the performance record:

Character data

Up to 16 characters can be entered. Character data containing blanks must be enclosed in single quotation marks. The quotation marks cannot be part of the comparison value.

Hexadecimal data

This value or constant represents a hexadecimal value, such as 0001D0F2. Up to 8 characters can be specified. Do not enclose the value in quotation marks.

Fullword binary

This data is converted into a fullword binary value. It is limited to the maximum value allowed in a fullword field.

Halfword binary

This data is converted into a halfword binary value. It is limited to the maximum value allowed in a halfword field.

Example using the FIELD command

This example requests the following:

- The name attributed to this field comparison is QXSELECT.
- The IFCID of the input record where this comparison is made is 003.
- The data section containing the data to be compared is defined by the second self-defining section.
- The data to be compared is at decimal offset 008 in the data section.
- The comparison is to be made in all occurrences of that data section (0).
- The data to be compared is a fullword binary value (F).
- The comparison is made as to whether the data in the data section is greater than the comparison value.
- The comparison value is 100 in fullword binary format.

```
FIELD (  
    QXSELECT,  
    003,  
    02,  
    008,  
    0,  
    F,  
    GT,  
    100  
)
```

GLOBAL command

This section provides an overview of the Auxiliary command GLOBAL.

Usage

Use the GLOBAL command to filter input data and set default values for subcommand options. It is also used to define global processing options (such as DD statements for various data sets), the number of lines printed per page, whether the OMEGAMON XE for DB2 PE main internal sort occurs, and time zone adjustments for different locations.

Usage notes

The GLOBAL command affects the processing of each report set. It enables you to:

- Specify values for the FROM, TO, INCLUDE, and EXCLUDE subcommand options, which provide primary filtering of input records. Only those records that satisfy these options are available for further processing.
- Define global processing options:
 - Change the ddname for input data
 - Change the ddname for the System Parameters report
 - Change the ddname for the frequency distribution data set
 - Define the number of lines printed on each report or trace page
 - Specify a time adjustment for DB2 locations
 - Control whether the OMEGAMON XE for DB2 PE internal sort is performed
 - Control how the OMEGAMON XE for DB2 PE internal sort responds to possible incomplete spanned records in variable spanned input data sets (RECFM=VS or VBS)
- Define default values for the following REDUCE, REPORT, TRACE, and FILE subcommand options:
 - FROM and TO subcommand options: If you do not include either FROM or TO in a REDUCE, REPORT, TRACE, or FILE subcommand, the default dates and times specified in GLOBAL are applied.
 - INCLUDE and EXCLUDE subcommand options: If you do not provide either INCLUDE or EXCLUDE in a REDUCE, REPORT, TRACE, or FILE subcommand, the appropriate GLOBAL INCLUDE or EXCLUDE filters are applied as defaults.
 - INTERVAL subcommand option: If INTERVAL is not specified in a REDUCE subcommand, the default specified in GLOBAL is applied.
 - BOUNDARY subcommand option: If BOUNDARY is not specified in a REDUCE subcommand, the default specified in GLOBAL is applied.

There can be only one GLOBAL command in a OMEGAMON XE for DB2 PE execution.

Syntax

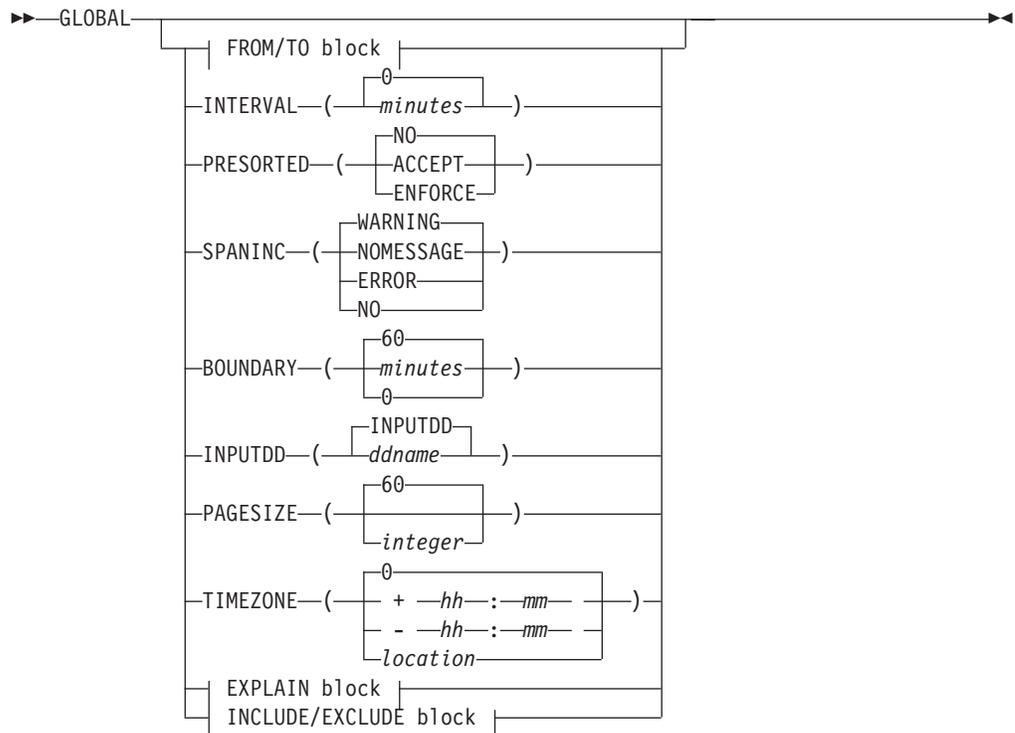


Figure 12. Syntax of the GLOBAL command

Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

FROM/TO

The range of record timestamps processed by the primary filtering function of GLOBAL.

For details, see “FROM/TO subcommand options” on page 30.

INTERVAL

Defines the time interval that data is summarized for.

For more information, see “INTERVAL subcommand option” on page 37.

PRESORTED

Controls the internal sort. It has the following values:

NO Does not disable the sort. This is the default.

ACCEPT

Disables the sort and accepts out-of-sequence records. By using this option, you can create complete and accurate Accounting reports (including records from several locations) from data sets that have not been sorted, for example SMF or GTF. The following limitations apply with this option:

- Duplicate input data is not checked. Therefore, do not concatenate two data sets that contain the same trace records, such as when trace data has been collected in two data sets during the same period.

- Only one location is reported per trace. To report more, specify several TRACE subcommands with INCLUDE on LOCATION.
- Trace entries might not be printed in time sequence. If a trace contains entries that are out of sequence, a message appears at the end showing the number.
- When INTERVAL(0) is in effect (the default), the interval times appearing in a report heading might not be accurate. Avoid ordering reports by interval when the interval is zero.

