

5.5

*IBM OMEGAMON for Db2 Performance
Expert on z/OS
ISPF Client User's Guide*



2024-07-18 edition

IBM Confidential. This edition applies to the IBM® OMEGAMON for DB2® Performance Expert on z/OS VNext Sponsor User Program (product number 5655-W37) only.

© **Copyright International Business Machines Corporation 2005, 2022.**

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

© **Rocket Software Inc. 2016, 2022.**

Contents

- About this information..... V**

- Chapter 1. Overview..... 1**
 - Service updates and support information..... 1
 - How to read syntax diagrams..... 1
 - Conventions..... 2
 - Terminology..... 3
 - Accessibility features..... 4

- Chapter 2. About ISPF Online Monitor..... 5**
 - Required authority..... 5
 - Accessing the ISPF Online Monitor..... 5
 - Main menu..... 5
 - Options..... 6
 - Control Exception Processing..... 6
 - Collect Report Data - General..... 6
 - Collect Report Data - For Buffer Pool Analysis..... 6
 - Explain..... 6
 - Options..... 7
 - Db2 Subsystem windows..... 7
 - History Defaults window..... 9
 - Monitor Output window..... 10
 - Auto Display window..... 10
 - Session options..... 11
 - Current SQLID window..... 12
 - Moving between panels..... 12
 - Values not updated..... 12
 - Leaving the ISPF Online Monitor..... 12
 - Function key default settings..... 13
 - Issuing Db2 commands..... 16
 - Using the COLLECT command..... 17
 - Using the OPTIONS command..... 17
 - Correlation ID translation..... 18
 - The default translation..... 18
 - Changing the default translation..... 18
 - How large and missing values are displayed..... 19
 - Elapsed time formats..... 20
 - Date formats..... 20
 - Displaying long names and values in scrollable fields..... 20

- Chapter 3. Monitoring and tuning..... 23**
 - Changed access path to the ISPF Online Monitor..... 24
 - Performance data generation and online monitoring..... 24
 - The Db2 operating environment..... 25
 - Performance objectives and exception processing..... 25
 - Monitoring performance..... 25
 - Monitoring using reports..... 26
 - Observing performance trends..... 26
 - Determining performance problems..... 26
 - Problems detected in periodic exception processing..... 26

Problems detected in exception event processing.....	27
Problems detected in exception reports.....	27
Viewing past data.....	28
Exception processing.....	29
Exception processing modes.....	30
Exception Notification window.....	31
How to define exception threshold values.....	32
How to start exception processing.....	32
How to restart exception processing (REINIT command).....	36
Examining exception messages (LOOK command).....	36
Exception processor output data sets.....	41
Stopping exception processing.....	42
Monitoring the access path with Explain.....	42
Authorizations required to use Explain.....	43
Performance considerations.....	43
What to look for in Explain information.....	43
Plan table considerations.....	44
Explain an SQL statement.....	45
Explain Menu.....	46
Explain output (Db2 Explain Output panel).....	52
Online Monitor EXPLAIN command.....	54
ISPF Online Monitor Source Explain.....	54
Db2 System Parameters panel.....	57
System Parameters commands.....	58
Problem identification examples.....	59
Example of exception processing.....	59
Using the ISPF Online Monitor in deadlock resolution.....	68
Application development and capacity planning.....	69
Chapter 4. Understanding common errors.....	71
STEPLIB missing from IRF-generated JCL.....	71
Chapter 5. Additional Information.....	73
Layout of the Exception Log File output record.....	73
Product legal notices.....	77
Index.....	81

About this information

IBM OMEGAMON for Db2 Performance Expert on z/OS (also referred to as OMEGAMON for Db2 Performance Expert) is a performance analysis, monitoring, and tuning tool for Db2 on z/OS® environments.

The document is part of the OMEGAMON for Db2 Performance Expert documentation library which provides instructions for installing, configuring, and using OMEGAMON for Db2 Performance Expert and is designed to help database administrators, system programmers, application programmers, and system operators perform these tasks:

- Plan for the installation of OMEGAMON for Db2 Performance Expert
- Install and operate OMEGAMON for Db2 Performance Expert
- Customize your OMEGAMON for Db2 Performance Expert environment
- Diagnose and recover from OMEGAMON for Db2 Performance Expert problems
- Design and write applications for OMEGAMON for Db2 Performance Expert
- Use OMEGAMON for Db2 Performance Expert with other DB2 products

Chapter 1. Overview

IBM OMEGAMON for Db2 Performance Expert on z/OS (OMEGAMON for Db2 Performance Expert) enables you to monitor, analyze, and tune the performance of your Db2 subsystems and Db2 applications.

Service updates and support information

Service updates and support information for this product, including software fix packs, PTFs, frequently asked questions (FAQs), technical notes, troubleshooting information, and downloads, are available from the web.

To find service updates and support information, see the following website:

<https://www.ibm.com/support/pages/omegamon-xe-db2-pepm-web-based-delivery-and-updates-windows-and-unix-based-components>

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.



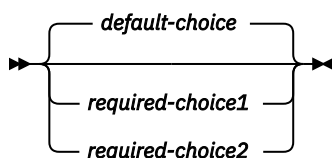
If the repeat arrow contains a number in parenthesis, the number represents the maximum number of times that the item can be repeated.



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.



Conventions

These conventions are used throughout the documentation.

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={<u>YES</u> 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.

Terminology

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.

<i>Table 1. Product names and their short names</i>	
Product name	Short name
IBM OMEGAMON for Db2 Performance Expert on z/OS	OMEGAMON for Db2 Performance Expert
IBM zSystems Monitoring Configuration Manager	Configuration Manager

Accessibility features

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use a software product successfully.

The major accessibility features in this product enable users to perform the following activities:

- Use assistive technologies such as screen readers and screen magnifier software. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.
- Customize display attributes such as color, contrast, and font size.
- Operate specific or equivalent features by using only the keyboard. Refer to the following publications for information about accessing ISPF interfaces:
 - *z/OS ISPF User's Guide, Volume 1*
 - *z/OS TSO/E Primer*
 - *z/OS TSO/E User's Guide*

These guides describe how to use the ISPF interface, including the use of keyboard shortcuts or function keys (PF keys), include the default settings for the PF keys, and explain how to modify their functions.

Chapter 2. About ISPF Online Monitor

When you invoke the ISPF Online Monitor by selecting option 3 (Trace collection, Explain, and Admin functions) on the OMEGAMON® for Db2 PE main menu, the Online Monitor Main Menu is displayed.

The Online Monitor Main Menu provides access to functions through sets of menus and panels.

```
03/25/22 17:13          Online Monitor Main Menu          PM06D861 D861  V10
Select one of the following.

  1. Options
  2. Control Exception Processing
 3a. Collect Report Data - General
 3b. Collect Report Data - For Buffer Pool Analysis
  4. Explain

Command ===>
  F1=Help   F2=Split  F3=Exit   F9=Swap   F12=Cancel
```

Figure 1. Online Monitor Main Menu

Considerations

Real time monitoring is provided by the Classic User Interface.

Required authority

The information in this section applies only if you do not use the user authorization exit.

If the user exit is active, the description about the user authorization exit provided in [IBM Db2 for z/OS in the IBM Knowledge Center](#) applies.

To use the Online Monitor, you need the following authorities:

- Access to a Db2 subsystem
- EXECUTE authority on the Online Monitor plan. The plan name is KO2PLAN.
- MONITOR1 privilege
- Display trace privilege

You also need the Db2 authority for any Db2 commands you issue.

The following authorities are recommended:

- MONITOR2 or SYSADM privilege if you need access to the currently executing SQL statement
- Db2 privilege to start or stop a Db2 trace, if you want to collect report data

Important: Granting MONITOR2 privilege enables access to potentially sensitive data, for example, the SQL statement being executed.

Accessing the ISPF Online Monitor

To access the ISPF Online Monitor, start your IBM OMEGAMON for Db2 Performance Expert on z/OS session from the TSO/ISPF environment (Start EXEC FPEJINIT).

Main menu

This section outlines the privileges and traces required to use the ISPF Online Monitor, shows how to change parameters that affect the behavior of an Online Monitor session, describes the online help, the

default function key settings, how to move between Online Monitor panels, how to issue Db2 commands, global commands, and how to treat some common errors.

Options

You can use options to change the parameters affecting the behavior of your own Online Monitor session. For example, you might want to specify the following options:

- Select the Db2 subsystem to be monitored.
- Display your installation-specific history parameters (only available if the data collector is installed at your site).
- Control the writing of monitor records to an output data set.
- Set the default auto-display interval.
- Set several Interactive Report Facility (IRF) session options.
- Set the current SQLID for Explain processing.

Control Exception Processing

Use Online Monitor exception processing to identify DB2 thread and statistics fields that contain values outside limits you have specified. This gives you better management of service levels by identifying problems in the DB2 subsystem and threads causing performance problems.

Collect Report Data - General

Use this option to start and stop Db2 traces either manually or automatically by specified triggers. The resulting trace data can be directed to a data set for immediate input to batch reporting. The Db2 traces required are determined by the reports that you want to produce.

See the [Reporting User's Guide](#) for more information.

Collect Report Data - For Buffer Pool Analysis

Use this option to collect data for the buffer pool analysis function. The collected data is used to report buffer pool efficiency, and to simulate the effects buffer pool tuning actions before altering a buffer pool's characteristics.

See the [Buffer Pool Analyzer User's Guide](#) for more information.

Explain

Use this option to examine the access path method chosen by Db2 for a given SQL statement in an easy-to-read format. The SQL statement you want to explain can be an existing entry in a specific PLAN_TABLE, an SQL statement from a previously bound plan or package, or a dynamically entered SQL statement.

You can view detailed information about packages, DBRMs, tables, indexes, and the SQL text. The SQL statement can be modified for online tuning.

You can also explain an SQL statement from within the thread activity function to explain a currently executing SQL statement, or from within an ISPF/PDF editor to explain an SQL statement imbedded in a source program or SPUFI input.

See [“Monitoring the access path with Explain” on page 42](#) for more information.

Options

Use the Options panels to change parameters affecting the behavior of your Online Monitor session. To display the Options menu, select option 1 (Options) from the Online Monitor Main Menu.

```
DGOMDPMN 13:13          Options          PM01DL0C DSN1 V10
Select one of the following.
--  1.  DB2 Subsystem
    2.  History Defaults
    3.  Monitor Output
    4.  Auto Display
    5.  Session Options
    6.  Current SQLID

Command ==>
F1=Help    F2=Split    F3=Exit    F9=Swap    F12=Cancel F16=Look
F17=Collect
```

Figure 2. Options menu

You can select one of the following options from this menu:

- Select option 1 (DB2 Subsystem) to display either the Db2 Subsystems List window or Db2 Subsystem window, where you can select the Db2 subsystem to be monitored.
- Select option 2 (History Defaults) to display the History Defaults window, where you can display the installation-specified history defaults. This option is available only if a data collector is active for the Db2 subsystem you are monitoring.
- Select option 3 (Monitor Output) to display the Monitor Output window, where you can control the writing of monitor records to an output data set.
- Select option 4 (Auto Display) to display the Auto Display window, where you can set the default interval for the AUTO command.
- Select option 5 (Session Options) to display the Session Options window, where you can set several options controlling the environment of your IRF session.
- Select option 6 (Current SQLID) to display the Current[®] SQLID window, where you can specify a different SQL authorization ID for qualifying the unqualified tables in the statements being explained.

Db2 Subsystem windows

Use the Db2 Subsystems List window to select the Db2 subsystem to be monitored.

Only the Db2 subsystems with data collectors started within the MVS™ system are listed on the Db2 Subsystems List window.

To display this window, select **Option 1 (DB2 Subsystem)** from the Options menu.

```

DGOMDWSL                DB2 Subsystems List                ROW 1 TO 3 OF 3

Select the Data Collector to use or specify a DB2 subsystem ID

DB2 Subsystem  DSN1

   Data Collector DB2 Location          DB2 Release
-   DSNB          PM01DB11             V10
-   DSNC          PM01DC11             V10
***** BOTTOM OF DATA *****

Command ==>
F1=Help   F2=Split  F7=Up      F8=Down   F9=Swap   F12=Cancel
F16=Look  F17=Collect

```

Figure 3. Db2 Subsystems List window

The following columns are shown:

Data Collector

The data collector identifier.

Db2 Location

The location as obtained from the trace records for the data collector.

Db2 Release

The Db2 release level.

To monitor a Db2 subsystem that has an active data collector running, enter any character in the input field beside a data collector Db2 location and press Enter.

To monitor a Db2 subsystem that does not have an active data collector running, enter the name of the Db2 subsystem in the DB2 Subsystem field and press Enter.

If no data collectors are active on your MVS system, the Db2 Subsystem window is displayed instead.

Note: The version and release numbers of the Db2 load library allocated to the Online Monitor session must match the version and release numbers of the Db2 subsystems being monitored.

```

DGOMDWS1                DB2 Subsystem

Enter the subsystem ID of the DB2 subsystem to monitor and
press Enter.

DB2 Subsystem  ____

F1=Help   F2=Split  F9=Swap   F12=Cancel

```

Figure 4. Db2 Subsystem window

If the data collector is not installed at your site or the monitored Db2 subsystem does not have an active data collector running, the following Online Monitor functions are not available:

- Viewing past data
- Periodic exception processing while you are not logged on
- Exception event processing
- Collection of parallel tasks for query CP parallelism

Changing Db2 subsystems terminates all asynchronous tasks that were started using the Online Monitor. If you change Db2 subsystems while an asynchronous task is active, you are notified by the Asynchronous Task Termination panel.

If you work in split screen mode and you change the subsystem in one session, asynchronous tasks are terminated even if they were activated in the other session.

```

DGOMGP03      Asynchronous Task Termination

Press Enter to change Subsystem or request Cancel to restore the
Subsystem ID to its previous value.

Changing the Subsystem results in the termination of all active
tasks as well as the stopping of all the DB2 traces started by the
Collect Report Data facility.

Exception/Collect tasks                               Status
-----
Periodic Exception                                   Active
Exception Event                                     Not active

Collect data for acct/io/stats/audit                 Active
CICS problem                                         Active
Collect Task C                                       Not active
Collect Task D                                       Not active

Command ==> -----
  F1=Help      F2=Split    F3=Exit      F9=Swap      F12=Cancel   F16=Look
  F17=Collect
  
```

Figure 5. Asynchronous Task Termination panel

To change Db2 subsystems, press Enter. All asynchronous tasks are terminated.

To return to the previous panel without changing Db2 subsystems, press F3 (Exit) or F12 (Cancel). All active asynchronous tasks remain active.

History Defaults window

This function is no longer supported. The menu entry for it will be removed in a future maintenance update.

Use the History Defaults window to view the history parameters that are currently in effect.

To display this window, select option 2 (History Defaults) from the Options menu.

```

DGOMDWH0      History Defaults      ROW 1 TO 6 OF 6

History Collected . . . . . : YES
History Interval . . . . . : 10

History From . . . . . : 04/05/08 13:30:36
History To . . . . . : Present

Data for which History is stored :
Statistics
System Parameters
Thread SQL Statement
Thread Summary
Thread Detail
Locked Resources
***** BOTTOM OF DATA *****

Command ==> -----
  F1=Help      F2=Split    F3=Exit      F7=Up
  F8=Down      F9=Swap     F12=Cancel   F16=Look
  F17=Collect
  
```

Figure 6. History Defaults window

By using the HISTORY command, you can view thread activity, statistics, and system parameters data previously gathered by the data collector. See [“Viewing past data” on page 28](#) for a complete description about how to view past performance data.

The History Collected field shows either YES or NO to indicate if any past data is available.

The `History Interval` field indicates how often DB2 instrumentation data is being gathered by the data collector. This field is a numeric value in seconds.

The `History From` field shows the date and time of the earliest history data available.

The `History To` field shows the date and time of the most recent history data available.

The `Data for which History is stored` field lists the types of data gathered by the data collector. This list includes the data types specified at data collector startup and subsequent changes made by the administrator users. Statistics history keeps subsystem-wide statistical information for a Db2 subsystem. For thread activity you can collect thread history individually in a summary format, in detail format, or in detail format with locking information, SQL statement, or both. System Parameters history keeps an overview of DSNZPARM values for a Db2 subsystem.

For more information about viewing past data, see [“Viewing past data” on page 28](#). For more information about data collector installation options, see [IBM Db2 for z/OS in the IBM Knowledge Center](#).

Monitor Output window

Use the Monitor Output window to enable or disable the writing of monitor output to a data set.

The main use of this option is to gather data for problem determination. Records in DPMOUT format that correspond to data shown in the Online Monitor panels are written to a data set whenever the display is updated. Output from the Statistics and System Parameters panels can be used as input to the Batch Statistics and System Parameter reports respectively. The output from the Thread Display panels can only be processed by a batch Record trace.

To display this window, select option 3 (`Monitor Output`) from the Options menu.

```
DGOMDWMO                Monitor Output

To write monitor output, enter 1 and dataset name and disposition

Write output . . . . . 2 1=yes 2=no

Dataset name . . . . . _____
Disposition . . . . . 1 1=append 2=overwrite 3=new

F1=Help   F2=Split   F9=Swap   F12=Cancel
```

Figure 7. Monitor Output window

Use the `Write output` field to specify whether you want monitor output written to a data set. Enter 1 in this field to enable the writing of monitor output to a data set.

The `Dataset name` field requires the name of the data set to which the data is to be written.

If you specify a disposition of 3 (new), the data set is dynamically allocated with the following attributes:

```
RECFM:
VBS
LRECL:
32 756
BLKSIZE:
6 233
```

Auto Display window

Use the Auto Display window to set the auto-display refresh interval for data shown in the Online Monitor panels. When auto display is active, the values displayed in the current panel are updated periodically, as specified in the Auto Display window. If you activate auto-display mode while viewing past data, the panel is updated with subsequent past records at every auto-display interval. The time difference between these records is determined by the history interval when the data was collected, and not by the current auto-display interval.

To display this window, select option 4 (Auto Display) from the Options menu.

```
DGOMDWAD                Auto Display
Update values then Enter
Units . . . . . 1  1=seconds
                   2=minutes
Interval . . . . . 10  1-7200 seconds
                   1-120  minutes
F1=Help   F2=Split  F9=Swap  F12=Cancel
```

Figure 8. Auto Display window

Use the Units field to specify the unit of time that is to be associated with the value in the Interval field.

Use the 1 field to specify how frequently the panels are refreshed when the AUTO command is issued.

To start auto-display mode, type AUTO on any Thread Activity or Statistics panel command line and press Enter. To stop auto-display mode, press the Attention key. You cannot perform any other actions with your Online Monitor session while auto-display mode is running.

Session options

Use the Session Options window to set several options controlling the environment of your Interactive Report Facility (IRF) session.

To display this window, select option 5 (Session Options) from the Options menu, or type OPTIONS on any command line and press Enter.

```
DGOF0PTS                Session Options
Update fields as required, then press Enter.
Confirmation display . . . . . 1  1=yes  2=no
Initial menu choice . . . . . 1  1-6 or blank
Execution mode . . . . . 1  1=Background
                           2=Foreground
                           3=Prompt
DPMPARMS data set . . . -----
F1=Help   F2=Split  F9=Swap  F12=Cancel
```

Figure 9. Session Options window

The field values shown in [Figure 9 on page 11](#) are the default settings.

Use **Confirmation display** to specify whether or not you want to be prompted each time you issue a delete, replace, or reset request.

Use **Initial menu choice** to specify the panel you want to start on.

Use **Execution mode** to specify how to execute batch jobs. Use the 1 field to specify the DPMPARMS data set to be used by the Online Monitor. The DPMPARMS data set contains information about exception thresholds, customized report layouts, time zone specifications, correlation ID translation, and the MAINPACK identifier. Only the correlation ID translation part is used by the Online Monitor.

