

5.5

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



2024-07-18 edition

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

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

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

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

Contents

About this information.....	vii
Chapter 1. Overview.....	1
Service updates and support information.....	1
How to read syntax diagrams.....	1
Conventions.....	2
Terminology.....	3
Accessibility features.....	4
Chapter 2. About the Performance Expert Client.....	5
Online monitoring and reporting.....	5
Performance Warehouse management.....	5
Buffer pool analysis.....	6
Integration with IBM Db2 AI.....	6
PE Client considerations.....	6
Chapter 3. Getting started.....	7
Configuring AT-TLS encryption.....	7
Starting Performance Expert Client.....	8
Working with the System Overview window.....	9
Defining the Db2 systems that must be monitored.....	10
Defining a Db2 subsystem.....	10
Defining a Db2 Connect gateway.....	12
Changing the properties of a Db2 subsystem.....	13
Changing the properties of a Db2 Connect gateway.....	13
Organizing the DB2 systems.....	14
Specifying global settings.....	14
Logging on to DB2 systems.....	14
Logging on to one DB2 subsystem or Db2 Connect gateway.....	14
Logging on to several DB2 subsystems.....	15
Starting several monitoring functions in one or more DB2 subsystems.....	15
Starting a monitoring function from the DB2 Control Center.....	15
Issuing DB2 commands.....	16
Integrating IBM Db2 Administration Foundation for z/OS.....	16
Getting help.....	17
Logging off from DB2 systems.....	18
Deleting a DB2 system.....	18
Leaving Performance Expert Client.....	18
Chapter 4. Online monitoring and reporting.....	19
Monitoring threads.....	19
Viewing all connected threads.....	19
Viewing details of a selected thread.....	19
Sorting threads.....	20
Filtering threads.....	21
Canceling threads.....	22
Creating and viewing SQL activity traces.....	23
Resolving locking conflicts.....	24
Viewing all threads involved in a locking situation.....	24
Viewing the threads associated with a specific resource.....	24

Viewing DB2 Statistics.....	25
Statistics values.....	25
Processing modes.....	26
Viewing Statistics information.....	27
Viewing the SQL statements in the dynamic SQL cache.....	28
Refreshing buffer pool Statistics.....	29
Viewing system health.....	29
Getting started.....	30
Scenario 1: Viewing current system health.....	31
Scenario 2: Viewing past system health.....	32
Scenario 3: Creating a data view.....	33
Changing the properties of a data view.....	33
Freezing the display of snapshots.....	39
Displaying data views as performance graphs in the System Overview window.....	40
Reusing data-view and data-group configurations.....	40
Gaps and sudden drops or rises in charts.....	40
Monitoring activities of Db2 Connect gateways and connections of DCS applications.....	40
Viewing all Db2 Connect gateways connected to a DB2 subsystem.....	42
Viewing Statistics details on a specific Db2 Connect gateway.....	42
Viewing information about all DCS applications connected to any DB2 subsystem.....	42
Viewing information about the DCS application related to a thread.....	43
The gateway snapshot time.....	43
Viewing DB2 system parameters.....	43
Working with exception processing.....	44
Starting exception processing.....	44
Processing event exceptions.....	45
Processing periodic exceptions.....	47
Working with the exception user exit.....	49
Launching IBM Db2zAI System Assessment Tool.....	49
Monitoring data sharing groups.....	49
Monitoring the threads of a data sharing group.....	50
Viewing Statistics on a data sharing group.....	50
Viewing the system health of a data sharing group.....	50
Viewing history data.....	51
Using Explain products to analyze and improve SQL statements.....	51
Displaying the access path while monitoring an application or thread.....	52
Displaying the access path while viewing Statistics on SQL statements.....	52
Exporting snapshot data for saving and printing.....	53
Working with windows displaying snapshots.....	54
Customizing table columns.....	54
Viewing details on a table entry in a contents pane.....	54
Viewing long performance counter values.....	54
Working in different time zones.....	55
Chapter 5. Working with Performance Warehouse.....	57
Running Performance Warehouse in a data-sharing environment.....	57
Logging on to Performance Warehouse.....	58
Logging on to a Performance Warehouse automatically.....	59
Logging on to a Performance Warehouse manually.....	59
Managing Performance Warehouse systems.....	59
Adding a Performance Warehouse system.....	59
Deleting a Performance Warehouse system.....	60
Defining and running processes.....	60
Report types.....	61
Scenario 1: Transferring performance data into the Performance Warehouse database.....	62
Scenario 2: Defining Accounting and Statistics reports using an existing input data set.....	65
Scenario 3: Collecting performance data for the generation of Accounting and Statistics reports...68	

Scenario 4: Collecting performance data regularly.....	69
Hints and tips.....	70
Using rules of thumb to analyze performance data.....	70
Managing rule-of-thumb groups.....	71
Managing rule-of-thumb clusters and rules of thumb.....	72
Analyzing the performance data.....	75
Defining and running queries.....	79
Managing query groups.....	79
Managing queries.....	80
Running a query.....	81
Chapter 6. Overview of the Performance Warehouse database.....	83
The Performance Warehouse data tables.....	83
The LOADLOG table.....	84
Chapter 7. Accessibility features in Performance Expert Client.....	87
Operating all features by using the keyboard.....	87
Changing the system settings for colors, fonts, and accessibility options.....	89
Product legal notices.....	91
Index.....	95

About this information

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

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

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

Chapter 1. Overview

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

Service updates and support information

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

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

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

How to read syntax diagrams

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

Arrow symbols

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



Two right arrows followed by a line indicate the beginning of a statement.



One right arrow at the end of a line indicates that the statement syntax is continued on the next line.



One right arrow followed by a line indicates that a statement is continued from the previous line.



A line followed by a right arrow and a left arrow indicates the end of a statement.

Conventions

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


Required items

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




Optional items

Optional items appear below the main path.




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




Multiple required or optional items

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




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




Repeatable items


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



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




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



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

Default keywords

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



Conventions

These conventions are used throughout the documentation.

Symbols

The following symbols might appear in command syntax:

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

Notation conventions

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

hilev

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

For example:

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

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

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

Terminology

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

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

Accessibility features

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

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

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

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

Chapter 2. About the Performance Expert Client

This topic introduces the main components and features of Performance Expert Client, which is one of the graphical interfaces of OMEGAMON for Db2 Performance Expert.

With Performance Expert Client you can monitor and analyze the performance of an application online and handle exception notifications. The Performance Warehouse provides a process-oriented view of the performance analysis tasks. It helps you manage and analyze buffer pools more efficiently by providing specific recommendations to adjust buffer pool sizes and threshold values.

Online monitoring and reporting

Performance Expert Client enables online monitoring and reporting, including exception notification handling.

You can use the online monitoring and reporting component to:

- Determine the total DB2 system performance and efficiency.
- Measure the performance of an application.
- Review the resources used by an application.
- Evaluate an application's impact on other applications and the DB2 system.
- Identify the cause of performance problems by checking certain thread-activity and Statistics fields for exceptions and listing the occurrence of specific events.
- Determine the tuning requirements for DB2.
- View performance data graphically.
- Analyze and improve SQL statements by using Explain products.

For more information refer to [“Using Explain products to analyze and improve SQL statements”](#) on page 51.

- Monitor DB2 Connect gateways, including system-related information.

This component is similar to the host-based (ISPF) Online Monitor. It has some advantages over the host-based Online Monitor, such as:

- The graphics improve the representation of the relationship between counters and the trends and evolution of counters.
- You can monitor several DB2 systems at the same time.
- When DB2 performance is not satisfactory, you receive tuning recommendations.

Performance Warehouse management

Performance Warehouse provides a process-oriented view of the performance analysis tasks.

It allows you to automate tasks that previously required user interaction, such as loading DB2 data into the Performance Warehouse database or generating reports.

Using the Performance Warehouse window, you can define, schedule, and run processes that automate the following tasks:

- Generating DB2 performance data
- Loading this data into the Performance Warehouse database
- Creating reports on the data in the Performance Warehouse database
- Collecting report data

You can then use rules of thumb and queries to analyze and evaluate the data in the Performance Warehouse database.

Samples of processes, rules of thumb, and queries are provided to help you get started.

Buffer pool analysis

Buffer pool analysis helps you manage buffer pools more efficiently by providing specific recommendations to adjust buffer pool sizes and threshold values.

If you are interested in background information about current buffer pool behavior, you can get in-depth statistical data, comprehensive reports, and easy-to-read graphic information for all buffer pools and their objects.

Buffer pool analysis helps you tune DB2 buffer pools and get the most out of your system resources. It lets you optimize memory usage by specifically recommending which buffer pools to change and by how much.

It allows you to:

- Save central processing unit (CPU) cycles
- Improve response times and elapsed times
- Simulate buffer pool behavior

This information does not cover buffer pool analysis. For information about how to use this component, see [Buffer Pool Analyzer User's Guide](#).

Integration with IBM Db2 AI

The integration of OMEGAMON for Db2 Performance Expert and IBM Db2 AI for z/OS (Db2zAI) enables you to launch an IBM Db2zAI System Assessment from the OMEGAMON for Db2 Performance Expert periodic and event exception detail displays, and in doing so helps you to determine the root cause of problems.

Note: Db2zAI integration is valid only for Db2 13 and Db2 12 subsystems.

To configure a Db2zAI server, see [“Defining a Db2 subsystem” on page 10](#).

To launch an IBM Db2zAI System Assessment when viewing periodic and event exception details, click the **Analyze Root Cause** button.

PE Client considerations

Review the following information before using the PE Client.

Effects of recycling a Db2 subsystem on snapshot history processing

If a Db2 subsystem is recycled at the same time a current delta or interval request is executed for snapshot history subsystem statistics data, then most of the delta or interval counters will overflow. This is caused by the second snapshot being taken after the Db2 subsystem restarts, at which time there are smaller values than when the first snapshot was taken (before the restart happened). To resolve this situation, reset the interval timestamp to a point after the Db2 subsystem was restarted.

Chapter 3. Getting started

This section describes the tasks you need to perform before you can start working with OMEGAMON for Db2 Performance Expert.

Before you start working with OMEGAMON for Db2 Performance Expert, ensure that you have the correct authorization and DB2 privileges to perform monitoring tasks.

To find out what authorization and DB2 privileges you need, see [Configuration and Customization](#).

Before you start working with OMEGAMON for Db2 Performance Expert, ensure that all DB2 traces are running. If you do not get any monitoring data, check if one or more DB2 traces are switched off.

Configuring AT-TLS encryption

You can use Application Transparent Transport Layer Security (AT-TLS) to secure communications between the OMEGAMON Db2 Collector PE Server subtasks and Performance Expert Client (PE Client).

Before you begin

It is highly recommended that you enable end-to-end encryption for PE Client communication with the OMEGAMON Db2 Collector PE Server on z/OS.

Important: When an instance of a PE Client is configured for encryption, it is assumed that all of the connections to the OMEGAMON Db2 Collectors on z/OS are encrypted. It is not possible to configure encryption for only a subset of OMEGAMON Db2 Collectors or selected Db2 subsystems that an instance of a PE Client wants to connect to.

The configuration consists of two parts:

z/OS Communication Server side

Implement AT-TLS policy, generate server certificates for the OMEGAMON Db2 Collector, and assign a keyring to the OMEGAMON Db2 Collector STC.

PE Client side

Create a Java keystore (JKS) containing the certification authority (CA) certificate, which is used in the certification chain for the OMEGAMON Db2 Collector server certificate (or alternatively containing the self-signed server certificate).