ENFORCE

Disables the sort but terminates processing if out-of-sequence records are present. Use this option only when the input data set has been sorted, for example when reprocessing data from the DPMOUT data set.

Note: Do not specify PRESORTED for Parallel I/O or Sysplex query parallelism, the results are unpredictable. Use the default (PRESORTED(NO)) if you see message FPEA4534I.

SPANINC

Controls how OMEGAMON XE for DB2 PE's underlying sort facility responds to incomplete spanned records in variable-length, unblocked, spanned (RECFM=VS) or variable-length, blocked, spanned (RECFM=VBS) records.

Sort facilities like DFSORT provide options to control the action to be taken if incomplete spanned records are detected. For DFSORT, the option SPANINC=*value*, where *value* is RC0, RC4, or RC16, controls the return code to be set, the message to be issued, and whether to terminate if necessary. For more details, see the *DFSORT Application Programming Guide*.

By default, OMEGAMON XE for DB2 PE passes option SPANINC=RC4 to the underlying sort facility. For DFSORT, this option sets a return code of 4, issues message ICE197I once, eliminates all incomplete spanned records, and continues with valid records being recovered. Usually, if no errors are detected, this process is transparent.

The following GLOBAL SPANINC subcommand options can be used to control the underlying sort facility if the default is not sufficient. The described behavior only applies to DFSORT and functionally equivalent sort facilities.

NOMESSAGE

Corresponds to SPANINC=RC0. Sets a return code of 0, issues message ICE197I (once), and eliminates all incomplete spanned records it detects.

WARNING

Corresponds to SPANINC=RC4. Sets a return code of 4, issues message ICE197I (once), and eliminates all incomplete spanned records it detects.

ERROR

Corresponds to SPANINC=RC16. Sets a return code of 16, issues message ICE204A, and terminates if an incomplete spanned record is detected.

NO No SPANINC option is passed to the underlying sort facility. Use this subcommand option if you run a sort facility that does not support the SPANINC option or the listed option values.

BOUNDARY

Controls the alignment of the intervals used to summarize records in the reduction process.

For more information, see "BOUNDARY subcommand option" on page 26.

INPUTDD

The ddname of the input data set. The default is INPUTDD.

SYSPRMDD

The ddname for the System Parameters report. The default is SYSPRMDD.

PAGESIZE

The number of lines printed per page. Specify an integer in the range 50 to 999. The default for PAGESIZE is 60.

Some reports have a fixed number of lines per page. PAGESIZE is ignored for I/O Activity summary reports, multi-page records in long Record Trace reports, and Explain reports.

TIMEZONE

The time adjustment applied to record timestamps during OMEGAMON XE for DB2 PE processing. By using TIMEZONE, you can process data from locations in different time zones based on the local time of a single location. All further processing is based on the adjusted time. The timestamp used in FROM and TO, printed on reports and traces, and recorded in the file and save data sets is the adjusted value. The DPMOUT data set contains both the original and the adjusted value.

Note:

When combining newly reduced data with restored data, make sure that the TIMEZONE specifications for the new data match those for the restored data. Misleading results can occur if the time adjustments are different.

If you change the reference location for data, the time adjustments in old save or file data might be incompatible with newly processed data.

This does not apply to the DPMOUT data, as time adjustments are recalculated when the data is read from INPUTDD.

You can specify time adjustments as follows:

location

The location of the reference time zone. The time is adjusted relative to the specified location. The adjustment is the difference between the time zone of the reference location and the CPU clock of the reported location. The LOCDATA member of the DPMPARMS data set must contain an entry for this location.

This is the recommended method.

*+hh:mm**-hh:mm*

The time difference between the reference time zone and Greenwich Mean Time (GMT). The time difference is specified as $\pm hh:mm$, where *hh* is hours in the range 00 to 23, and *mm* is minutes in the range 00 to 59. You can specify any value in the range -12:00 to +12:00. Use + for local times west of Greenwich, and - for local times east of Greenwich. The time is adjusted relative to the specified time zone difference.

The data for calculating the required adjustments is stored in the LOCDATA member of the DPMPARMS data set. It is stored and edited by using the IRF.

Note:

1. If TIMEZONE is not specified, no timestamps are adjusted. If the location supplied in the TIMEZONE option does not have a corresponding entry in the LOCDATA member, a message is generated and execution is terminated.
2. During OMEGAMON XE for DB2 PE processing, if data is encountered for a location whose time zone adjustment values have not been supplied in the LOCDATA member, the time adjustment default is applied to all records for that location. The default adjustment is stored in LOCDATA under a location name of *. If LOCDATA does not contain a default entry, no adjustment is applied.
3. If several systems that should have the same CPU clock time have synchronization errors, you can use TIMEZONE to correct the times for OMEGAMON XE for DB2 PE processing. Change the CPU clock time for the nonsynchronized locations in LOCDATA to reflect the errors. For example, if SAN_JOSE_LAB is one minute late in synchronizing with SANTA_TERESA_LAB, edit LOCDATA and add a minute to the CPU clock value for SAN_JOSE_LAB.
4. Synchronization can only be adjusted to the nearest minute. Any adjustments made to LOCDATA might become invalid if the CPU clock time at any of the locations is reset.

For more information about TIMEZONE, see *Monitoring Performance from ISPF*.

EXPLAIN

Defaults for EXPLAIN options can be specified with the GLOBAL command. These defaults are overridden by specifying different values in the EXPLAIN command, see "GLOBAL command with EXPLAIN option."

INCLUDE/EXCLUDE

Includes or excludes data associated with specific OMEGAMON XE for DB2 PE identifiers.

Table 12 on page 156 shows which OMEGAMON XE for DB2 PE identifiers can be used with the different INCLUDE and EXCLUDE subcommand options for GLOBAL.

GLOBAL command with EXPLAIN option

This section describes the GLOBAL command with the EXPLAIN option.

Usage

Some OMEGAMON XE for DB2 PE Explain options can be made globally valid, if specified in the GLOBAL command. An option value specified in a specific EXPLAIN command takes precedence over the corresponding GLOBAL option value.

