

5.5

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



**2025-07-04 edition**

This edition applies to IBM® OMEGAMON for Db2® Performance Expert on z/OS® (product number 5655-W37) and to any subsequent releases until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

Specific changes are indicated by a vertical bar to the left of a change. A vertical bar to the left of a figure caption indicates that the figure has changed. Editorial changes that have no technical significance are not noted.

© **Copyright International Business Machines Corporation 2005, 2025.**

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

© **Rocket Software Inc. 2016, 2025.**

---

# Contents

<b>About this information.....</b>	<b>V</b>
<b>Chapter 1. Overview.....</b>	<b>1</b>
Service updates and support information.....	1
How to read syntax diagrams.....	1
Conventions.....	2
Terminology.....	3
Accessibility features.....	4
<b>Chapter 2. About ISPF Online Monitor.....</b>	<b>5</b>
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
Monitor Output window.....	9
Session options.....	10
Current SQLID window.....	10
Moving between panels.....	10
Leaving the ISPF Online Monitor.....	11
Function key default settings.....	11
Issuing Db2 commands.....	14
Using the COLLECT command.....	16
Using the OPTIONS command.....	16
Correlation ID translation.....	16
The default translation.....	17
Changing the default translation.....	17
How large and missing values are displayed.....	17
Elapsed time formats.....	18
Date formats.....	19
Displaying long names and values in scrollable fields.....	19
<b>Chapter 3. Monitoring and tuning.....</b>	<b>21</b>
Changed access path to ISPF Online Monitor.....	22
Performance data generation and online monitoring.....	22
The Db2 operating environment.....	22
Performance objectives and exception processing.....	23
Monitoring performance.....	23
Monitoring using reports.....	24
Observing performance trends.....	24
Determining performance problems.....	24
Problems detected in periodic exception processing.....	24
Problems detected in exception event processing.....	25
Problems detected in exception reports.....	25
Exception processing.....	26

Exception processing modes.....	26
Exception Notification window.....	27
How to define exception threshold values.....	28
How to start exception processing.....	28
How to restart exception processing (REINIT command).....	33
Examining exception messages (LOOK command).....	33
Exception processor output data sets.....	37
Stopping exception processing.....	38
Exception processing exit.....	38
Monitoring the access path with Explain.....	40
Authorizations required to use Explain.....	40
Performance considerations.....	40
What to look for in Explain information.....	40
Plan table considerations.....	41
Explain an SQL statement.....	42
Explain Menu.....	43
Explain output (Db2 Explain Output panel).....	49
Online Monitor EXPLAIN command.....	51
ISPF Online Monitor Source Explain.....	51
Problem identification examples.....	55
Example of exception processing.....	55
Using the ISPF Online Monitor in deadlock resolution.....	58
Application development and capacity planning.....	59
<b>Chapter 4. Understanding common errors.....</b>	<b>61</b>
STEPLIB missing from IRF-generated JCL.....	61
<b>Chapter 5. Additional Information.....</b>	<b>63</b>
Layout of the Exception Log File output record.....	63
<b>Notices.....</b>	<b>67</b>
<b>Index.....</b>	<b>71</b>

## About this information

---

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

The document is part of the OMEGAMON for Db2 PE documentation library which provides instructions for installing, configuring, and using OMEGAMON for Db2 PE 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 PE
- Install and operate OMEGAMON for Db2 PE
- Customize your OMEGAMON for Db2 PE environment
- Diagnose and recover from OMEGAMON for Db2 PE problems
- Design and write applications for OMEGAMON for Db2 PE
- Use OMEGAMON for Db2 PE with other Db2 products



---

# Chapter 1. Overview

IBM OMEGAMON for Db2 Performance Expert on z/OS (OMEGAMON for Db2 PE) allows 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 [OMEGAMON for Db2 PE: Web-based delivery and updates for 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:



Table 1. Symbol usage	
Symbol	Usage
	<p>The <b>or</b> symbol is used to denote a choice. You can use the argument on the left or the argument on the right. For example:</p> <pre>YES   NO</pre> <p>In this example, you can specify YES or NO.</p>
()	<p>Denotes optional arguments. Arguments that are not enclosed in square brackets are required. For example:</p> <pre>APPLDEST DEST (ALTDEST)</pre> <p>In this example, DEST is a required argument and ALTDEST is optional.</p>
{ }	<p>Some documents use braces to denote mandatory arguments, or to group arguments for clarity. For example:</p> <pre>COMPARE {workload} - REPORT={SUMMARY   HISTOGRAM}</pre> <p>In this example, the workload variable is mandatory. The REPORT keyword must be specified with a value of SUMMARY or HISTOGRAM.</p>
–	<p>Default values are underscored. For example:</p> <pre>COPY infile outfile - [COMPRESS={YES   NO}]</pre> <p>In this example, the COMPRESS keyword is optional. If specified, the only valid values are YES or NO. If omitted, the default is YES.</p>

## Notation conventions

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

### *hilev*

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

For example:

- *thilev* refers to the high-level qualifier for your target data set.
- *rhilev* refers to the high-level qualifier for your runtime data set.

For members in target libraries, the high-level qualifier is *thilev* rather than *rhilev*.

- *shilev* refers to the SMP/E library high-level qualifier.

## 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.

Table 2. Product names and their short names	
Product name	Short name
IBM OMEGAMON for Db2 Performance Expert on z/OS	OMEGAMON for Db2 PE
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

To invoke the ISPF Online Monitor, select 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 enhanced 3270 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 IBM Documentation](#) 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 40](#) 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 z/OS 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 ==> _____ Scroll ==> CSR
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 z/OS 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 are not available:

- 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.