Procedure

1. Create a public/private key pair (private key plus certificate) for the OMEGAMON Db2 Collectors. The OMEGAMON Db2 Collectors present this server certificate in the TLS handshake process to the PE Client requesting a connection to the OMEGAMON Db2 Collector.
 - The certificate can be self-signed or signed by a CA according to your company's standards and rules. You can create and sign the certificate on z/OS with RACF or equivalent security products, or outside of z/OS. If a certificate is created outside of z/OS you must import it into z/OS.
 - Create a SAF keyring on z/OS or a file-based keystore on z/OS, and import the server certificate into it. Assign the keyring to the owner of the OMEGAMON Db2 Collector STC.
 - If you created a self-signed certificate, export it in text format and make it available to the client computers on which PE Client is running. If your server certificate is signed by a CA, make sure that the root CA or the certification chain is available at the client computers on which PE client is running.
2. Create an AT-TLS policy for encryption. To manage encrypted connections between the PE Client and the OMEGAMON Db2 Collector, it is necessary to create AT-TLS rules on every LPAR on which OMEGAMON Db2 Collectors are running. Define policy rules with the following characteristics:
 - Direction is "Inbound"

- The port range needs to cover the TCP/IP ports on which the OMEGAMON Db2 Collector is listening for incoming requests from PE Clients
- Allow TLSv1.2 and TLSv1.3 connections
- Specify the name of the SAF keyring you created in the previous task
- When the server certificate you created for the OMEGAMON Db2 Collector is not the default certificate for the STC user, specify the label of the certificate which AT-TLS should use for the TLS handshake.

Refer to IBM “z/OS Communications Server: IP Configuration Guide” for set up instructions.

- Policy Agent and policy applications: <https://www.ibm.com/docs/en/zos/2.5.0?topic=reference-policy-agent-policy-applications>
 - TLS/SSL security: <https://www.ibm.com/docs/en/zos/2.5.0?topic=guide-tlssl-security>
3. On the PE Client side, a Java truststore is required that must contain the server certificates (if you are working with self-signed certificates), or the root CA certificate or the certificate chain, if you signed the server certificate.

You can use Java keytool to create a password protected truststore and import certificates into it.

Important: The truststore must contain the certificates for all LPARs you want to monitor with that instance of the PE Client.

Continue to the PE Client for setup.

4. Start the PE Client and proceed to the **Monitor > Configurations** menu and select the **Encryption** tab.
5. In this screen, select the **Enable end-to-end encryption** checkbox to enable TLS encryption for the PE Client. Specify the truststore location and the password for the truststore.
6. Verify that all of the information provided is correct, and click **Apply**. Restart the PE Client. If there are no errors, all communication between the PE Client and z/OS is now encrypted. If the information provided is invalid (for example an incorrect truststore, or an incorrect password, or both), the Configuration window will not close and the information will not be saved.

Starting Performance Expert Client

You can start Performance Expert Client directly or from the DB2 Control Center.

If you installed a firewall on your workstation, ensure that Performance Expert Client has correct access permissions. Otherwise, Performance Expert Client does not work correctly.

You can start Performance Expert Client directly in one of the following ways:

- Double-click the **IBM DB2 Performance Expert for zOS** icon on your desktop.
- Click **Start > All Programs > IBM DB2 Performance Expert for zOS > IBM DB2 Performance Expert for zOS**.

To start Performance Expert Client from the DB2 Control Center, click the  toolbar button.

If you want to start Performance Expert Client more than once on a workstation and on the same operating system, Performance Expert Client must run under different user IDs. For example, if you want to start Performance Expert Client three times in such an environment, Performance Expert Client must run under three different user IDs or on different workstations.

After you start Performance Expert Client, the System Overview window is displayed, which is the main window.

Working with the System Overview window

The System Overview window shows the main performance information about the DB2 systems that you monitor and to which you are logged on. From this window, you can access all monitoring functions that are supported by Performance Expert.

You can view performance information for all DB2 systems or for a specific DB2 system. If a member is in DSG, you can view the information for each member. Figure 1 on page 9 shows the KPIs page of the System Overview window after you logged on to a DB2 subsystem.

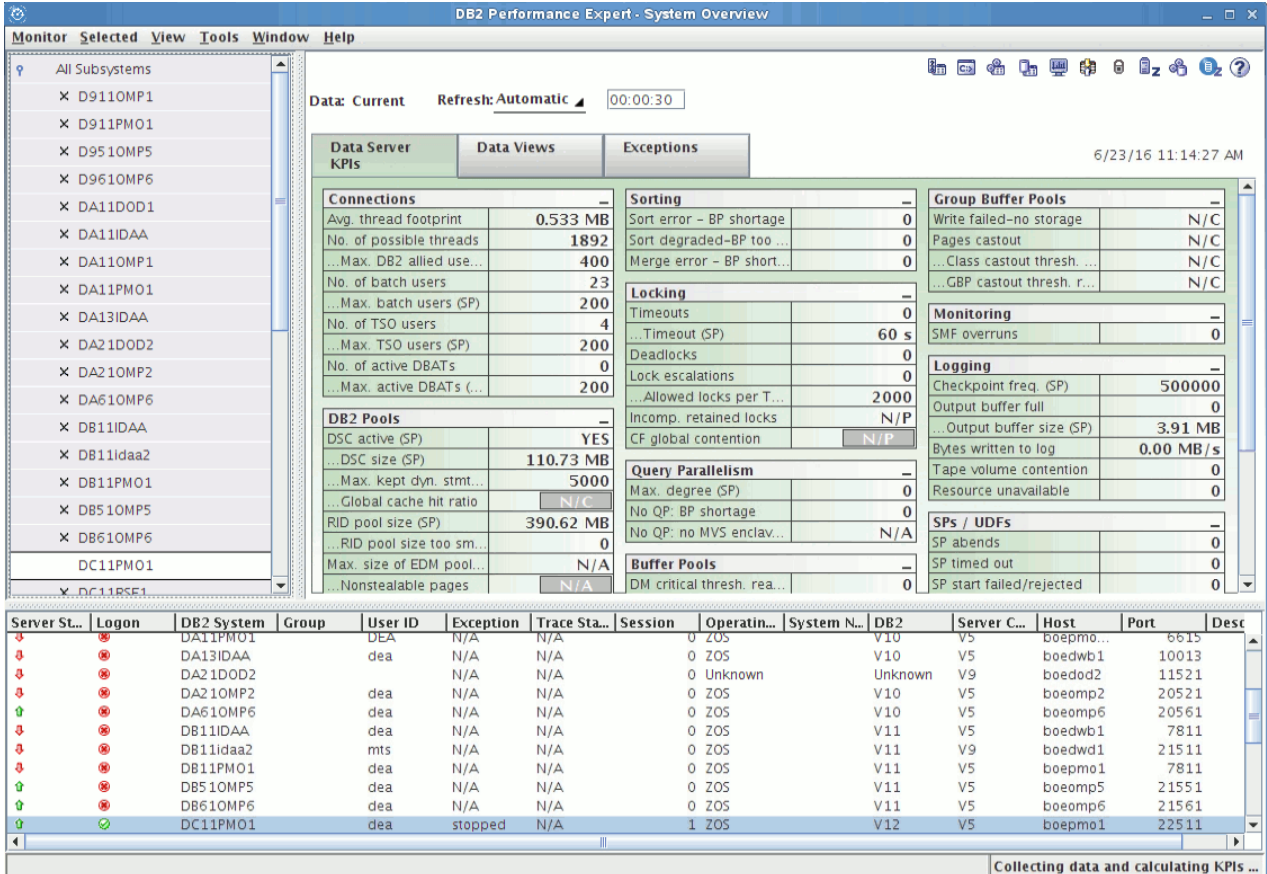


Figure 1. System Overview window

The folders pane on the left lets you select the information level. The contents pane on the right shows the performance information for the selected information level. The performance information is shown on the following tabbed pages:

- The KPIs page shows the counters that are key performance indicators (KPIs) for the selected information level. Many of these counters have cumulative values, where the delta between the values of the last two snapshots is calculated.

When you move the cursor over a KPI, hover help is displayed, which contains a short description of the KPI. Double-clicking a KPI takes you to the pane or page of the Statistics Details or System Parameters window that contains this counter. The counter is then highlighted in green.

- The Data Views page shows the performance graphs for the selected information level. Performance graphs show Statistics information in graphical format. You create these performance graphs as data views in the System Health window and make them available in the System Overview window. When you click and move the cursor over a data view, the values for the counters that are represented by the data view are shown.

To open the pane or page of the Statistics Details that contains the counters in the data view, right-click the data view and then select Show in Statistics Details. To open the System Health window, in which you created the data view, right-click the data view and then select Show in System Health.

- The Exceptions page shows the last 10 periodic exceptions and the last 10 event exceptions for the selected information level.

To view more details on an event exception or a periodic exception, do one of the following:

- Click the exception. The exception details are shown at the bottom of the page.
- Double-click the exception. The appropriate Details window opens.
- Delete link removes the entry from the table.
- New exceptions are shown in red. Exceptions that you have double-clicked to display details are shown in grey.

The toolbar buttons at the top of the window take you directly to the monitoring functions.

Specific icons and colors indicate whether the value of a counter falls below or above a warning or problem threshold. Event exceptions are also indicated by an icon. You are only informed of a warning or problem if the warning or problem threshold is active and if you activated the threshold set that contains this warning or problem threshold. For more information, refer to [“Processing periodic exceptions” on page 47](#).

The data in the System Overview window is refreshed at the rate that you specify for a DB2 system.

Defining the Db2 systems that must be monitored

To monitor a Db2 subsystem or Db2 Connect gateway, you must enable Performance Expert Client to access the performance data of this DB2 system which is described in the following sections.

Note: [“Changing the properties of a Db2 subsystem” on page 13](#) and [“Changing the properties of a Db2 Connect gateway” on page 13](#) show how to change the properties of a Db2 system.

Defining a Db2 subsystem

Each OMEGAMON Collector monitors the Db2 subsystem on which it is installed. The OMEGAMON Collector collects the performance data of the subsystem and stores it in its DB2PM database.

To monitor a Db2 subsystem from Performance Expert Client, you must establish a connection between Performance Expert Client and the OMEGAMON Collector installed on the Db2 subsystem so that Performance Expert Client can access the performance data of this subsystem. You do this by defining the Db2 subsystem:

1. On the System Overview window, click **Monitor > New Db2 System**.

The New Db2 System wizard is displayed.

2. On the Object tab, click **Db2 on z/OS, OS/390 systems**, and then click **Next**.

3. On the Data Collector tab, specify the connection to the Db2 subsystem on which OMEGAMON Collector is installed:

Host

The MVS host name.

Port

Specify the port number to which the Performance Expert Client can connect when communicating with the OMEGAMON Collector. The port number specifies the TCP/IP port on which the server listens for incoming sessions of Performance Expert Client for the specified Db2 subsystem.

If you do not know this information, contact your database administrator (DBA).

4. Click **Next**.

5. On the Db2 Subsystem tab, do one of the following:

- Click **Retrieve** to use the Db2 subsystem information defined on the host.

The Logon window is displayed.

- In the **User ID** field, type your MVS host user ID. In the **Password** field, type your password. Then click **OK**.

The Db2 Subsystem page of the New Db2 System window is displayed. For the **Name** field, OMEGAMON for Db2 Performance Expert supplies the same value as for the **Database alias** field. Ensure that the name is unique throughout Performance Expert Client.

- Type the values for the following fields:

Name

Type a name for the Db2 subsystem that is unique throughout all Db2 subsystems. If you clicked Retrieve, this field shows the same value as the Database alias field. Ensure that this value is unique throughout Performance Expert Client.