Syntax of the GLOBAL EXPLAIN options

EXPLAIN Block:

Table 12. OMEGAMON XE for DB2 PE identifiers used with INCLUDE and EXCLUDE subcommand options for GLOBAL

OMEGAMON XE for DB2 PE identifiers	GLOBAL subcommands			
	REDUCE	TRACE	REPORT	FILE
ACE (Agent control element address)	X	X	X	X
CLASS (DB2 trace class)	X	X	X	X
CONNECT (Connection ID)	X	X	X	X
CONNTYPE (Connection type)	X	X	X	X
CORRNAME (Correlation name)	X	X	X	X
CORRNMBR (Correlation number)	X	X	X	X
DATABASE (Database name)	X	X	X	X
DATASET (Data set name)	X	X	X	X
ENDUSER (End user ID)				
FIELD (Comparison with data in a record field)	X	X	X	X
GROUP (Group name)	X	X	X	X
IFCID (Instrumentation Facility Component Identifier)	X	X	X	X
INSTANCE (Instance number)	X	X	X	X
LOCATION (Location name)	X	X	X	X
MAINPACK (Main package)				
MEMBER (Member name)	X	X	X	X
ORIGAUTH (Original authorization ID)	X	X	X	X
PACKAGE (Package information) or PROGRAM (Program information)				
PAGESET (Page set name)	X	X	X	X
PLANNAME (Plan name)	X	X	X	X
PRIMAUTH (Primary authorization ID) or AUTHID (Authorization ID)	X	X	X	X
REQLOC (Requester location)	X	X	X	X
RESOURCE TYPE (Resource type)				
RMID (Resource manager identifier)	X	X	X	X
SQLCODE				
SUBSYSTEMID (Subsystem ID)	X	X	X	X
THREADTYPE (Thread type)				
TRANSACTION (End user transaction name)				
TYPE (Event type)				
WSNAME (End user workstation name)				

GROUP command

This section provides an overview of the Auxiliary command GROUP.

Usage

Use the GROUP command to collect several OMEGAMON XE for DB2 PE identifier values under one name. When you request a report and specify this name by using the INCLUDE or EXCLUDE subcommand option, the events for all individual items are consolidated into one.

Usage notes

- With the GROUP command you can define a named set of OMEGAMON XE for DB2 PE identifier values. A set name can contain values or lists of values for a particular identifier.
- When you request a report and specify this name in INCLUDE or EXCLUDE, the events for all individual items are consolidated into one. For example, you might request that all PRIMAUTHs used by the accounting department be reported under the set name ACCTS. Thus, the entire department is reported in one entry rather than in individual entries for each PRIMAUTH.
- The GROUP command is only used to collect OMEGAMON XE for DB2 PE identifier values. It is not related to DB2 data sharing *groups*, and should not be confused with the GROUP keyword of the SCOPE subcommand option, or the GROUP OMEGAMON XE for DB2 PE identifier.
- Sets can be used in the REDUCE and REPORT subcommands of the following report sets:
 - Accounting
 - I/O Activity
 - Locking
 - SQL Activity
 - Utility Activity
- Sets can also be used in the following report sets and commands, however, OMEGAMON XE for DB2 PE treats them as lists:
 - Statistics
 - Audit
 - Record Traces
 - GLOBAL in all report sets
 - FILE and TRACE in all report sets

General rules regarding the use of GROUP

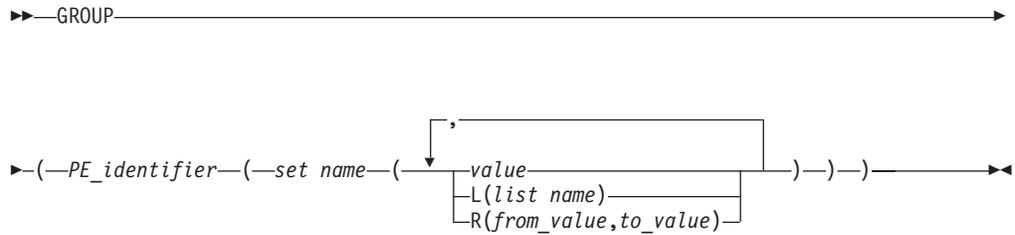
- Each GROUP command defines one set.
- You can use as many GROUP commands as you want.
- All values associated with a set must be for the same OMEGAMON XE for DB2 PE identifier.
- You can use more than one GROUP command for the same OMEGAMON XE for DB2 PE identifier.
- The combination of set name and OMEGAMON XE for DB2 PE identifier must be unique in the job step.

Rules applying to the use of GROUP with INCLUDE and EXCLUDE

- All set names that are referenced with INCLUDE or EXCLUDE must be defined by a GROUP command in the same job step.
- Set definitions are searched in the order that they are specified in.
- If more than one set name is used in INCLUDE or EXCLUDE, and the sets contain common values, the data is assigned to the first set that the value is

found in. When OMEGAMON XE for DB2 PE has found a value in a set, it does not attempt to find other occurrences of the same value.

Syntax



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

PE_identifier

The OMEGAMON XE for DB2 PE identifier that the set contains. Identifiers valid for the different report sets are shown in the following table.

Table 13. OMEGAMON XE for DB2 PE identifiers used with GROUP

OMEGAMON XE for DB2 PE Identifier	Accounting	I/O Activity	Locking	Utility Activity	SQL Activity
CONNECT	●	●	●	●	●
CONNTYPE	●	●	●	●	●
CORRNAME	●	●	●	●	●
CORRNMBR	●	●	●	●	●
DATABASE		●	●		
DATASET		●			
ENDUSER	●				
MAINPACK	●				
ORIGAUTH	●	●	●	●	●
PACKAGE	●				
PAGESET		●	●		
PLANNAME	●	●	●	●	●
PRIMAUTH or AUTHID	●	●	●	●	●
REQLOC	●	●	●	●	●
THREADTYPE	●				
TRANSACTION	●				
WSNAME	●				

set name

The name of the set. It must be unique within the job step. Set names can

consist of up to eight alphanumeric characters (A — Z, _, #, \$, @, 0 — 9) with no embedded spaces. Set names must begin with a character in the range A to Z.

value

A value for the specified OMEGAMON XE for DB2 PE identifier.

Identifier values must consist of the following characters: A — Z, #, \$, @, >, <, or 0 — 9. If the value you want to include contains a character that is not in this list, use an asterisk in its place or place the string in quotes.

Except for the CONNTYPE identifier, a value can be specified in generic form. Place an asterisk (*) in the value to indicate that any value is valid. For example:

- * processes any value.
 - ABCD* processes any value starting with ABCD.
 - *BCDE processes any value in the first character position where the second, third, fourth, and fifth character positions contain BCDE.
 - **CDE processes any value in the first and second character positions where the third, fourth, and fifth character positions contain CDE.
 - A*CDE processes any value in the second character position where the first character position contains A, and the third, fourth, and fifth character positions contain CDE.
- L** Indicates that the following value is a list name. You cannot use the generic form for list names. For example, L(ABC*E) cannot be used.
- R** Denotes a *range* of values beginning with *from_value* and ending with *to_value*. The *from_value* must be less than *to_value*. The generic form can be used only in the last character position in range values. For example, R(AUTH1*,AUTH2*) is acceptable, but R(AUTH*1,AUTH*2) is not.