Current SQLID window

Use the Current SQLID window to specify a different SQL authorization ID for qualifying the unqualified tables in the statements being explained if they do not belong to a plan or package. The Current SQLID window is also used to qualify the plan table to be accessed.

To display this window, select option 6 (Current SQLID) from the Options menu.

```
DGOMDWSS                               Current SQLID
                                         Specify the current SQLID, then press Enter.

Current SQLID . . . . . USERT001

F1=Help      F2=Split
F9=Swap      F12=Cancel
```

Figure 10. Current SQLID window

If the 1 field in this window is not specified, it defaults to your user ID.

Moving between panels

When you have entered a command on the command line, or have entered the information required to complete a panel, press Enter.

To go back to a previous panel or menu, press F3 (Exit).

Each time you press Enter or F3 (Exit), data entered in that panel is validated. If an error is detected, an error message is displayed and the cursor is positioned on the field in error. If no error is detected, processing continues.

To leave a panel without saving the entries and return to the previous panel, or to cancel all windows, press F12 (Cancel).

You can move between various Online Monitor options by typing =x where x is the Online Monitor Main Menu option number on the command line of any Online Monitor panel:. For example, type =1, =2...

Values not updated

Values are not updated if:

- Any OMEGAMON for Db2 PE or ISPF command is entered on the command line
- Windows are selected for display
- Function keys are used

Note: The values shown in the Thread Summary and Threads Holding Resource panels are refreshed automatically each time the panels are displayed.

Leaving the ISPF Online Monitor

To exit the ISPF Online Monitor from the Online Monitor Main Menu, press F3 (Exit) or F12 (Cancel). From any other Online Monitor panel, type =X on the command line and press Enter.

If any asynchronous tasks are active when you exit the Online Monitor, you are notified by one of the Asynchronous Task Termination panels. If the data collector is not active, the tasks are terminated when you exit the Online Monitor. If the data collector is active, you have the choice of leaving the selected tasks active or not when you exit the Online Monitor.

```

FPEMGP01           Asynchronous Task Termination
Command ==>> -----

Press Enter to exit the Online Monitor or request Exit or Cancel to
return to the Online Monitor Main Menu.

Select an Exception task to keep it active. All non-selected tasks
will be terminated when leaving the Online Monitor as well as all the
DB2 traces started by the Collect Facility will be stopped.

Exception/Collect tasks           Status
-----
Periodic Exception                Not active
Exception Event                   Not active

run 3pm today                     Active
run at 10 today                   Active
Collect Task C                    Not active
Collect Task D                    Not active
Collect Task for BPA              Not active
F1=Help   F2=Split   F3=Exit   F9=Swap   F12=Cancel
F17=Collect

```

Figure 11. Asynchronous Task Termination panel

To exit the Online Monitor from this window, press Enter. If the data collector is active, any tasks selected in this window remain active.

To return to the Online Monitor Main Menu, press F3 (Exit) or F12 (Cancel). All active tasks remain active.

Function key default settings

Use the function key default settings to view help, move between panels, or access certain Online Monitor and ISPF functions with a single keystroke.

Note:

- The Online Monitor is an ISPF application, and all normal ISPF behaviors apply.
- The default function key settings in help panels follow the standard conventions for help in the ISPF environment.

Table 2 on page 13 describes the Online Monitor function keys, their default settings, their functions, and the Online Monitor panels on which they are available.

Table 2. Function key default settings			
Function key	Default setting	Function	Online Monitor panels
F1	HELP	Used to view help information for a panel or field.	All panels
F2	SPLIT	Used to divide the display into two logical displays separated by a horizontal line and starts another ISPF session, or changes the location of the horizontal line.	All panels
F3	EXIT	Used to validate and save the data entered in a panel, exit the panel, and return to the previous panel.	All panels that contain a command line
F4	PROMPT	Used to view a list of possible values for a field. The prompt fields are followed by a plus sign (+).	Some Collect Report Data panels and Exception Threshold Field Details panel

Table 2. Function key default settings (continued)

Function key	Default setting	Function	Online Monitor panels
F5	AUTO	Used to refresh the data shown in the current panel periodically.	All Thread and Statistics panels
	ADD	Used to add a new exception threshold entry.	Exception Threshold Field Details panel
F6	HISTORY	Used for viewing historical data in panels. Note: If your installation has installed the Online Monitor without the data collector, this function is not available.	All Thread Activity, Statistics, and System Parameters panels
	DELETE	Used to delete an exception threshold entry.	Exception Threshold Field Details panel
F7	UP	Used to scroll toward the top of the data.	All scrolling panels
F8	DOWN	Used to scroll toward the bottom of the data.	All scrolling panels
F9	SWAP	Used to switch between ISPF sessions.	All panels
F10	QUALIFY	Used to filter the threads listed in the Thread Summary panel.	Thread Summary panel
	DELTA	Used to begin delta processing mode that calculates the statistics values between the last two times you pressed Enter.	All Statistics panels
	PREVIOUS	Used to display the previous exception threshold entry for a field.	Exception Threshold Field Details panel
F11	SORT	Used to sort the threads listed in the Thread Summary panel.	Thread Summary panel
	INTERVAL	Used to establish a base point in time from which statistics are to be calculated.	All Statistics panels
	NEXT	Used to display the next exception threshold entry for a field.	Exception Threshold Field Details panel
F12	CANCEL	Used to exit a panel without saving the entries, and return to the previous panel canceling all related panels.	All panels
F14	PURGE	Used to purge a thread currently processing in the Db2 subsystem you are monitoring. You require the Db2 privilege to perform CANCEL THREAD.	All Thread Activity panels
F16	LOOK	Used to view the following exception information and authorization failures: <ul style="list-style-type: none"> • Periodic exceptions • Periodic exceptions messages • Display exceptions • Authorization failure summary • Exception event summary • Exception event messages 	All panels (except LOOK and related panels)

<i>Table 2. Function key default settings (continued)</i>			
Function key	Default setting	Function	Online Monitor panels
F17	COLLECT	Used to display the Collect Report Data panel, where you can collect specific Db2 instrumentation data and direct this data to a data set.	All panels (except Collect Report Data panels)
F18	EXPLAIN	Used to explain the access path methods chosen by Db2 for a given SQL statement.	All Thread panels that show the SQL statement
F19	LEFT	Used to scroll toward the left.	Thread Summary panel
F20	RIGHT	Used to scroll toward the right.	Thread Summary panel
F21	EXPAND	Used to expand a field content that is too long to fit in a panel into a separate window where the entire field content can be shown.	All panels that contain information of a length that cannot be shown in the available panel space.
F22	LEFT	Used to scroll through a field content that is too long to fit in a panel, if the cursor is positioned on such a field.	All panels that contain information of a length that cannot be shown in the available panel space.
F22	RIGHT	Used to scroll through a field content that is too long to fit in a panel, if the cursor is positioned on such a field.	All panels that contain information of a length that cannot be shown in the available panel space.

You can use the following ISPF commands to alter the function key settings:

KEYLIST

Changes the function key settings.

FKA

Alternates between the function key long display format, short display format, and no function key display at all.

PFSHOW

Toggles on and off the display of function key settings.

PFSHOW TAILOR

Specifies how function keys are displayed.

All function keys correspond to a command entered on the command line. For example, typing the command CANCEL on the command line is the same as pressing the F12 (Cancel) key.

You can use the command line with the function keys to enter function parameters. To do this, type the parameters on the command line and press the function key to activate the command.

You can use the ISPF PANELID command to choose whether to display a panel ID on the top line of the panels. To toggle the panel ID display, type PANELID on the command line and press Enter.

Issuing Db2 commands

Use the command Db2 to execute any Db2 command during your Online Monitor session. You can enter this command on the command line of any Online Monitor panel.

The command syntax is:




Figure 12. Syntax of the Db2 command

The Db2 command accepts the following parameter:

command text

The Db2 command you want to execute.

For example:

```
DB2 DISPLAY THREAD(*)
```

Type DB2 and the *command text* on the command line and press Enter. The Db2 Command Output window is overlaid on the current panel, where you can view the Db2 command output. An example of the Db2 Command Output window is shown in Figure 13 on page 16.

```

ROW 1 TO 5 OF 5
PM01DLOC          DSN1 V10          GROUP001 MEMBER01

  DGMCDCO          DB2 Command Output          ROW 1 TO 12 OF 12
  ---
  DSNV401I _ DISPLAY THREAD REPORT FOLLOWS -
  DSNV402I _ ACTIVE THREADS -
  NAME      ST A  REQ ID          AUTHID  PLAN    ASID  TOKEN          N/P
  DB2CALL   T    3  USER030      USER030 DB2PMOM 002B  25             N/P
  DB2CALL   T    5  USER001      USER001 DB2PMOM 002E  10             N/P
  DB2CALL   T    6  USER023      USER023 DB2PMOM 0031  10
  DB2CALL   T   10  USER027      USER027 DB2PMOM 0082  20
  DB2CALL   T    5  USER040      USER040 DB2PMOM 0027  15
  DB2CALL   T   50  USER009      USER009 DB2PMOM 002D  15
  DB2CALL   T    4  USER009      USER009 DB2PMOM 002D  22
  DISPLAY ACTIVE REPORT COMPLETE
  DSN9022I _ DSNVDT '-DISPLAY THREAD' NORMAL COMPLETION
  ***** Bottom of data *****

Command ==>
  F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
  F12=Cancel F16=Look   F17=Collect

Command ==> DB2 DISPLAY THREAD(*)
  F1=Help   F2=Split  F3=Exit   F5=Auto   F6=History F7=Up
  F8=Down   F9=Swap   F10=Qualify F11=Sort  F12=Cancel F16=Look
  F17=Collect F19=Left  F20=Right  F22=Purge

```

Figure 13. Db2 Command Output window

If you issue the command Db2 without any parameters, the Db2 Command window is displayed.

```

DGOMCCMD          DB2 Command
Enter DB2 Command below:
START TRACE(MON) _____
-----
-----
-----

Place cursor on choice and press Enter to retrieve command

=> start trace(perfm) class(6)
=> display thread(*)
=> display trace
=>
=>
=>
=>
=>
=>
=>

Command ==> _____
F1=Help      F2=Split    F3=Exit      F9=Swap      F12=Cancel   F16=Look
F17=Collect

```

Figure 14. Db2 Command window

This window displays command lines on which you can issue Db2 commands. The window shows a list of the last ten Db2 commands that you entered.

You can retrieve a command that is stored in the list by positioning the cursor under the command and pressing Enter. Then, the command is shown on the command line, where you can edit and then submit the command by pressing Enter again.

If a new or modified command is submitted, the command is added to the top of the list. If the command was not edited before submission, or the command was not submitted at all, the list is not updated.

Note: You cannot abbreviate the command Db2.

Using the COLLECT command

Use the COLLECT command to display the Collect Report Data panel, where you can collect specific Db2 instrumentation data and direct this data to a data set.

The command syntax is:

```
▶▶ COLLECT ◀◀
```

To issue the COLLECT command, type COLLECT on the command line and press Enter. COLLECT can be truncated to a minimum of COL.

For more information about collecting report data, see the *Reporting User's Guide*.

Using the OPTIONS command

Use OPTIONS to display the Session Options window, where you can set several options controlling the environment of your Interactive Report Facility (IRF) session.

The command syntax is:

```
▶▶ OPTIONS ◀◀
```

Figure 15. Syntax of the OPTIONS command

To issue the `OPTIONS` command, type `OPTIONS` on any command line and press Enter. `OPTIONS` can be truncated to a minimum of `OPT`.

Correlation ID translation

The correlation ID is a Db2 field that identifies the task executed by Db2.

The correlation ID contains:

Batch jobs

Jobname

TSO applications

Original authorization ID (the logon user ID)

Applications using the Db2 call attachment facility

Original authorization ID (the logon user ID)

CICS® transactions

Connection type, thread type, thread number, and the transaction ID

IMS applications

PST number and PSBNAME of the application

RRS applications

The character string provided by the application during signon

Particularly for CICS and IMS it is useful to break the correlation ID into several parts, so that you can easily distinguish the transaction ID (for CICS threads) from the PSBNAME (for IMS threads).

The default translation

OMEGAMON for Db2 PE breaks the correlation ID into parts by translating the correlation ID into two separate identifiers, the *correlation name* and the *correlation number*. Unless it was changed in your installation, this translation is based on the connection type of the thread.

Connection Type	1	2	3	4	5	6	7	8	9	10	11	12
Batch	Correlation name: job name								Correlation number: blank			
TSO, Db2 call attach	Correlation name: original authorization ID								Correlation number: blank			
CICS	Correlation number: pool thread				Correlation name: transaction ID							
IMS	Correlation number: application PSBNAME				Correlation name: application PST							
RRS	Correlation name: the first 8 characters of the correlation ID provided by the application during signon								Correlation number: the remaining 4 characters			

Changing the default translation

The correlation translation information is kept in the member `CORRDATA` of the `DPMPARMS` data set.

Each record in the `CORRDATA` member specifies the translation that is to be used for a specific connection ID. The connection ID is used here, not the connection type.

The translation is expressed as:

- Offset where the correlation name starts
- Length of the correlation name

- Offset where the correlation number starts
- Length of the correlation number

If OMEGAMON for Db2 PE does not find the connection ID for a given thread in the CORRDATA member, the default translation is used.

To activate the tailored correlation translation, issue the OPTIONS command and specify the DPMPARMS data set in the Session Options window.

How large and missing values are displayed

Values shown in panels are either total values or average values.

If there is insufficient space to display a value in a field, a rounded value is shown followed by one of the following letters to indicate magnitude:

K

thousand (kilo - 10^3)

M

million (mega - 10^6)

G

billion (giga - 10^9)

TT

trillion (tera - 10^{12})

The letter is displayed directly after the number, without blank spaces. There can, however, be decimal places, as in the following examples:

- 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 in 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.

Elapsed time formats

This section shows the different time formats in which time values are presented.

- *dd hh:mm:ss.ffffff*, where:

dd

represents days

hh

represents hours

mm

represents minutes

ss

represents seconds

ffffff

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.ffffff*, where:

sssssss

represents seconds

ffffff

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.

Date formats

The date format for the Online Monitor is taken from your ISPF environment, and is configured during ISPF installation.

Displaying long names and values in scrollable fields

Certain identifiers, such as authorization IDs and program and collection names, can be up to 128 characters long. In general, if an identifier or any applicable field value is too long for the space available in a panel, the Online Monitor displays the information in a so-called scrollable field. To view the non-visible section of a scrollable field, you can place the cursor on the field and use an assigned function key to scroll through the information.

Figure 16 on page 21 shows a fictitious example of a panel with possible variations of scrollable fields. Relevant parts of the panel are in italics. The string *abcdefghijklmnopqrstuvwxy* represents a hypothetical field value that is too long to fit into the available space (of whatever field).

- Variation 1 shows that the string is left-aligned and truncated at the right side (*uvwxyz* is not shown). A plus sign (+) follows the string, which indicates that more of the field's content is available on the right side.
- Variation 2 shows that the string is right-aligned and truncated at the left side (*abcdef* is not shown). A minus sign (-) follows the string, which indicates that more of the field's content is available on the left side.
- Variation 3 shows that the string is centered and truncated at the left and right sides (*abc* and *xyz* are not shown). A plus sign and a minus sign (-+) follow the string, which indicates that more of the field's content is available on both sides.

```

DGOMTWPK                               Current Package

                                           More:    +
Location . . . . . : PM01D711
Collection ID . . . . . : abcdefghijklmnopqrst +
Program name . . . . . : ghijklmnopqrstuvwxyz -
Consistency token . . . . . : defghijklmnopqrstuvw -+
----- QPACPKID+0 -----
| ISPEXPND                               Line 1 of 1 |
| abcdefghijklmnopqrstuvwxyz                |
| |                                         |
| Command ==>                               Scroll ==> CSR |
| F1=Help   F2=Split   F3=Exit   F5=Rfind   F6=Rchange   F7=Up   |
| F8=Down   F9=Swap    F10=Left  F11=Right  F12=Cancel  |
-----
F7=Up      F8=Down    F9=Swap    F12=Cancel  F14=Purge
F16=Look   F17=Collect F21=Expand F22=Left   F23=Right

```

Figure 16. Fictitious panel with scrollable field variations

The plus and minus signs next to fields act as *scroll indicators*.

- If shown, they indicate that the content of a field is shown in part. No scroll indicator is shown if the complete field content can be presented.
- Whether more of a field's content is available on the left, on the right, or on both sides. Minus (-) associates the left side, plus (+) associates the right side.

Initially, all field values are shown left-aligned, which means that only the plus sign is shown. The minus sign is shown after you scrolled through a field by means of the function keys described next.

To view non-visible sections of an appropriate field, you need to position the cursor on the field and use the following function keys (the default function key settings are assumed).

- F21 (Expand) shows the field's content in a separate window, which is overlaid on the current panel. If the content exceeds the available width of this window, it continues on the next line. Pressing F3 (Exit) or F12 (Cancel) closes the window.
- F22 (Left) moves the begin of the shown field content toward the start of the actual content. The increment is the length of the available space, until the shown field content is left-aligned in the available space.
- F23 (Right) moves the end of the shown field content toward the end of the actual content. The increment is the length of the available space, until the shown field content is right-aligned in the available space.

These keys can be used in any meaningful order, provided the cursor is positioned on a scrollable field. Otherwise, pressing a key has no effect.

The mechanism described so far is applicable to *input fields* and *output fields*. When you type information in a field, you can use the same function keys to position the content of a field. However, a subtle distinction should be noted: In output fields potentially existing trailing space characters are treated as nonexistent. This means, you cannot accidentally scroll to a non-visible section of the field content. In input fields trailing space characters are considered valid. This means, when you scroll through an input field with a long sequence of space characters, the field might appear empty but in fact is not.

Finally, you can get help about the use of scrollable fields by positioning the cursor on the scroll indicator area (reserved for the -+ indicators) and pressing F1 (Help).

Chapter 3. Monitoring and tuning

This section describes how to monitor a Db2 subsystem, determine performance problems by using batch reports and the Online Monitor functions of OMEGAMON for Db2 PE, and how-to and tuning tasks with the ISPF Online Monitor. Such tasks include viewing past data and statistics, monitoring exceptions and thread activities, explaining SQL statements, and evaluating Db2 system parameters.

You can use OMEGAMON for Db2 PE for continuous or periodic monitoring of a Db2 subsystem and for determining specific performance problems in Db2. OMEGAMON for Db2 PE can help you in:

- Determining how an application will perform or is performing over a period of time
- Indicating where there are tuning opportunities in your system
- Distinguishing among subsystem and application problems
- Monitoring an application in a detailed manner so you can identify problem areas
- Analyzing constraints acting on an application
- Determining the performance effects of any adjustments made within the Db2 subsystem
- Determining whether performance objectives are being met

OMEGAMON for Db2 PE offers different ways to monitor your subsystem depending on whether you want to see current or past activity. Use the Online Monitor to monitor an active subsystem as well as to view events that happened in the recent past. Use the Batch reports to examine performance problems in the more distant past and trends over a period of time.

The amount of data generated for monitoring a Db2 subsystem is vast, so limiting the amount of data to show only potential problem areas is essential. OMEGAMON for Db2 PE offers several ways of reducing the amount of data that needs to be examined, the most important is exception processing. Exception processing makes it easy for you to focus on possible performance problems by highlighting data that is outside defined thresholds. It is available in Batch reporting and the Online Monitor.

A good approach is to monitor an active Db2 subsystem using Online Monitor exception processing (you do not need to be logged on to have exception processing running) and to regularly generate Statistics and Accounting exception reports. Online Monitor exception processing alerts you to performance problems as soon as they occur, and the Accounting and Statistics reports give you a detailed picture of application and system performance over a period of time.