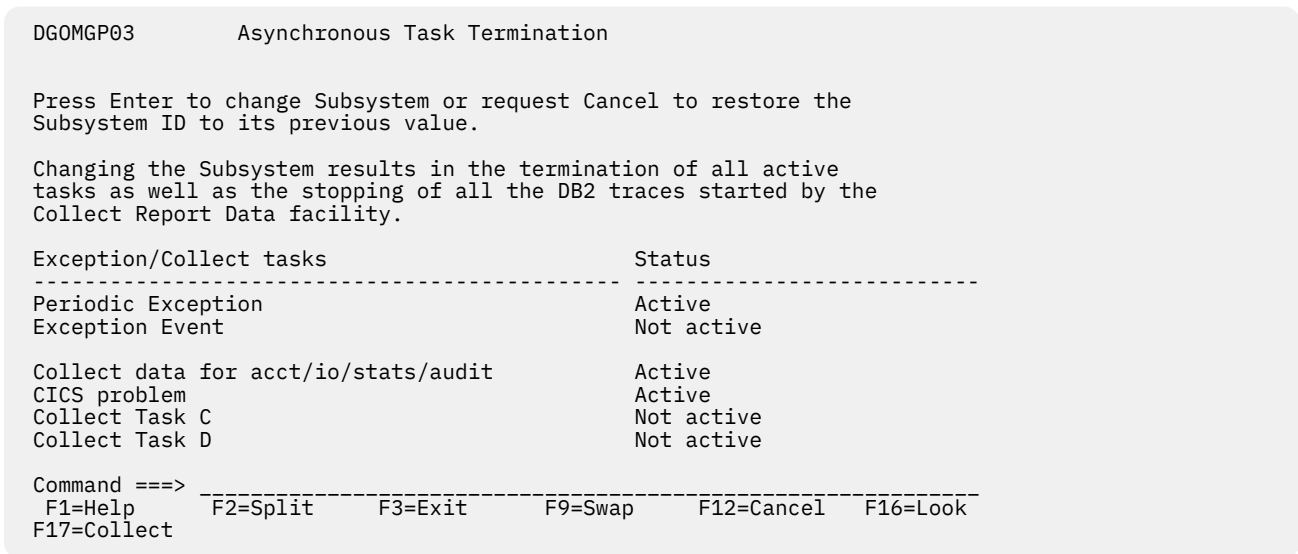


Figure 5. Asynchronous Task Termination panel

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.

## 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.

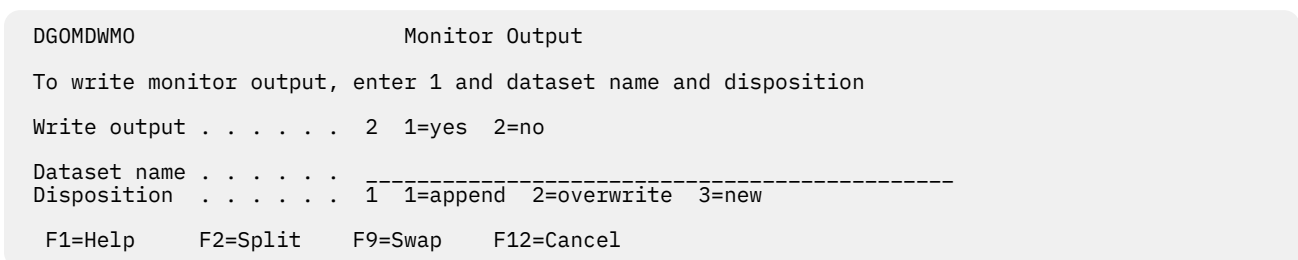


Figure 6. Monitor Output window

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

## 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.

```
DGOFOPTS                Session Options

Update fields as required, then press Enter.

Confirmation display . . . . . 1  1=yes 2=no
Initial menu choice . . . . . _  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 7. Session Options window

The field values shown in [Figure 7 on page 10](#) 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 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 8. 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...

## 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 9. 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 3 on page 12 describes the Online Monitor function keys, their default settings, their functions, and the Online Monitor panels on which they are available.

Table 3. 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
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	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

Table 3. Function key default settings (continued)			
Function key	Default setting	Function	Online Monitor panels
F16	LOOK	Used to view the following exception information and authorization failures: <ul style="list-style-type: none"> <li>• Periodic exceptions</li> <li>• Periodic exceptions messages</li> <li>• Display exceptions</li> <li>• Authorization failure summary</li> <li>• Exception event summary</li> <li>• Exception event messages</li> </ul>	All panels (except LOOK and related 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 10. 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 11 on page 15](#).

```

ROW 1 TO 5 OF 5

PM01DL0C          DSN1 V10          GROUP001 MEMBER01

DGOMCDCO          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 11. 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 12. 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 13. 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.

Table 4. The 12-Byte Correlation ID field and the default translation												
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 14 on page 19 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.

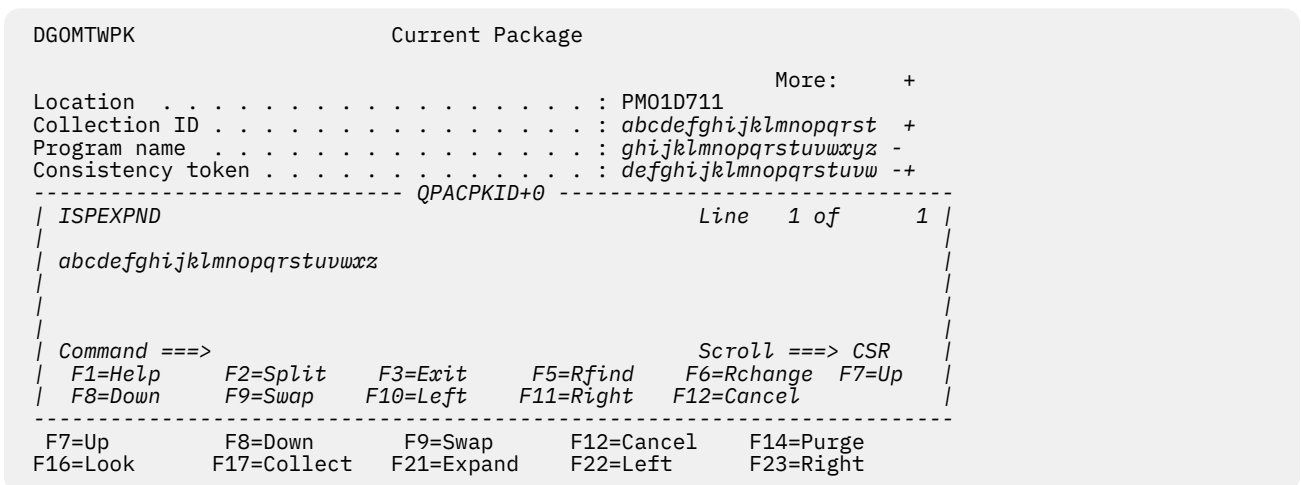


Figure 14. Fictitious panel with scrollable field variations

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 determine performance problems by using ad hoc Db2 trace collection and by creating and running batch reports in OMEGAMON for Db2 PE. Such tasks include monitoring exceptions and explaining SQL statements.