Database alias

Type the alias of the Db2 subsystem location. If you clicked Retrieve, this field shows the alias of the Db2 subsystem location as defined on the host. The specified alias will appear in the folders pane of the System Overview window.

MVS system

The name of the MVS system.

Data Sharing Group

If the Db2 subsystem is a member of a data sharing group, type the group name here.

Data Sharing Member

If the Db2 subsystem is a member of a data sharing group, type the member name here.

Description

Use this field for a description of the Db2 subsystem. This field is optional.

Local DB alias

If you want to use Query Workload Tuner, type the local database alias of the Db2 subsystem.

6. Click **Next**.

7. On the Performance Warehouse tab, you can associate a Performance Warehouse system with the Db2 subsystem you are defining. You must specify a Performance Warehouse system if you want to create SQL activity traces.

In the **Database alias** field, type the alias of the Performance Warehouse database that is maintained by the OMEGAMON Collector that monitors this Db2 subsystem. This is the database alias that you specified when cataloging the DB2PM database of the OMEGAMON Collector.

For further information about cataloging the DB2PM database, see [Configuration and Customization](#).

8. Click **Next**.

9. On the DSM Information tab, you specify the settings for the Data Server Manager (DSM) Server:

Protocol

The transfer protocol to be used for the Data Server Manager (DSM) Server. This can be HTTP or HTTPS.

DSM Server

Shows the host name of the DSM Server. When the Protocol field has been set to HTTPS, this value should be matched to the host name specified in the public key stored in TrustStore.

DSM Port

The port number for the Data Server Manager.

TrustStore

The path to the TrustStore file. TrustStore should have the public key for the host name in the DSM Server field. If the Protocol is HTTP, this field is greyed out.

DSM user

The user ID for logging into DSM.

DSM user password

The user password for logging into DSM.

DSM connection

The profile name for the monitored subsystem. For example, VB6A_RS29. When the SSL Secured checkbox is checked, this profile name is the Database connection name in the DSM Database Connections, and this profile should have SSL connectivity as JDBC security.

SSL Secured

This checkbox enables SSL Secured connections. Note: there must be a predefined DSM connection in order to use SSL.

JDBC URL

The URL for the monitored Db2 subsystem. If the SSL Secured box is checked, this field is greyed out.

Explain SQLID

An optional field that tells the tuning service to use the Explain tables pertaining to a specific SQLID. If this field is blank, the PE Client logon user ID will be used as the Explain SQLID.

10. Click **Next**.

11. On the Db2zAI tab, specify the configuration properties for the IBM Db2 AI for z/OS (Db2zAI) integration.

Note: Db2zAI integration is valid only for Db2 13 and Db2 12 subsystems.

Protocol

The transfer protocol (HTTPS only) to be used for the Db2zAI Server. This field is display only (cannot be edited).

Server

The host name of the Db2zAI Server.

SA Port

The Db2zAI System Assessment HTTPS port number. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSGKMA_1.2.0.1/src/ai/ai_configurenetworkport.html

UI Port

The Db2zAI User Interface port number. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSGKMA_1.2.0.1/src/ai/ai_configurenetworkport.html

User

The user ID for logging into Db2zAI. This ID must have proper authorization on the Db2zAI connection related to the respective Db2 subsystem connection.

Password

The password for logging into Db2zAI.

Start time (-)

(Default = 6) The number of hours before the exception timestamp at which point the Db2zAI assessment is started. For example, if an exception occurs at 5pm, and the **Start time (-)** is 2h and **End time (+)** is 1h, the assessment period would have a start time of 3pm and an end time of 6pm.

End time (+)

(Default = 1) The number of hours after the exception timestamp at which point the Db2zAI assessment is ended. For example, if an exception occurs at 5pm, and the **Start time (-)** is 2h and **End time (+)** is 1h, the assessment period would have a start time of 3pm and an end time of 6pm.

12. Click **Finish**.

The Db2 subsystem is shown in the **System Overview** window under the **z/OS** folder.

Defining a Db2 Connect gateway

If an OMEGAMON Collector is registered for a Performance Expert Agent for DB2 Connect Monitoring, it receives information about the DB2 Connect server that is installed on the same system as Performance

Expert Agent for DB2 Connect Monitoring. The DB2 Connect server serves as Db2 Connect gateway. The information is stored in the DB2PM database of the OMEGAMON Collector.

To monitor a Db2 Connect gateway from Performance Expert Client, you must establish a connection between Performance Expert Client and the OMEGAMON Collector collecting the information about the activities of a DB2 Connect server that serves as Db2 Connect gateway. You do this by defining the Db2 Connect gateway:

1. On the System Overview window, click **Monitor > New DB2 System**.

The New DB2 System wizard is displayed.

2. On the Object page, click **DB2 Connect/Gateways** and then click **Next**.

3. On the DB2 Connect/Gateways page, do the following:

- a. From the **DB2 system alias** list, select the DB2 subsystem monitored by the OMEGAMON Collector that receives data from one or more Performance Expert Agents for DB2 Connect Monitoring.
- b. Click **Retrieve** to retrieve the names of all Db2 Connect gateways with Performance Expert Agents for DB2 Connect Monitoring that send data to the selected OMEGAMON Collector.

The Logon window is displayed.

- c. In the **User ID** field, type your user ID. In the **Password** field, type your password. Then click **OK**.

The DB2 Connect/Gateways page of the New DB2 System window is displayed. The DB2 Connect gateway name consists of the following parts: the IP address, the workstation name, and the name of the DB2 subsystem containing Performance Expert Agent for DB2 Connect Monitoring. The fields under **Gateway Information** are shown.

- d. Select a Db2 Connect gateway name.

This name is also displayed in the **Gateway alias** field.

- e. Optionally you can specify another name in the **Gateway alias** field that is unique throughout Performance Expert Client. This name is shown in the folders pane of the System Overview window.

- f. Optionally you can specify a description for the selected Db2 Connect gateway.

4. Click **Finish** to return to the System Overview window.

The name of the Db2 Connect gateway is displayed in the System Overview window under the **DB2 Connect/Gateways** folder.

Changing the properties of a Db2 subsystem

This topic describes how to change the properties of a Db2 subsystem.

To change the properties of a Db2 subsystem:

1. Log off from the OMEGAMON Collector that monitors this Db2 subsystem. Refer to [“Logging off from DB2 systems” on page 18](#) for more information.
2. Click **Selected > Properties**.

The Db2 System Properties window is displayed.

You can define your changes on the various pages.

Changing the properties of a Db2 Connect gateway

This topic describes how to change the properties of a Db2 Connect gateway.

To change the properties of a Db2 Connect gateway:

1. Log off from the OMEGAMON Collector that receives information from this Db2 Connect gateway. Refer to [“Logging off from DB2 systems” on page 18](#) for more information.
2. Click **Selected > Properties**.

The DB2 System Properties window is displayed.

On the DB2 Connect/Gateways page, you can change the **Gateway alias** and **Description** fields.

Organizing the DB2 systems

When you specify or import a new DB2 subsystem or DB2 Connect gateway for monitoring from Performance Expert Client, this DB2 system is added to the corresponding folder of the System Overview window. Within these folders, the DB2 systems are listed in alphabetical order. You can structure and sort the DB2 systems according to your needs using the My Shortcuts folder.

You can copy one DB2 system and paste it as a shortcut to the **My Shortcuts** folder. In this folder, you can create a hierarchical structure of subfolders and copy the shortcuts to one or more subfolders.

You can change a DB2 system or shortcut in one of the following ways:

- Rename a DB2 system and thus all of its shortcuts. To do this, you must use the DB2 System Properties window (see [“Changing the properties of a Db2 subsystem” on page 13](#) and [“Changing the properties of a Db2 Connect gateway” on page 13](#)).
- Delete a DB2 system and thus all of its shortcuts: click the DB2 system that you want to delete in the **All Subsystems**, **All Data Sharing Group**, or **DB2 Connect Gateways** folder and then click **Selected > Delete**.
- Remove a shortcut from a folder: click the shortcut in the appropriate folder and then click **Selected > Delete** to delete only this shortcut.

Specifying global settings

Click **Monitor > Configuration** to specify preferences for global settings that are used throughout OMEGAMON for Db2 Performance Expert, such as your Web browser of choice, diagnostics settings, and accessibility settings. The Configuration window is displayed.

- On the Preferences page, you can:
 - Specify the location of your Web browser.
 - Permanently save your logon passwords for all DB2 systems.
 - Reactivate the display of message DGOK3638 (Monitor switches are not set) if you suppressed this message.
- On the Diagnostics page, you can specify e mail settings for sending diagnostics reports to IBM support.
- On the Accessibility page, you can specify high-contrast settings for your screen.

Logging on to DB2 systems

To monitor a DB2 system, you must log on to the OMEGAMON Collector that collects the performance data for this DB2 system.

If you use the same password for several OMEGAMON Collectors, you can log on to them at once. During such a multiple logon, you can start monitoring functions.

Logging on to one DB2 subsystem or Db2 Connect gateway

You log on to one DB2 subsystem or Db2 Connect gateway by logging on to the OMEGAMON Collector that collects the performance data for this DB2 system.

To log on to one DB2 subsystem or Db2 Connect gateway:

1. In the System Overview window, click the DB2 system and then click **Selected > Logon**.
The Logon window is displayed.
2. In the **User ID** field, type your MVS host user ID.
3. In the **Password** field, type your password.

4. You can type your TSO group ID in the **Group ID** field. This is optional.
5. Click **OK**.

Logging on to several DB2 subsystems

You log on to several DB2 subsystems at once by logging on to the OMEGAMON Collectors that collect the performance data for these DB2 subsystems and that are accessible with the same password.

To log on to several DB2 subsystems at once:

1. In the System Overview window, click **Monitor > Multiple Logon**.

The Multiple Logon window is displayed.

2. Select the DB2 subsystems to which you want to log on. You can access the OMEGAMON Collector of these DB2 subsystems with the same password. You can also select the monitoring functions that you want to start.
3. Click **OK**. The Logon window is displayed.
4. Fill in the fields as described under [“Logging on to one DB2 subsystem or Db2 Connect gateway” on page 14](#) and click **OK**.

The Log window is opened telling you whether your logon requests were successful. If you selected one or more monitoring functions, the appropriate windows are also opened.

Starting several monitoring functions in one or more DB2 subsystems

You can start several monitoring functions in one or more DB2 subsystems to which you are logged on.

To start several monitoring functions in one or more DB2 subsystems:

1. In the System Overview window, click **Monitor > Multiple Launch**.

The Multiple Launch window is displayed.

2. Select the DB2 subsystems and the monitoring functions that you want to start and click **OK**.

The Log window is opened telling you whether your launch requests were successful. In addition, the appropriate monitoring windows are displayed.

Starting a monitoring function from the DB2 Control Center

When you are working with DB2 you might want to start a specific performance-monitoring function from DB2. You can do this using the DB2 Control Center.

The DB2 system for which you want to view the performance data must have already been defined for monitoring from Performance Expert Client (see [“Defining the Db2 systems that must be monitored” on page 10](#)). However, Performance Expert Client need not be open when you select the monitoring function from the Control Center.

To start a monitoring function from the Control Center, do the following:

1. In the folders pane, expand the system folder and right-click the DB2 subsystem for which you want to view performance data. If you are not connected to the selected DB2 subsystem, you are prompted to log on. You can also select **DB2 Performance Expert** to open a submenu containing the monitoring functions available for this DB2 subsystem.
2. Select the monitoring function you want to work with.
3. If the DB2 subsystem location alias of a DB2 subsystem is not found, a window is opened where you can select the OMEGAMON Collector monitoring this DB2 subsystem.