Note:

1. Quoted values can also be specified. Any quoted string is accepted, provided that it passes length and format checking.
2. Range cannot be specified for CONNTYPE.

Grouping records

This section describes how to group records.

Usage

The record key for reduced data contains a number of OMEGAMON XE for DB2 PE identifiers. The identifiers contained in the key vary according to the report set.

When records are grouped, the set name of the records relating to grouped items is substituted for the original value of the specified identifier in the record key. When the substitution has been made, the records are not available for processing by using the original key values.

The following examples illustrate how sets are processed.

Example 1:

The following records are processed:

Table 14. Record set processing

Location	Connection ID	Correlation Name	Correlation Number	Plan Name	Primary Authorization ID
LOCATION_1	TSO	USER_1	0	PLAN_1	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_1	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_2	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_3	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_4	ACCOUNTS

The following OMEGAMON XE for DB2 PE command stream example groups three of the four plans:

```
⋮
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
⋮
EXEC
```

This establishes a set named PLANGRP that you can use with INCLUDE or EXCLUDE on REDUCE and REPORT.

The stage of processing when the key value substitution takes place depends on where you include the set:

- You can use sets on GLOBAL(INCLUDE) to include records for grouped items. No substitution takes place during preprocessing, but the input records for the grouped items are included as if you entered each item.

Note: When using sets on GLOBAL, remember that the GLOBAL INCLUDE or EXCLUDE specification becomes the default for all other commands in the job step. GLOBAL INCLUDE is illustrated in the following example:

```
⋮
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
GLOBAL
INCLUDE (
  PLANNAME (G(PLANGRP))
)
```

```

:
EXEC

```

The following records are available for further processing:

Table 15. Records grouped using GLOBAL(INCLUDE)

Location	Connection ID	Correlation Name	Correlation Number	Plan Name	Primary Authorization ID
LOCATION_1	TSO	USER_1	0	PLAN_1	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_1	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_2	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_3	ACCOUNTS

- When you use the set name on REDUCE(INCLUDE), the set name is substituted for the original value of the identifier in the key during REDUCE processing. The reduced data contains only the grouped records. You cannot process data by using the original key values in REPORT, and only the records for the set are stored in the save data set.

```

:
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
GLOBAL
INCLUDE (
  PLANNAME (G(PLANGRP))
)
ACCOUNTING
REDUCE
  INCLUDE (PLANNAME (G(PLANGRP)))
SAVE
:
EXEC

```

- When you use the set name on GLOBAL(INCLUDE), it acts as the default for subcommands without an INCLUDE specification. The result for the following example is the same as the result for the preceding example. The records are grouped during REDUCE processing by using the GLOBAL(INCLUDE) default.

```

:
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
GLOBAL
INCLUDE (
  PLANNAME (G(PLANGRP))
)
ACCOUNTING
REDUCE

```

```

      SAVE
      :
      EXEC

```

The reduced data and the save data set contain the following record:

Table 16. Records grouped during REDUCE processing using the GLOBAL(INCLUDE)

Location	Connection ID	Correlation Name	Correlation Number	Plan Name	Primary Authorization ID
LOCATION_1	TSO	USER_1	0	PLANGRP	ACCOUNTS

- When you use the set name on REPORT(INCLUDE), the records for grouped items are consolidated during report processing. In the following examples, the REDUCE subcommand specification overrides the GLOBAL default. All records that pass data filtering are included in the save data set because they are not grouped during REDUCE processing.

```

      :
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
GLOBAL
  INCLUDE (
    PLANNAME (G(PLANGRP))
  )
ACCOUNTING
  REDUCE
    INCLUDE (PLANNAME(*))
  REPORT
    INCLUDE (PLANNAME(G(PLANGRP)))
  SAVE
      :
      EXEC

```

In the following example, the set name on GLOBAL(INCLUDE) acts as the default for the REPORT subcommand. The result is the same as for the previous example.

```

      :
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
GLOBAL
  INCLUDE (
    PLANNAME (G(PLANGRP))
  )
ACCOUNTING
  REDUCE
    INCLUDE (PLANNAME(*))
  REPORT
  SAVE

```

```

:
EXEC

```

The Accounting report contains the following entry:

Table 17. Accounting report set

Location	Connection ID	Correlation Name	Correlation Number	Plan Name	Primary Authorization ID
LOCATION_1	TSO	USER_1	0	PLANGRP	ACCOUNTS

The save data set contains the following records:

Table 18. Save data for the accounting report set

Location	Connection ID	Correlation Name	Correlation Number	Plan Name	Primary Authorization ID
LOCATION_1	TSO	USER_1	0	PLAN_1	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_1	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_2	ACCOUNTS
LOCATION_1	TSO	USER_1	0	PLAN_3	ACCOUNTS

Example 2

In this example, two reports are generated. The first presents set items individually. In the second report, records are grouped. The REDUCE subcommand specification again overrides the GLOBAL default, so all records that pass preprocessing are included in the save data set.

```

:
//SYSIN DD *
GROUP (
  PLANNAME (
    PLANGRP (
      PLAN_1
      PLAN_2
      PLAN_3
    )
  )
)
GLOBAL
INCLUDE (
  PLANNAME (G(PLANGRP))
)
ACCOUNTING
REDUCE
  INCLUDE (PLANNAME(*))
REPORT
  INCLUDE (PLANNAME(*))
REPORT
  INCLUDE (PLANNAME(G(PLANGRP)))
SAVE
:
EXEC

```

In the following example, the set name on GLOBAL(INCLUDE) acts as the default for the second REPORT subcommand. The result is the same as for the previous example.

```

:
//SYSIN DD *
GROUP (

```

```

    PLANNAME (
      PLANGRP (
        PLAN_1
        PLAN_2
        PLAN_3
      )
    )
  )
GLOBAL
  INCLUDE (
    PLANNAME(G(PLANGRP))
  )
ACCOUNTING
  REDUCE
    INCLUDE (PLANNAME(*))
  REPORT
    INCLUDE (PLANNAME(*))
  REPORT
  SAVE
  :
EXEC

```

LIST command

This section provides an overview of the Auxiliary command LIST.

Usage

Use the LIST command to define a named list of values for a OMEGAMON XE for DB2 PE identifier, and to use the list name in INCLUDE or EXCLUDE instead of entering each list item.