Deadlock and timeout participant details are available online through exception event processing. Consider generating deadlock and timeout traces regularly, because in this way information is available to help you investigate any locking problems in detail. There is no significant performance overhead on the Db2 side in collecting the data for these reports.

The best way to investigate performance trends is by producing Accounting and Statistics reports that are ordered by interval.

To detect problems as they occur, use the Online Monitor periodic exception processing. When you detect poor thread performance, you can examine the comprehensive performance data that is shown in Thread and Statistics panels. If you assume the problem is caused by SQL, you can analyze the access path using the online explain function.

If the problem occurred in the recent past, you can use the Online Monitor HISTORY command to view the events surrounding the problem without having to re-create it.

If the panels do not provide enough information to solve the problem, you can use the Online Monitor to collect instrumentation data for batch reports. You can specify the data collection to be triggered by exception thresholds; in this way you can minimize the time high-volume and high-cost traces are active and ensure that the data needed is collected at the right time.

OMEGAMON for Db2 PE provides a comprehensive set of reports with different levels of detail and for different areas of performance. This represents a top-down approach to problem determination: the most

generic reports indicate the problem area and, if necessary, more details can be shown to narrow down the cause of the problem.

Changed access path to the ISPF Online Monitor

When you start IBM Tivoli® OMEGAMON for Db2 Performance Expert on z/OS from ISPF (Start EXEC FPEJINIT), the main menu is displayed.

```
FPEFMENU      IBM OMEGAMON for Db2 Performance Expert on z/OS

Select one of the following.

--  1. Create and execute reporting commands
    2. View online DB2 activity - Classic Interface
    3. Trace collection, Explain, and Admin functions
    4. Maintain parameter data sets
    5. Customize report and trace layouts
    6. Exception profiling

Command ===> _____
F1=Help   F2=Split  F3=Exit   F9=Swap   F12=Cancel
```

Figure 17. Main menu

Note: The OMEGAMON for Db2 PE options shown on this menu are described in detail in [Monitoring Performance from the OMEGAMON Classic Interface](#).

Performance data generation and online monitoring

Db2 generates trace data about its own performance and events in Db2 subsystems. However, Db2 does not provide any reporting facilities for analyzing this data. The Online Monitor provides you with the capability to view an active Db2 subsystem and identify performance problems online.

The Online Monitor displays subsystem-wide performance information, such as processor times, buffer pool usage, locking, log, and I/O activity. For an individual thread, the Online Monitor displays information such as the elapsed time, the time spent in Db2, the time it was suspended, the read and write activity involved, the locks obtained, and the SQL statements executed.

The Online Monitor displays subsystem-wide Db2 performance information in a comprehensive form that is easy to understand and analyze. You can use the Online Monitor to:

- Determine total Db2 system performance and efficiency
- Measure an application's performance and resource use
- Evaluate an application's effect on other applications and the system
- Analyze and improve SQL statements
- Identify potential problems
- Determine tuning requirements for Db2.

When changes are made to an application or to the Db2 subsystem, the Online Monitor can help you determine the effects. This is very important for determining whether the changes increased or decreased performance.

When Db2 performance is not satisfactory, the Online Monitor can help you identify areas where tuning is required to optimize the performance of Db2. The Online Monitor can log Db2 activities and events and provide this information for later viewing to assist you in determining the cause of potential problems.

For a long-term view of Db2 performance, your needs are best served by the batch reporting capabilities.

The Db2 operating environment

The performance of a Db2 transaction or query is dependent not only on the performance of the Db2 subsystem, but also on the performance of the transaction manager, such as CICS or IMS, and the MVS system itself. Therefore the environment in which the Db2 subsystem is operating should be tuned before Db2 is tuned.

For example, if the MVS system is overloaded, tuning a Db2 subsystem is unlikely to improve Db2 system performance. Db2 performance can only be improved by reducing or balancing the load of the MVS system.

Specialized tools are available to monitor the different system components:

- CICSplex System Manager for CICS
- IMS/VS DC Monitor or IMSPARS for IMS
- RMF for MVS

The relationship between the different systems and performance tools is complex, however, it is not within the scope of this documentation. For more information, see [IBM Db2 for z/OS in the IBM Knowledge Center](#). It is assumed that the environment in which the Db2 subsystem is operating is well tuned.

Performance objectives and exception processing

Before you can start monitoring the system, define your performance objectives on the basis of the business needs, the workload for the system, and the resources available. Typically, the objectives would include acceptable response times, average throughput, and system availability.

These objectives are usually formalized in service-level agreements between the users and the data processing groups in an organization. The agreements can include expectations of query response times and transaction throughput.

You can monitor how well these objectives are being met.

The most efficient way to do this is to set limits, exception thresholds, for key fields that reflect your performance objectives using exception processing.

Exception profiling can assist you in establishing exception thresholds. This facility sets exception thresholds automatically based on your application configuration. For reports, the Accounting TOP subcommand option is also useful in determining Accounting exception thresholds.

For example, you can monitor response times by setting exception thresholds for class 1 and class 2 elapsed times to reflect the acceptable response times for your environment. Class 1 elapsed time shows the thread time (from thread creation to thread termination) and class 2 time shows the time Db2 spent processing SQL statements.

Monitoring performance

The key to effective performance monitoring is in identifying unusual situations and thereby limiting the amount of data that needs to be examined. In addition to exception processing, OMEGAMON for Db2 PE offers several other ways of filtering the data and highlighting potential problems.

The following options are available in Batch reporting:

- Filtering data by date and time (FROM and TO subcommand options)
- Filtering data by identifiers such as user ID (INCLUDE and EXCLUDE subcommand options)
- Filtering data by resource usage (TOP subcommand option)
- Ordering data on reports by interval (INTERVAL subcommand option)
- Summarizing and sorting data on SQL reports
- Tailoring report layouts for your own needs. The User-Tailored Reporting feature (UTR) and its use is described in detail in the [Reporting User's Guide](#).

The following options are available in the Online Monitor:

- Qualify and sort functions
- History data collection with qualifications

Monitoring using reports

You can filter the data shown in reports by using the INCLUDE and EXCLUDE subcommands to show, for example, information only for certain plans, authorization IDs, or locations. You can use the FROM and TO subcommands to filter data within specified times.

Another way to limit the data that needs to be examined is to use the TOP subcommand option to obtain a high water mark type of reporting on resource usage. The TOP lists, printed at the end of an Accounting report or trace, can identify the threads or users that have required the most use of the resources specified in the TOP subcommand option. Alternatively, you can use the TOP subcommand option with the ONLY keyword to filter entries based on resource usage and produce a report that shows only entries with the highest resource usage.

You can summarize data for certain periods by ordering your Accounting and Statistics reports by interval. For example, you can summarize data for the peak periods during the day.

If you produce SQL Activity reports and traces, potential problems can be highlighted by sorting and summarizing the information within the report and trace entries by several criteria. For example, a problem cursor can be identified by summarizing SQL activity by cursor and ordering the cursors by TCB times. The sorted entries provide an easy way to identify SQL statements that might be causing performance problems.

You can tailor your own report layouts and trace layouts with the User-Tailored Reporting feature (UTR). Use UTR to control the volume, contents, and layout of your traces and reports. The User-Tailored Reporting feature (UTR) and its use is described in detail in the [Reporting User's Guide](#).

Observing performance trends

You can use reports and graphs to summarize data over periods of several days, weeks, or months to observe trends in performance. Pay special attention to peak periods of activity, for new applications and for the system as a whole. During peak periods, constraints and response-time problems are most evident.

Some trends to look for are:

- Increases in response times, number of I/Os, resource contention, and processor usage
- Changing workload patterns over a period
- Changes in the transaction distribution and frequency
- Changes in the SQL activity pattern

Determining performance problems

When you find that there are performance problems when you are monitoring the system, you can use several panels and reports to investigate the cause of the problems.

Problems detected in periodic exception processing

To view the past events surrounding the problem online, you can use the HISTORY command. The amount of available historical data is determined by installation-defined options. Note, however, that the batch reports are more comprehensive than the Online Monitor panels. Therefore, if the panels do not provide enough information to solve a problem, the Online Monitor Collect Report Data function should be used to gather information in a data set for input to the batch reports.

Problems detected in exception event processing

If you detect a problem using the online exception event processing, collect the appropriate trace data to produce Locking reports, I/O Activity reports, Audit reports, or Record Trace reports. This section describes how to proceed with specific exception events.

Deadlock or Timeout

If deadlocks or timeouts occur too often, generate a Lockout report to see which applications and objects are affected. Having identified the objects and applications causing the deadlocks or timeouts, use EXPLAIN to understand the locking behavior of the SQL statements or consider reorganizing the database.

EDM Pool Full

First check online or batch statistics to obtain more information about the EDM Pool situation. More details are provided in the I/O Activity EDM Pool report.

Authorization Failure

If authorization failures occur too often, generate an Audit authorization failure report for details.

Thread Commit Indoubt

Run a Record trace on the Statistics Class 4 IFCIDs to see details of communication problems. These are likely to be either VTAM® or Db2 internal problems.

Coupling Facility Rebuild

Use the COLLECT command to automatically start tracing IFCID 268 (CF rebuild end) when a coupling facility rebuild starts, and run a Record trace for this IFCID. You can get more details about coupling facility behavior from RMF reports.

Problems detected in exception reports

If you use an Accounting exception report to monitor your system, produce this report using the TOP or the INTERVAL subcommand option so that you can immediately focus on potential problem areas.

If you use an Statistics exception report to monitor your system, produce this report using the INTERVAL subcommand option so that you can immediately focus on potential problem areas

Sometimes, however, you require more detailed reports to determine the exact cause of a problem.

System problems

If exception processing indicates problems in system-wide resource usage and a Statistics trace does not clarify the reason for the problem, but points to EDM pool or logging activity, consider running I/O Activity reports. Or, if the Statistics trace indicates a problem with binds, generate Utility Activity reports. If the number of deadlocks is high, run Locking reports. If none of these report sets offer adequate information to determine the cause of the problem, you can run a Record trace to format the individual instrumentation records.

Application problems

If exception processing indicates an application-related problem, which is typically indicated as an elapsed time problem for an application or a user, use Explain reports to determine the access path of the suspected plan. In addition, if necessary, generate Accounting traces by using the TOP subcommand option.

- If the Explain reports and the Accounting traces do not identify the reason for poor SQL performance, use SQL Activity reports for detailed information about the specific statement, such as scans or I/O per page set or sort specifics.
- If the Accounting traces indicate a locking problem, run Locking reports.
- If the problem seems to be with binds or Db2 utilities, run Utility Activity reports.

As with system-related problem determination, you can run a Record trace if none of these report sets provide adequate information.

Viewing past data

You can recall and view statistics, thread activity, and system parameters data previously collected by the data collector. The data is gathered at installation-defined intervals.

You can view past data by issuing the HISTORY command or pressing F6 (History) in any panel that can display past data.

When viewing past data, the word HISTORY is displayed on the line following the heading of the current panel.

```
DGOMTPSM 11:57          Thread Detail          PM01DL0C DSN1 V10
                        GROUP001 MEMBER01 HISTORY 05/30/08 11:45:00
For details, place any character next to heading, then press Enter.
```

Figure 18. Thread Detail panel showing history active

The date and time following the word HISTORY indicate when the data being displayed was collected.

Note: You can view past data only if the data collector is installed at your site.

Your installation can choose to collect data at the following levels:

- Thread Summary
- Thread Detail
- Thread Detail with locking information
- Thread Detail with SQL statement
- Thread Detail with locking information and SQL statement
- Statistics
- System Parameters

Historical data is kept in either a VSAM data set or data space. The data is gathered by the data collector. When the data set or data space is full, the data is written to the beginning again, writing over the earliest data gathered.

The availability of the data is limited by the collection rate, the size of the data space or data set used, the amount of historical data to be maintained, and the thread history qualification definitions (all of these are defined in the data collector startup parameters).

The data collector parameters can also be modified by the administrator user.

The syntax of the HISTORY command is:




Figure 19. Syntax of the HISTORY command

The HISTORY command accepts the following parameters:

OFF

Returns the display to the current time.

BACK

Displays data for the previous interval. If no more data is available, data from the earliest available interval is displayed.

FORWARD

Displays information for the next available interval. If no data is available, data from the current time is displayed and processing continues as if HISTORY is OFF.

date

The date of the required information. If no date is specified, today's date is assumed.

hh:mm:ss

The time of the required information. If no time is specified, the default is the time of the earliest available history record for the specified date. The seconds do not need to be specified.

Note: HISTORY can be truncated to a minimum of HIS.

If no data is available for the specified date and time, the next available record is displayed.

When a panel is refreshed while viewing past data, the panel is updated with data from the next stored record, just as if you had entered HISTORY FORWARD on the command line.

If no parameters are specified, the History window is displayed.

```
DGOMHWHY                History
                          05/09/08  14:03
Date . . . . . 03/16/08
Time . . . . . 14:03:24
F1=Help   F2=Split   F9=Swap
F12=Cancel
```

Figure 20. History window

The current system date and time are the defaults for the History window.

Note: When monitoring with HISTORY OFF, it is possible to view data that cannot be viewed by using the HISTORY command. This occurs when the “current moment” does not coincide with the history collection interval. For example, when the history interval is 10 minutes in length and data is collected at 8:30 and 8:40, the data retrieved at 8:32 from DB2 is not subsequently retrievable using the HISTORY command.

A sample JCL is provided in library RK02SAMP (DG0MMJCL) that can be used to unload history VSAM data sets created by the data collector. After unloading, you can use the data as input to the batch Record trace function for problem determination. The sample JCL might need to be modified.

Exception processing

Exception processing is the most effective way of identifying performance problems. The first step in monitoring your Db2 system should always be to start thread and statistics exception processing.

Use Online Monitor exception processing to identify Db2 thread and statistics fields with values outside defined thresholds. This allows better management of service levels by identifying problems in the Db2 subsystem and threads causing performance problems.

By using display exception processing and periodic exception processing, you can monitor and identify:

- Threads that might be experiencing problems
- Subsystem-wide Db2 conditions that might be causing performance problems

By using exception event processing, you can monitor the following events:

- Deadlock
- Timeout
- EDM pool full

- Authorization failure
- Thread commit indoubt
- Coupling Facility (CF) rebuild/alter start
- CF rebuild/alter end
- Global trace started

The threshold values for thread activity and statistics fields are specified in the Exception Threshold data set. When exceptions are detected during your monitoring session, you are notified so that appropriate action can be taken. In addition, you can activate a user exit that can automatically trigger any immediate reaction, for example, issue an alert to NetView®.

You can view information about any of the exceptions using the Online Monitor **LOOK** command.

With exception processing, you can create the following two data sets for later analysis:

- The Exception Log data set (to print a list of exceptions or to load exception data into Db2)
- The Exception DPMOUT data set (records that had exceptions)

Exception processing modes

This section describes the three basic types of exception processing available: display exception processing, periodic exception processing, and exception event processing.

Display exception processing

Use display exception processing to monitor the occurrence of a specific exception. This exception processing mode operates in the foreground of Online Monitor processing. With display exception processing you can view thread activity exceptions, and statistics exceptions in interval or delta processing mode.

With display exception checking, fields shown in the current panel are checked whenever the display is refreshed with new or historical data. Fields with exception conditions are shown in reverse video, and the color of the field indicates the level of the exception. *Warning* level exceptions are highlighted in yellow, while *Problem* level exceptions are highlighted in red. These colors can differ if you have changed your ISPF default colors. Selection fields in the detail panels are also shown in reverse video if any of their lower level windows contain fields in exception status. The exception notifications are stored and can be examined using the LOOK command.

Periodic exception processing

Use periodic exception processing to periodically monitor thread activity fields and statistics fields for exception conditions. This processing mode runs in the background of Online Monitor processing.

With periodic exception processing, data is tested for exception whenever the interval that you specified has elapsed (whether or not you are viewing the relevant data). If any *Problem* level exceptions are detected, the Exception Notification window is displayed. If only *Warning* level exceptions are detected, a message is displayed to notify that a periodic exception has occurred. This window is overlaid on the current panel and shows the number of periodic exceptions detected during the interval. The exception notifications are stored and can be examined using the LOOK command.

If the data collector is active, periodic exception processing can continue when you exit the Online Monitor. When you use the Online Monitor the next time, any periodic exceptions that were detected while you were not logged on are displayed. Periodic exception processing is not terminated until you stop it or until the data collector itself is terminated. If the data collector is not active, periodic exception processing terminates when you exit the Online Monitor.

Exception event processing

Use exception event processing to monitor the Db2 subsystem for the occurrence of particular events. This processing mode runs in the background of Online Monitor processing. The following events can be monitored:

- Deadlock
- Timeout
- EDM pool full
- Authorization failure
- Thread commit indoubt
- CF rebuild/alter start
- CF rebuild/alter end
- Global trace started
- Data set extension
- Unit of recovery problem
- Log space shortage

The events must be specified with the EXCEPTIONEVENT data collector parameter or in the Data Collector Parameters window before exception event processing can be activated from the Exception Processor panel. For more information, see the *IBM Db2 for z/OS in the IBM Knowledge Center*.

When an exception event is detected, the Exception Notification window is displayed to notify you that an exception event has occurred. Exception event notifications are stored and can be examined using the LOOK command.

Past data and exceptions

Display exception processing operates with past data just as if the data had been retrieved directly from Db2. Exceptions are shown in reverse video and logged as normal.

Periodic exception processing and exception event processing do not report on past data. They report on current data, whether or not past data is currently being displayed.

Exception Notification window

The Exception Notification window is displayed whenever periodic exception or exception event processing is active and either a problem level exception or an event exception is detected. This window is overlaid on the current panel and shows the number of periodic *problem* and *warning* level exceptions, and the number of exception events since the last exception notification or since exception processing was started.

If exception event processing is active, the Exception Notification window is shown, as in the following figure.

```

DGOMEPE1                               Exception Notification
Time . . : 04/14/08 12:42:32

Periodic Exceptions
Problem . . . . . : 0
Warning . . . . . : 0

Exception Events
Deadlock . . . . . : 0
Timeout . . . . . : 0
EDM Pool Full . . . . . : 0
Authorization Failure . . . . . : 1
Thread Commit Indoubt . . . . . : 0
CF rebuild/alter start . . . . . : 0
CF rebuild/alter end . . . . . : 0
Global trace started . . . . . : 0
Data set extension . . . . . : 0
Unit of Recovery problem . . . . . : 0
Log space shortage . . . . . : 0

F1=Help      F2=Split      F9=Swap
F12=Cancel

```

Figure 21. Exception Notification window (exception event processing active)

If exception event processing is not active, the Exception Notification window is shown, as in the following figure.

```

DGOMEPEX                               Exception Notification
Time . . : 04/04/08 11:53:26

Periodic Exceptions
Problem : 1
Warning : 2

F1=Help      F2=Split      F9=Swap
F12=Cancel

```

Figure 22. Exception Notification window (exception event processing inactive)

To exit the Exception Notification window and return to the panel you were viewing, press Enter or F12 (Cancel).

The Periodic Exceptions section of this window shows the number of Problem and Warning level periodic exceptions since the last time you were notified of an exception.

The Exception Events section of this window shows the number of exception events that occurred for various events since the last time you were notified of an exception. This field categorizes the seven different classes of exception events that can occur.

How to define exception threshold values

This information shows where exception thresholds are defined.

This information has been consolidated in the [Reporting User's Guide](#) for consistency reasons.

How to start exception processing

Exception processing is started from the Exception Processor panel. All exception processing functions are activated from this panel.

Figure 23 on page 33 shows the Exception Processor panel.

Alternatively, exception processing can automatically be started when the OMEGAMON Collector is started.