If OMEGAMON Collector is not already open, it is started. You can then view the requested performance data for the selected DB2 subsystem in a window.

Issuing DB2 commands

You can issue any DB2 command for which you have authority from Performance Expert Client, except **–START DB2**. This command can be issued only from an MVS terminal.

To issue a DB2 command:

1. You can select one of the following ways:

- In the System Overview window, select the DB2 subsystem in the folders pane and then click the **DB2 Commands** toolbar button ().
- You can also use the menu bar or pop-up menu: click the DB2 subsystem in the folders pane, then click **Selected > Monitor > DB2 Commands** on the menu bar or **Monitor > DB2 Commands** on the pop-up menu.

If you are not connected to the selected DB2 subsystem, you are prompted to log on.

The DB2 Commands window is displayed.

2. In the **Command input** field, type your command or click the arrow to select a previously issued command from the list.
3. Click **Send** to send the command to your DB2 subsystem.

The command is added to the **Command input** list. The DB2 response is displayed in the **DB2 response is** box.

To remove the DB2 response, right-click the **DB2 response is** box and then click **Clear**.

Integrating IBM Db2 Administration Foundation for z/OS

Administration Foundation integration for Performance Expert Client allows you to submit explain jobs for SQL Statements captured in the Dynamic Statement Cache under the **Statistics Details** and **Thread Summary** panels. It is a replacement for Data Studio Manager integration.

For more information about Administration Foundation, refer to [IBM Db2 Administration Foundation for z/OS](#).

Prerequisites and setup

To use Administration Foundation in conjunction with Performance Expert Client, a Zowe Virtual Desktop environment is needed with Administration Foundation installed and configured. Additionally, users that will be accessing the system are to be added, and if required, alternate explain SQLID's setup, if different EXPLAIN tables are to be used.

To be able to submit tuning jobs, the required subsystems need to be registered in the Administration Foundation interface.

For further information on how to set up Administration Foundation, refer to [Installing Db2 Administration Foundation](#).

For secure communication, Administration Foundation integration for Performance Expert Client utilizes encrypted communication via certificates. To acquire it, export the certificate in your browser from the Zowe environment with Administration Foundation, and convert it to *.jks keystore.

For example, using Google Chrome (Firefox requires additional intermediary conversions) export the certificate as a single-chain *.pem/*.crt file and use command line utility **keytool** to export it to Java Keystore *.jks.

Note the files location and password you provided.

```
Convert it using the command line: KEYTOOL -IMPORTCERT -FILE CERTIFICATEFROMBROWSER.CRT  
-KEYSTORE CERTIFICATEFORPEC.JKS
```


If using multi factor authentication setup, make sure that RACF profiles accessing the system are setup with AZFTOTP1 and AZFPTKT1 factors configured. For an example setup, refer to [IBM Z Multi-Factor Authentication 2.3 - Download PDFs](#).

Note: Userids are to be connected to groups IZUUSER (for zOSMF which Zowe needs) or MVDUSER for Zowe itself.

Take note of the host name (or IP address) and port of the Zowe host for Performance Expert Client side setup, as it might differ from regular web access.

Configuration

From within Performance Expert Client, right-click the selected subsystem select "Properties". Under the "Admin Foundation" tab, fill in the parameters acquired in the previous setup section. These include:

- Host name (or IP address) of the Zowe host
- Port of the Zowe host
- Certificate store saved on your computer
- Certificate store's password
- Username for Zowe Administration Foundation access
- Password for the user accessing Zowe Administration Foundation
- Select checkbox "Use MFA" if MFA setup is being used (password field will be greyed out)
- Alternate explain SQLID, if used

Click **OK** to save the Administration Foundation configuration.

If Administration Foundation is configured with a given subsystem, said subsystem will always take precedence using Administration Foundation over DSM, regardless of DSM's status. If Administration Foundation is not setup, all explain functionality will try to use DSM. This applies for each subsystems configuration individually.

Certificate store and its password combination is validated upon clicking **OK** in the configuration window.

Usage


Administration Foundation Explain integration functionality can be accessed from two locations:

- Under the **Statistics Details > Dynamic SQL Statements > SQL Statements** pane, for any of the selected SQL statements (after clicking **Receive statement cache information**, taking note of any filters that might be present), and clicking **Explain** in the individual SQL statement details view.
- Under **Threads > Thread Summary**, double-click a thread to open the **Threads Details** window. On the **SQL Statement** pane, click **Explain**.

Getting help

You can display help for all elements in a window or wizard and for the windows and wizards themselves.

To get help on:

- A window or wizard, click the Help button or toolbar button () or **Help > General Help**.
- A pane of a window, click the name of the pane in the folders pane and then press F1.
- A tabbed page of a window, click the tab and then press F1.
- A toolbar button, move your mouse pointer over this button.
- A field in a window or wizard, click in the field and then press F1 or display help for the window or wizard (general help) and click Field help. Field help also covers the buttons displayed between and next to the fields.

- A command button at the bottom of a window or wizard, move the input focus to this button and press F1 or display help for the window or wizard (general help) and click Button help.
- A table in a window or pane, click in the table and then press F1.

Logging off from DB2 systems

To log off from a DB2 subsystem or Db2 Connect gateway you must log off from the OMEGAMON Collector that collects the performance data for this DB2 system.

In the System Overview window, click the DB2 system and then click **Selected > Logoff**. If one or more monitoring activities are still running on your DB2 system, you are asked whether you want to continue or stop them.

The exceptions that occurred for the DB2 system from which you logged off are removed from the 100 Most Recent Exceptions pane of the System Overview window even if exception processing continues running. After your next logon, the exceptions that were collected during your previous logon and logoff are shown again.

Deleting a DB2 system

If you do not want to monitor a DB2 subsystem or Db2 Connect gateway from Performance Expert Client anymore, you can delete it.

To delete a DB2 system:

1. If you are still logged on to the OMEGAMON Collector that collects the performance data of the DB2 system that you want to delete, you must log off. See [“Logging off from DB2 systems” on page 18](#) for more information.
2. In the System Overview window, click the DB2 system that you want to delete and then click **Selected > Delete**.

A message appears asking you whether you also want to delete the DB2 catalog entries for this DB2 system. If you decide to leave the catalog entries and a Performance Warehouse system is associated with this DB2 system, you can continue working with this Performance Warehouse.

3. Make your selection, then click **Yes**.

The DB2 system and all its shortcuts are removed from the folders pane of the System Overview window.

Leaving Performance Expert Client

To leave Performance Expert Client, click **Monitor > Exit**.

A message or a window appears:

- If you have already logged off from all DB2 systems or stopped all monitoring activities on the systems to which you are still logged on, a message is displayed asking you whether you want to leave Performance Expert Client. Monitoring activities that you did not stop when logging off from a DB2 system continue running.
- If you are still logged on to DB2 systems on which monitoring activities are running, a window is displayed asking you whether you want to stop or continue all activities.

Chapter 4. Online monitoring and reporting

These topics describe how to monitor the performance of DB2 subsystems and Db2 Connect gateways and how to create reports and traces.

Monitoring threads


You can view the current and recent activity of all active threads connected to a DB2 subsystem. You can get an overview of all connected threads and detailed information about a specific thread. You can also view the SQL activity of a single thread and create traces on this activity.

This section describes how to get the overview and the detailed information, customize the displayed windows, and create SQL activity traces.

Viewing all connected threads

You can display all active threads connected to a DB2 subsystem or a member of a data sharing group.

To do this, you can choose one of the following ways:

- In the System Overview window, select the DB2 subsystem or member of the data sharing group in the folders pane and then click the **Thread Summary** toolbar button ().
- You can also use the menu bar or pop-up menu: click the DB2 subsystem or member of the data sharing group in the folders pane, then click **Selected > Monitor > Thread Summary** on the menu bar or **Monitor > Thread Summary** on the pop-up menu.

If you are not connected to the selected DB2 subsystem, you are prompted to log on.

The Thread Summary window is displayed. It contains the key performance counters for all connected threads. If a thread has children, the number shown in the Parallel Threads column includes the parent thread.

To refresh the thread activity data, in the toolbar, select **Automatic** in the **Refresh** field. By default, values are updated every six seconds. To change the interval, click in the time field next to the **Refresh** field and overwrite the value according to your requirements.

Related information:

- If the Thread Summary window lists too many threads, you can filter them. For more information refer to [“Filtering threads” on page 21](#).
- You can sort the threads in the window. For more information, see [“Sorting threads” on page 20](#).
- The Thread Summary window shows the key performance counters in a table. One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns” on page 54](#).
- If you are monitoring a member of a data sharing group, refer to [“Monitoring data sharing groups” on page 49](#) for more information.
- If you are working in different time zones, refer to [“Working in different time zones” on page 55](#) for more information.

Viewing details of a selected thread

You can get more details of a thread by double-clicking it in the Thread Summary window to open the Thread Details window.

The Overview pane of the Thread Details window shows the main counters of the selected thread. The folders pane lists additional counter categories. You can get more details on a counter category by clicking it in the folders pane.

If a thread has children, the children are listed on the Parallel Threads pane. You can get more details on a child by double-clicking it.

The folders pane of the Thread Details window lists the counter categories. The contents pane shows the counters for the selected counter category. You can get more details on a counter category by clicking it in the folders pane.

Some DB2 counters are also shown as percentages in bar charts, stacked to 100%. Each bar in the chart represents a snapshot and reflects several counters. The percentages of these counters are shown proportionately and accumulated to add up to 100%. Click the bar chart to see which counters are represented. You can increase or decrease the number of snapshots that are shown in a bar chart by using the plus (+) and minus (-) signs next to the x-axis of the chart.

For one or more DB2 counters in the Thread Details window an icon appears next to the counter affected if:

- You run periodic exception processing
- You specified a warning threshold, problem threshold, or both
- A counter value falls below or above a threshold

When you move your mouse pointer over such a counter, you see the timestamp of the exception. The icon disappears after the counter value falls below or above the threshold or after periodic exception processing is stopped and you refresh the data in the window. No icon is shown for counters listed in a table.

On the SQL Statement pane, you see the text of the SQL statement that was recently performed by the thread. You can view the access plan for this statement in graphical format by clicking Explain. Refer to [“Using Explain products to analyze and improve SQL statements” on page 51](#) for more information. If the statement text is much longer than the space provided and, therefore, difficult to read, you can view it in another window. Click View Statement in New Window.

Related information:

- The Thread Details window shows part of its performance information in tables. In contents panes containing a table:
 - You can filter the entries listed. For more information refer to [“Filtering threads” on page 21](#).
 - One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns” on page 54](#).
 - You can choose to view details on a table entry in another window or on a notebook page. For more information, refer to [“Viewing details on a table entry in a contents pane” on page 54](#).
- Performance counter values that are too long are truncated. For information about how to view the entire value, refer to [“Viewing long performance counter values” on page 54](#).
- If you are working in different time zones, refer to [“Working in different time zones” on page 55](#) for more information.

Sorting threads

You can sort threads by a specific performance counter (simple sort) or by using several performance counters (hierarchical sort).

In windows that list threads, the entries are sorted by a performance counter (simple sort). By default, this is the leftmost column in the window. To change the order, click the name of another column.

In the Thread Summary window, you can create a specific sort hierarchy. You can select up to five performance counters as sort criteria; the first one selected has the highest sort criterion. You can sort the rows in ascending or descending order. You can also view them in a different order for each DB2 subsystem.