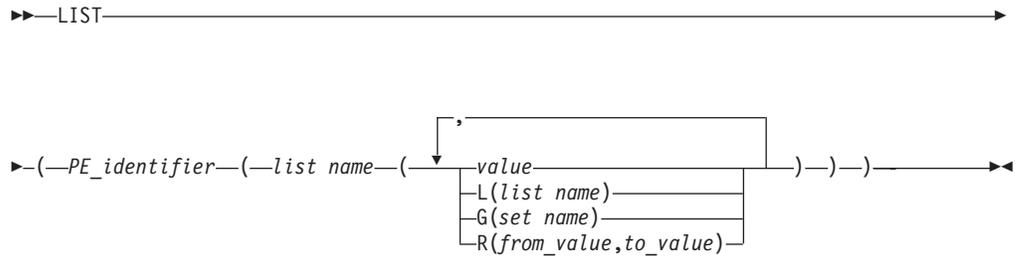
Usage notes

- The LIST command is available in all report sets.
- A list can contain values or lists of values for a particular identifier. Items in a list are treated independently, as if they were entered individually in INCLUDE or EXCLUDE.
- You can assign a name to the list and specify this name by using the INCLUDE or EXCLUDE subcommand option. In this way, you do not need to enter each list item individually.
- Each LIST command defines one list.
- You can use as many LIST commands as you want.
- All values itemized in a list must be for the same OMEGAMON XE for DB2 PE identifier.
- You can use more than one LIST command for the same OMEGAMON XE for DB2 PE identifier.
- The combination of list name and OMEGAMON XE for DB2 PE identifier must be unique in the job step.

Rules applying to the use of LIST with INCLUDE and EXCLUDE

- List definitions are searched in the order that they were specified for INCLUDE or EXCLUDE.
- All list names that are referenced in INCLUDE or EXCLUDE must be defined by a LIST command in the same job step. Those list names that are not defined by a LIST command are ignored during INCLUDE and EXCLUDE processing.

Syntax



Subcommand options

The syntax diagram shows the options that are available with this subcommand. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions of these options. The following list gives additional or specific descriptions of selected options, where appropriate.

PE_identifier

The OMEGAMON XE for DB2 PE identifier that the list contains. Any identifier can be used with the LIST command except FIELD.

list_name

The name of the list. The combination of list name and OMEGAMON XE for DB2 PE identifier must be unique in the job step. List names can consist of up to eight alphanumeric characters (A — Z, _, #, \$, @, 0 — 9) with no embedded spaces. List names must begin with a character in the range A to Z.

set_name

The name of the set. It must be unique within the job step. Set names can consist of up to eight alphanumeric characters (A — Z, _, #, \$, @, 0 — 9) with no embedded spaces. Set names must begin with a character in the range A to Z.

value

A value for the specified OMEGAMON XE for DB2 PE identifier.

Identifier values must consist of the following characters: A — Z, _, #, \$, @, >, <, or 0 — 9. If the value you want to include contains a character that is not in this list, use an asterisk in its place.

Except for the INSTANCE identifier, a value can be specified in generic form. Place an asterisk (*) in the value to indicate that any value is valid. For example:

- * processes any value.
- ABCD* processes any value starting with ABCD.
- *BCDE processes any value in the first character position where the second, third, fourth, and fifth character positions contain BCDE.
- **CDE processes any value in the first and second character positions where the third, fourth, and fifth character positions contain CDE.
- A*CDE processes any value in the second character position where the first character position contains A, and the third, fourth, and fifth character positions contain CDE.

- L** Denotes that the following value is a list name. You cannot use the generic form for list names. For example, L(ABC*E) cannot be used.

The L(list name) option cannot specify the list name for this LIST command.

- G** Denotes that the following value is a set name. You cannot use the generic form for set names. For example, G(ABC*E) cannot be used.
- R** Denotes a *range* of values beginning with *from_value* and ending with *to_value*. The *from_value* must be less than *to_value*. The generic form can be used only in the last character position in range values. For example, R(AUTH1*,AUTH2*) is acceptable, but R(AUTH*1,AUTH*2) is not.

Example using the LIST command

This example requests the following:

- The name of the list is AUTHLST5.
- The primary authorization IDs (PRMAUTH) associated with this list definition are:
 - USER01 through USER05
 - USER11.

```
//SYSIN DD *
LIST
      (PRMAUTH(AUTHLST5(R(USER01,USER05),USER11)))
ACCOUNTING
      TRACE
      INCLUDE (PRMAUTH(L(AUTHLST5)))
EXEC
```

Chapter 9. Troubleshooting commands

This section provides an overview of the Troubleshooting commands.

Usage

This section describes the DUMP command used to dump records from an input data set, and the TAPECOPY command used to copy records from an input data set to an output data set. The options for these commands are almost identical. DUMP and TAPECOPY can each occur once in a job step. The description of the subcommand options is intentionally brief, to avoid recurrences. See Chapter 6, “OMEGAMON XE for DB2 PE subcommand options,” on page 25 for comprehensive descriptions.

The following topics provide additional information:

- “DUMP command”
- “TAPECOPY command” on page 169

DUMP command

This section provides an overview of the Troubleshooting command DUMP.

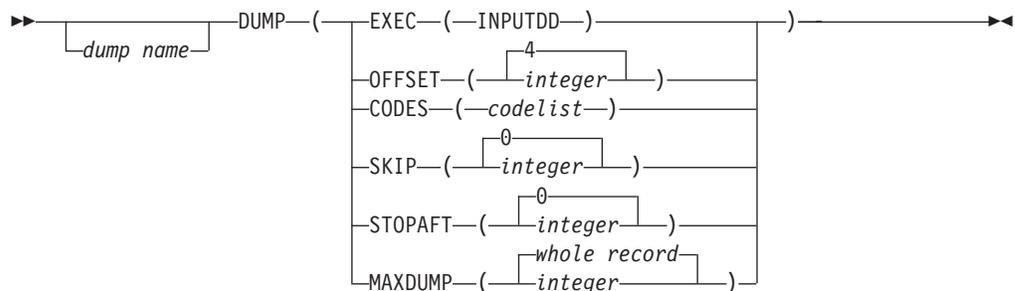
Usage

The DUMP command is a tool used for diagnosis. It provides, in dump format, a printout of an input data set. All records in the input data set, a selected range of records, or specific record types can be dumped.

Usage notes

- The DUMP command defines options for the record dump module.
- Dumps are written to SYSPRINT in a hexadecimal dump format.
- “Sample JCL for DUMP and TAPECOPY commands” on page 170 shows sample JCL for DUMP and TAPECOPY commands.

Syntax



Subcommand options

dump name

A user-defined name printed on the dump. If this field is omitted, the records are identified with the name DUMP0001.

Ensure that the specified name is not the same as a OMEGAMON XE for DB2 PE command keyword or abbreviation.

EXEC

Specifies INPUTDD as the ddname of the input data set. INPUTDD is the only valid ddname for this option and must be specified.