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. Exception Processing 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 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 ISPF Online Monitor

When you start IBM 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 15. 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 z/OS

system itself. Therefore the environment in which the Db2 subsystem is operating should be tuned before Db2 is tuned.

For example, if the z/OS 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 z/OS 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 z/OS

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 IBM Documentation\*](#). 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

Use reports or traces to diagnose and identify the root cause of 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.

## Exception processing

---

Exception processing is the most effective way to identify 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 that are outside defined thresholds. This allows better management of service levels by identifying problems in the Db2 subsystem and threads that are causing performance problems.

By using 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, you can 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 two basic types of exception processing available: periodic exception processing and exception event processing.

### 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. The next time you use the Online Monitor, 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 EVENTOBSERVATION 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 IBM Documentation\*](#).

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.

## 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.

The Exception Notification 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.

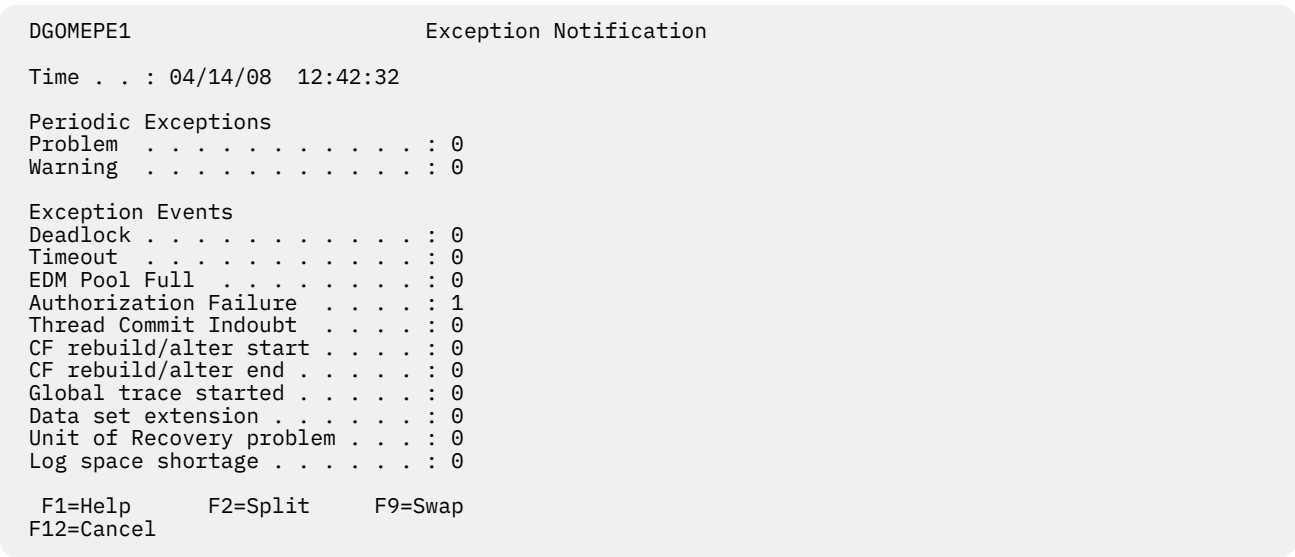


Figure 16. Exception Notification window (exception event processing active)

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

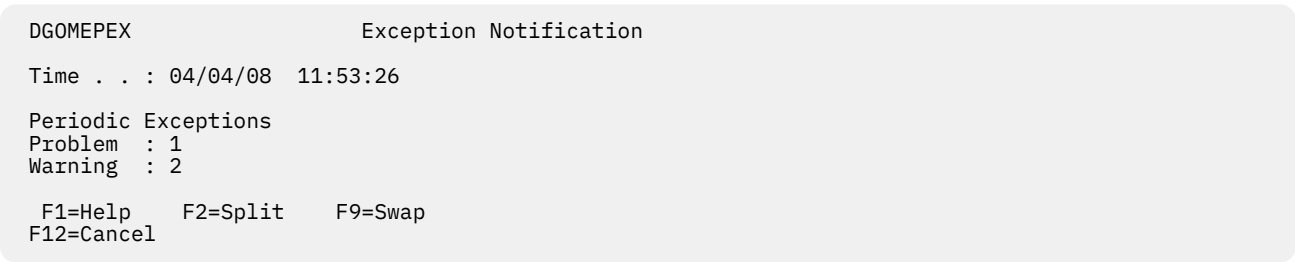


Figure 17. Exception Notification window (exception event processing inactive)

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

Use this information to understand where exception thresholds are defined.

For information about defining exception threshold values, see the [Reporting User's Guide](#).

How to start exception processing

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

The Exception Processor panel is shown in [Figure 18 on page 29](#).

```

DG0MEP02                Exception Processor                PM01DL0C DSN1 V10

For any field enter any character to activate

Activate/Deactivate Exception Processing
- Periodic                >      User Exit
- Exception event notification

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 18. 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.

Depending 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 using this method, 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 **EVENTOBSERVATION** 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 method for 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 to ensure that the new exception criteria is recognized.

For more information, see “How to restart exception processing (REINIT command)” on page 33.

**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. The Exception Processor panel is shown in [Figure 19 on page 30](#).

```

DG0MEP02                               Exception Processor                               PM01DL0C DSN1 V10
For any field enter any character to activate
Activate/Deactivate Exception Processing
- Periodic                               >      User Exit
- Exception event notification
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 19. Exception Processor panel

Use the 1 field to specify whether the 1 field value is displayed 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* level 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. For more information, see “Exception Output Data Sets window” on page 31.

**Tip:** 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 IBM Documentation](#).

### **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 20. 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 37](#) for more information.

For the layout of the Exception Log File data set, see [“Layout of the Exception Log File output record” on page 63](#).

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 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 21 on page 32](#).