To sort the threads:

1. Click **View > Sort**.

The Sort window is displayed.

2. In the Available columns box, click the column name to be the highest sort criterion and then click Add (▶).

The selected column name is removed from the Available columns box and displayed in the Sort columns box.

You can continue selecting additional sort criteria until you have reached the limit of five column names. Your selection is always hierarchical. This means, the column name you select first is the highest sort criterion, the one you select second is the second highest criterion, and so on. For example, if you select Plan as the highest sort key and there are several threads with the same Plan, you might want to sort these threads by another column name, such as Elapsed Class 2, and would thus select Elapsed Class 2 as second sort key.

3. Within a sort criterion, the threads are, by default, sorted in ascending order. In the Sort columns box, click the column name for which you want to specify descending order and then click Descending to sort the threads in descending order.
4. You can use the Move Up and Move Down buttons to change the hierarchy of the sort criteria.
5. At any time during or after your selection, you can remove a sort criterion from your selection. In the Sort columns box, click the column name to be removed and then click Remove (◀). The selected column name is removed from the Sort columns box and displayed in the Available columns box at its previous position.
6. When you have finished specifying your sort criteria, click OK.

The Thread Summary window is displayed. The rows are sorted according to your specifications. The Sort (↓^AZ) icon is displayed in the right-hand corner of the status bar to show that your sort criteria are active. The current sort criteria remain valid for this DB2 subsystem until you change them again or disable them. If you want to keep your sort criteria but temporarily disable them to view the default order, right-click the Sort (↓^AZ) icon and then click Disable. The Sort (↓^AZ) icon is dimmed and the threads are displayed in the default order.

To reactivate your sort criteria, right-click the dimmed Sort (↓^AZ) icon and then click Enable. The threads are sorted according to your specifications again.

Filtering threads

You can filter threads so that only those meeting your specifications are displayed. This minimizes network traffic and can decrease response times considerably. For example, if you suspect that threads with a certain CPU Class 1 are causing problems, you can specify your filter to display only threads within CPU Class 1. You can define a different filter for each DB2 subsystem.

You start filtering on the Thread Summary window. If you have large amounts of data, you can also set filter criteria in the System Overview window.

To filter threads:

1. Click **View > Filter/Qualify**.

The Filter/Qualify window is displayed.

2. On the Identifier Type/Status page, under Column name, click the column name for which you want to specify one or more values.
3. Type a value in the Value field. You can use a wildcard character (*) in the value. For example, if you want to cover all values that start with PM, you would specify PM*. Characters following the wildcard are ignored.
4. Click Add (▶).

The selected column name is displayed in the Filter table, with its value and the include (☑) symbol.

The include symbol means that all threads having the specified value for the selected column name are included in the list of threads.

5. At any time during or after your selection, you can change your criterion:

- You can change the include symbol to an exclude (☒) symbol for any previous specification in order to exclude threads having the specified value:

Click the column name and then click Exclude under Qualifier criteria.

- To narrow the filter, you can select the same column name several times to specify more values. For example, you might want all threads starting with PM but exclude those starting with PMD. Up to 10 values are allowed for the same column name.

6. You can continue specifying additional filter criteria. All criteria and values that you specify are evaluated in the sequence in which you specify them and are combined. This means, the more criteria you specify the narrower the filter and the fewer the threads that are displayed.

7. At any time during or after your selection, you can remove filter criteria from the table.

In the Filter table, click the criterion to be removed and then Remove (✖). If you remove all criteria of a column name, this column name is removed from the Filter table.

8. Click the Additional Qualification tab to specify the thread status and type.

On the Additional Qualification page:

- By default, all check boxes are selected, except for the Exclude threads with no work in process check box.
- If the In DB2 check box is selected, all items in the list are automatically selected. You can deselect the ones that you do not need.
- If the In check box is selected, the field next to this check box and the check box of Exclude threads with no work in process can be selected. By default, Application is selected and the check box of Exclude threads with no work in process is not selected.

9. Click OK.

The Thread Summary window is displayed, which is refreshed according to your specifications. The

Filter/Qualify (☑) icon is displayed in the right-hand corner of the status bar to show that your filter is active. The current filter remains valid until you change it again or disable it.

If you want to keep your filter but temporarily view the threads unfiltered, right-click the Filter (☑) icon and then click Disable. The Filter (☑) icon is dimmed and the data is displayed in the default order.

To make your filter active again, right-click the dimmed Filter (☑) icon and then click Enable.

The window is refreshed according to your specifications.

Canceling threads

If you have the appropriate DB2 authority, you can cancel threads.

You can do this from the Thread Summary or the Thread Details window:

1. Do one of the following:

- In the Thread Summary window, click the thread you want to cancel and then click **Selected > Cancel Thread**.
- In the Thread Details window, click **Thread Details > Cancel Thread**.

2. Click **Yes** to confirm the cancelation.

Creating and viewing SQL activity traces

If you have MONITOR1 and MONITOR2 privileges, you can create traces on the DB2 system-wide SQL activity that takes place during the processing of a DB2 application. You can create such traces only if you are not viewing history data.

The trace results are displayed in a Web browser window. To ensure correct display, use the latest version of a web browser.

You can select a preferred browser go to **Monitor > Configuration**, and specify a default HTML browser in the **Preferences** window.

A trace can only be created if:

- The Performance Warehouse server is running.
- A Performance Warehouse system has been associated with the monitored DB2 subsystem. Verify that the Performance Warehouse page of the DB2 System Properties window contains the alias of the Performance Warehouse database maintained by the OMEGAMON Collector that is monitoring the DB2 subsystem. This alias is equal to the database alias that you specified when cataloging the DB2PM database by using the Configuration Assistant of DB2. For information about how to display the DB2 System Properties window, refer to [“Defining a Db2 subsystem” on page 10](#).

To create a trace, start in the Thread Summary or Thread Details window. When you select a thread, all threads with the same LUWID as the selected one, are included in the trace.

1. Do one of the following:

- In the Thread Summary window, select the thread for which you want to create the trace and then click **Selected > SQL Activity Tracing**.
- In the Thread Details window, click **Thread Details > SQL Activity Tracing**.

The SQL Activity Report Generation window is displayed.

Here, you can specify the criteria for the data collection to stop and the SQL activity events you want to see in the trace report.

2. In the **Elapsed time** and **Number of records collected** fields, you can leave the default specifications or specify both or only one of the criteria.

If you specify both criteria, data collection stops as soon as the first criterion is met. For example, if you specify an elapsed time of five minutes (00:05:00) and a number of records of 5 and this number is reached after a minute, the trace stops and disregards the remaining four minutes.

3. Optionally, select the SQL statements with a negative return code and the workload blocks that are shown in trace report.

4. Click **OK**.

If the message `Data set not cataloged` is displayed, set the OMEGAMON Collector startup parameter `USERAUTHORIZATION` to YES.

For more information about how to set this parameter, refer to [Configuration and Customization](#).

The collection of the performance data and the generation of the SQL activity trace are started and the SQL Activity Tracing window is displayed.

When the trace report is complete and the data sets are downloaded, the results are shown in a Web browser window.

The SQL activity trace results are downloaded as temporary files and are deleted when you leave OMEGAMON for Db2 Performance Expert. To save these files, use the save functions provided by your Web browser.

Resolving locking conflicts

When several threads try to access the same data at the same time, DB2 uses locks to control these concurrent attempts. A lock associates a DB2 resource with an application process in a way that affects how threads can access this resource.

You can view the threads involved in a locking situation where at least one thread is suspended. You can view the threads of a single DB2 subsystem or an entire data sharing group.

This section describes how to display the threads independently of the associated resource and to view only the threads associated with a specific resource.


Related information:

- Some windows show the key performance counters in a table. One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns” on page 54](#).

Viewing all threads involved in a locking situation

This topic describes how to view all threads for an entire DB2 subsystem or data sharing group.

To view all threads for an entire DB2 subsystem or data sharing group you can do one of the following:

- In the System Overview window, select the DB2 subsystem in the folders pane and then click the **Threads in Lock Conflicts** toolbar button ().
- You can also use the menu bar or pop-up menu: click the DB2 subsystem in the folders pane, then click **Selected > Monitor > Threads in Lock Conflicts** on the menu bar or **Monitor > Threads in Lock Conflicts** on the pop-up menu.

If you are not connected to the selected DB2 subsystem, you are prompted to log on.

The Threads in Lock Conflicts window shows all threads that are involved in a locking situation, independent of the resource with which they are associated. They are grouped by the thread that holds the lock (Holder) and the threads that wait for the resource to be freed (Waiter).

This window shows all threads involved in a locking situation, independent of the resource that they are associated with. They are grouped by the thread holding the lock (Holder) and the threads waiting for the resource to be freed (Waiter).

To view more details on a thread, double-click the thread, or click the thread and then click **Selected > Thread Details**.


Related information:

- If you are monitoring a member of a data sharing group, refer to [“Monitoring data sharing groups” on page 49](#) for more information.

Viewing the threads associated with a specific resource

This topic describes how to view the threads associated with a specific resource.

To view the threads associated with a specific resource:

1. Open the Locking Conflicts window in one of the following ways:
 - In the System Overview window, select the DB2 subsystem in the folders pane and then click the **Locking Conflicts** toolbar button ().
 - You can also use the menu bar or pop-up menu: click the DB2 subsystem in the folders pane, then click **Selected > Monitor > Locking Conflicts** on the menu bar or **Monitor > Locking Conflicts** on the pop-up menu.

This window lists all resources involved in a locking situation.

2. In the Locking Conflicts window, double-click the resource, or click the resource for which you want to view the associated threads and then click **Selected > Threads in Lock Conflicts**.

The Threads in Lock Conflicts window for the selected resource type is displayed.

This window contains detailed information about the selected resource and a list of the associated threads. The resource information changes with the selected resource. The threads are grouped by the thread holding the lock (Holder) and the threads waiting for the resource to be freed (Waiter).

3. To view more details on a thread, double-click the thread, or click the thread and then click **Selected > Thread Details**.

Viewing DB2 Statistics

You can view important Statistics and ratios of a DB2 subsystem in various levels of detail and processing modes.

This section explains some Statistics terms and the processing modes. It also describes how to get the Statistics information and how to refresh buffer pool Statistics.

Statistics values

Most of the DB2 Statistics values accumulate while DB2 is active. When DB2 is started, accumulation begins and continues until DB2 is stopped. There are, however, certain fields that are not cumulative.

Some are reported as *snapshot* values (*gauge* values in a workgroup environment), and some are reported as *high-water mark* values. A snapshot value is a current value. It is updated each time the Statistics values are displayed. A high-water mark is the highest value reached since startup. It is updated each time the Statistics values are displayed.