Starting exception processing when the OMEGAMON Collector is started

When the OMEGAMON Collector is started, its startup parameters determine whether exception events are to be processed. Dependent on which exception events are specified, the appropriate traces are started. If an Exception Threshold data set with exception criteria is already available, it can be specified in the OMEGAMON Collector startup parameters. By this means, exception processing automatically starts with predefined exception thresholds whenever the OMEGAMON Collector is started.

For more information about the Exception Threshold data set, a sample data set provided with the product, its data set attributes, and how its content can be modified, see the *Reporting User's Guide*.

For more information about OMEGAMON Collector startup parameters, see the topic about configuring OMEGAMON Collector for exception processing in *Monitoring Performance from the OMEGAMON Classic Interface*. The **EXCEPTIONEVENT** startup parameter determines the events, and the **AUTOEXCPTHNAME** startup parameter specifies the name of the Exception Threshold data set that is to be used at startup.

If this means of starting exception processing is used, and if the content of the Exception Threshold data set is modified while the OMEGAMON Collector is running, remember to refresh the environment if you want the new exception criteria to be recognized.

Refer to “How to restart exception processing (REINIT command)” on page 36 for more details.

Starting exception processing from the Exception Processor panel

Use the Exception Processor panel to activate and deactivate various exception processes by selecting exceptions under Activate/Deactivate Exception Processing.

To display this panel, select option 2 (Control Exception Processing) from the Online Monitor Main Menu.

```

DGOMEPO2                Exception Processor                PM01DLOC DSN1 V10
For any field enter any character to activate
                                                                More:  - +
Activate/Deactivate Exception Processing
- Display thread summary
- Display thread detail
- Display statistics detail > Periodic
- Exception event notification
                                                                > User Exit

Options
Periodic units . . . . . 2          1=Seconds
                                                                2=Minutes
Periodic interval . . . . . 10     1-7200 Seconds
                                                                1-120 Minutes

> Disable auto-display for problem exceptions
> Sound alarm for exception warnings
- Log file data set output needed
- DPMOUT data set output needed

Exception threshold data set
Name . . . . . -----

Command ==>
  F1=Help    F2=Split  F3=Exit  F7=Up    F8=Down  F9=Swap
  F12=Cancel F16=Look   F17=Collect
```

Figure 23. Exception Processor panel

Use the 1 field to specify whether the 1 field value is in seconds or minutes. This field is only required for periodic exception processing.

The 1 field specifies how often Db2 instrumentation data is gathered, examined for exceptions, and reported if an exception condition occurs. All periodic exception messages are gathered periodically as specified by this field, and reported when the display is refreshed. After the elapsed period, if any periodic exceptions have occurred, the Exception Notification window is displayed and shows the number of warning and problem exceptions. This field is only required for periodic exception processing.

You can request a Log File data set or DPMOUT data set to be written by selecting `Log file data set output needed` or `DPMOUT data set output needed` from the Exception Processor panel, respectively. When you request a log or DPMOUT file, the Exception Output Data Sets window is displayed. See [“Exception Output Data Sets window”](#) on page 34 for details.

Previously selected fields are indicated by a greater than symbol (>).

The Exception threshold data set field contains the name of the Exception Threshold data set used as input to the exception processor. The data set is built using the Exception Threshold Field Details panel. The Online Monitor exception processor uses the information in this data set to test fields for exception conditions.

Note: A valid Exception Threshold data set needs to be specified in the Exception Processor panel before activating display exception, periodic exception, or exception event processing.

Exception processing user exit

OMEGAMON for Db2 PE supports a user exit to enhance the exception processing capabilities. This exit can handle periodic exceptions and event exceptions. Therefore, you can start the exception user exit for periodic processing, event processing, or both, depending on which kind of exception processing you have activated in the Exception Processor panel.

You can activate the user exit together with periodic exception processing, event exception processing, or both. But you can also activate it when periodic exception processing, event exception processing, or both are already active.

To activate the user exit, type any character in the `User Exit` field. To deactivate the user exit, enter a blank in this field. The `User Exit` field is displayed only if the data collector is active.

Several users can invoke periodic exception processing or event exception processing simultaneously using different threshold data sets. The users can start or stop this user exit independently. The user exit routine can check for the user ID, the exception field name, the field value, or other characteristic items to select individual paths of processing.

OMEGAMON for Db2 PE provides a sample of the exception processing exit, called `DGOMUPXT`, which issues a message to the operator. The message text varies depending on whether the situation is an event exception or a periodic exception. For a description of the different formats of this message, see message number `DGOV0100I` in [Messages](#).

You can modify `DGOMUPXT` according to your needs, for example, to examine the type of exception and perform any action necessary to handle the situation.

For more information, see [IBM Db2 for z/OS in the IBM Knowledge Center](#).

Exception Output Data Sets window

The Exception Output Data Sets window is displayed whenever you have requested a Log File or DPMOUT data set in the Exception Processor panel.


```

DGOMEPO3          Exception Output Data Sets

Specify the data set(s) to be used, and press Enter

Display log file data set
Name . . . . .
Disposition . . 1 1=append 2=overwrite 3=new

Display DPMOUT data set
Name . . . . .
Disposition . . 1 1=append 2=overwrite 3=new

Periodic log file data set
Name . . . . .
Disposition . . 1 1=append 2=overwrite 3=new

Periodic DPMOUT data set
Name . . . . .
Disposition . . 1 1=append 2=overwrite 3=new

F1=Help  F2=Split  F7=Up    F8=Down  F9=Swap  F12=Cancel

```

Figure 24. Exception Output Data Sets window

Use this window to enter the names of the Exception Log data set and DPMOUT data set you want the exception processor to write to. If the data collector is active for your current subsystem, make sure that it is authorized to write to these data sets. Different data sets are used for display exceptions and periodic exceptions.

For data sets selected in the Exception Processor panel, a valid name is required. Do not specify a name for those data sets that were not selected. Therefore, if a LOG data set was selected but not a DPMOUT data set, then you must enter a valid name for the Exception Log data set, but not for the DPMOUT data set.

You can enter the name of an output Log File data set where exception conditions are written. An entry is written to the data set for each exception condition detected by the exception processor. This data set can subsequently be used as input to the Db2 LOAD utility, or printed using the Exception Log print utility.

See [“Printing the Exception Log File data set” on page 41](#) for more information.

For the layout of the Exception Log File data set, see [“Layout of the Exception Log File output record” on page 73](#).

You can enter the name of an output DPMOUT data set where DPMOUT-formatted records that contain exception conditions are written. This data set can be used as input to a batch Record trace or Statistics trace for a more detailed analysis of exception conditions. See the [Report Reference](#) for information about the layout of DPMOUT-formatted records.

The Disposition fields specify how the data is to be written to the data sets. Enter 1 (append) to append the data, 2 (overwrite) to overwrite the old data, or 3 (new) to dynamically allocate a new data set.

Press Enter to initialize exception processing and to activate the selected functions. If there are any errors during this process, a panel is displayed that shows the errors.

Starting display exception processing

Display exception processing is started from the Exception Processor panel under Activate/Deactivate Exception Processing.

The Exception Processor panel is shown in see [Figure 23 on page 33](#).

Enter any character in the following fields to activate display exception processing for the appropriate displays:

- Display thread summary
- Display thread detail
- Display statistics detail

Starting periodic exception processing

Periodic exception processing is started from the Exception Processor panel under Activate/Deactivate Exception Processing.

The Exception Processor panel is shown in [Figure 23 on page 33](#).

Enter any character in the 1 field and specify the periodic exception units and interval under the Options field to activate periodic exception processing.

Starting exception event notification

Exception event notification is started from the Exception Processor panel under Activate/Deactivate Exception Processing.

The Exception Processor panel is shown in [Figure 23 on page 33](#).

Enter any character in the Exception event notification field to activate exception event notification.

The events must be specified with the EXCEPTIONEVENT data collector parameter or from the Data Collector Parameters window of the administrator user dialog before exception event processing can be activated.

How to restart exception processing (REINIT command)

If the values in the Exception Threshold data set are changed during an Online Monitor session, the exception processor needs to be restarted (reinitialized) to load and use the new values.

Exception processing *initialization* occurs when you start display or periodic exception processing using the options in the Exception Processor panel as shown in [Figure 23 on page 33](#). You can *reinitialize* exception processing by stopping and starting exception processing from the Exception Processor panel.

You can also reinitialize exception processing by using the REINIT command. The command syntax is:

➤ REINIT ➤

Figure 25. Syntax of the REINIT command

Type REINIT (or a valid abbreviation, beginning with a minimum of REI) on any command line (except on the Asynchronous Task Termination panel) as long as exception processing is active, and press Enter.

Examining exception messages (LOOK command)

This section describes how to examine exception messages issued during exception processing using the LOOK command.

Whenever an exception occurs, an exception message is written to the appropriate exception list where it can be examined using the LOOK command.

The command syntax is:




Figure 26. Syntax of the LOOK command

The LOOK parameters 1 to 6 correspond to the options on the Look Selections menu in [Figure 27](#) on page 37.

To display the Look Selections menu, type LOOK and press Enter on any command line, or press F16 (LOOK).

```

DGOMLMSP 03:35    Look Selections

Subsystem: PM01DL0C          DSN1 V10

Select one of the following displays

--  1.  Periodic Exceptions
    2.  Periodic Exceptions Messages
    3.  Display Exceptions
    4.  Authorization Failure Summary
    5.  Exception Event Summary
    6.  Exception Event Messages

Command ==>
F1=Help    F2=Split    F3=Exit    F6=History    F9=Swap    F12=Cancel
F17=Collect

```

Figure 27. Look Selections menu

Use the Look Selections menu to reach panels that display exception messages and the status of exception processes. The list panels display the last 500 exceptions and authorization failures that occurred during your Online Monitor session. When this limit is reached, the oldest entries in the list are discarded as new entries are added. When these windows are first displayed, the bottom of the list is displayed and shows the most recent exceptions that occurred.

The following topics describe the windows that can be accessed from the Look Selections menu.

Examining the Periodic Exceptions

To display the Periodic Exceptions List window, select option 1 (Periodic Exceptions) from the Look Selections menu or type LOOK 1 and press Enter on any command line.

```

DGOMLAXP                Periodic Exceptions List                Row 499 to 500 of 500

Periodic Interval started . . . . . : 03/30/08 08:52:04.10
Last Interval . . . . . : 03/30/08 08:53:35.85

   Time      Location      Group  Subsystem  Member  Corrname
   Field     Reqloc          Primauth Planname  Connect Corrnmbr
   Descr
-----
- 08:35:36 PM01DLOC      GROUP001 N/P      MEMBER01
   N/P      N/P      N/P
SLRSUSP 100 > 0 Problem Total
TOTAL ALL SUSPENSIONS

- 13:21:28 PM01DLOC      GROUP001 N/P      MEMBER01
   N/P      N/P      N/P
ADRECETT 0.0 < 100 Problem Commit
ELAPSED TIME IN APPLICATION (CLASS 1)
***** Bottom of data *****

Command ==>
F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel

```

Figure 28. Periodic Exceptions List window

Use this window to view the most recent periodic exceptions that have occurred. You can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

Select any of the entries listed in this window to display either the Thread Detail panel or Db2 Statistics Detail panel as appropriate, where you can examine the field causing the exception. Pressing F3 (Exit) or F12 (Cancel) returns you to the Periodic Exceptions List window. If the selected exception was caused by a thread that is no longer active, past data is automatically retrieved if available to display the thread.

For more information about history, see “Viewing past data” on page 28.

This list is cleared if you exit the Online Monitor without an active data collector running.

Examining the Periodic Exceptions Messages

To display the Periodic Exception Messages window, select option 2 (Periodic Exceptions Messages) from the Look Selections menu or type LOOK 2 and press Enter on any command line.

```

DGOMLAMP                Periodic Exception Messages                ROW 1 TO 2 OF 2

Message
DGOM944 Periodic Exception Processor started at 03/30/08 08:52:04.100
DGOM945 Periodic Exception Processor stopped at 03/30/08 08:53:07.290
***** BOTTOM OF DATA *****

Command ==>
F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel

```

Figure 29. Periodic Exception Messages window

Use this window to view messages issued by the periodic exception processor. You can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

This list is cleared when you exit the Online Monitor.

Examining the Display Exceptions

To display the Display Exception List window, select option 3 (Display Exceptions) from the Look Selections menu or type LOOK 3 and press Enter on any command line.

```

DGOMLSXP                Display Exception List                Row 499 to 500 of 500

Time      Location      Group  Subsystem  Member  Corname
Field     Reqloc              Prmauth Planname  Connect Corrnbr
Descr     Value              Compare Threshold Type      By
-----
08:35:36 PM01DL0C      GROUP001 N/P      MEMBER01
          N/P              N/P      N/P
SLRSUSP  100                >        0        Problem Total
TOTAL ALL SUSPENSIONS

13:21:28 PM01DL0C      GROUP001 N/P      MEMBER01
          N/P              N/P      N/P
ADRECETT 0.0                <        100     Problem Commit
ELAPSED TIME IN APPLICATION (CLASS 1)

***** Bottom of data *****

Command ==>
F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel

```

Figure 30. Display Exception List window

Use this window to view information about the latest display exceptions that have occurred during your Online Monitor session. You can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

This list is cleared when you exit the Online Monitor.

Examining the Authorization Failure Summary

To display the Authorization Failure Summary window, select option 4 (Authorization Failure Summary) from the Look Selections menu or type LOOK 4 and press Enter on any command line.

```

DGOMLAFS                Authorization Failure Summary                ROW 1 TO 3 OF 3

Reporting Started . . . . . : 03/30/08 08:52:04
Last Interval . . . . . : 03/30/08 09:01:23

For details, place any character next to date and press Enter

  Date      Time      Authorization ID
_ 03/30/08  08:52:04  USER300
_ 03/30/08  08:52:34  USER300
_ 03/30/08  08:53:07  USER300

***** BOTTOM OF DATA *****

Command ==>
F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel

```

Figure 31. Authorization Failure Summary window

Use this window to view a list of authorization failures. You can select any item from the list to display the Authorization Failure Detail window, where the authorization failure can be examined in greater detail. You can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

Examining the Exception Event Summary

To display the Exception Event Summary window, select option 5 (Exception Event Summary) from the Look Selections menu or type LOOK 5 and press Enter on any command line.

```

DGOMLEXP                Exception Event Summary                ROW 1 TO 4 OF 4

Reporting Started . . . . . : 12/07/09 08:52:04
Last Interval . . . . . : 12/07/09 09:01:23

   Date      Time      IFCID
- 12/07/09   08:52:50   172 Deadlock
- 12/07/09   08:53:06   267 CF rebuild/alter started
- 12/07/09   08:54:11   090 Global trace started
- 12/07/09   08:54:12   359 Index split
- 12/07/09   08:54:13   337 Lock escalation
***** BOTTOM OF DATA *****

Command ==>
F1=Help   F2=Split   F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel

```

Figure 32. Exception Event Summary window

Use this window to view the most recent exception events that occurred during your Online Monitor session. You can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

Select any event listed on the Exception Event Summary window for closer examination. When an exception event is selected, one of the following windows is displayed as determined by the type of event:

- EDM Pool Full Data window
- Deadlock Data window
- Timeout Data window
- Coupling Facility (CF) Rebuild/Alter Start Event window
- CF Rebuild End Event window
- CF Alter End Event window
- Unformatted Record panel (thread commit indoubt)
- Global trace started
- Data Set Extension Data window
- Unit of Recovery Inflight or Indoubt Data window
- Active Log Space Shortage Data window

For IFCID 337 (Lock Escalation) an IFCID 359 (Index Split) exception events no details are available because the ISPF Online Monitor runs in Db2 10 toleration mode only.

Examining the Exception Event Messages

To display the Exception Event Messages window, select option 6 (Exception Event Messages) from the Look Selections menu or type LOOK 6 and press Enter on any command line.

```

DGOMLEMP                Exception Event Messages                ROW 1 TO 4 OF 4
-----
DGOM953 Exception event processor started at      03/30/08 08:52:04.100
DGOM954 Exception event processor stopped at     03/30/08 08:52:10.100
***** BOTTOM OF DATA *****

Command ==>
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down     F9=Swap
F12=Cancel

```

Figure 33. Exception Event Messages window

Use this window to view messages issued by the exception event processor. You can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

Exception processor output data sets

Exception processing output data can be written to data sets. Use these data sets for further analysis of exception conditions.

Note: Always write monitor, trace, Exception Log File, and exception DPMOUT data to separate data sets.

Exception Log File data set

You can specify a particular Exception Log File data set where information about exceptions is written. The contents of the Exception Log File data set can be either printed or loaded into a Db2 table for further investigation.

To retain log file data on the same data set across multiple Online Monitor sessions, specify APPEND for the disposition.

To dynamically allocate a new Exception Log File data set with the following attributes, specify NEW for the disposition:

```

RECFM:
  VB
LRECL:
  512
BLKSIZE:
  4 096

```

The data set contains an entry for each field found in exception status.

See [“Layout of the Exception Log File output record” on page 73](#) for the layout of the Exception Log File data set.

Printing the Exception Log File data set

You can print the contents of the Exception Log File data set using the Exception Log File print utility.

To use this utility, submit the sample member DGOMEJCL found in the RKO2SAMP library. This member is provided as an example and can be modified as required.

Performance Database

Exception data can be loaded into OMEGAMON for Db2 PE’s Performance Database.

You can find CREATE TABLE DDL, LOAD, CREATE VIEW statements, and sample SQL queries in the RKO2SAMP library, in the following members:

```

CREATE TABLE statement:
  DGOECFIL

```

LOAD utility control statement:

DGOELFIL

Sample CREATE VIEW statement:

DGOEVFIL

Sample SQL query:

DGOEQFIL

Exception DPMOUT data set

You can specify a particular exception DPMOUT output data set. This data set holds the Db2 instrumentation records that contain at least one field in exception status, in DPMOUT format. You can use this data set as input to a batch Record trace or Statistics trace for a more detailed analysis of exception conditions.

To retain DPMOUT data on the same data set across multiple Online Monitor sessions, specify APPEND for the disposition.

To dynamically allocate a new DPMOUT data set with the following attributes, specify NEW for the disposition:

RECFM:

VBS

LRECL:

32 756

BLKSIZE:

6 233

See the *Report Reference* for information about the layout of the DPMOUT record.

Stopping exception processing

The different types of exception processing can be stopped using the Exception Processor panel by deselecting the appropriate fields under Activate/Deactivate Exception Processing.

Figure 23 on page 33 shows the Exception Processor panel.

If an exception task is active when you exit the Online Monitor, you are notified by one of the Asynchronous Task Termination panels. You can either keep selected tasks active (if the data collector is active) and exit the Online Monitor, or return to the Online Monitor Main Menu keeping all asynchronous tasks active.

See [“Leaving the ISPF Online Monitor”](#) on page 12 for more information.

If an exception task is active when you change Db2 subsystems, you are notified by one of the Asynchronous Task Termination panels. You can either change Db2 subsystems and terminate all asynchronous tasks, or return to the previous panel keeping all asynchronous tasks active.

See [“Db2 Subsystem windows”](#) on page 7 for more information.

If you work in split screen mode and you change the subsystem in one session, periodic exceptions are terminated even if they were turned on in the other session.

Monitoring the access path with Explain

Online Monitor Explain provides a real-time analysis of the access path methods that are chosen by Db2 for a given SQL statement. You can also modify the SQL statement text and reexplain it. This provides you with the ability to tune a Db2 subsystem while you are online.

This topic describes how to select the SQL statement you want to explain from within the Online Monitor or from within an ISPF/PDF editor, the authorization required to use Explain, performance considerations, tuning advice, and the Db2 Explain Output panel.

Online Monitor Explain is based on DB2's SQL Explain function. However, it translates the more important PLAN_TABLE codes into full English sentences, so you do not need to know the values and columns of the PLAN_TABLE. In addition, Online Monitor Explain provides direct access to related catalog information (the catalog statistics of the tables and indexes that are used in the SQL statement).