DGOMEPO2	Exception Processor	PM01DL0C DSN1 V10
For any field enter any character to activate		
		More: - +
Activate/Deactivate Exception Processing		
- Periodic	>	User Exit
- Exception event notification		
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
F12=Cancel	F16=Look	F17=Collect
	F7=Up	F8=Down
		F9=Swap

*Figure 21. Exception Processor panel*

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 22 on page 32](#).

DGOMEPO2	Exception Processor	PM01DL0C DSN1 V10
For any field enter any character to activate		
		More: - +
Activate/Deactivate Exception Processing		
- Periodic	>	User Exit
- Exception event notification		
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
F12=Cancel	F16=Look	F17=Collect
	F7=Up	F8=Down
		F9=Swap

*Figure 22. Exception Processor panel*

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

The events must be specified with the EVENTOBSERVATION 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 19 on page 30](#). 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 23. 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:

➤ LOOK ➤

1
2
3
4
5
6

*Figure 24. Syntax of the LOOK command*

The LOOK parameters 1 to 6 correspond to the options on the Look Selections menu in [Figure 25 on page 34](#).

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 25. 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     Corrnmb
   Descr      Value      Compare   Threshold     Type        By
-----
- 08:35:36 PM01DL0C      GROUP001  N/P          MEMBER01
   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 26. 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.

Pressing F3 (Exit) or F12 (Cancel) returns you to the Periodic Exceptions List window.

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 27. 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 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 28. 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

You can view details about exception events from the Exception Event Summary window.

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 29. 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

You can view messages issued by the exception event processor from the Exception Event Messages window.

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 30. 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 63 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.

The Exception Processor panel is shown in [Figure 19 on page 30](#).

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. For more information, see [“Leaving the ISPF Online Monitor” on page 11](#).

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. For more information, see [“Db2 Subsystem windows” on page 7](#).

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.

## Exception processing exit

Exception processing exits are user exits that handle periodic exceptions and event exceptions.

You can activate or deactivate exception processing exits from the Exception Processor panel with the Host Online Monitor (as described in [IBM Db2 for z/OS in IBM Documentation](#)) or from the Exception Processing window with the Workstation Online Monitor.

The exception processing exit is implemented as an asynchronous task. It is provided as a load module named DGOMEPT, which contains the modules DGOMEPT, DGOMUPXT, and DGOVQUEU.

Module DGOMEPTX is posted whenever the periodic exception task or the event exception task detects an exception condition. It receives the queued exception log file records using the services of module DGOVQQUEU and calls the user exit routine DGOMUPXT.

In RKO2SAMP, 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 change DGOMUPXT according to your needs, for example, to examine the type of exception and perform an action necessary to handle the situation.

The exception processing exit runs in the Data Collector address space as a subfunction of periodic exception processing. Therefore, the exit remains active even if the administrator exits OMEGAMON for Db2 PE for as long as periodic exception processing remains active in the Data Collector address space.

## Generating a NetView alert

You can generate a NetView generic alert for the exception message by using the sample CLIST DGOMSPXT, which is provided in RKO2SAMP.

### Procedure

1. Update the NetView automation table and add the CLIST to a NetView DSICLD concatenated data set.

The following example shows how you can update the NetView automation table:

```
*****
* ENTRY FOR DB2 PM EXCEPTION MESSAGE PROCESSING *
*****
SYN %AOFINGOP1% = 'AUTO';
IF MSGID='DGOV0100I' & TEXT=MESSAGE
  THEN EXEC (CMD('DGOMSPXT' MESSAGE)
    ROUTE (ONE %AOFINGOP1%));
```

2. Modify the sample CLIST DGOMSPXT according to your installation requirements.

## Customizing and activating the exception processing exit

You can customize and activate the exception processing exit by completing the steps in the procedure.

### Procedure

1. Using Assembler language, change the user exit routine DGOMUPXT according to your installation requirements. For more information about running this user exit, such as the layout of the input area or the parameter list that is passed along in register 1, refer to the description in the source code of the exit.

**Tip:** Rather than using your own JCL, it is recommended that you use DGOMJPXT, which is supplied in RKO2SAMP, to assemble and link-edit the user exit routine you have modified. The execution of DGOMJPXT includes SMP/E RECEIVE and APPLY. This method allows you to keep track of your modifications, which is helpful in case future changes of the product affect this function. During the assembly of the user exit routine, macro DGOFLGD is accessed, which is also supplied in RKO2SAMP. View the expansion of this macro as it contains the field names used to map the exception log file record.

2. Stop and restart the Data Collector to activate the new load module.
3. Activate the exception processing exit using the Exception Processor panel.

## 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 42](#) 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 42](#) 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 DSNDDB04. 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 . . . . . DSNDDB04
Table space name . . . . .

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

Figure 31. 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 32. Upgrade Plan Table window

## Explain an SQL statement

You can explain an SQL statement by various actions.

- 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 51](#) for more information.

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

[Figure 33 on page 43](#) shows the flow of the Explain process.






Figure 33. 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 SPUFI 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 51.

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 49.

## 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 34. 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 . . . . . DG0@TPG1_____ +
Version . . . . . _____ +

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

```

Figure 35. 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-----+-----+
_      1011001          DG0@TPG1      OMPE_FINAL    2011021615243488
_      1011001          DG0@TPG1      OMPE_FINAL    2011030811011901
_      1011001          DG0@TPG1      0510_PM32647C 2011092016072777
_      1011001          DG0@TPG1      0510_PM32647C 2012012413093911
_      1011001          DG0@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 36. 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 . . . . . DG0@TPG*_____+
Version . . . . . _____+

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

```

Figure 37. 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 . . . . . : PM0DB2A          SDA2  V10
Current server . . . . . : PM0DB2A          SDA2  V10

S   Collection      Package      Version
----+-----1----+-----2-----+-----1-----+-----2-----+
-   K02EX510        DG0@TPG1      OMPE_FINAL
-   K02EX510        DG0@TPG1      0510_PM32647C
-   K02EX510        DG0@TPG2      OMPE_FINAL
-   K02EX510        DG0@TPG2      0510_PM32647C
-   K02EX510        DG0@TPG3      0510_PM32647C
-   K02EX510        DG0@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 38. 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 39. 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 40. 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 41. 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 . SYSDDUMMY1
-   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 42. 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 43. 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 44. 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  DSNKXX01

- 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 45. 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 49](#). 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 46. 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 49](#).