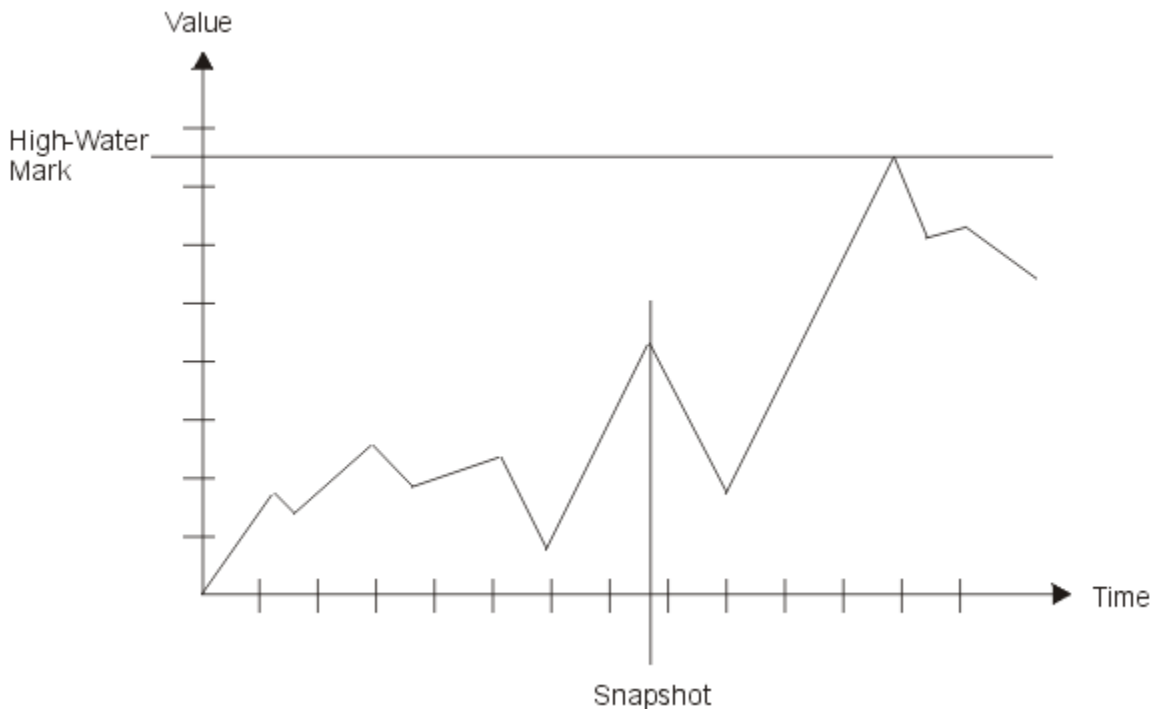


Figure 2. Snapshots and high-water marks

The values in fields other than snapshot or high-water mark fields are *accumulated* or *delta* values. They are calculated by the Statistics display and derived differently for each of the Statistics processing modes.

Processing modes

You can display the Statistics values in different modes: regular, interval, and delta.

Interval and *delta* modes determine the time period over which delta values are accumulated.

In all modes, Statistics values are updated each time a refresh occurs.

Regular processing

Regular processing is the default processing mode whenever you open a Statistics window.

Figure 3 on page 26 shows that the values displayed are accumulated from the start of DB2.

Regular Mode




Figure 3. Regular processing

No calculations are performed on the Statistics values in this mode.

Interval processing

The interval mode allows you to isolate DB2 Statistics activity beginning at a specific time.

Figure 4 on page 26 shows the interval processing mode.




Figure 4. Interval processing

In this mode, all Statistics windows show the delta values accumulated since the start of interval processing. In Figure 4 on page 26, interval processing was requested when the snapshot showed a value of 9. In this case, accumulation is set to zero and all future snapshots are accumulated from this point. When you stop interval processing, all following snapshots are shown in regular mode again, that is, as accumulated since the start of DB2.

Delta processing

Delta processing allows you to isolate DB2 Statistics activity between two snapshot refreshes.

Figure 4 on page 26 shows delta processing. It can help you discover a problem when running exception processing.





Figure 5. Delta processing

In this mode, all Statistics windows show the delta between the values of two consecutive snapshots. This means, accumulation starts at zero after each snapshot. When you stop delta processing, all following snapshots are shown in regular mode again, that is, as accumulated since the start of DB2.

Viewing Statistics information

This topic describes how to view Statistics information for a DB2 subsystem.

To view Statistics information for a DB2 subsystem:

- In the System Overview window, select the DB2 subsystem in the folders pane and then click the **Statistics Details** toolbar button ().
- You can also use the menu bar or pop-up menu: click the DB2 subsystem in the folders pane, then click **Selected > Monitor > Statistics Details** on the menu bar or **Monitor > Statistics Details** on the pop-up menu.

If you are not connected to the selected DB2 subsystem, you are prompted to log on.

The Statistics Details window is displayed.

The folders pane of the Statistics Details window lists the counter categories. The contents pane shows the counters for the selected counter category. You can get more details on a counter category by clicking it in the folders pane.

Some DB2 counters are also shown as percentages in bar charts, stacked to 100%. Each bar in the chart represents a snapshot and reflects several counters. The percentages of these counters are shown proportionately and accumulated to add up to 100%. Click the bar chart to see which counters are represented. You can increase or decrease the number of snapshots that are shown in a bar chart by using the plus (+) and minus (-) signs next to the x-axis of the chart.

An icon appears next to the counter affected if:

- You run periodic exception processing.
- You specified a warning threshold, problem threshold, or both for one or more DB2 counters in the Statistics Details window.
- A counter value falls below or above a threshold.

When you move your mouse pointer over such a counter, you see the timestamp of the exception. The icon disappears after the counter value falls below or above the threshold or after periodic exception processing is stopped and you refresh the data in the window. No icon is shown for counters listed in a table.

On the SQL Statement pane, you see the text of the SQL statement that was recently performed by the thread. You can view the access plan for this statement in graphical format by clicking Explain. Refer to “Using Explain products to analyze and improve SQL statements” on page 51 for more information. If the statement text is much longer than the space provided and, therefore, difficult to read, you can view it in another window. Click View Statement in New Window.

Related information:

- The Statistics Details window shows part of its performance information in tables. In contents panes containing a table:
 - One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns”](#) on page 54.
 - You can choose to view details on a table entry in another window or on a notebook page. For more information, refer to [“Viewing details on a table entry in a contents pane”](#) on page 54.
- Performance counter values that are too long are truncated. For information about how to view the entire value, refer to [“Viewing long performance counter values”](#) on page 54.
- If you are monitoring a member of a data sharing group, refer to [“Monitoring data sharing groups”](#) on page 49 for more information.
- If the OMEGAMON Collector and Performance Expert Client are in different time zones, refer to [“Working in different time zones”](#) on page 55 for more information.


Viewing the SQL statements in the dynamic SQL cache

This topic describes how to view the SQL statements in the dynamic SQL cache.


You can view detailed information about the SQL statements in the dynamic SQL cache as follows:

1. In the folders pane of the Statistics Details window, click **Dynamic SQL Statements** and then click **SQL Statements**.


The SQL Statements pane is displayed without any entries.

2. To display the SQL statements that are in the cache:
 - a. Click the **Filter** () toolbar button to check the filter values and set the appropriate values for your subsystem. For more information about defining filter criteria, see step [“5”](#) on page 28.
 - b. Select the **Receive Statement Cache information** check box.

The list is refreshed whenever you click **View > Refresh** or when an automatic refresh occurs.

3. The SQL statements in the cache might be shown with the value 0 for all counters because no trace data has been collected yet. To get all values for the counters, click the **Start** () toolbar button to start the collection of the specified counter values.

After the next refresh, the counter values of all SQL statements displayed reflect the values collected since the trace started.


Trace data collection continues until you click the **Stop** () toolbar button. Because tracing can increase the overhead, stop tracing when you have finished working with the SQL statements. You can only stop the traces that you started.

4. To view the text of an SQL statement, double-click the statement.

The SQL Statements *text* page opens, where *text* stands for the first couple of characters of the statement text.

Note: If the text is much longer than the space provided and therefore difficult to read, click **View Statement in New Window** to view the text in another window.

This page contains the text and the counters of the SQL statement. If you have an Explain product installed and the statement contains a SELECT, UPDATE, INSERT, DELETE, or VALUES clause, you can click **Explain** to view the access path. For more information about Explain products, refer to [“Using Explain products to analyze and improve SQL statements”](#) on page 51.

5. To reduce the list of SQL statements:
 - a. Return to the SQL Statements pane of the Statistics Details window.
 - b. Click the **Filter** () toolbar button.

The Dynamic Statement Cache Filter Criteria window is displayed.

c. Choose your filter method:

- If you want to list only those SQL statements where a specific counter exceeds a certain value:
Click **Statements with value higher than**, specify a value, and then select one of the counters listed under **Fields**.
- If you want to display the SQL statements sorted by the values of a counter in descending order:


Click **Statements with highest value** and then select one of the counters listed under **Fields**.

This option is useful, for example, if there are more SQL statements in the cache than can be displayed (up to 1 MB). For example, if there are 5 MB of SQL statements in the cache and you are mainly interested in those with high Getpage counter values, ensure that these SQL statements are displayed by using this option.

d. Click **OK** to return to the SQL Statements pane.

The list of SQL statements is changed according to your filter criteria. The filter remains valid until you change it.

Restriction: If you are viewing history data, you cannot define a filter. For more information about history data, refer to [“Viewing history data” on page 51](#).

6. To change the order of the table columns shown or hide the table columns that you are not interested in, click the **Customize Columns** () toolbar button. The Customize Columns window is displayed.

For further information about how to customize columns, see [“Customizing table columns” on page 54](#).

Refreshing buffer pool Statistics

The contents of a window or contents pane is normally refreshed when an automatic refresh occurs or when you click **View > Refresh**.

This does not apply, however, to buffer pool Statistics at dataset level.

Buffer pool Statistics at dataset level are received using IFCID 199. If I/O activity has taken place (at least one I/O per second), DB2 updates the list of buffer pool data sets every five minutes or according to the time interval set in the Online data set Statistics time interval (minutes) parameter in the Tracing Parameters (DSNTIPN) pane of the System Parameters window.

To view the update:


1. On the Data Sets pane of the Buffer Management page of the Statistics Details window, select the **Receive data set information** check box.
2. Click **View > Refresh** when the time interval has elapsed.

Viewing system health

Statistics information for a DB2 subsystem can be viewed as numerical values and percentages on various Statistics Details windows or as snapshots in graphical format.

[“Viewing DB2 Statistics” on page 25](#) explains the Statistics Details windows.

To view Statistics information in graphical format you can do one of the following:

- In the System Overview window, select the DB2 subsystem in the folders pane and then click the System Health toolbar button ().
- You can also use the menu bar or pop-up menu: click the DB2 subsystem in the folders pane, then click **Selected > Monitor > System Health** on the menu bar or **Monitor > System Health** on the pop-up menu.

If you are not connected to the selected DB2 subsystem, you are prompted to log on.

The System Health window is displayed, split into the following panes:

- The folders pane on the left lists one or more of the following data groups:
 - The System Overview data group, which contains links to data views in data groups that you have defined. These linked data views are displayed as performance graphs on the Data Views page of the System Overview window. They are refreshed at the refresh rate that you specified for the DB2 subsystem in the System Overview window.
 - Any data groups and data views that you have defined.
- The contents pane on the right displays the contents of the selected data group.

You can arrange one or more counter categories, such as Hit ratios or Physical locks (P-locks), as data views in a data group. Each data view consists of a graphical representation, a legend, and a report of the counters selected for the counter category. Each data group can contain a different selection and arrangement of data views.

This section uses scenarios to help you understand and get started with system health. In addition, it describes the tasks that you are most likely to perform in system health.

Related information:

- If you are monitoring a member of a data sharing group, refer to [“Monitoring data sharing groups” on page 49](#) for more information.
- If you are working in different time zones, refer to [“Working in different time zones” on page 55](#) for more information.

Getting started

The scenarios in this section are based on a data view that you create with the following steps. The options that you select and any alternatives are explained later in the section.