Authorizations required to use Explain

To explain an SQL statement, you need Db2 EXECUTE privilege on the OMEGAMON for Db2 PE Explain plan defined at installation. The name of this plan is KO2EXPL.

The following Db2 PLAN_TABLE privileges might also be required:

- Explain an entry in a plan table.

You require Db2 SELECT privilege on the PLAN_TABLE.

- Explain an SQL statement from a previously bound plan or package.

You require Db2 SELECT privilege on the plan or package owner's PLAN_TABLE. If you do not have the privilege, a dynamic Explain of the statement is performed if you have UPDATE privilege on your own PLAN_TABLE.

- Explain an SQL statement entered in the panel.

You require Db2 UPDATE privilege on your own PLAN_TABLE.

- Explain an SQL statement from a Thread Activity window.

You require Db2 SELECT privilege on the plan or package owner's PLAN_TABLE. If you do not have the privilege, a dynamic Explain of the statement is performed if you have UPDATE privilege on your own PLAN_TABLE.

- CREATE or ALTER a plan table.

You require a minimum of CREATETAB authority on the database that the PLAN_TABLE is created in, or ALTER authority to modify the PLAN_TABLE.

Performance considerations

The execution time to explain a plan or package depends on the number of rows in the accessed PLAN_TABLE. If the table has many rows, you can improve performance by creating an index on PROGNAME, QUERYNO. Alternatively, delete unnecessary rows from your PLAN_TABLE.

What to look for in Explain information

When an SQL statement is explained, it is important that the explaining takes place on the production Db2 subsystem, or at least on a Db2 subsystem where the catalog statistics have been updated to reflect the real production system in terms of table size, available indexes, and other key values. You can specify the "current server" on which Explain is to be executed. This option makes it possible that you are connected to a Db2 test subsystem while you execute the Explain on the remote production system.

The following Explain information can be useful in determining why an application does not perform as expected:

- Access path chosen

Table space scans and nonmatching index scans should be avoided, unless you intend to access all rows in a given table or the table is very small. If the table has one or more indexes, try to reconstruct the SQL statement in such a way that Db2 chooses a better access path. If there is no index, consider creating one.

- Index-Only-Access

When you only select a few column values, consider the possibility of including these few columns in the column list of one of the indexes. In this way, all requested data can be found in the index. The access path message informs you if you succeed in doing so. Likewise, if you select a maximum value,

consider building a descending index on that column (or an ascending index, if you select a minimum value). In this way, you can even avoid the scanning of leaf pages in the index structure.

- Clustering versus clustered

If a clustering index has been chosen by Db2, ensure that the actual index is clustered. In the Index Information window, if the clustered value is NO, or if the cluster ratio is less than 95%, the table space might need a reorganization to bring the data rows into clustering sequence.

- Number of matching columns

On the Plan Table Data panel, if Db2 has selected a matching index scan, you should verify in the Index Information window that the number of columns used in the index is what you expect.

- Active pages versus pages with rows

Verify that the number of pages with rows is approximately the same as active pages, especially if you are performing table space scans. The value shown in the 1 field in the Table Information window should be as close as possible to 100 percent.

- Number of tables per table space

On the Table Space Information window, you should monitor the Tables field. This field shows the number of tables located in the table space. If the access path is Tablespace scan and the table space is not segmented, there should be only one table in the table space. In a nonsegmented table space, all tables are scanned, not only the selected table.

- Host variable definitions versus column definitions

An inconsistent definition of host variables shown in the Host Variable Definition window, compared to the corresponding column definitions shown in the Key Column Selection window, can indicate an inefficient access path selection, resulting from a possible disqualification of index usage. If, for example, an index column is defined as 3 characters, and that column is being compared in a WHERE-clause with a host variable defined as 4 characters, then Db2 does not base its access path selection on the mentioned index. You should verify that a column and a host variable being compared in a WHERE-clause have compatible definitions.

Plan table considerations

Whenever an SQL statement is explained, the result is written as an entry into a plan table. When the Explain function performs a dynamic Explain, the result is written to the PLAN_TABLE of the current SQLID. The current SQLID is either the same ID as the one used for your Online Monitor session or the ID that you specify in the Explain Menu. You can then examine the explained SQL statement residing in the plan table by using the Db2 Explain Output panel and its associated windows.

If you request a dynamic Explain of an SQL statement, and if you do not have an up-to-date plan table, one of these windows is displayed:

- If a plan table currently does not exist, the [“Create Plan Table window”](#) on page 44 prompts you to create a plan table.
- If a plan table exists, but does not match your current version of Db2, the [“Upgrade Plan Table window”](#) on page 45 prompts you to upgrade the plan table.

Create Plan Table window

If a plan table does not exist when a dynamic Explain is required, the Create Plan Table window is displayed.

Use the Create Plan Table window to specify the database name (optional), and table space name of the table to be created (optional). If neither are entered, a plan table is created in the default database DSNDB04. If this database does not exist or you do not have sufficient authority, an error message is displayed.

```

DGOMYWQC                Create Plan Table

You do not have a plan table. Update the database name and table
space name as required, then press Enter to create the plan table.

Database name . . . . . DSNDB04
Table space name . . . . .

F1=Help   F2=Split   F9=Swap   F12=Cancel

```

Figure 34. Create Plan Table window

Upgrade Plan Table window

If a plan table exists but does not match your current version of Db2, the Upgrade Plan Table window is displayed. Use this window to confirm that the table definition be upgraded to match your current version of Db2.

```

DGOMYWQU                Upgrade Plan Table

Your plan table definition does not include all the columns
required by DB2PM.

To upgrade the plan table, press Enter.

F1=Help   F2=Split   F9=Swap   F12=Cancel

```

Figure 35. Upgrade Plan Table window

Explain an SQL statement

You can explain an SQL statement by various actions.

- Entering the EXPLAIN command or pressing F18 (Explain) in any Thread Activity panel in either the SQL Statement and Package window or SQL Statement and DBRM window.

The current SQL statement is then explained.

- Selecting option 4 (Explain) from the Online Monitor Main Menu.

The Explain Menu is displayed, where you can select one of four options for explaining an SQL statement.

- Marking a range of lines to be explained from within an ISPF/PDF editor and issuing the EXPLAIN command from the editor command line.

See [“ISPF Online Monitor Source Explain” on page 54](#) for more information.

The SQL statement is then explained and the result displayed in the Db2 Explain Output panel.

[Figure 36 on page 46](#) shows the flow of the Explain process.




Figure 36. Flow of Explain process

Selecting Explain from the Online Monitor Main Menu

You can access Explain by selecting option 4 (Explain) from the Online Monitor Main Menu. See [Figure 1](#) on page 5 for an example of the Online Monitor Main Menu.

The Explain Menu is displayed, where you can select one of four methods for specifying the SQL statement you want explained.

Explaining an SQL statement from an ISPF/PDF editor

You can explain an SQL statement that is imbedded in a source program or SPUIFI input from an ISPF/PDF editor. To do this, specify a range of source lines to explain using the E line prefix command, then type EXPLAIN on the command line and press Enter. Any valid SQL statements within the specified range can then be selected from a list for Explain processing. For more information about source Explain, see [“ISPF Online Monitor Source Explain”](#) on page 54.

If Explain executed successfully, the Db2 Explain Output panel is displayed. For more information about the Db2 Explain Output panel, see [“Explain output \(Db2 Explain Output panel\)”](#) on page 52.

Explain Menu

Use the Explain Menu to specify the Db2 subsystem where the Explain request is to be performed (optional), a different SQL authorization ID for qualifying the unqualified tables in the statements being explained, and to select an option for explaining an SQL statement. To display the Explain Menu, select option 4 (Explain) from the Online Monitor Main Menu.

```

DGOMYPMN                                Explain Menu

Local location . . . . . : PMODB2A          SDA2  V10
Current server . . . . . : PMODB2A          SDA2  V10

Change current server if required, then select one of the following.

- 1. Explain an existing entry in the plan table
  2. Explain a package's SQL statement
  3. Explain a DBRM's SQL statement
  4. Enter an SQL statement to be explained

Current server . . . . . PMODB2A_____ +
Current SQLID . . . . . PMDEV51_____ +

Command ==> _____
F1=Help   F2=Split  F3=Exit   F9=Swap   F12=Cancel F21=Expand F22=Left
F23=Right

```

Figure 37. Explain Menu

Use the `Current server` field to specify the Db2 subsystem where you want the Explain request to be performed. This field makes it possible that you are connected to a Db2 test subsystem while you execute the Explain on the remote production system. The default value is the local location.

Use the `Current SQLID` field to specify a different SQL authorization ID for qualifying the unqualified tables in the SQL statements to be explained. The `Current SQLID` field can also be used to qualify the plan table to be accessed. The default value is your user ID.

Whenever package or DBRM SQL statements are explained, the package or plan qualifier is used to qualify unqualified SQL statements. The package or plan owner is used to qualify the plan table.

From this window, you can select one of the following options:

- Select option 1 (Explain an existing entry in the plan table) to display the Plan Table Entry window, where you can interpret an existing entry in your plan table, or in another user's plan table if you have the required Db2 authority.
- Select option 2 (Explain a package's SQL statement) to display the Package SQL Statement window, where you can explain an SQL statement for a previously bound package.
- Select option 3 (Explain a DBRM's SQL statement) to display the DBRM SQL Statement window, where you can explain an SQL statement for a previously bound plan.
- Select option 4 (Enter an SQL statement to be explained) to display the SQL Text Entry window, where you can explain a dynamic SQL statement. After you specified the query number, the SQL Text Edit panel is displayed, where you can enter and explain a dynamic SQL statement.

Use these options, and related windows, to select the SQL statement you want to explain. The SQL statement is then explained, and the result displayed in the Db2 Explain Output panel.

Explain existing entry in Plan Table

Use the Plan Table Entry window to explain an entry in your plan table or, if you have the required Db2 authority, an entry in another user's plan table. Regardless of the user specified, a Db2 plan table needs to exist for that user.

To display this window, select option 1 from the Explain Menu.

```

DGOMYWPE                               Plan Table Entry

Local location . . . . . : PMODB2A          SDA2  V10
Current server . . . . . : PMODB2A          SDA2  V10

Specify the owner of the plan table. Then either specify a
query number, or leave query number blank and enter plan,
program and version details for a list of query numbers.

Owner of plan table . . . . . PMDEV51_----- +
Query number . . . . . -----

Plan name . . . . . -----
Program name . . . . . : DGO@TPG1         +
Version . . . . . ----- +

Command ==>
F1=Help   F2=Split   F3=Exit   F9=Swap   F12=Cancel F21=Expand
F22=Left  F23=Right

```

Figure 38. Plan Table Entry window

Enter a user ID in the Owner of plan table field, or leave this field blank to specify your own plan table. Then specify the Query number of the plan table entry you want to explain.

If the 1 field is left blank, all entries are listed that match the plan name, program name, and version that you specified. Leaving any of these criteria fields blank is equivalent to specifying an asterisk (*), which matches all entries. The Plan Table Entry List panel is then displayed, which lists all matching entries in the plan table.

After selecting an entry in a plan table, the entry is processed and the result is shown in the Db2 Explain Output panel.

Plan Table Entry List panel

Use the Plan Table Entry List panel to select an entry from the list of entries belonging to the plan table. If many table entries are listed, you can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

```

DGOMYWPL                               Plan Table Entry List                               Row 1 to 5 of 5

Local location . . . . . : PMODB2A          SDA2  V10
Current server . . . . . : PMODB2A          SDA2  V10

Table . . . . . : PMDEV51.PLAN_TABLE      +

S      Query  Plan      Program      Version      Timestamp
-----+-----1-----+-----1-----+
-      1011001  DGO@TPG1    OMPE_FINAL   2011021615243488
-      1011001  DGO@TPG1    OMPE_FINAL   2011030811011901
-      1011001  DGO@TPG1    0510_PM32647C 2011092016072777
-      1011001  DGO@TPG1    0510_PM32647C 2012012413093911
-      1011001  DGO@TPG1    0510_PM32647C 2012012413101054
***** Bottom of data *****

Command ==>
F1=Help   F2=Split   F3=Exit   F7=Up      F8=Down    F9=Swap
F12=Cancel F16=Look   F17=Collect F21=Expand F22=Left   F23=Right

```

Figure 39. Plan Table Entry List panel

After making a selection, the entry is processed and the result is shown in the Db2 Explain Output panel.

Explain an SQL statement of a package

Use the Package SQL Statement window to explain an SQL statement of a previously bound package.

To display this window, select option 2 from the Explain Menu.

```

DGOMYWKS                Package SQL Statement
Collection identifier . . . . K02EX510_-----+
Package name . . . . . DGO@TPG*-----+
Version . . . . . -----+

Command ==> -----
F1=Help      F2=Split    F3=Exit     F9=Swap     F12=Cancel
F16=Look     F17=Collect F21=Expand  F22=Left    F23=Right

```

Figure 40. Package SQL Statement window

To explain an SQL statement from this window, type the collection identifier, package name, and version into the appropriate fields and press Enter. Wildcards are accepted in these fields. Empty fields are equivalent to asterisks (*), which is the default. The Package Version List panel is then displayed, where you can select a package from a list of all matching packages.

Package Version List panel

The Package Version List panel shows a list of packages that match the fields specified in the Package SQL Statement window. If many packages are listed, you can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

```

DGOMYWKV                Package Version List                Row 1 to 8 of 8
Local location . . . . . : PMODB2A                SDA2 V10
Current server . . . . . : PMODB2A                SDA2 V10

S  Collection          Package          Version
----+-----1-----2  ----+-----1-----2  ----+-----1-----2-----+
-  K02EX510            DGO@TPG1        OMPE_FINAL
-  K02EX510            DGO@TPG1        0510_PM32647C
-  K02EX510            DGO@TPG2        OMPE_FINAL
-  K02EX510            DGO@TPG2        0510_PM32647C
-  K02EX510            DGO@TPG3        0510_PM32647C
-  K02EX510            DGO@TPG4        0510_PM32647C
***** Bottom of data *****

Command ==> -----
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down     F9=Swap
F12=Cancel   F16=Look    F17=Collect F21=Expand  F22=Left    F23=Right

```

Figure 41. Package Version List panel

From this window, you can select a package to display the SQL Statement List panel, where you can examine a list of SQL statements belonging to the package.

SQL Statement List panel (Packages)

The SQL Statement List panel shows a list of SQL statements for the selected package. For each item listed, the statement number and the first 60 characters of the SQL statement text are shown. If many

SQL statements are listed, you can use the scrolling keys F7 (Up) and F8 (Down) to browse the list. Use the key F21 (Expand) to browse the SQL statement.

```

DGOMYWKQ                SQL Statement List                Row 1 to 7 of 7

Local location . . . . . : PMODB2A                SDA2  V10
Current server . . . . . : PMODB2A                SDA2  V10

Collection . . . . . : K02EX510                +
Package . . . . . : DG0@TPG1                +
Version . . . . . : 0510_PM32647C            +

S   Statement  SQL Statement Text
-----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6
-      36  DECLARE DGO_SYSPACKAGE TABLE ( LOCATION VARCHAR ( 128 ) NOT
-      1011001 SELECT LOCATION , COLLID , NAME , CONTOKEN , OWNER , CREATOR
-      1011002 SELECT LOCATION , COLLID , NAME , CONTOKEN , OWNER , CREATOR
-      1011003 SELECT DEGREE INTO :HV_SYSPACKAGE_DEGREE FROM DGO_SYSPACKAGE
-      1011051 SELECT LOCATION , COLLID , NAME , CONTOKEN , OWNER , CREATOR
-      1011052 SELECT LOCATION , COLLID , NAME , CONTOKEN , OWNER , CREATOR
-      1011053 SELECT DEGREE INTO :HV_SYSPACKAGE_DEGREE FROM DGO_SYSPACKAGE
***** Bottom of data *****

Command ==>
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down    F9=Swap
F12=Cancel   F16=Look    F17=Collect F21=Expand  F22=Left   F23=Right

```

Figure 42. SQL Statement List panel

From this window, you can select the SQL statement to be explained. The SQL statement is then processed and the result is shown in the Db2 Explain Output panel.

Explain an SQL statement of a DBRM

Use the DBRM SQL Statement window to explain an SQL statement of a previously bound plan.

To display this window, select option 3 from the Explain Menu.

```

DGOMYWDS                DBRM SQL Statement

DBRM name . . . . . DSN*_____ +
Plan name . . . . . _____

Command ==>
F1=Help      F2=Split    F3=Exit     F9=Swap     F12=Cancel
F16=Look     F17=Collect F21=Expand  F22=Left    F23=Right

```

Figure 43. DBRM SQL Statement window

To explain an SQL statement from this window, type the DBRM name and plan name into the appropriate fields and press Enter. Wildcards are accepted in these fields. Empty fields are equivalent to asterisks (*), which is the default. The DBRM and Plan List panel is then displayed, where you can select a DBRM from a list of all matching DBRMs and plans.

DBRM and Plan List panel

The DBRM and Plan List panel shows a list of DBRMs and plans matching the fields specified in the DBRM SQL Statement window. If many items are listed, you can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

```

DGOMYWDL                                DBRM and Plan List                                Row 1 to 4 of 4
Local location . . . . . : PMODB2A                SDA2  V10
Current server . . . . . : PMODB2A                SDA2  V10

S   DBRM                                Plan Name
----+-----1-----+-----2
-   DSNTIAD                             DSNTIA81
-   DSNTIAD                             DSNTIAD
-   DSNTIAUL                             DSNTIAUL
-   DSNTIAUL                             DSNTIB81
***** Bottom of data *****

Command ==> -----
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down     F9=Swap
F12=Cancel   F16=Look    F17=Collect F21=Expand  F22=Left    F23=Right
  
```

Figure 44. DBRM and Plan List panel

From this window, you can select any DBRM listed to display the SQL Statement List panel, where you can examine a list of SQL statements belonging to the DBRM.

SQL Statement List panel (DBRMs)

The SQL Statement List panel shows a list of SQL statements for the selected DBRM. This window shows the statement number and the first 60 characters of the SQL statement text for each item listed. If many SQL statements are listed, you can use the scrolling keys F7 (Up) and F8 (Down) to browse the list. Use the key F21 (Expand) to browse the SQL statement.

```

DGOMYWDQ                                SQL Statement List                                Row 1 to 4 of 59
Local location . . . . . : PMODB2A                SDA2  V10
Current server . . . . . : PMODB2A                SDA2  V10

DBRM . . . . . : DSNTIAUL                        +
Plan name . . . . . : DSNTIAUL

S   Statement  SQL Statement Text
----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6
-   724  DECLARE MSEL STATEMENT
-   935  WHENEVER SQLERROR GO TO EXECER11
-   8064  SELECT CURRENT SQLID INTO :CRNAME FROM SYSIBM . SYSDUMMY1
-   9162  ROLLBACK
***** Bottom of data *****

Command ==> -----
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down     F9=Swap
F12=Cancel   F16=Look    F17=Collect F21=Expand  F22=Left    F23=Right
  
```

Figure 45. SQL Statement List panel

From this window, you can select an SQL statement to be explained. The SQL statement is then processed and the result is shown in the Db2 Explain Output panel.

Explain dynamic SQL statement

Use the SQL Text Entry window and its associated SQL Text Edit panel to enter an SQL statement to be explained.

To display this window, select option 4 from the Explain Menu.

```

DGOMYWES                                SQL Text Entry

Enter query number if needed, then press Enter to display the SQL Text
Edit panel.

  Query number . . . . . 999735912

  Set current degree to ANY

F1=Help   F2=Split   F3=Exit   F9=Swap   F12=Cancel

```

Figure 46. SQL Text Entry window

In this window, enter the query number you want to assign to the SQL statement. The query number identifies in the plan table the SQL statement that you are entering. If no query number is specified, the default is 999 735 912.