## 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 47. 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 . . . . . -----

7 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 48. 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.

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

## SQL Statement Selection panel

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
	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

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 49.

## Source EXPLAIN command

The command syntax looks like this:



54 ISPF Client User's Guide

## 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 51. 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 52. 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 53. 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.

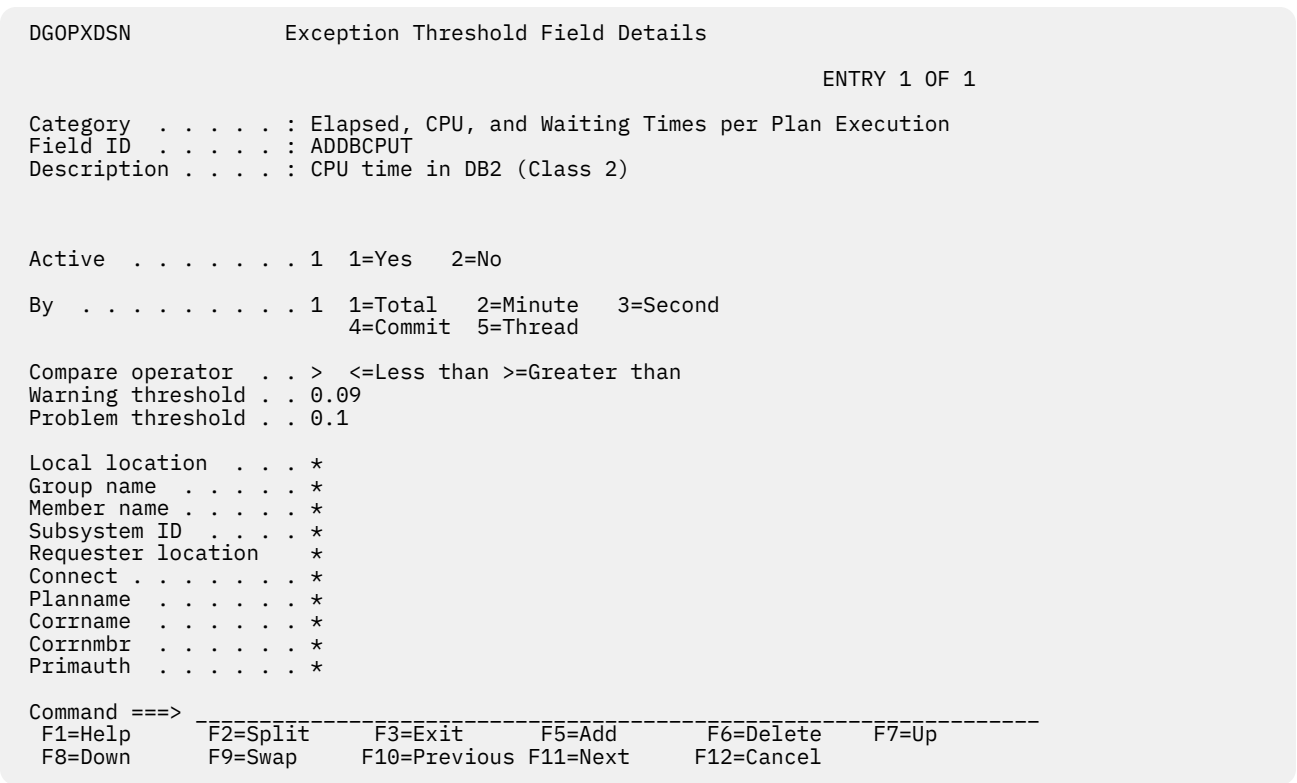


Figure 54. Exception Threshold Field Details panel

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 28.

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.

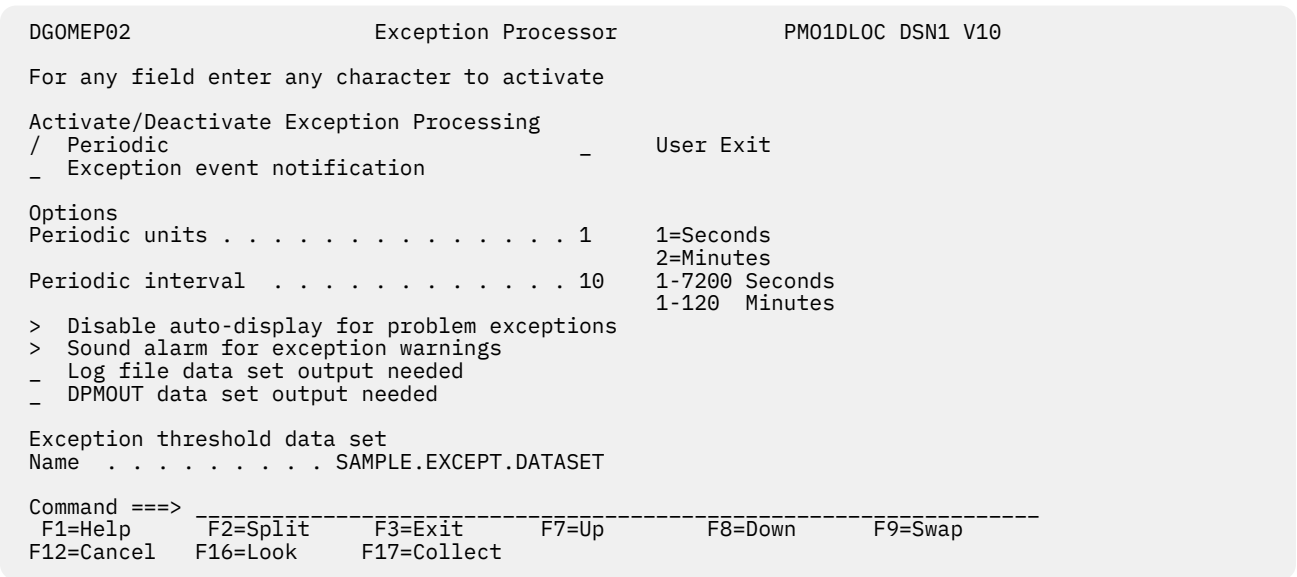


Figure 55. Exception Processor panel

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 28](#).