OFFSET

The offset of the record code into the record. The record code is a 1-byte field at position offset-plus-1. For example, OFFSET(4) defines a record code in the fifth byte of the record. The offset must be a numeric value less than the actual length of the record. The maximum value is 999 999 999. The default is 4.

CODES

The code values for records to be processed. Each code is a 2-digit hexadecimal number. You can specify either of the following:

- A list of values, for example '01,02,03'
- A range of values, for example '01-03,05-07'

Each entry must be separated by a comma. Enclose the code list in quotes if more than one value is specified.

If this option is omitted, all record codes (00-FF) are processed.

SKIP

The number of records to be skipped before processing begins. The maximum value is 999 999 999. The default is 0, which means that processing begins with the first record.

STOPAFT

The number of records to be processed, starting after the number of records to be skipped (**SKIP** option). The maximum value is 999 999 999. The default is 0, which causes all records (after skipping, if specified) to be processed.

MAXDUMP

The length of the dump in bytes, starting from the beginning of the record. The default is the full length of the record. You can enter any integer in the range of 1 to 99 999. For example, if you specify MAXDUMP(128), only the first 128 bytes of input records are dumped.

Note: Some IFCID records can be up to 32 KB in length. If you use the default for MAXDUMP (the entire record), very large reports can be produced.

Example using DUMP

In this example:

- The DUMP is named DUMPSTAT.
- The ddname of the input data set is INPUTDD (the GLOBAL default).
- The offset of 4 defines a record code in the fifth byte of the record.
- Only records with a value of 01 or 02 in the fifth byte are dumped.
- The first 125 records of the input data set are skipped.
- The next 10 records that meet the specifications are dumped.
- Only the first 1 000 bytes of each record are dumped.

```
DUMPSTAT DUMP (  
EXEC (INPUTDD)  
OFFSET (4)  
CODES ('01,02')  
SKIP (125)  
STOPAFT (10)  
MAXDUMP (1000))
```

Note: See also “Sample JCL for DUMP and TAPECOPY commands” on page 170.

TAPECOPY command

This section provides an overview of the Troubleshooting command TAPECOPY.

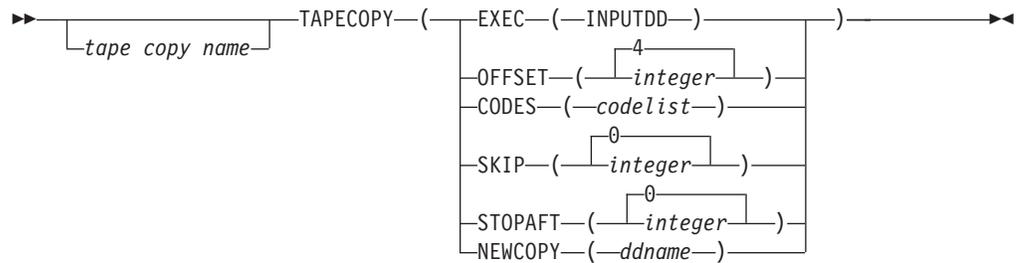
Usage

The TAPECOPY command is a utility tool that you can use to select a specified subset or all records from an input data set, and copy them to an output data set. The selection logic is identical to that used for the DUMP command.

Usage notes

- The TAPECOPY command defines options for the data set copy module.
- Copies of selected portions of the input data set are produced on a user-specified output data set.

Syntax



Subcommand options

tape copy name

A user-defined name identifying the records copied. If this field is omitted, the records are identified by the name COPY0001.

Ensure that the specified name is not the same as a OMEGAMON XE for DB2 PE command keyword or abbreviation.

EXEC

Specifies INPUTDD as the ddname of the input data set. INPUTDD is the only valid ddname for this option and must be specified.

OFFSET

The offset of the record code into the record. The record code is a 1-byte field at position offset-plus-1. For example, OFFSET(4) defines a record code in the fifth byte of the record. The offset must be a numeric value less than the actual length of the record. The maximum value is 999 999 999. The default is 4.

CODES

The code values for records to be processed. Each code is a 2-digit hexadecimal number. You can specify either of the following:

- A list of values, for example '01,02,03'
- A range of values, for example '01-03,05-07'

Each entry must be separated by a comma. Enclose the code list in quotes if more than one value is specified.

If this option is omitted, all record codes (00-FF) are processed.

SKIP

The number of records that are skipped before processing begins. The maximum value is 999 999 999.

The default is 0. If 0 is specified, processing begins with the first record.

STOPAFT

The number of records to be processed, starting after the number of records to be skipped (SKIP option). The maximum value is 999 999 999.

The default is 0. This causes all records (after skipping, if specified) to be processed.

NEWCOPY

The ddname of the output data set. The default is TAPECOPY.

Example using TAPECOPY

In this example:

- The TAPECOPY is named COPYSTAT.
- The ddname of the input data set is INPUTDD (the GLOBAL default).
- The offset of 4 defines a record code in the fifth byte of the record.
- Only records with a value of 01 or 02 in the fifth byte are copied.
- The first 50 records of the input data set are skipped.
- The next 10 records that meet the specifications are copied.
- The ddname of the output data set is OUTDATA.

```
COPYSTAT TAPECOPY (  
EXEC      (INPUTDD)  
OFFSET   (4)  
CODES    ('01,02')  
SKIP     (50)  
STOPAFT  (10)  
NEWCOPY  (OUTDATA))
```

Sample JCL for DUMP and TAPECOPY commands

The following figure shows a sample JCL for the DUMP and TAPECOPY commands.

Note: The command syntax shown below is not appropriate in all circumstances. You must modify it to meet your requirements and system setup.

```

//TPCDUMP1 JOB    (INSTALLATION DEPENDENCIES)
//*
//*
//*****
//*
//*
COPY INPUT DATA FILE TO SYSPRINT OR AN OUTPUT DATA FILE
//*
//*****
//*
// PEMAIN EXEC PGM=FPECMMAIN
//STEPLIB DD DSN=FPE.FPELIB.RKANMOD,DISP=SHR
//SYSPRINT DD SYSOUT=X
//SYSOUT DD SYSOUT=X
//*
//*
--- INPUT DATA SET ---
//*
//INPUTDD DD DSN=FPE.FPELIB.INPUT1,DISP=SHR
//*
//*
--- OUTPUT DATA SET ---
//*
//OUTDATA DD DSN=FPE.FPELIB.OUTPUT1,
//          DISP=(NEW,CATLG),
//          UNIT=SYSDA,SPACE=(TRK,(10,10),RLSE),
//          DCB=(RECFM=VB,LRECL=4092,BLKSIZE=4096)
//*
//SYSIN DD *
*
*
--- TAPECOPY COMMAND EXAMPLE ---
*
COPY0001 TAPECOPY (EXEC(INPUTDD),SKIP(125),STOPAFT(10),NEWCOPY(OUTDATA))
*
*
*
--- DUMP COMMAND EXAMPLE ---
*
DUMP0001 DUMP (EXEC(INPUTDD),SKIP(125),STOPAFT(10),MAXDUMP(1000))
EXEC
*
*

```

Figure 13. Sample JCL for the DUMP and TAPECOPY commands

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Deutschland GmbH
Dept. M358
IBM-Allee 1
71139 Ehningen
Germany

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com/legal/us/en/copytrade.shtml) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/us/en/copytrade.shtml>.