The 1 field specifies whether the SQL statement is eligible for query parallelism. If you enter a forward slash (/) in this field, the current degree of parallelism is set to ANY. Otherwise, the current degree of parallelism is set to 1.

After you press Enter, the Edit SQL Text panel is displayed, where you can enter the SQL statement text you want to explain.

SQL Text Edit panel

Use this panel to view or modify the full text of an SQL statement for purposes of online tuning, or to create an SQL statement to be explained.

This panel is displayed by pressing Enter in the SQL Text Entry window, or by selecting the 1 field from the Db2 Explain Output panel.

```

ISREDDE2   XRK.PM04.SPFTEMP1.CNTL           Columns 00001 00072
*****  ***** Top of Data *****
000001 SELECT COUNT ( * ) INTO :HV_LOC10_COUNT FROM DGO_SYSPACKAGE WHERE
000002 LOCATION = :HV_LOC10_LOCATION AND ( COLLID LIKE :HV_LOC10_COLLID OR
000003 COLLID LIKE :HV_LOC10_COLLID2 ) AND ( NAME LIKE :HV_LOC10_NAME OR NAME
000004 LIKE :HV_LOC10_NAME2 ) AND VERSION LIKE :HV_LOC10_VERSION QUERYNO
000005 001011010
*****  ***** Bottom of Data *****

Command ==>
F1=Help      F2=Split    F3=Exit     F5=Rfind    F6=Rchange  F7=Up
F8=Down      F9=Swap     F10=Left   F11=Right   F12=Cancel  F13=Help
F14=Split    F15=End     F16=Return F17=Rfind   F18=Rchange F19=Up
F20=Down     F21=Swap    F22=Left   F23=Right   F24=Cretriev

Scroll ==> CSR_

```

Figure 47. SQL Text Edit panel

From this panel, you can enter a new (or modify an existing) SQL statement to assess the effect on access path selection. The statement can be adjusted and then explained again to dynamically assess the effect.

You can enter SQL text that exceeds the display depth by using the scrolling keys F7 (Up) and F8 (Down).

To process and explain the entered SQL text, press F3 (Exit). The Db2 Explain Output panel is displayed with a confirmation message.

To ignore the changes and return to the previous panel, press F12 (Cancel).

Explain output (Db2 Explain Output panel)

Use the Db2 Explain Output panel and Catalog Table windows to investigate the access path methods of an explained SQL statement. Information regarding packages, DBRMs, table spaces, tables, indexes, and column attributes can also be viewed.

Use the Db2 Explain Output panel to examine the results of an explain request.

When an SQL statement has been chosen for explanation, the following processing depends on the origin of the statement:

Explain existing entry in the plan table

The specified PLAN_TABLE entry is shown as a result of the explain request.

Explain a package's SQL statement

The plan table of the package owner is searched to check whether the SQL statement has been explained during BIND. If the search is successful, this PLAN_TABLE entry is taken. If the search is not successful, a dynamic explain is performed, that is, SQL EXPLAIN is invoked for the statement. Then the result of this explain is displayed.

The plan table is searched by using the bind time of the package. There might be multiple occurrences of the package with different bind times. For example, the Db2 system catalog might contain the latest package that is created by the Db2 command BIND or REBIND. It might also contain a former package version that is activated by the Db2 command REBIND SWITCH.

Explain a DBRM's SQL statement

The plan table of the DBRM owner is searched to check whether the SQL statement has been explained during BIND. If the search is successful, this PLAN_TABLE entry is taken. If the search is not successful, a dynamic explain is performed, that is, SQL EXPLAIN is invoked for the statement. Then the result of this explain is displayed.

SQL statement to be explained has been entered

A dynamic explain is performed, that is, SQL EXPLAIN is invoked for the statement, and the result is displayed.

SQL statement is modified in Db2 Explain Output panel

A dynamic explain is performed, that is, SQL EXPLAIN is invoked for the modified statement and the new result is displayed.

Only the Db2 Explain Output panel for packages (DGOMYPKM) is shown in this section.

```

DGOMYPKM                DB2 Explain Output
Local location . . . . . : PMODB2A          SDA2  V10
Current server . . . . . : PMODB2A          SDA2  V10

- Package . . . . . : K02EX510.DG0@TPG1      +
  Version . . . . . : 0510_PM32647C         +
  Explain executed at . . . : 12/01/25 18:33:56.63

----- SQL Text -----
- SELECT LOCATION , COLLID , NAME , CONTOKEN , OWNER , CREATOR , TIMESTAMP
  , BINDTIME , QUALIFIER , PKSIZE , AVGSIZE , SYSENTRIES , VALID ,
  Status . : Compiled-bound using defaults for input variables
  Isolation: Uncommitted Read

- Host variable definitions

----- Access path summary for query block 1 step 1 -----
Matching index scan with scan of referenced data pages
Number of matching columns: 4. The index has 4 columns
Non clustered index scan will be used
Page range scan will not be used

- Table  SYSIBM  SYSPACKAGE
- Index  SYSIBM  DSNKX01

- PLAN_TABLE details for step

FPEM762 This statement was explained at bind time
Command ==> -----
F1=Help      F2=Split      F3=Exit      F7=Up        F8=Down      F9=Swap
F12=Cancel   F16=Look      F17=Collect  F21=Expand   F22=Left     F23=Right

```

Figure 48. Db2 Explain Output panel (packages)

The Db2 Explain Output panel is divided into three sections:

- The first section of the panel shows the plan name, and information about the package or DBRM that contains the SQL statement. You can select this section for further investigation of the package, DBRM, or plan name associated with the SQL statement.
- The second section of the panel, titled `SQL Text`, shows the first 150 characters of the SQL statement text. You can view or modify the full SQL statement text by selecting the input field shown beside the SQL text. The SQL Text Edit panel is then displayed. For more information about the SQL Text Edit panel, see [“SQL Text Edit panel” on page 52](#). After the SQL text is modified, the SQL statement is reexplained, and the new result is shown in the Db2 Explain Output panel. From this section you can also reach windows that provide information about the host variable definitions.
- The third section of the panel, titled `Access path summary`, shows information about the access path methods chosen by Db2 for each of the individual query blocks and steps required to execute the SQL statement. Each step also shows a list of the tables and indexes accessed for that step. You can select any table or index listed in a step for further investigation. You can also display the raw `PLAN_TABLE` data for a step. This also provides online access to help information for `PLAN_TABLE` columns, similar to the information in *Db2 SQL Reference* for the EXPLAIN statement.

If the execution of the SQL statement contains many steps, you can use the scrolling keys F7 (Up) and F8 (Down) to browse through the steps.

Online Monitor EXPLAIN command

Use the EXPLAIN command to explain the currently executing SQL statement in the Thread Activity panels.

- SQL Statement and Package
- SQL Statement and DBRM

Note: If the Thread Detail panel has been modified to show the SQL statement text, you can use the EXPLAIN command from the Thread Detail panel as well.

If the explain request was successful, the Db2 Explain Output panel is displayed.

The command syntax is:

►► EXPLAIN ◀◀

Figure 49. Syntax of the EXPLAIN command

Note: EXPLAIN can be truncated to a minimum of EXP.

For a description of the Db2 Explain Output panel, see [“Explain output \(Db2 Explain Output panel\)” on page 52](#).

ISPF Online Monitor Source Explain

Use Source Explain to explain SQL statements that are embedded in a source program or SPUFI input. Source explain is performed from within the ISPF/PDF editor.

The following languages are supported:

- Assembler
- C
- COBOL
- FORTRAN
- PL/I
- SPUFI

Note: Ensure that Source Explain is installed at your site. Before you activate Source Explain, make sure that the Db2 load library is allocated to your TSO ISPF session.

To explain an SQL statement while editing a source program (or SPUFI input), specify the lines you want to have explained by using the ISPF/PDF editor line prefix command E, type EXPLAIN on the command line, and press Enter.

You can use the line prefix command E like this:

E

To explain a single line, type E in the prefix area of the line to be scanned for SQL statements.

EE

To explain a range of lines, type EE in the prefix area of the first line and the last line of the range of lines to be scanned for SQL statements.

E[n]

To explain a specific number of lines, type E[n] on the first line of the area to be scanned, where *n* is the number of lines to be scanned for SQL statements.

When you enter the EXPLAIN command, Source Explain scans the specified range for valid SQL statements. If a range is not specified, the entire source is scanned. The valid SQL statements are then listed on the SQL Statement Selection panel.

The following figure shows how to explain an SQL statement while editing COBOL source code.

To explain the SQL statements in the source code from line 3040 to 3160:

1. In the line prefix area of lines 3040 and 3160, type EE.
2. On the command line, type explain and press Enter.

```

EDIT ---- SYS1.DSN610.SDSNSAMP(DSNBBC3) - 01.00 ----- COLUMNS 001 072
003010      *** CURSOR LISTS ALL EMPLOYEE NAMES WITH A PATTERN (%) OR (_)
003020      *** FOR LAST NAME
003030
EE3040          EXEC SQL DECLARE TELE2 CURSOR FOR
003050              SELECT *
003060              FROM   VPHONE
003070              WHERE  LASTNAME LIKE :LNAME-WORK
003080              AND    FIRSTNAME LIKE :FNAME-WORK
003090              END-EXEC.
003100
003110      *** CURSOR LISTS ALL EMPLOYEES WITH A SPECIFIC
003120      *** LAST NAME
003130
003140          EXEC SQL DECLARE TELE3 CURSOR FOR
003150              SELECT *
003160              FROM   VPHONE
003170              WHERE  LASTNAME = :LNAME
003180              AND    FIRSTNAME LIKE :FNAME-WORK
003190              END-EXEC.
003200      /
003210      /*****
003220      * FIELDS SENT TO MESSAGE ROUTINE      *
003230      *****/
003240      01 MAJOR                                PIC X(07) VALUE 'DSNBBC3'.
003250
003260      01 MSGCODE                               PIC X(4).
COMMAND ==> explain                                SCROLL ==> CSR
F1=HELP      F2=SPLIT      F3=END      F4=RETURN      F5=RFIND      F6=RCHANGE
F7=UP        F8=DOWN      F9=SWAP      F10=LEFT      F11=RIGHT      F12=RETRIEVE

```

Figure 50. Source Explain example

If you previously selected the **Always display this window** field in the Source Explain Options window as shown in [Source Explain Options window](#), the Source Explain Options window is displayed. If you did not select this field, the [SQL Statement Selection panel](#) is displayed.

Source Explain Options window

You can display the Source Explain Options window by specifying a command in the ISPF/PDF editor or on the SQL Statement Selection panel.

- In the command line of the ISPF/PDF editor, type EXPLAIN OPTIONS and press Enter.

If you have previously selected the **Always display this window** field in the Source Explain Options window, you can type EXPLAIN in the command line of the ISPF/PDF editor and press Enter.

- In the command line of the SQL Statement Selection panel, type OPTIONS and press Enter.

You can use the Source Explain Options window to specify the Source Explain processing options, for example, the language of the source code you are editing, the subsystem ID of the explaining Db2, the SQLID to be used, or the degree of parallelism.

```
DGOMYWSO                Source Explain Options
Update fields as required, then press Enter

Local DB2 subsystem . . . . . ----
Current server location . . . . . -----
Current SQLID . . . . . -----
Query number . . . . . -----

/ Set current degree to ANY
/ Always display this window

Define source language . . . . . _  1. Assembler
                                       2. C/370
                                       3. COBOL
                                       4. FORTRAN
                                       5. PL/I
                                       6. SPUFI

F1=Help   F2=Split   F3=Exit   F9=Swap   F12=Cancel
```

Figure 51. Source Explain Options window

On the Source Explain Options window, you can specify the following options:

Local Db2 Subsystem

The local Db2 subsystem to which you want to connect.

Current Server Location

The Db2 subsystem where you want the source SQL statements to be explained.

In this field, you can specify the Db2 test subsystem you want to connect to while you are executing the explain on the remote production system. If you do not specify a subsystem, the local Db2 subsystem is used.

Current SQLID

A different SQL authorization ID for qualifying the unqualified tables in the SQL statements being explained.

A different SQL authorization ID is also used to qualify the plan table to be accessed. If you do not specify an SQL authorization ID, your user ID is used.

Query number

A number that identifies the Explain statement.

If you do not specify a number, the value 999 735 912 is assigned to this field.

Set current degree to ANY

To specify whether the SQL statement is eligible for query parallelism.

By default, the current degree of parallelism is set to 1. To set the current degree of parallelism to ANY, type a slash (/) in this field.

Always display this window

To control the display of the Source Explain Options window.

By default, the Source Explain Options window is not displayed on each Source Explain request. To display the Source Explain Options window on each Source Explain request, type a slash (/) in this field.

Define source language

In this field, you must specify the language of the source that contains the SQL statements. This field cannot be left blank.

After you have specified the Source Explain options to be used, press Enter to proceed to the SQL Statement Selection panel.

SQL Statement Selection panel

Use the SQL Statement Selection panel to view a summarized list of the valid SQL statements within the specified source area. If there are many SQL statements listed, you can use the scrolling keys F7 (Up) and F8 (Down) to browse the list.

```
DGOMYWSS                SQL Statement Selection                Row 1 to 2 of 2

This panel summarizes the SQL statements found in the following source
module that can be explained.

Source . . : PMDEV51.COMM.COBSRCE(PGMEXP01)

Select an SQL statement, then press Enter

S   Line No  SQL Statement Text
7      80  DECLARE CUR1 CURSOR WITH HOLD FOR S
7      110  SELECT C1,C2,C3,C4,C5,C6 INTO :DCLTAB1 FROM TAB1 WHERE C1 =
--- End of List ---

Command ===>
F1=Help  F2=Split  F3=Exit  F7=Up  F8=Down  F9=Swap  F12=Cancel
```

Figure 52. SQL Statement Selection Panel

You can select any SQL statement that is listed on the SQL Statement Selection panel for Explain processing by typing a slash (/) next to the line to be explained as shown in the figure above. The SQL statement is then processed even if the SQL statement text exceeds the specified range. Then the Online Monitor is accessed. With the Online Monitor, you can view the Explain output on the Db2 Explain Output panel. For more information, see [“Explain output \(Db2 Explain Output panel\)”](#) on page 52.

From the SQL Statement Selection panel, you can access the Source Explain Options window by using the OPTIONS command. In the Source Explain Options window, you can change the Source Explain processing options.

Source EXPLAIN command

Use the EXPLAIN command to explain an SQL statement from within an ISPF/PDF editor, or to specify Source EXPLAIN options.

The command syntax looks like this:




Figure 53. Syntax of the Source EXPLAIN command

Db2 System Parameters panel

Use the Db2 System Parameters panel to view a summary of important DSNZPARM values, maximum concurrent users, and IRLM thresholds. Use the scrolling keys F7 (Up) and F8 (Down) to view all information in the Db2 System Parameters panel.

To display the Db2 System Parameters panel, select option 3 (Display System Parameters) from the Online Monitor Main Menu.

The following figure shows the Db2 System Parameters panel.

```
DG0ZPSM          DB2 System Parameters
Command ===>

For details, type any character next to heading, then press Enter.

PM05D851          D851  V10

- Data (DSNTIPA1, DSNTIPA2)
- Data Sharing Parameters (DSNTIPK)
- Sizes Panel 2(DSNTIP7)
- Thread Management(DSNTIPE)
- Buffer Pools (DSNTIP1, DSNTIP2, DSNTIP6)
- Tracing Parameters (DSNTIPN)
- Operator Functions (DSNTIPO)
- Application Programming Defaults 1(DSNTIPF)
- Application Programming Defaults 2(DSNTIP4)
- Performance and Optimization (DSNTIP8)
- IRLM (DSNTIPI, DSNTIPJ)
- IRLM Processing Parameters
- Protection (DSNTIPP)
- MVS Parmlib Updates (DSNTIPM)
- Active Log Data Set Parameters (DSNTIPL, DSNTIPH)
- Archive Log Data Set Parameters (DSNTIPA, DSNTIPH)
- Databases and Table Spaces (DSNTIPS)
- Distributed Data Facility (DSNTIPR, DSNTIP5)
- Routine Parameters (DSNTIPX)
- Data Definition Control Support (DSNTIPZ)
- Storage Sizes(DSNTIPC)
- Group Buffer Pools
- ZParm Online Change Auditing
- Other Parameters

F1=Help      F2=Split    F3=Exit     F6=History  F7=Up       F8=Down
F9=Swap      F12=Cancel  F16=Look    F17=Collect
```

Figure 54. Db2 System Parameters panel

Note: The Db2 System Parameters panel you see might have been tailored by your installation and consequently might not necessarily match the panel examples shown in this information.

From the Db2 System Parameters panel, you can reach a set of windows that provide system parameters information in greater detail. To reach a window, type any character in the space provided beside the heading.

Several windows can be accessed simultaneously by selecting multiple fields in the Db2 System Parameters panel. To leave the current window and proceed with the next window, press F3 (Exit). To cancel all windows and return to the Db2 System Parameters panel, press F12 (Cancel).

System Parameters commands

The following command applies to System Parameters panels. You can type the command on any System Parameters panel command line, or invoke it using the appropriate function key. A list of active function keys is displayed at the bottom of each panel. You can abbreviate commands down to three characters in length. For example, the HISTORY command can be abbreviated to HIS.

Problem identification examples

This section provides examples of how to discover and identify various system problems by using exception processing, viewing past data, and analyzing the access paths chosen by Db2 to process an SQL statement.

Example of exception processing

This example demonstrates how you can use exception processing to discover a possible class 2 CPU (processing) time problem, and identify the cause by viewing past data and analyzing the access paths chosen by Db2 to process an SQL statement. The data collector needs to be installed at your site to view past data.

First, set the exception threshold for class 2 CPU time to a value applicable to your environment. To do this, select option 4 from the OMPE menu. This displays the Data Set Maintenance Menu.

```
DGOPMENU                Data Set Maintenance Menu
Select one of the following.
1  1. Maintain exception thresholds
   2. Maintain correlation translations
   3. Maintain time zone information
   4. Maintain MAINPACK definitions

Exception data set
'SAMPLE.EXCEPT.DATASET'

DPMPARMS data set
-----

Command ==> -----
F1=Help    F2=Split    F3=Exit    F6=History  F9=Swap    F12=Cancel
F17=Collect
```

Figure 55. Data Set Maintenance Menu

Ensure that you have specified an Exception Threshold data set name in the 1 field. Select option 1 (Maintain exception thresholds) to display the Exception Threshold Category Selection panel.

The Exception Threshold Category Selection panel is displayed.

```

DGOPXDS1      Exception Threshold Category Selection

Select one or more categories, then press Enter. Overtyping with space to
deselect any category. Request EXIT when complete.

  Category
/ Elapsed, CPU and Waiting Times per Plan Execution
- Elapsed, CPU and Waiting Times per Program Execution
- CPU Times per Address Space
- SQL Statements per Plan Execution
- SQL Statements per Program Execution
- SQL Statements per System
- Subsystem Events per Plan Execution
- Subsystem Events per System
- Locking Activity per Plan Execution
- Locking Activity per System
- RID List Processing per Plan Execution
- RID List Processing per System
- Query Parallelism per Plan Execution
- Query Parallelism per System
- Buffer Pools Activity per Plan Execution
- Buffer Pools Activity per System
- Distributed Activity per Location per Plan Execution

Command ==> _____
F1=Help   F2=Split  F3=Exit   F7=Up    F8=Down  F9=Swap  F12=Cancel

```

Figure 56. Exception Threshold Category Selection panel

Select the category that contains the exception field you want. For this example, select the Elapsed, CPU and Waiting Times per Plan Execution field.