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	Corrname
Field	Reqloc	Primauth	Planname	Connect	Corrnmb
Descr	Value	Compare	Threshold	Type	By
-----					
- 14:27:41	PM01DL0C	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 56. 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.

## 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 28](#).

### Deadlock resolution

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



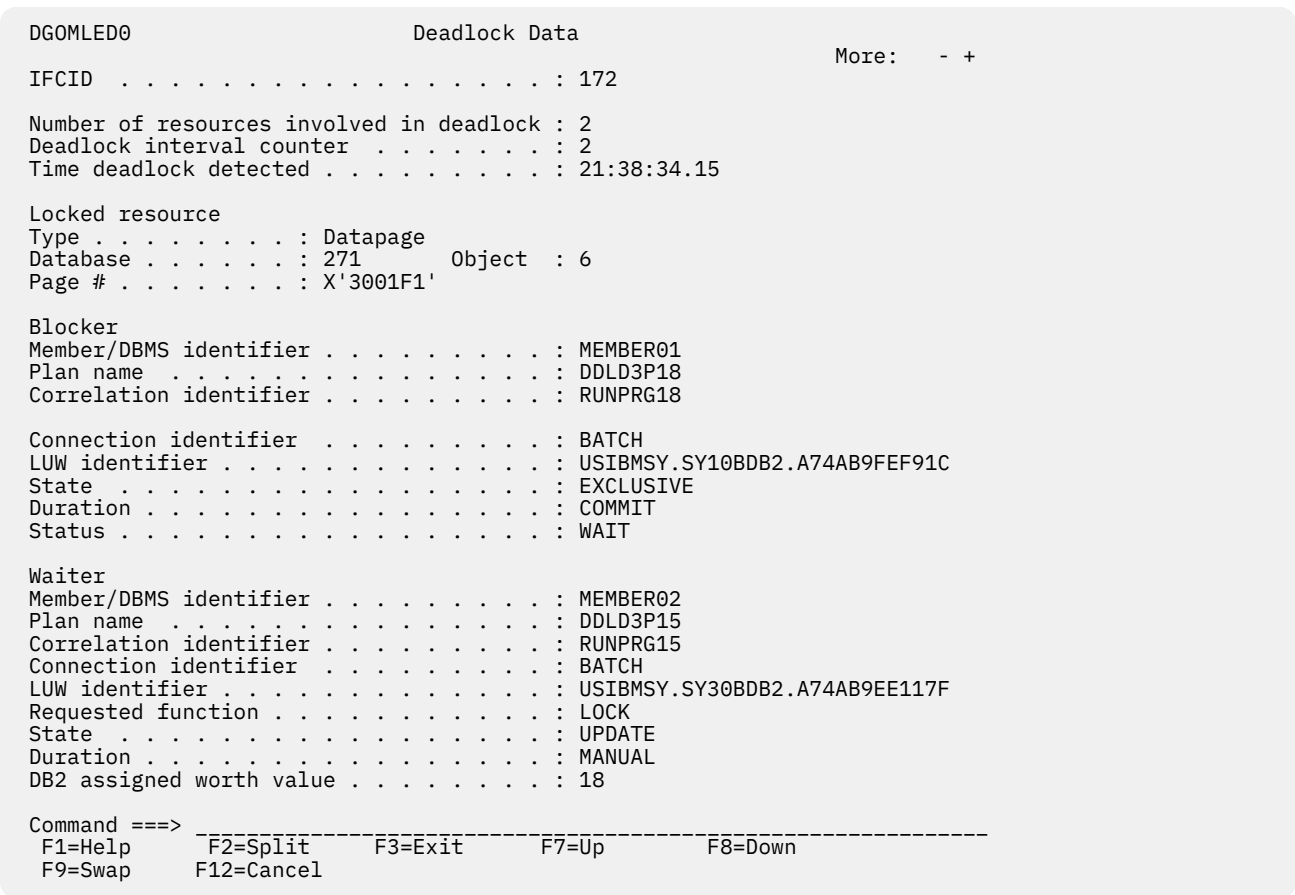


Figure 57. Deadlock Data window

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 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

Table 5. Layout of the Exception Log File output record			
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"><li>• F = Field-level exception record</li><li>• M = Exception message (Online Monitor only)</li></ul>
1	1	CHAR	Processing origin ID: <ul style="list-style-type: none"><li>• B = Record written by batch exception processor</li><li>• S = Record written by online display exception processor</li><li>• A = Record written by online periodic exception processor</li></ul>
2	4	CHAR	Per: <ul style="list-style-type: none"><li>• SYST = System</li><li>• PROG = Program</li><li>• PLAN = Plan name</li></ul>
6	16	CHAR	Local location
22	2	CHAR	Reserved
24	8	CHAR	Primary authorization ID

Table 5. 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 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"> <li>• Blank = Normal field</li> <li>• B = Buffer pool field</li> <li>• D = DDF field</li> <li>• G = Group buffer pool field</li> <li>• C = Global group buffer pool statistics field</li> <li>• P = Package</li> </ul>
218	54	CHAR	Field name qualifier used in the Exception Threshold data set
272	2	CHAR	Reserved

Table 5. 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"> <li>• Buffer pool ID for buffer pool records</li> <li>• Remote location for DDF records</li> <li>• Package location, connection ID, and program name for package records</li> <li>• Blank for all other records</li> </ul>
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"> <li>• V = Total value</li> <li>• M = By minute</li> <li>• S = By second</li> <li>• C = By commit</li> <li>• T = By thread</li> </ul>
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"> <li>• P = Problem</li> <li>• W = Warning</li> </ul>
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.





## Notices

---

This information was developed for products and services offered in the US. This material might 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, MD-NC119  
Armonk, NY 10504-1785 US

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

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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, MD-NC119  
Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 actual people or business enterprises 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 shown below:

© (*your company name*) (*year*).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. (*enter the year or years*).

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 IBM OMEGAMON for Db2 Performance Expert on z/OS.

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

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

This publication also documents information that is NOT intended to be used as Programming Interfaces of IBM OMEGAMON for Db2 Performance Expert on z/OS. 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 IBM OMEGAMON for Db2 Performance Expert on z/OS.