Perform the following steps:

1. Define a data group in which you will place the data view:
 - a. Click **Data Groups** and then click **Selected > New**. The New Data Group window is displayed.
 - b. Type My Locking in the **New name** field and then click **OK**.The new group appears in the folders pane of the System Health window.
2. Click **Data views** and then click **Selected > New** to define a data view. The New Data View wizard appears.
3. On the Category page:
 - In the **Data view name** field, type Locking as the name for this data view.
 - In the **Data view category** list, select **Locking**. The categories in the list correspond to the Statistics categories for counters and data elements.
 - Click **Next** to display the Counter page.
4. On the Counters page:
 - In the **Counter selection** list, select the **Lock requests** and **Unlock requests** check boxes to specify the counters for the data view.
 - Leave **Not stacked** selected as type of counter display. The counters are then displayed independently of one another as separate bars in a bar chart or lines in a line chart.
 - Click **Next** to display the Thresholds page.
5. On the Thresholds page:
 - Next to **Static scale**, change the default 1000 units to 30 and select **Second**.
 - Click **Next** to display the Graphics page.

6. On the Graphics page, leave the default selections and values. Note that **Bar chart** is selected to display the counters as bars.

Click **Finish**.

When you work with System Health for the first time, data is refreshed every 6 seconds and the new My Locking data view looks similar to this one:

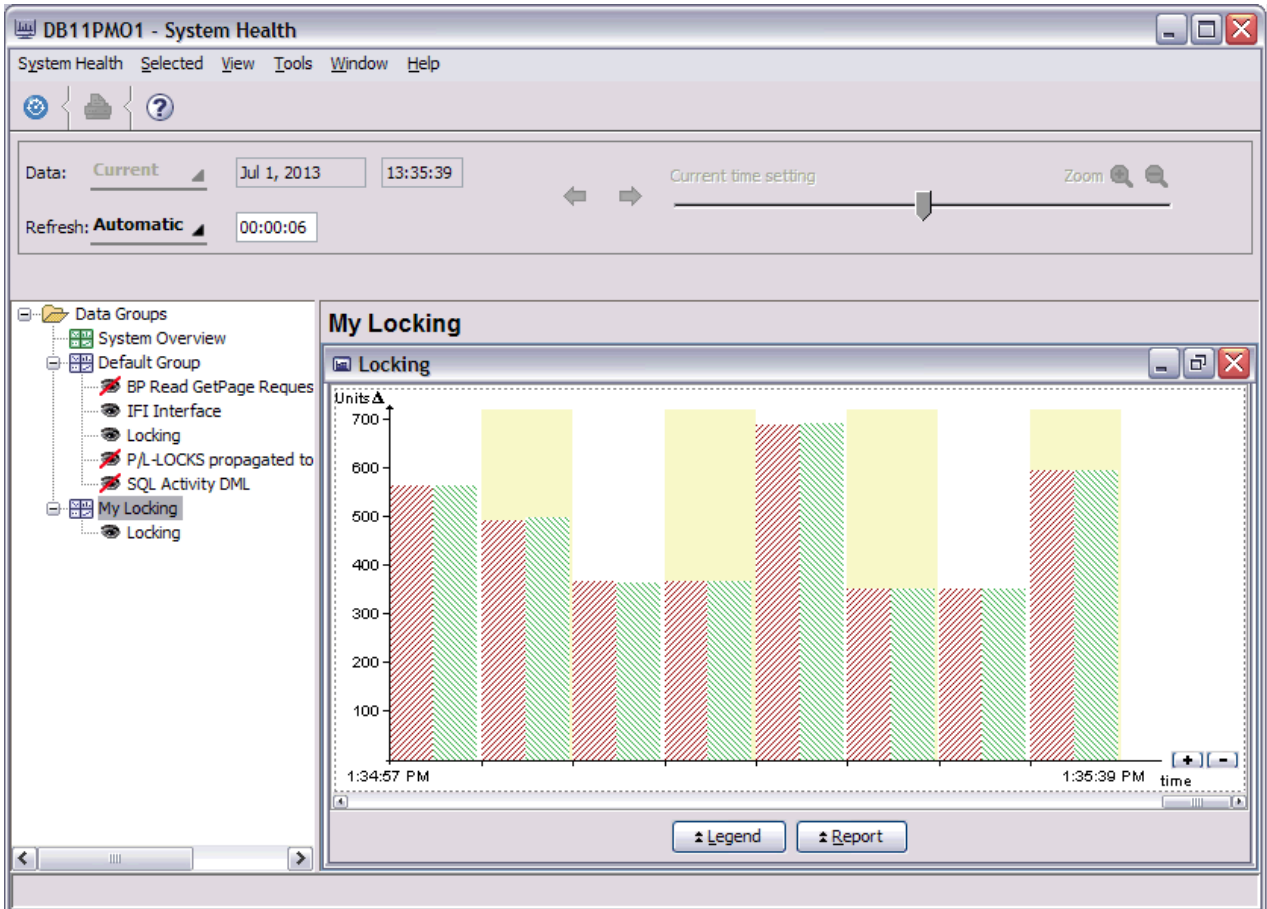


Figure 6. System Health window containing new data group and data view

Scenario 1: Viewing current system health

The following scenario shows how to collect and display snapshots and helps you understand the chart and the values displayed. As an example, it uses the data view that you previously created.

Perform the following steps:

1. To start the collection and display of snapshots, in the toolbar, select **Automatic** in the **Refresh** field.
2. To specify how often a snapshot is to be collected and shown in the data view: click in the time field next to the **Refresh** field and overwrite the displayed time with 0 : 00 : 06, which is 6 seconds. Then press Enter.

The data in the data view is refreshed every 6 seconds, which is the minimum refresh time.

By default, the latest eight snapshots are displayed in a bar chart. When this number is exceeded, a scroll bar appears at the bottom of the chart for you to scroll through the snapshots collected since the start of the collection. By default, up to 100 snapshots are kept in memory and can be viewed. Any older snapshots are stored in the history data set and can be viewed in history mode. See [“Scenario 2: Viewing past system health”](#) on page 32.

The background color of a snapshot alternates for you to see when a new snapshot starts. This is useful if you selected several counters for a data view and, therefore, a snapshot consists of several bars.

The y-axis of the chart ends slightly above unit 700 because the highest visible value of one of the counters is just under 700. The y-axis changes when the highest value of one of the counters increases or decreases.

The values shown in the chart represent the delta between two subsequent snapshots. For example, assume that each minute 30 lock and unlock requests are performed, which is one every two seconds. Then a delta value of 3 is shown for snapshots that are taken every 6 seconds. The delta is smaller if you decrease the refresh time and is higher if you increase the refresh time.

3. To see which counters are represented in the data view and in which colors, click **Legend**.
4. To view the values for the individual counters, you can do one of the following:
 - Click a snapshot in the data view, to view the values for a specific snapshot.
 - Click **Report** to get an overview of several snapshots.

The scale of a chart depends on the threshold that you set for the counters in the data view.

5. To change the threshold for the data view:
 - a. Click **Selected > Properties**. The Properties window is displayed. Click the **Thresholds** tab.

A threshold for a counter is specified as units per time interval. The **Static scale** field shows 30 units per second, which you specified when creating the data view. This means that the threshold is reached if a lock or an unlock request is performed 30 times per second.
 - b. Change 30 to 60 and then click **Finish**.

The y-axis of the data view reflects the threshold change. Because the data is refreshed every 10 seconds, the threshold is reached if a lock or an unlock request is performed, 600 times during this period. Therefore, the red horizontal (threshold) line is now at unit 600 and the y-axis of the chart ends slightly above it.

Scenario 2: Viewing past system health

If your system is not performing well and you want to know when this started, you are interested in snapshots that were taken earlier and are probably no longer displayed to you. The following scenario tells you how to view past system health using the data view that you created and helps you understand the chart and the values displayed.

For the purpose of this scenario, it is assumed that you have collected locking data for at least a couple of minutes and you selected a refresh time of 10 seconds and a threshold of 60 units per second.

Perform the following steps:

1. To start the display of past system health, in the toolbar, select **History** in the **Data** field.

The history slider is activated and the **Refresh** field is disabled.

2. Move the slider to a time during which snapshots were collected and wait a couple of seconds to receive the data.

The chart has changed as follows:

- The chart shows the eight snapshots collected from the time that the slider indicates. The time of the first and last snapshots displayed is shown below the chart.

Note that the collection of the current snapshot continues in the background. So when you leave history mode, the most recent snapshots are displayed.

- The scale of the y-axis, the place of the threshold line, and the size of the bars changed:

When you view the current system health you can specify how often the snapshots are refreshed. History snapshots are, however, displayed at the interval at which the OMEGAMON Collector collects data, which is 60 seconds for the recording interval. Because the refresh rate is now 60 seconds (instead of the 10 seconds you previously specified), the y-axis of the chart now ends above 3500 and contains a (threshold) line at unit 3600 and the size of the bars is adjusted accordingly.

For more information about how to select history snapshots in different ways, see [“Viewing history data” on page 51](#).

The number of snapshots available is limited by the size of the history file on the host. If their number exceeds this size, the oldest snapshot is deleted. Your chart reflects this change and the timestamp to the lower left of the chart is adapted.

Scenario 3: Creating a data view

Scenario 3 describes how to create a data view.

You can create a data view in one of the following ways:

- Create a data view from scratch as shown in [“Getting started” on page 30](#).
- Copy an existing data view and then change it according to your needs. You can also copy a data view that is delivered with OMEGAMON for Db2 Performance Expert and contained in the Default Group data group.
- Use the templates delivered with OMEGAMON for Db2 Performance Expert. To create a data view using a template click **Data views** under the data group that is to contain the new data view and then click **Selected > Open Predefined Data View**. The Open Predefined Data View window appears, which lists the templates that you can select.

Changing the properties of a data view

This topic describes how to change the properties of a data view.

Most of your tasks in system health are changes to the system health chart. They are explained in the following sections using the data view you created in [“Getting started” on page 30](#). To change this data view:

Click the data view in the folders pane and then click **Selected > Properties**. The Properties window is displayed, which is similar to the New Data View wizard.

Changing the method for displaying the counters

In many cases, it is useful to display the counters as lines instead of bars. Or, you might want to see the ratio between the counters.

OMEGAMON for Db2 Performance Expert provides the following chart types:

- [“Non-stacked bar chart” on page 33](#)
- [“Non-stacked line chart” on page 34](#)
- [“Stacked bar chart” on page 35](#)
- [“Stacked area chart” on page 36](#)
- [“Stacked-to-100% bar chart” on page 37](#)

To view these charts, you must change your selections under **Type of counter display** on the Counters page of the Properties window and under **Chart type** on the Graphics page of this window.

If a data view contains counters represented by numeric values and counters expressed as percentages, the chart has two scales, one for the numeric values and one for the percentages.

Non-stacked bar chart

This section describes the non-stacked bar chart.

A sample is shown in [Figure 7 on page 34](#):

- Characteristics and purpose:

The counters are displayed independently of one another as separate bars in a bar chart. They can be a mixture of numeric and percentage values. This chart type is useful if you want to know the absolute

values of each counter displayed and specify thresholds for individual counters (see [“Specifying and adjusting single thresholds”](#) on page 38).

- Selections on Counters and Graphics pages:

On the Counters page, select **Not stacked** as type of counter display.

On the Graphics page, select **Bar chart** as chart type.