The Exception Threshold Field Selection panel is displayed.

```

DGOPXDS2      Exception Threshold Field Selection

Select only one field, then press Enter. Overtyping with space to
deselect any field. Request EXIT when complete.

Field category . . : Elapsed, CPU, and Waiting Times per Plan Execution

  Field      Description
- ADRECETT   Elapsed time in application (Class 1)
- ADCPUT     CPU time in application (Class 1)
- ADDB2ETT   Elapsed time in DB2 (Class 2)
/ ADDBCPUT   CPU time in DB2 (Class 2)
- ADTWTDB    Total wait time in DB2 (Class 2)
- ADTWTAP    Total wait time in application (Class 1)
- ADTSUST    Total Class 3 suspensions time
- ADTSUSC    Total Class 3 suspensions
- QWACAWTL   Lock/latch suspensions time (Class 3)
- ADLLSUSC   Lock/latch suspensions (Class 3)
- QWACAWTI   Synchronous I/O susp. time (Class 3)
- ADIOSUSC   Synchronous I/O suspensions (Class 3)
- QWACAWTR   Other read I/O susp. time (Class 3)

Command ==> _____
F1=Help   F2=Split  F3=Exit   F7=Up    F8=Down  F9=Swap  F12=Cancel

```

Figure 57. Exception Threshold Field Selection panel

This panel shows all available fields within the specified category. Select the ADDBCPUT field and press Enter.

The Exception Threshold Field Details panel is displayed.

```

DGOPXDSN                Exception Threshold Field Details

                                                                ENTRY 1 OF 1

Category . . . . . : Elapsed, CPU, and Waiting Times per Plan Execution
Field ID . . . . . : ADDBCPUT
Description . . . . : CPU time in DB2 (Class 2)

Active . . . . . 1 1=Yes 2=No

By . . . . . 1 1=Total 2=Minute 3=Second
                4=Commit 5=Thread

Compare operator . . > <=Less than >=Greater than
Warning threshold . . 0.09
Problem threshold . . 0.1

Local location . . . *
Group name . . . . . *
Member name . . . . . *
Subsystem ID . . . . *
Requester location . *
Connect . . . . . *
Planname . . . . . *
Corrname . . . . . *
Corrnbr . . . . . *
Primauth . . . . . *

Command ==>-----
F1=Help      F2=Split      F3=Exit      F5=Add      F6=Delete    F7=Up
F8=Down      F9=Swap       F10=Previous F11=Next    F12=Cancel

```

Figure 58. Exception Threshold Field Details panel

Use the Exception Threshold Field Details panel to specify the threshold criteria for the exception field selected from the Exception Threshold Field Selection panel. For this example, enter a 1 in the Active field, set the By field to Total (1), the Compare operator field to a greater than symbol (>), and the Problem threshold field to a value of 0.1 seconds.

Note: Choose a threshold value that is applicable to your environment. For information about defining exception threshold values, see [“How to define exception threshold values”](#) on page 32.

Now that the exception threshold has been defined, select option 2 (Control Exception Processing) from the Online Monitor Main Menu to display the Exception Processor panel.

```

DGOMEPO2                Exception Processor                PM01DLOC DSN1 V10

For any field enter any character to activate

Activate/Deactivate Exception Processing
- Display thread summary
- Display thread detail
- Display statistics detail
/ Periodic                - User Exit
- Exception event notification

Options
Periodic units . . . . . 1      1=Seconds
                                   2=Minutes
Periodic interval . . . . . 10  1-7200 Seconds
                                   1-120 Minutes

> Disable auto-display for problem exceptions
> Sound alarm for exception warnings
- Log file data set output needed
- DPMOUT data set output needed

Exception threshold data set
Name . . . . . SAMPLE.EXCEPT.DATASET

Command ==> -----
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down     F9=Swap
F12=Cancel   F16=Look     F17=Collect

```

Figure 59. Exception Processor panel

Use the Exception Processor panel to activate periodic exception processing. Enter the Exception Threshold data set name you specified on the Data Set Maintenance Menu and type a forward slash (/) in the Periodic field to activate periodic exception processing. In this example, the exception processor has been set to check for periodic exceptions every 10 seconds.

For more information about activating exception processing, see [“How to start exception processing” on page 32](#).

When a periodic exception occurs, the Online Monitor notifies you by displaying the Exception Notification window. This window is overlaid on the current panel.

```

DGOMTPSM 14:27                               Thread Detail                               PM01DLOC DSN1 V10
                                                GROUP001 MEMBER01
For details, place any character next to heading, then press Enter.
More: - +
- Thread Identification
  Primauth . . . . . : TFA6
  Planname . . . . . : CICSEMP      DGOMEPEX Exception Notification
  Connection ID . . . : APPCICP3
  Requesting Location: PM01DLOC   Time . . : 05/03/08 14:27:41
- Current Package . . . . .
- Times
  Class 1 . . . . . : Periodic Exceptions
  Class 2 . . . . . : Problem : 1
  Class 3 . . . . . : Warning : 0
  Class 7 . . . . . : F1=Help      F2=Split      F9=Swap
  Class 8 . . . . . : F12=Cancel
- Locking Activity
  Timeouts . . . . . : 0
  Deadlocks . . . . . : 0
  Suspensions . . . . : 1
  Lock escalations . . : 0
  Maximum page locks held . . . . : 9
- Locked Resources
- RID List Processing
  Unsuccessful - any reason . . . . . : 0
- SQL Activity, Commits and Rollbacks
  DML . . . . . : 4 Commit . . . . . : 0
  DCL . . . . . : 0 Rollback . . . . . : 0
  DDL . . . . . : 0 Changes/Commit . . . : 0.0
- Buffer Manager Activity
  Getpage requests . . . . . : 60
  Buffer updates . . . . . : 0
  Prefetch reads . . . . . : 1
  Synchronous I/O . . . . . : 0
  SQL Statement and Package . . . . . : EMPSRCH
- Distributed Data
  Requester elapsed time . . . . . : N/P
- IFI (Class 5) and Data Capture
- Query Parallelism Data
- Data Sharing Locking Activity
  Suspensions . . . . . : N/A
- Group Buffer Pools Activity
- Stored Procedures

Command ==> -----
F1=Help      F2=Split    F3=Exit     F5=Auto     F6=History  F7=Up
F8=Down     F9=Swap     F12=Cancel  F16=Look    F17=Collect F22=Purge

```

Figure 60. Exception Notification window (overlaid on Thread Detail panel)

The Exception Notification window indicates that a periodic exception has occurred. You can examine the exception in greater detail using the LOOK command. Type LOOK 1 on any command line and press Enter to display the Periodic Exceptions List window.

```

DGOMLAXP                Periodic Exceptions List                Row 499 to 500 of 500
Periodic Interval started . . . . . : 05/03/08 11:43:26:10
Last Interval . . . . . : 05/03/08 14:27:41:22

   Time      Location      Group  Subsystem  Member  Corname
   Field     Reqloc      Primauth  Planname  Connect  Corinmbr
   Descr
-----
- 14:27:41  PM01DLOC    GROUP001 DSN1      MEMBER01 TESTDC6
   'BLANK'  TFA6        CICSEMP   APPCICP3 'BLANK'

   ADDBCPUT 0.108004    >      0.1      Problem  Total
   CPU TIME IN DB2 (CLASS 2)

***** Bottom of data *****

Command ==>
F1=Help    F2=Split   F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel

```

Figure 61. Periodic Exceptions List window

From this window you can examine a list that contains the last 500 periodic exceptions that have occurred. The most recent exceptions are displayed at the bottom of the list.

In this example the thread causing the exception has ended. However, you can still select the exception to be examined from the list displayed in the Periodic Exceptions List window. If history and data collector are available, the Online Monitor retrieves past data and displays the Thread Detail panel where you can examine the thread causing the exception.

Some fields can have values that are longer than the space available to show them in this panel. When this happens, the last character in the field is replaced with an asterisk (*) to indicate that the reported value is truncated.

For details, place any character next to heading, then press Enter.

```

_ Thread Identification
Primauth . . . . . : TFA6          Correlation name . . . . . : TESTDC6
Planname . . . . . : CICSEMP       Connection type . . . . . : CICS
Connection ID . . . : APPCICP3     Type . . . . . : ALLIED
Requesting Location: PM01DLOC     Status . . . . . : DB2
_ Current Package . . . . . : EMPSRCH
_ Times
Class 1 . . . . . : 1:09:58.64058   Elapsed           CPU
Class 2 . . . . . : 11.563398       0.593096
Class 3 . . . . . : 6.186118        N/A
Class 7 . . . . . : N/P             N/P
Class 8 . . . . . : N/P             N/A
_ Locking Activity
Timeouts . . . . . : 0
Deadlocks . . . . . : 0
Suspensions . . . . . : 1
Lock escalations . . . . . : 0
Maximum page locks held . . . . . : 9
_ Locked Resources
RID List Processing
Unsuccessful - any reason . . . . . : 0
_ SQL Activity, Commits and Rollbacks
DML . . . . . : 4 Commit . . . . . : 0
DCL . . . . . : 0 Rollback . . . . . : 0
DDL . . . . . : 0 Changes/Commit . . . . . : 0.0
_ Buffer Manager Activity
Getpage requests . . . . . : 60
Buffer updates . . . . . : 0
Prefetch reads . . . . . : 1
Synchronous I/O . . . . . : 0
/ SQL Statement and Package . . . . . : EMPSRCH
Distributed Data
Requester elapsed time . . . . . : N/P
_ IFI (Class 5) and Data Capture
_ Query Parallelism Data
_ Data Sharing Locking Activity
Suspensions . . . . . : N/A
_ Group Buffer Pools Activity
_ Stored Procedures

Command ==> -----
F1=Help      F2=Split    F3=Exit    F5=Auto    F6=History  F7=Up
F8=Down      F9=Swap     F12=Cancel F16=Look   F17=Collect
F22=Purge

```

Figure 62. Thread Detail panel

The word HISTORY is displayed under the heading line of the Thread Detail panel to indicate that you are viewing past data. The date and time displayed next to HISTORY represent the time just after the exception occurred.

This is where you can use past data to examine the data on a thread that has ended. In this way, you can examine the thread before and after the exception occurred by moving backward and forward through time. See [“Viewing past data” on page 28](#) for more information about viewing past data.

Because class 2 time is the time spent within Db2, the excessive class 2 time causing the exception could be caused by a possible SQL statement problem.

From the Thread Detail panel, select SQL Statement and Package to display the SQL Statement and Package window.

```

DGOMTWCS                SQL Statement and Program
                        HISTORY 00/05/08 14:27:41
                        More: +
_ SQL Statement : SELECT * FROM SYSIBM.SYSTABLES WHERE CREATOR='POD'

Location . . . . . : PM01D711
Collection ID . . . . . : DSNESPCS
Program name . . . . . : DSNESM68
Nested activity name . . . . . : N/P
Program type . . . . . : Package
Consistency token . . . . . : X'149EEA901A79FE48'
Version
N/P
Statement type . . . . . : CLOSE
Statement number . . . . . : 0
Current SQL ID . . . . . : POD
Bind type . . . . . : Dynamic
Cached dynamic SQL identifier . . . : N/A
Thread status . . . . . : In DB2
Database name . . . . . : N/P
Page set name . . . . . : N/P
Page number . . . . . : N/P
Elapsed time . . . . . : N/P
CPU time . . . . . : N/P
_ Used Buffer Pools
  Getpages . . . . . : N/P
  Synch read I/O . . . . . : N/P

Command ==>
F1=Help      F2=Split    F3=Exit     F5=Auto     F6=History
F7=Up        F8=Down     F9=Swap     F12=Cancel  F16=Look
F17=Collect  F18=Explain F22=Purge

```

Figure 63. SQL Statement and Package window

This window displays the SQL statement executing at the time the exception occurred. To obtain details about the access paths chosen by Db2 to process the SQL statement, press F18 (Explain) to explain the SQL statement and display the Db2 Explain Output panel.

```

DGOMYPKM                DB2 Explain Output
                        More - +

Local location . . . . . : PM01DL0C          DSN1 V10
Current server . . . . . : PM01DL0C          DSN1 V10

_ Package . . . . . : EMPLOYEE.EMPSRCH
Version . . . . . : EMPSRCH_PACKAGE_VERSION1

Explain executed at . . . . : 13:10:52 04/17/08

----- SQL Text -----
_ SELECT EMPNO, LASTNAME, WORKDEPT, BIRTHDATE FROM DSNB710.EMP WHERE
EMPNO = :EMPLOYEE_NUMBER
Status : Compiled-REPORT specified. Rebound at exec.time using values
Isolation: Cursor stability

_ Host variable definitions

----- Access path summary for query block 1 step 1 -----
Table space scan - no index will be used
Standard sequential prefetch will be performed
Lock mode is share lock for the page
Page range scan will not be used

/ Table DSNB710 EMP

_ PLAN_TABLE details for step

----- Access path summary for query block 1 step 2 -----
Command ==>
F1=Help      F2=Split    F3=Exit     F7=Up        F8=Down     F9=Swap
F12=Cancel   F16=Look    F17=Collect

```

Figure 64. Db2 Explain Output panel (packages)

The access path summary in the Db2 Explain Output panel shows that a table space scan is being used to access the table. This could indicate the reason for the excessive class 2 time.

Select the table from the Db2 Explain Output panel to display the Table Information window, where you can see if any indexes have been defined for the table.

```

DGOMYWTI                Table Information
Local location . . . . . : PM01DLOC          DSN1 V10
Current server  . . . . . : PM01DLOC          DSN1 V10
More:           - +
Table name . . . . . : DSNB710.EMP
Type . . . . . : TABLE
/ Indexes . . . . . : YES
Database name . . . . . : DSNBD31A
_ Table space name . . . . . : DSNBS31E
Table identifier . . . . . : 14
Columns . . . . . : 14
Rows . . . . . : 32
Status . . . . . : HAS PRIMARY INDEX
Maximum record length . . . . . : 107
Pages . . . . . : 1
Percentage of pages used . . . . . : 2
Last RUNSTATS . . . . . : 04/13/08 10:53:27
EDIT procedure name . . . . . : DSNBEAE1
VALIDATION procedure . . . . . : NONE
Auditing . . . . . : NONE
DATA CAPTURE option . . . . . : NONE
Row compression . . . . . : 0
Alias table or view . . . . . :
Check constraints . . . . . : 0
Encoding scheme . . . . . : EBCDIC

Command ==>
F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel F16=Look  F17=Collect

```

Figure 65. Table Information window

From this window you can see that an index has been defined for the table. Select Indexes to display the Index Selection window where you can select the index to be examined. In this example, the index selected is DSNB710.XEMP1. This displays the Index Information window.

```

DGOMYWII                Index Information
Local location . . . . . : PM01DLOC          DSN1 V10
Current server  . . . . . : PM01DLOC          DSN1 V10
More:           - +
_ Index name . . . . . : DSNB710.XEMP1
Index space name . . . . . : XEMP1
Table name . . . . . : DSNB710.EMP
Database name . . . . . : DSNBD31A
Buffer pool . . . . . : BPO
/ Key columns . . . . . : 1
Subpage size (bytes) . . . . . : 512
Unique rule . . . . . : PRIMARY - UNIQUE
Clustering index . . . . . : YES
Currently clustered . . . . . : YES
Cluster ratio . . . . . : 100
Full key card . . . . . : 32
First key card . . . . . : 32
Levels . . . . . : 1
Leaf pages . . . . . : 1
Maximum piece size (KB) . . . . . : 256
Close rule . . . . . : LEAVE OPEN
Last RUNSTATS . . . . . : 04/13/08 10:53:27
Allocated space . . . . . : 0
Erase rule . . . . . : NO
Index type . . . . . : 2
Time of CREATE INDEX . . . . . : 2008-01-27-08.12.59.000000
Time of last ALTER INDEX . . . . . : 2008-01-27-10.22.30.000000

Command ==>
F1=Help   F2=Split  F3=Exit   F7=Up     F8=Down   F9=Swap
F12=Cancel F16=Look  F17=Collect

```

Figure 66. Index Information window

From the Index Information window, you can select Key columns to display the Key Column Information window.

```

DGOMYWCI                Key Column Information

Local location . . . . . : PM01DLOC          DSN1 V10
Current server  . . . . . : PM01DLOC          DSN1 V10
                                           More:  - +

Column name . . . . . : EMPNO
Table name . . . . . : DSNB710.EMP
Index name . . . . . : DSNB710.XEMP1
Position . . . . . : 6
Sequence . . . . . : Ascending
Type . . . . . : CHAR
Length . . . . . : 6
Scale . . . . . : 0
Key cardinality . . . . . : 123
Null value . . . . . : No
Second highest value . . . . . : 200330
Second lowest value . . . . . : 000020
Last RUNSTATS . . . . . : 2008-04-13-13.10.33.103784

_ Show Key Distribution values

Command ==>
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down
F9=Swap     F12=Cancel  F16=Look   F17=Collect

```

Figure 67. Key Column Information window

Use this window to examine the characteristics of the key column on which the index has been defined.

One of the possible reasons for an index not being used is that the host variable defined in the program does not match the characteristics of the column as defined in the table. To determine whether this is the reason for the index not being used, you can return to the Db2 Explain Output panel where the Host Variable Definition window can be selected.

```

- Host Variable Definition -
DGOMYWHV

Local location . . . . . : PM01DLOC          DSN1 V10
Current server  . . . . . : PM01DLOC          DSN1 V10

Name                                     Type                               Length
EMPLOYEE_NUMBER                          FIXED CHARACTER                      8

Command ==>
F1=Help      F2=Split    F3=Exit     F7=Up       F8=Down    F9=Swap
F12=Cancel  F16=Look   F17=Collect

```

Figure 68. Host Variable Definition window

In this particular example, the Host Variable Definition window shows that the definitions do not match, thereby disqualifying use of the index.

The host variable defined in the source program should be changed to match the table definition so that Db2 can use the index and thereby reduce the class 2 time.

Using the ISPF Online Monitor in deadlock resolution

A deadlock is an error condition where processing cannot continue because of an unresolved contention for a particular resource. A deadlock is signaled when the deadlock detection cycle time has been reached. Ensure that the resource timeout limit is at least twice that of the deadlock detection cycle so that deadlocks are detected before a timeout occurs.

Deadlock detection

Start the data collector to collect deadlock events.

Start exception event processing from the Exception Processor panel to be notified when an exception event occurs. See [“How to start exception processing”](#) on page 32.

Deadlock resolution

On detection of a deadlock, use LOOK to display the participants in the deadlock.

```

DGOMLED0                                Deadlock Data                                More:  - +
IFCID . . . . . : 172

Number of resources involved in deadlock : 2
Deadlock interval counter . . . . . : 2
Time deadlock detected . . . . . : 21:38:34.15

Locked resource
Type . . . . . : Datapage
Database . . . . . : 271      Object : 6
Page # . . . . . : X'3001F1'

Blocker
Member/DBMS identifier . . . . . : MEMBER01
Plan name . . . . . : DDL3P18
Correlation identifier . . . . . : RUNPRG18

Connection identifier . . . . . : BATCH
LUW identifier . . . . . : USIBMSY.SY10BDB2.A74AB9FEF91C
State . . . . . : EXCLUSIVE
Duration . . . . . : COMMIT
Status . . . . . : WAIT

Waiter
Member/DBMS identifier . . . . . : MEMBER02
Plan name . . . . . : DDL3P15
Correlation identifier . . . . . : RUNPRG15
Connection identifier . . . . . : BATCH
LUW identifier . . . . . : USIBMSY.SY30BDB2.A74AB9EE117F
Requested function . . . . . : LOCK
State . . . . . : UPDATE
Duration . . . . . : MANUAL
DB2 assigned worth value . . . . . : 18

Command ==>
  F1=Help      F2=Split    F3=Exit      F7=Up        F8=Down
  F9=Swap      F12=Cancel

```

Figure 69. Deadlock Data window

If the 1 field is Datapage, ensure that all applications (Holder and Waiter) access the resource in the same order.