This publication also documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM OMEGAMON for Db2 Performance Expert on z/OS. 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 marks 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 <https://www.ibm.com/legal/copyright-trademark>.

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.

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

## 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 <https://www.ibm.com/us-en/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and 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 [59](#)
- Asynchronous Task Termination panel [11](#)
- authorization
  - changing SQL authorization [10](#)
  - Explain function [40](#)
  - requirements [5](#)
- authorization failure
  - listing [35](#)
- Authorization Failure Summary window [35](#)
- AUTOEXCPHNAME
  - OMEGAMON Collector startup parameter [29](#)

## C

- capacity planning [59](#)
- COLLECT command [16](#)
- command
  - COLLECT [16](#)
  - DB2 command [14](#)
  - EXPLAIN [51](#)
  - LOOK [33](#)
  - OPTIONS [16](#)
  - PANELID [14](#)
  - REINIT [33](#)
  - source EXPLAIN [54](#)
  - understanding syntax diagrams [1](#)
- command not found [61](#)
- Command Output window [14](#)
- Command window [15](#)
- common errors [61](#)
- conventions [2](#)
- cookie policy [67](#), [69](#)
- correlation translation
  - default translation [17](#)
- Create Plan Table window [42](#)
- Current SQLID window [10](#)

## D

- data set
  - exception processing output [37](#)
- Data Set Maintenance menu
  - defining exception thresholds [55](#)
- database request module (DBRM), explaining [47](#)
- date format
  - online monitoring [19](#)
- DB2 command
  - description [14](#)
- DB2 Command Output window [14](#)
- DB2 Command window [15](#)
- DB2 Explain Output panel [49](#)
- DB2 subsystem windows [7](#)
- Db2 subsystems, monitoring

- Db2 subsystems, monitoring (*continued*)
  - exception processing [21](#)
- DB2 subsystems, monitoring
  - controlling the writing of monitor output [9](#)
  - exception profiling [23](#)
  - exception report [25](#)
  - identifying unusual situations [23](#)
  - performance objectives [23](#)
  - tuning the operating environment [22](#)
- DBRM (database request module)
  - explaining [47](#)
- DBRM (database request module), explaining [47](#)
- DBRM and Plan List panel [47](#)
- DBRM SQL Statement window [47](#)
- DGOFPTS panel [10](#)
- DGOMCCMD panel [15](#)
- DGOMCDCO panel [14](#)
- DGOMDPMN panel [7](#)
- DGOMDWMO panel [9](#)
- DGOMDWSS panel [10](#)
- DGOMEPO2 panel [29](#), [57](#)
- DGOMEPO3 panel [31](#)
- DGOMEPE1 panel [27](#)
- DGOMGP01 panel [11](#)
- DGOMLAFS panel [35](#)
- DGOMLAMP panel [35](#)
- DGOMLAXP panel [34](#)
- DGOMLED0 panel [58](#)
- DGOMLEMP panel [36](#)
- DGOMLEXP panel [36](#)
- DGOMLMSP panel [33](#)
- DGOMYPKM panel [49](#)
- DGOMYPMN panel [43](#)
- DGOMYWDL panel [47](#)
- DGOMYWDQ panel [48](#)
- DGOMYWDS panel [47](#)
- DGOMYWES panel [48](#)
- DGOMYWKQ panel [46](#)
- DGOMYWKS panel [45](#)
- DGOMYWKV panel [46](#)
- DGOMYWPE panel [44](#)
- DGOMYWPL panel [45](#)
- DGOMYWQC panel [42](#)
- DGOMYWQU panel [42](#)
- DGOMYWSS panel [54](#)
- DGOPMENU panel
  - defining exception thresholds [55](#)
- DPMOUT data set
  - exception processing [38](#)
  - specifying exception processor name [31](#)
- DPMPARMS data set
  - CORRDATA member [17](#)

## E

- elapsed time format
  - online monitoring [18](#)

- errors
  - command not found [61](#)
  - STEPLIB missing from IRF-generated JCL [61](#)
- EVENTOBSERVATION
  - OMEGAMON Collector startup parameter [29](#)
- exception DPMOUT data set
  - attributes [9](#), [38](#)
- Exception Event Messages window [36](#)
- exception event processing
  - description [27](#)
  - listing messages [36](#)
  - starting [32](#)
  - viewing events [36](#)
- Exception Event Summary window [36](#)
- Exception Log File data set
  - attributes [37](#)
  - description [37](#)
  - output record layout [63](#)
  - printing [37](#)
  - sample files [37](#)
  - specifying processor name [31](#)
- Exception Log File print utility [37](#)
- Exception Notification window [27](#)
- Exception Output Data Sets window [31](#)
- exception processing
  - activating and deactivating [29](#), [57](#)
  - description [26](#)
  - exception event
    - description [27](#)
    - starting [32](#)
    - viewing [36](#)
  - modes [26](#)
  - output data set [37](#)
  - performance problem identification
    - example [55](#)
  - periodic exception
    - description [26](#)
    - listing messages [35](#)
    - starting [32](#)
    - viewing [34](#)
  - restarting [33](#)
  - starting [28](#), [33](#)
  - starting from Exception Processor panel [29](#)
  - starting with OMEGAMON Collector [29](#)
  - stopping [38](#)
  - viewing messages and status [33](#), [37](#)
- exception processing user exit [30](#)
- Exception Processor panel [29](#), [57](#)
- exception profiling
  - performance problem identification [23](#)
- exception report
  - monitoring [25](#)
- exception threshold
  - defining [28](#)
- Exception Threshold data set
  - building [30](#)
  - reinitializing [29](#)
  - specifying in Exception Processor panel [30](#)
  - specifying thresholds in [28](#)
  - threshold values in [26](#)
  - using at OMEGAMON Collector startup [29](#)
- EXEC
  - FPEJINIT [5](#)
- Expand function key for scrollable fields [20](#)

- explain
  - authorization [40](#)
  - DBRM [47](#)
  - description [40](#)
  - method [42](#)
  - output [49](#)
  - performance considerations [40](#)
  - plan table considerations [41](#)
  - plan table entry [44](#)
  - process overview [42](#)
  - specifying the SQL statement [48](#)
  - SQL statement [42](#)
  - viewing access path method [49](#)
- EXPLAIN command [51](#), [54](#)
- Explain menu [43](#)
- Explain Output panel [49](#)