Adobe is either a registered trademark or a trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Bibliography

You can order many IBM publications such as product manuals or IBM Redbooks online at the IBM Publications Center website.

You can also order by telephone by calling one of the following numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968

In other countries, contact your software account representative to order Tivoli publications.

IBM Tivoli OMEGAMON XE for DB2 Performance Expert publications

The product library for Version 5 Release 2 covers the following information units:

OMEGAMON XE for DB2 PE and OMEGAMON XE for DB2 PM

- *Configuration and Customization*, GH12-6998
- *Parameter Reference*, SH12-6999
- *Monitoring Performance from ISPF*, SH12-6996
- *Monitoring Performance from the OMEGAMON Classic Interface*, SH12-6994
- *Monitoring Performance from Performance Expert Client*, SH12-6995
- *Report Command Reference*, SH12-6992
- *Report Reference*, SH12-6991
- *Reporting User's Guide*, SH12-6997
- *Messages*, GH12-6993
- *Program Directory for Performance Monitor*, GI19-5009
- *Program Directory for Performance Expert*, GI19-5007
- *Quick Start Guide for the end-to-end SQL monitoring function*, GH12-6990

Buffer Pool Analyzer

- *Buffer Pool Analyzer Configuration Guide*, SH12-7030
- *Buffer Pool Analyzer User's Guide*, SH12-7029
- *Program Directory for IBM DB2 Buffer Pool Analyzer for z/OS*, GI19-5010

InfoSphere Optim Performance Manager for Linux, UNIX, and Windows

- *InfoSphere® Optim™ Performance Manager Installation Guide*, GC19-2934

The documentation is provided in PDF and htm format in the:

- *Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS information center*
- *Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS information center*

IBM Tivoli Monitoring publications

For the most current list of publications, see the Tivoli Monitoring information center.

IBM DB2 publications

For the most current list of publications, see the IBM Information Management Software for z/OS Solutions information center.

Other IBM publications

For IBM publications that are not directly related to OMEGAMON XE for DB2 PE and PM, see Related information in the MVS™ Programming Hiperbatch Guide.

Index

A

- accessibility features xi
- ACCOUNTING command 53
 - with FILE subcommand 71
 - with REDUCE subcommand 62
 - with REPORT subcommand 55
 - with RESTORE subcommand 66
 - with SAVE subcommand 64
 - with TRACE subcommand 67
- ACCTYPE option
 - EXPLAIN command 83
- ACFILDD1 21
- ACMEM01 15
- ACRPTDD 22
- ACRSTDD 22
- ACSAVDD 23
- ACTRCDD1 23
- ACWORK 24
- ALL option
 - EXPLAIN command 83
- AUDIT command 73
 - with FILE subcommand 80
 - with REDUCE subcommand 77
 - with REPORT subcommand 74
 - with TRACE subcommand 79
- AUDWORK 24
- AUFILDD1 21
- AURPTDD 22
- AUTRCDD1 23
- AUWORK 24
- auxiliary command 145
 - CASE command 145
 - FIELD command 147
 - GLOBAL command 150
 - GROUP command 157
 - LIST command 164
 - troubleshooting command 167

B

- bibliography 177
- BLKSIZE 15
- BOUNDARY subcommand option 26

C

- CASE command 145
- command
 - understanding syntax diagrams viii
- command overview 52
- command stream 11
- comments, sending xii
- conventions used vi

D

- DATEFORMAT subcommand option 27
- DBRM option
 - EXPLAIN command 83

- DD statements 15
 - ACFILDD1 21
 - ACMEM01 15
 - ACRPTDD 22
 - ACRSTDD 22
 - ACSAVDD 23
 - ACTRCDD1 23
 - ACWORK 24
 - AUDWORK 24
 - AUFILDD1 21
 - AURPTDD 22
 - AUTRCDD1 23
 - AUWORK 24
 - BLKSIZE 15
 - DPMLOG 15
 - DPMOUTDD 16
 - DPMPARMS 17
 - EXCPTDD 17
 - EXFILDD1 18
 - EXPRODDA 18
 - EXTRCDD1 19
 - INPUTDD 19
 - IORPTDD 22
 - IOTRCDD1 23
 - IOWORK 24
 - JOBSUMDD 19
 - JSSRSDD 20
 - LOFILDD1 21
 - LORPTDD 22
 - LOTRCDD1 23
 - LOWORK 24
 - LRECL 15
 - RECFM 15
 - RTFILDD1 21
 - RTRPTDD 22
 - RTTRCDD1 23
 - RTWORK 24
 - SQLWORK 24
 - SQRPTDD 22
 - SQTRCDD1 23
 - STFILDD1 21
 - STRPTDD 22
 - STRSTDD 22
 - STSAVDD 23
 - STTRCDD1 23
 - STWORK 24
 - SYSIN 20
 - SYSOUT 21
 - SYSPRMDD 21
 - UTRPTDD 22
 - UTTRCDD1 23
 - UTWORK 24
- DDNAME subcommand option 29
- DEGREE option
 - EXPLAIN command 84
- DFSORT, SPANINC option 152
- DPMLOG 15
- DPMOUTDD 16
 - INCLUDE and EXCLUDE 16
- DPMPARMS 17
- DUMP command 167

E

- EXCEPTION subcommand option 30
- EXCLUDE and INCLUDE
 - DPMOUTDD 16
- EXCLUDE subcommand option 33
 - with GLOBAL command 155
- EXCPTDD 17
- EXEC command 82
- EXFILDD1 18
- EXPLAIN command 82
 - with EXPLAIN PACKAGE specifications 88
 - with EXPLAIN PLAN specifications 87
 - with EXPLAIN QMFQUERY specifications 90
 - with EXPLAIN QUERYNO specifications 92
 - with EXPLAIN SQLSTMT specifications 91
- EXPLAIN option 154
- EXPLAIN PACKAGE specifications 88
- EXPLAIN PLAN specifications 87
- EXPLAIN QMFQUERY specifications 90
- EXPLAIN QUERYNO specifications 92
- EXPLAIN SQLSTMT specifications 91
- EXPRODDA 18
- EXTRCDD1 19