If the deadlock concerns an index with high insert activity (the Locked resource Type field is Index Page), set the SUBPAGES parameter for the index to 1 or consider converting the index to type 2. If the problem is not related to SUBPAGES, it could be that the commit frequency for an update application is too low. The ratio of updates per commit can be viewed in the Thread Detail panel.

Application development and capacity planning

You can also use OMEGAMON for Db2 PE in application development and capacity planning.

The Explain and the SQL Activity reports can help you anticipate how an application will perform. The Online Monitor Explain and Source Explain functions can assist you in tuning your application’s SQL online.

Accounting and Statistics reports that are ordered by INTERVAL are useful in capacity planning. Use these reports to identify peak periods in system activity as well as the need to adjust system resources. monitoring Db2 subsystems

Chapter 4. Understanding common errors

This section describes problems often encountered when you use OMEGAMON for Db2 PE. If you come across a problem, read this section before contacting IBM service.

STEPLIB missing from IRF-generated JCL

Symptom

The STEPLIB DD statement is missing from the Interactive Report Facility-generated JCL.

Probable Cause

Wrong option selected while running FPEJVARs.

Explanation

FPEJVARs is an EXEC provided with OMEGAMON for Db2 PE that displays a panel for setting up certain defaults. One of the options in this panel specifies whether the OMEGAMON XE for Db2 PE load library is in the system LNKST concatenation. If this option is selected, no STEPLIB DD statement is in the IRF-generated JCL.

Command not found

Symptom

Message IKJ56500I Command xxx not found is issued during the execution of OMEGAMON for Db2 PE.

Probable Cause

The Program Control Facility (PCF) was not updated with the list of OMEGAMON for Db2 PE modules that are started as TSO command processors.

Explanation

The PCF performs a security check on all commands to examine if they are included in the PCF command list. See your system administrator to ensure that all necessary updates have been made. After the list has been updated, the commands can be used.

If you are running the ACF2 security program, the Command Limiting Table needs to be updated by adding FPEFMAIN, FPEMMAIN, and FPEF000.

STEPLIB missing from IRF-generated JCL

Symptom: The STEPLIB DD statement is missing from the Interactive Report Facility-generated JCL.

Probable Cause: Wrong option selected while running FPEJVARs.

Explanation: FPEJVARs is an EXEC provided with OMEGAMON for Db2 PE that displays a panel for setting up certain defaults. One of the options in this panel specifies whether the OMEGAMON for Db2 PE load library is in the system LNKST concatenation. If this option is selected, no STEPLIB DD statement is in the IRF-generated JCL.

Chapter 5. Additional Information

These topics provide additional information that can be helpful when you work with OMEGAMON for Db2 PE.

Layout of the Exception Log File output record

The table in this section outlines the format of the Exception Log File data set output record. Use these terms to interpret the table.

Offset

Length from the beginning of the record to the start of the field.

Length

Length of the field, in bytes.

Db2 Data Type

The format of the value in the field:

- SMALLINT
- INTEGER
- CHAR
- TIMESTAMP
- DOUBLE PRECISION

Offset (decimal)	Length (bytes)	Db2 Data Type	Description
N/A	2	SMALLINT	Length of record (LL)
N/A	2	SMALLINT	Binary zeros
0	1	CHAR	Log record type: <ul style="list-style-type: none">• F = Field-level exception record• M = Exception message (Online Monitor only)
1	1	CHAR	Processing origin ID: <ul style="list-style-type: none">• B = Record written by batch exception processor• S = Record written by online display exception processor• A = Record written by online periodic exception processor
2	4	CHAR	Per: <ul style="list-style-type: none">• SYST = System• PROG = Program• PLAN = Plan name
6	16	CHAR	Local location
22	2	CHAR	Reserved
24	8	CHAR	Primary authorization ID

Table 4. Layout of the Exception Log File output record (continued)

Offset (decimal)	Length (bytes)	Db2 Data Type	Description
32	8	CHAR	Original authorization ID
40	8	CHAR	Plan name
48	8	CHAR	Connection ID
56	8	CHAR	Connection type
64	12	CHAR	Untranslated correlation ID
76	8	CHAR	Correlation name
84	8	CHAR	Correlation number
92	8	CHAR	Network ID
100	8	CHAR	VTAM LU name
108	12	CHAR	Instance number
120	2	SMALLINT	LUW sequence number
122	18	CHAR	Main package (MAINPACK)
140	16	CHAR	Requesting location
156	2	CHAR	Reserved
158	4	CHAR	Db2 subsystem ID
162	6	CHAR	OMEGAMON XE for Db2 PE VnRnMn
168	6	CHAR	Db2 VnRnMn
174	1	SMALLINT	Db2 release code
175	8	CHAR	Store clock timestamp
183	26	TIMESTAMP	Db2 timestamp
209	8	CHAR	Field name
217	1	CHAR	Field type: <ul style="list-style-type: none"> • Blank = Normal field • B = Buffer pool field • D = DDF field • G = Group buffer pool field • C = Global group buffer pool statistics field • P = Package
218	54	CHAR	Field name qualifier used in the Exception Threshold data set
272	2	CHAR	Reserved

Table 4. Layout of the Exception Log File output record (continued)

Offset (decimal)	Length (bytes)	Db2 Data Type	Description
274	54	CHAR	Actual field name qualifier: <ul style="list-style-type: none"> • Buffer pool ID for buffer pool records • Remote location for DDF records • Package location, connection ID, and program name for package records • Blank for all other records
328	2	CHAR	Reserved
330	1	CHAR	Reserved
331	1	CHAR	Field description (without qualifier)
371	40	CHAR	Compare basis: <ul style="list-style-type: none"> • V = Total value • M = By minute • S = By second • C = By commit • T = By thread
372	1	CHAR	Operator (> or <)
373	12	CHAR	Exception value (character)
385	12	CHAR	Threshold value (character)
397	1	CHAR	Exception level: <ul style="list-style-type: none"> • P = Problem • W = Warning
398	8	DOUBLE PRECISION	Exception value (floating point)
406	8	DOUBLE PRECISION	Threshold value (floating point)
414	8	CHAR	Group name
422	8	CHAR	Member name

Note: The exception and threshold values in floating-point form (offsets 398 and 406 respectively) can contain time or counts fields. For time fields, the value is stored in terms of microseconds. For example, if you entered 0.3 seconds for the threshold value, it is stored as 300 000 microseconds in the Exception Log File data set. For counts fields, the value is stored as the actual value.

Product legal notices

This information was developed for products and services offered in the U.S.A.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 give 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 (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. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, 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.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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 Director of Licensing IBM Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

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 information 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

This publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of OMEGAMON for Db2 Performance Expert.

This publication documents information that is NOT intended to be used as Programming Interfaces of OMEGAMON for Db2 Performance Expert.

This publication primarily documents intended Programming Interfaces that allow the customer to write programs to obtain the services of OMEGAMON for Db2 Performance Expert.

This publication also documents information that is NOT intended to be used as Programming Interfaces of OMEGAMON for Db2 Performance Expert. This information is identified where it occurs by an introductory statement to a topic or section.

This publication primarily documents information that is NOT intended to be used as Programming Interfaces of OMEGAMON for Db2 Performance Expert.

This publication also documents intended Programming Interfaces that allow the customer to write programs to obtain the services of OMEGAMON for Db2 Performance Expert. This information is identified where it occurs by an introductory statement to a topic or section.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.html>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux® is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions:

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and the section titled "Cookies, Web Beacons, and Other Technologies" in IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details>. Also, see the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Index

A

- accessibility
 - overview [4](#)
- application development [69](#)
- Asynchronous Task Termination panel [13](#)
- authorization
 - changing SQL authorization [12](#)
 - Explain function [43](#)
 - requirements [5](#)
- authorization failure
 - listing [39](#)
- Authorization Failure Summary window [39](#)
- Auto Display window [11](#)
- AUTOEXCPHNAME
 - OMEGAMON Collector startup parameter [33](#)

C

- capacity planning [69](#)
- COLLECT command [17](#)
- command
 - COLLECT [17](#)
 - DB2 command [16](#)
 - EXPLAIN [54](#)
 - LOOK [36](#)
 - OPTIONS [17](#)
 - PANELID [15](#)
 - REINIT [36](#)
 - source EXPLAIN [57](#)
 - system parameters [58](#)
 - understanding syntax diagrams [1](#)
- command not found [71](#)
- Command Output window [16](#)
- Command window [16](#)
- common errors [71](#)
- conventions [2](#)
- cookie policy [77](#), [79](#)
- correlation translation
 - default translation [18](#)
- Create Plan Table window [44](#)
- Current SQLID window [12](#)

D

- data set
 - exception processing output [41](#)
- Data Set Maintenance menu
 - defining exception thresholds [59](#)
- database request module (DBRM), explaining [50](#)
- date format
 - Online Monitor [20](#)
- DB2 command
 - description [16](#)
- DB2 Command Output window [16](#)
- DB2 Command window [16](#)
- DB2 Explain Output panel [52](#)

- DB2 subsystem windows [7](#)
- DB2 subsystems, monitoring
 - controlling the writing of monitor output [10](#)
 - exception processing [23](#)
 - exception profiling [25](#)
 - exception report [27](#)
 - identifying unusual situations [25](#)
 - performance objectives [25](#)
 - tuning the operating environment [25](#)
- DB2 System Parameters panel [57](#)
- DBRM (database request module)
 - explaining [50](#)
- DBRM (database request module), explaining [50](#)
- DBRM and Plan List panel [51](#)
- DBRM SQL Statement window [50](#)
- DGOFPTS panel [11](#)
- DGOMCCMD panel [16](#)
- DGOMCDCO panel [16](#)
- DGOMDPMN panel [7](#)
- DGOMDWAD panel [11](#)
- DGOMDWH panel [9](#)
- DGOMDWMO panel [10](#)
- DGOMDWSS panel [12](#)
- DGOMEPO2 panel [33](#), [61](#)
- DGOMEPO3 panel [34](#)
- DGOMEPE1 panel [31](#)
- DGOMGP01 panel [13](#)
- DGOMHWHY panel [29](#)
- DGOMLAFS panel [39](#)
- DGOMLAMP panel [38](#)
- DGOMLAXP panel [37](#)
- DGOMLED0 panel [69](#)
- DGOMLEMP panel [40](#)
- DGOMLEXP panel [40](#)
- DGOMLMSP panel [37](#)
- DGOMLSXP panel [39](#)
- DGOMYPKM panel [52](#)
- DGOMYPMN panel [46](#)
- DGOMYWDL panel [51](#)
- DGOMYWDQ panel [51](#)
- DGOMYWDS panel [50](#)
- DGOMYWES panel [51](#)
- DGOMYWII panel [67](#)
- DGOMYWKQ panel [49](#)
- DGOMYWKS panel [48](#)
- DGOMYWKV panel [49](#)
- DGOMYWPE panel [47](#)
- DGOMYWPL panel [48](#)
- DGOMYWQC panel [44](#)
- DGOMYWQU panel [45](#)
- DGOMYWSS panel [57](#)
- DGOMYWTI panel [67](#)
- DGOMZPSM panel [57](#)
- DGOMPENU panel
 - defining exception thresholds [59](#)
- Display Exception List window [39](#)
- display exception processing

display exception processing (*continued*)

description [30](#)

starting [35](#)

viewing exceptions [39](#)

DPMOUT data set

exception processing [42](#)

specifying exception processor name [34](#)

DPMPARMS data set

CORRDATA member [18](#)

DSNZPARM value

viewing summary [57](#)

E

elapsed time format

Online Monitor [20](#)

errors

command not found [71](#)

STEPLIB missing from IRF-generated JCL [71](#)

exception DPMOUT data set

attributes [10](#), [42](#)

Exception Event Messages window [40](#)

exception event processing

description [31](#)

listing messages [40](#)

starting [36](#)

viewing events [40](#)

Exception Event Summary window [40](#)

Exception Log File data set

attributes [41](#)

description [41](#)

output record layout [73](#)

printing [41](#)

sample files [41](#)

specifying processor name [34](#)

Exception Log File print utility [41](#)

Exception Notification window [31](#)

Exception Output Data Sets window [34](#)

exception processing

activating and deactivating [33](#), [61](#)

description [29](#)

display exception

description [30](#)

starting [35](#)

viewing [39](#)

exception event

description [31](#)

starting [36](#)

viewing [40](#)

modes [30](#)

output data set [41](#)

performance problem identification

example [59](#)

periodic exception

description [30](#)

listing messages [38](#)

starting [36](#)

viewing [37](#)

restarting [36](#)

starting [32](#), [36](#)

starting from Exception Processor panel [33](#)

starting with OMEGAMON Collector [33](#)

stopping [42](#)

viewing messages and status [36](#), [41](#)

exception processing (*continued*)

viewing past data [31](#)

exception processing user exit [34](#)

Exception Processor panel [33](#), [61](#)

exception profiling

performance problem identification [25](#)

exception report

monitoring [27](#)

exception threshold

defining [32](#)

Exception Threshold data set

building [34](#)

reinitializing [33](#)

specifying in Exception Processor panel [34](#)

specifying thresholds in [32](#)

threshold values in [30](#)

using at OMEGAMON Collector startup [33](#)

EXCEPTIONEVENT

OMEGAMON Collector startup parameter [33](#)

EXEC

FPEJINIT [5](#)

Expand function key for scrollable fields [21](#)

explain

authorization [43](#)

DBRM [50](#)

description [42](#)

method [45](#)

output [52](#)

performance considerations [43](#)

plan table considerations [44](#)

plan table entry [47](#)

process overview [45](#)

specifying the SQL statement [51](#)

SQL statement [45](#)

viewing access path method [52](#)

EXPLAIN command [54](#), [57](#)

Explain menu [46](#)

Explain Output panel [52](#)

F

filtering input [25](#)

FPEJINIT EXEC [5](#)

function key

commands to change the defaults [15](#)

defaults [13](#)

for scrollable fields [21](#)

H

history

exception processing [31](#)

HISTORY command [28](#)

History Defaults window [9](#)

History panel [29](#)

I

IFCID 337

on Exception Event Summary panel [40](#)

IFCID 359

on Exception Event Summary panel [40](#)

Index Selection window [67](#)

Index Split exception event [40](#)
Interactive Report Facility (IRF)
 session options [11](#)
IRF (Interactive Report Facility)
 session options [11](#)
ISPF Online Monitor
 source explain [54](#)

L

large value
 Online Monitor [19](#)
layout
 Exception Log File data set output [73](#)
Left function key for scrollable fields [21](#)
legal notices
 cookie policy [77, 79](#)
 notices [77](#)
 programming interface information [77, 78](#)
 trademarks [77–79](#)
limiting input [25](#)
Lock Escalation exception event [40](#)
LOOK command [36, 37](#)
Look Selections menu [37](#)

M

menu
 data set maintenance
 defining exception thresholds [59](#)
 Explain [46](#)
 Look Selections [37](#)
 Options [7](#)
message
 exception processing [36, 41](#)
 IKJ56500I [71](#)
migration
 plan table [45](#)
missing value
 Online Monitor [19](#)
Monitor Output window [10](#)
monitoring DB2 subsystems
 controlling the writing of monitor output [10](#)
 exception processing [23](#)
 exception profiling [25](#)
 exception report [27](#)
 identifying unusual situations [25](#)
 performance objectives [25](#)
 tuning the operating environment [25](#)
monitoring using reports [26](#)

N

notices [77, 78](#)

O

Online Monitor
 accessing [5](#)
 description [24](#)
 overview [5](#)
 using [5, 20](#)
Online Monitor options [7](#)

options [6](#)
OPTIONS command [17](#)
Options menu [7](#)
overview [1](#)

P

package
 explain an SQL statement [48](#)
Package SQL Statement window [48](#)
Package Version List panel [49](#)
panel
 Asynchronous Task Termination [13](#)
 Authorization Failure Summary [39](#)
 DB2 Explain Output [52](#)
 DB2 System Parameters [57](#)
 Exception Event Messages [40](#)
 Exception Event Summary [40](#)
 Exception Processor [33, 61](#)
 moving between Online Monitor panels [12](#)
 moving between panels [12](#)
 SQL Statement List (packages) [49](#)
 SQL Statement Selection [57](#)
 SQL Text Edit [52](#)
PANELID command [15](#)
past data
 exception processing [31](#)
performance problem identification
 controlling the writing of monitor output [10](#)
 deadlock resolution [68, 69](#)
 examples
 exception processing [59](#)
 exception processing [23](#)
 exception profiling [25](#)
 exception report [27](#)
 identifying unusual situations [25](#)
 performance objectives [25](#)
 tuning the operating environment [25](#)
Periodic Exception Messages window [38](#)
periodic exception processing
 description [30](#)
 listing messages [38](#)
 starting [36](#)
 viewing exceptions [37](#)
Periodic Exceptions List window [37](#)
plan table
 considerations [44](#)
 explaining existing entry [47](#)
 upgrading [45](#)
Plan Table Entry List panel [48](#)
Plan Table Entry window [47](#)
programming interface information [77, 78](#)

R

REINIT command [36](#)
Right function key for scrollable fields [21](#)
RKO2SAMP library
 DGOECFIL member [41](#)
 DGOELFIL member [42](#)
 DGOEQFIL member [42](#)
 DGOEVFIL member [42](#)
 DGOMEJCL member [41](#)

RKO2SAMP library (*continued*)
DGOMMJCL member [29](#)

S

sample file
 Exception Log File data set [41](#)
screen readers and magnifiers [4](#)
scrollable field
 examples [20](#)
scrollable field support [20](#)
service information [1](#)
Session Options window [11](#)
source explain
 ISPF Online Monitor [54](#)
source EXPLAIN command [57](#)
Source Explain Options window [55](#)
SQL statement
 explain function [42](#)
SQL Statement List panel (DBRMs) [51](#)
SQL Statement List panel (packages) [49](#)
SQL Statement Selection panel [57](#)
SQL Text Edit panel [52](#)
SQL Text Entry window [51](#)
SQLID window [12](#)
starting
 OMEGAMON for DB2 PE [5](#)
status
 exception processing [36](#), [41](#)
STEPLIB missing from IRF-generated JCL [71](#)
support information [1](#)
syntax [16](#), [17](#), [28](#), [36](#), [37](#), [57](#)
syntax diagram
 source EXPLAIN [57](#)
syntax diagrams [1](#)
System Parameters panel [57](#)
system parameters reports
 commands [58](#)

T

Table Information window [67](#)
terminology used [3](#)
time format
 Online Monitor [20](#)
toleration mode
 IFCID 337 [40](#)
 IFCID 359 [40](#)
trademarks [77–79](#)

U

Upgrade Plan Table window [45](#)

V

value
 large and missing [19](#)
 Online Monitor [19](#)
 refreshing
 Online Monitor values periodically [10](#)
viewing past data [28](#)

W

window
 Auto Display [11](#)
 Create Plan Table [44](#)
 Current SQLID [12](#)
 DB2 Command [16](#)
 DB2 Command Output [16](#)
 DBRM and Plan List [51](#)
 DBRM SQL Statement [50](#)
 Display Exception List [39](#)
 Exception Notification [31](#)
 Exception Output Data Sets [34](#)
 History [29](#)
 History Defaults [9](#)
 Index Selection [67](#)
 Monitor Output [10](#)
 Package SQL Statement [48](#)
 Package Version List [49](#)
 Periodic Exception Messages [38](#)
 Periodic Exceptions List [37](#)
 Plan Table Entry [47](#)
 Plan Table Entry List [48](#)
 Session Options [11](#)
 SQL Statement List (DBRMs) [51](#)
 SQL Text Entry [51](#)
 Table Information [67](#)
 Upgrade Plan Table [45](#)



Product Number: 5655-W37

SH12-7070