## F

- filtering input [23](#)
- FPEJINIT EXEC [5](#)
- function key
  - commands to change the defaults [13](#)
  - defaults [11](#)
  - for scrollable fields [20](#)

## I

- IFCID 337
  - on Exception Event Summary panel [36](#)
- IFCID 359
  - on Exception Event Summary panel [36](#)
- Index Split exception event [36](#)
- Interactive Report Facility (IRF)
  - session options [10](#)
- IRF (Interactive Report Facility)
  - session options [10](#)
- ISPF Online Monitor
  - source explain [51](#)

## L

- large value
  - Online Monitor [17](#)
- layout
  - Exception Log File data set output [63](#)
- Left function key for scrollable fields [20](#)
- legal notices
  - cookie policy [67](#), [69](#)
  - notices [67](#)
  - programming interface information [67](#), [68](#)
  - trademarks [67–69](#)
- limiting input [23](#)
- Lock Escalation exception event [36](#)
- LOOK command [33](#)
- Look Selections menu [33](#)

## M

- menu
  - data set maintenance
    - defining exception thresholds [55](#)
  - Explain [43](#)

- menu (*continued*)
  - Look Selections [33](#)
  - Options [7](#)
- message
  - exception processing [33](#), [37](#)
  - IKJ56500I [61](#)
- migration
  - plan table [42](#)
- missing value
  - online monitoring [17](#)
- Monitor Output window [9](#)
- monitoring Db2 subsystems
  - exception processing [21](#)
- monitoring DB2 subsystems
  - controlling the writing of monitor output [9](#)
  - exception profiling [23](#)
  - exception report [25](#)
  - identifying unusual situations [23](#)
  - performance objectives [23](#)
  - tuning the operating environment [22](#)
- monitoring using reports [24](#)

## N

- notices [67](#), [68](#)

## O

- Online Monitor options [7](#)
- online monitoring
  - accessing [5](#)
  - description [22](#)
  - overview [5](#)
  - using [5](#), [19](#)
- options [6](#)
- OPTIONS command [16](#)
- Options menu [7](#)
- overview [1](#)

## P

- package
  - explain an SQL statement [45](#)
- Package SQL Statement window [45](#)
- Package Version List panel [46](#)
- panel
  - Asynchronous Task Termination [11](#)
  - Authorization Failure Summary [35](#)
  - DB2 Explain Output [49](#)
  - Exception Event Messages [36](#)
  - Exception Event Summary [36](#)
  - Exception Processor [29](#), [57](#)
  - moving between Online Monitor panels [10](#)
  - moving between panels [10](#)
  - SQL Statement List (packages) [46](#)
  - SQL Statement Selection [54](#)
  - SQL Text Edit [49](#)
- PANELID command [14](#)
- performance problem identification
  - controlling the writing of monitor output [9](#)
  - deadlock resolution [58](#), [59](#)
  - examples
    - exception processing [55](#)

- performance problem identification (*continued*)
  - exception processing [21](#)
  - exception profiling [23](#)
  - exception report [25](#)
  - identifying unusual situations [23](#)
  - performance objectives [23](#)
  - tuning the operating environment [22](#)
- Periodic Exception Messages window [35](#)
- periodic exception processing
  - description [26](#)
  - listing messages [35](#)
  - starting [32](#)
  - viewing exceptions [34](#)
- Periodic Exceptions List window [34](#)
- plan table
  - considerations [41](#)
  - explaining existing entry [44](#)
  - upgrading [42](#)
- Plan Table Entry List panel [45](#)
- Plan Table Entry window [44](#)
- programming interface information [67](#), [68](#)

## R

- REINIT command [33](#)
- Right function key for scrollable fields [20](#)
- RKO2SAMP library
  - DGOECFIL member [37](#)
  - DGOELFIL member [38](#)
  - DGOEQFIL member [38](#)
  - DGOEVFIL member [38](#)
  - DGOMEJCL member [37](#)

## S

- sample file
  - Exception Log File data set [37](#)
- screen readers and magnifiers [4](#)
- scrollable field
  - examples [19](#)
- scrollable field support [19](#)
- service information [1](#)
- Session Options window [10](#)
- source explain
  - ISPF Online Monitor [51](#)
- source EXPLAIN command [54](#)
- Source Explain Options window [52](#)
- SQL statement
  - explain function [40](#)
- SQL Statement List panel (DBRMs) [48](#)
- SQL Statement List panel (packages) [46](#)
- SQL Statement Selection panel [54](#)
- SQL Text Edit panel [49](#)
- SQL Text Entry window [48](#)
- SQLID window [10](#)
- starting
  - OMEGAMON for Db2 PE [5](#)
- status
  - exception processing [33](#), [37](#)
- STEPLIB missing from IRF-generated JCL [61](#)
- support information [1](#)
- syntax [14](#), [16](#), [33](#), [54](#)
- syntax diagram

syntax diagram (*continued*)  
source EXPLAIN [54](#)  
syntax diagrams [1](#)

## T

terminology used [3](#)  
time format  
    online monitoring [18](#)  
toleration mode  
    IFCID 337 [36](#)  
    IFCID 359 [36](#)  
trademarks [67–69](#)

## U

Upgrade Plan Table window [42](#)

## V

value  
    large and missing [17](#)  
    online monitoring [17](#)

## W

window  
    Create Plan Table [42](#)  
    Current SQLID [10](#)  
    DB2 Command [15](#)  
    DB2 Command Output [14](#)  
    DBRM and Plan List [47](#)  
    DBRM SQL Statement [47](#)  
    Exception Notification [27](#)  
    Exception Output Data Sets [31](#)  
    Monitor Output [9](#)  
    Package SQL Statement [45](#)  
    Package Version List [46](#)  
    Periodic Exception Messages [35](#)  
    Periodic Exceptions List [34](#)  
    Plan Table Entry [44](#)  
    Plan Table Entry List [45](#)  
    Session Options [10](#)  
    SQL Statement List (DBRMs) [48](#)  
    SQL Text Entry [48](#)  
    Upgrade Plan Table [42](#)







Product Number: 5655-W37

SH12-7070