F

- FIELD command 147
 - options 148
- FILE subcommand
 - options 72, 134
 - with ACCOUNTING command 71
 - with AUDIT command 80
 - with LOCKING command 105
 - with RETRACE command 108
 - with STATISTICS command 134
 - with SYSPARMS command 137
- FIRST option
 - EXPLAIN command 84
- FORCE option
 - EXPLAIN command 84
- FORMAT option
 - EXPLAIN command 84
- FROM subcommand option 30

G

- GEN option
 - EXPLAIN command 84
- GLOBAL command 150
 - with EXCLUDE subcommand option 155
 - with EXPLAIN option 154
 - with INCLUDE subcommand option 155
- GROUP command 157

GROUP command (*continued*)
grouping records 159

H

HOSTVAR option
EXPLAIN command 85

I

IBM support assistant xi
identifiers 3
used with ORDER 46
INCLUDE and EXCLUDE
DPMOUTDD 16
INCLUDE subcommand option 33
with GLOBAL command 155
INDEX option
EXPLAIN command 85
INPUTDD 19
Interactive Report Facility (IRF) 11
INTERVAL subcommand option 37
IOACTIVITY command 93
with REDUCE subcommand 97
with REPORT subcommand 94
IORPTDD 22
IOTRCDD1 23
IOWORK 24
IRF (Interactive Report Facility) 11

J

job summary data
RESTORE subcommand 20
SAVE subcommand 20
JOBSUMDD 19
JSSRSDD 20

L

LAST option
EXPLAIN command 84
LAYOUT subcommand option 38
LEVEL option
EXPLAIN command 85
LEVEL subcommand option 45
LIMIT subcommand option 45
LIST command 164
LOCATION option
EXPLAIN command 86
LOCKING command 99
with FILE subcommand 105
with REDUCE subcommand 102
with REPORT subcommand 100
with TRACE subcommand 103
LOFILDD1 21
LORPTDD 22
LOTRCDD1 23
LOWORK 24
LRECL 15

M

MATCHING option
EXPLAIN command 83

message
ICE197I 152
ICE204A 152

N

NOEXCEPTION subcommand option 30
NONMATCH option
EXPLAIN command 83

O

options
FIELD command 148
ORDER subcommand option 46
output, report 13
overview, command 52
OWNER option
EXPLAIN command 86

P

PACKAGES option
EXPLAIN command 86
PACKLIMIT option
EXPLAIN command 86

R

RECFM 15
redbooks x
REDUCE subcommand
options 62, 98, 125
with ACCOUNTING command 62
with AUDIT command 77
with IOACTIVITY command 97
with LOCKING command 102
with STATISTICS command 125
with UTILITY command 142
report output 13
REPORT subcommand
options 75, 101, 112, 140
with ACCOUNTING command 55
with AUDIT command 74
with IOACTIVITY command 94
with LOCKING command 100
with SQLACTIVITY command 111
with STATISTICS command 121
with TOP subcommand 55
with UTILITY command 140
RESTORE subcommand
effect on job summary 20
options 130
with ACCOUNTING command 66
with STATISTICS command 130
RETRACE command 105
with FILE subcommand 108
with TRACE subcommand 107
RTFILDD1 21
RTRPTDD 22
RTTRCDD1 23
RTWORK 24

S

SAVE subcommand
effect on job summary 20
options 64, 129
with ACCOUNTING command 64
with STATISTICS command 128
SCOPE subcommand option 48
sending comments xii
service x
sort facility
incomplete spanned record 152
other than DFSORT 152
SORTBY subcommand option 48
SPANINC option, DFSORT 152
spanned record, incomplete 152
SPREADSHEETDD subcommand
option 48
SQLACTIVITY command 110
with TRACE subcommand 115
SQLACTIVITY subcommand
with REPORT command 111
SQLID option
EXPLAIN command 86
SQLWORK 24
SQRPTDD 22
SQTRCDD1 23
SSID option
EXPLAIN command 87
STATISTICS command 119
with FILE subcommand 134
with REDUCE subcommand 125
with REPORT subcommand 121
with RESTORE subcommand 130
with SAVE subcommand 128
with TRACE subcommand 131
STFILDD1 21
STRPTDD 22
STRSTDD 22
STSAVDD 23
STTRCDD1 23
STWORK 24
subcommand option
BOUNDARY 26
DATEFORMAT 27
DDNAME 29
EXCEPTION 30
EXCLUDE 33
FROM 30
INCLUDE 33
INTERVAL 37
LAYOUT 38
LEVEL 45
LIMIT 45
NOEXCEPTION 30
ORDER 46
SCOPE 48
SORTBY 48
SPEADSHEETDD 48
SUMMARIZEBY 49
TO 30
TOP 49
TYPE 50
WORKLOAD 50
subcommand options 25
SUMMARIZEBY subcommand
option 49
support home website x

- syntax diagrams viii
- SYSIN 20
- SYSOUT 21
- SYSPARMS command 135
 - with FILE subcommand 137
 - with TRACE subcommand 137
- SYSRMDD 21
- SYSRMDD card 135

T

- TABLE option
 - EXPLAIN command 87
- TABSCAN option
 - EXPLAIN command 83
- TAPECOPY command 169
- terminology online x
- terminology used vii
- time formats, elapsed 14
- TO subcommand option 30
- TOP subcommand option 49
- TRACE subcommand
 - with ACCOUNTING command 67
 - with AUDIT command 79
 - with LOCKING command 103
 - with RETRACE command 107
 - with SQLACTIVITY command 115
 - with STATISTICS command 131
 - with SYSPARMS command 137
 - with TOP subcommand 67
 - with UTILITY command 142
- troubleshooting command 167
 - DUMP command 167
 - TAPECOPY command 169
- TYPE subcommand option 50

U

- updates x
- UTILITY command 138
 - with REDUCE subcommand 142
 - with REPORT subcommand 140
 - with TRACE subcommand 142
- UTRPTDD 22
- UTTRCDD1 23
- UTWORK 24

W

- where to find information x
- WORKLOAD subcommand option 50

Readers' Comments — We'd Like to Hear from You

IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS
IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS
Report Command Reference
Version 5.2.0

Publication No. SH12-6992-00

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

Email address



Fold and Tape

Please do not staple

Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM Deutschland Research & Development GmbH
Dept. 0446
User Experience & Information Development
Schoenaicher Strasse 220
71032 Boeblingen
Germany

Fold and Tape

Please do not staple

Fold and Tape



Product Number: 5655-W37
5655-W38

SH12-6992-00