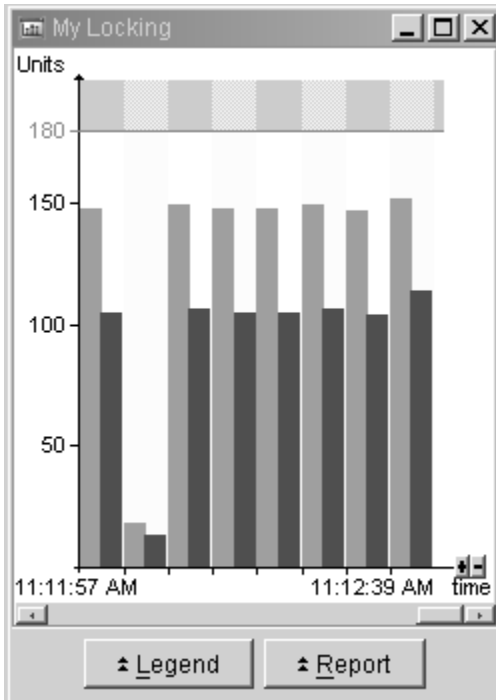


Figure 7. Non-stacked bar chart

Non-stacked line chart

This section describes the non-stacked line chart.

A sample is shown in [Figure 8 on page 35](#):

- Characteristics and purpose:

The counters are displayed independently of one another as separate lines in a line chart. They can be a mixture of numeric and percentage values. This chart type is useful if you want to get an overview of the absolute values of each counter displayed over a larger time period. You can specify thresholds for individual counters (see [“Specifying and adjusting single thresholds”](#) on page 38).

- Selections on Counters and Graphics pages:

On the Counters page, select **Not stacked** as type of counter display.

On the Graphics page, select **Line chart** as chart type.




Figure 8. Non-stacked line chart

Stacked bar chart

This section describes the stacked bar chart.

A sample is shown in [Figure 9 on page 36](#):

- Characteristics and purpose:

The counters are accumulated in a single bar. They can be numeric values or percentage values. This chart type is useful if you want to see the value that each snapshot adds up to.

- Selections on Counters and Graphics pages:

On the Counters page, select **Stacked** as type of counter display.

On the Graphics page, select **Bar chart** as chart type.




Figure 9. Stacked bar chart

Stacked area chart

This section describes the stacked area chart.

A sample is shown in [Figure 10 on page 37](#):

- Characteristics and purpose:

The counters are accumulated in an area. They can be numeric values or percentage values. This chart type is useful if you want to view, over a larger period, the value that each snapshot adds up to.

- Selections on Counters and Graphics pages:

On the Counters page, select **Stacked** as type of counter display.

On the Graphics page, select **Area chart** as chart type.

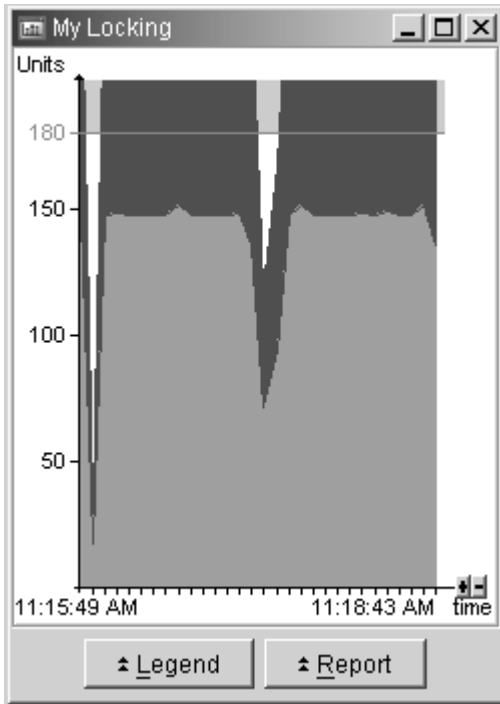


Figure 10. Stacked area chart

Stacked-to-100% bar chart

This section describes the stacked-to-100% bar chart.

A sample is shown in [Figure 11 on page 38](#):

- Characteristics and purpose:

The counters are shown proportionally in a bar chart. They are accumulated to add up to 100%. For example, if the counters have the values: a=2, b=5, c=7. the values shown are a=14%, b=36%, c=50%.

This chart type is comparable to a pie chart where single values are changed to percentages to fill 100%. It is useful if you want to see the ratio between the counter values and you are not interested in thresholds for individual counters.

- Selections on Counters and Graphics pages:

On the Counters page, select **Stacked to 100%** as type of counter display.

On the Graphics page, **Bar chart** is selected because this is the only choice.

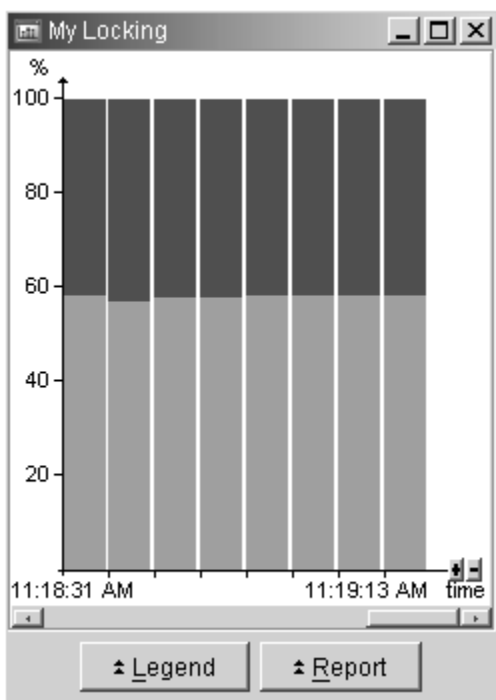


Figure 11. Stacked-to-100% bar chart

Specifying a global threshold

When defining a data view, you must specify a global threshold for all counters. This threshold enables you to see at which time a counter exceeds a certain value.

For example, when setting the threshold for the My Locking data view to 60 units per second, you want to see at a glance when a lock or an unlock request is performed, more than 60 times per second. The global threshold determines the scale of the chart and is represented as a red horizontal line. Step “5” on page 32 shows you how to define or change this global threshold.

If you are not sure which unit and time interval to specify, you can try a couple of values for the units and time interval. Or, to get you started, you can define the scale to automatically adapt to the highest visible value of a counter.

To adapt the scale to the highest visible value:

1. Display the Thresholds page of the Properties window.
2. Click **Dynamic scale** and then click **Finish**.

You can use this kind of counter display to calculate the units and time interval for the scale. However, it is also useful if you want to compare the actual values of the counters with each other or get a better view of the relation between two counters.

Specifying and adjusting single thresholds

In addition to the global threshold, you can also specify a threshold for each counter in your data view. These thresholds do not affect the scale of the chart. You can specify the thresholds if you view the counters in a non-stacked bar or line chart.

To specify thresholds for the lock and unlock request counters of the My Locking data view:

1. Display the Thresholds page of the My Locking Properties window.
2. Under **Single thresholds**, select the **Threshold** check box for each counter and specify a threshold:
 - For Lock requests, specify higher than 200 per Minute
 - For Unlock requests, specify higher than 100 per Minute

3. On the Graphics page, click **Line chart** as chart type and then click **Finish**.

You get a line chart with two additional threshold lines in the color of the respective counter. They are at about 20 units for lock requests and about 15 for unlock requests.

If you select a bar chart as chart type, you only see the threshold line if it is exceeded. In this case, the part of the bar exceeding this threshold is shaded.

Later you might want to change the global threshold (next to **Static scale** on the Thresholds page), which affects the scale of the chart. If you leave the single thresholds as they are, the proportions between the scale and the single thresholds are no longer correct.

To set the proportions right, perform the following steps on the Thresholds page of the Properties window:

1. Double the number next to **Static scale**: change 30 to 60.
2. The **Adjust single thresholds** button becomes available. Click this button.

Under **Single thresholds**, the thresholds for the counters are automatically doubled.

3. Click **Finish**.

Changing the appearance of the bars, lines, or areas in the chart

At any time, you can change the color of bar, line, and area charts and the pattern of bar charts. In addition, you can add a grid to your chart where a vertical line separates the snapshots. You perform these changes on the Graphics page of the Properties window.

Except for non-stacked bar charts, a grid separates each snapshot. In a non-stacked bar chart, however, two subsequent snapshots are distinguished by a different background. Therefore, the grid appears after every two snapshots.

Changing the number of visible snapshots in the chart

By default, a chart contains eight snapshots in a bar chart or 30 snapshots in a line or area chart. This is defined in the **Number of visible snapshots** field on the Graphics page of the Properties window. You might, however, want to examine a longer snapshot period in a bar chart or view fewer snapshots in a line chart. You can change the default to a number from 4 through 99 or the number in the **Maximum number of kept snapshots** field.

You can also use the plus (+) and minus (-) signs next to the x-axis of the chart to increase or decrease the number of snapshots displayed.

The **Maximum number of kept snapshots** field defines how many snapshots are kept in memory by the data view. When this limit is exceeded the oldest snapshot is deleted and the latest snapshot is added. You can change the number in this field according to your needs.

Freezing the display of snapshots

If you want to examine a snapshot in more detail or print the current chart, you must stop the continuous addition of new snapshots to your chart without canceling data collection. To achieve this freeze the display of snapshots. In the folders pane, click the data view that you want to freeze and then click **Selected > Freeze**.

The name of the data view is shown in *italics* in the folders pane for you to see that the view is frozen. Data collection continues in the background. To see the dynamic updates, click **Selected > Freeze** again.

Displaying data views as performance graphs in the System Overview window

You can display one or more data views as performance graphs on the Data Views page of the System Overview window. This data view must have been created in a data group that you defined and it must contain a link in the System Overview data group of the System Health window.

If you want to display the IFI interface data view in the System Overview window, click the IFI interface data view in the folders pane of the System Health window and then click **Selected > Display in System Overview**.

A link to the selected data view is shown in the System Overview data group and the linked data view immediately appears on the Data Views page of the System Overview window. The name of the link consists of the name of the original data group and the name of the original data view. The icon in front of the original data view has changed to indicate that this data view is shown in the System Overview window.

When you change a data view, you always change the original data view. The link is automatically updated.

The data views that are displayed on Data Views page of the System Overview are refreshed at the rate that you specified for the DB2 subsystem on the System Overview window. Snapshots continue to be collected and shown in the System Overview window even after you stop snapshot collection in the System Health window.

Reusing data-view and data-group configurations

You can copy data views and data groups within a DB2 subsystem and export and import entire data groups. This is useful, for example, if you want to share your configurations with your colleagues or use the same data views or data groups under several DB2 subsystems. Copying and moving data views between data groups also helps you organize your views.

To export or import a data group, click the data group in the folders pane and then click **Selected > Export** or **Selected > Import**. Use copy, cut, and paste to copy or move a data group or data view.

Gaps and sudden drops or rises in charts

This topic lists reasons for gaps and sudden drops or rises in charts.

If there is a gap in a line chart or between two bars in a bar chart, or two consecutive snapshots differ considerably from each other, the interval between the snapshots might have changed for one of the following reasons:

- DB2 has been restarted. Because the data view shows delta values, the delta between the last value before the DB2 stop and the first value after the DB2 restart would be a negative number. To reflect this situation, a gap appears in the chart and the Report section of the data view shows a value of *N/C*. To view the Report section click **Report**.
- The interval at which history data is collected might have been configured. A line chart shows a sudden drop or rise and a bar chart continues with much smaller or larger bars.

Monitoring activities of Db2 Connect gateways and connections of DCS applications

You can monitor the activities of DB2 Connect servers (DB2 subsystem) that serve as Db2 Connect gateways. You can also monitor the connections of Database Connection Services (DCS) applications within the Distributed Relational Database Architecture™ (DRDA) of DB2. You can use the information that you obtain through these monitoring activities to evaluate the health of a Db2 Connect gateway and to make changes necessary to tune the system.

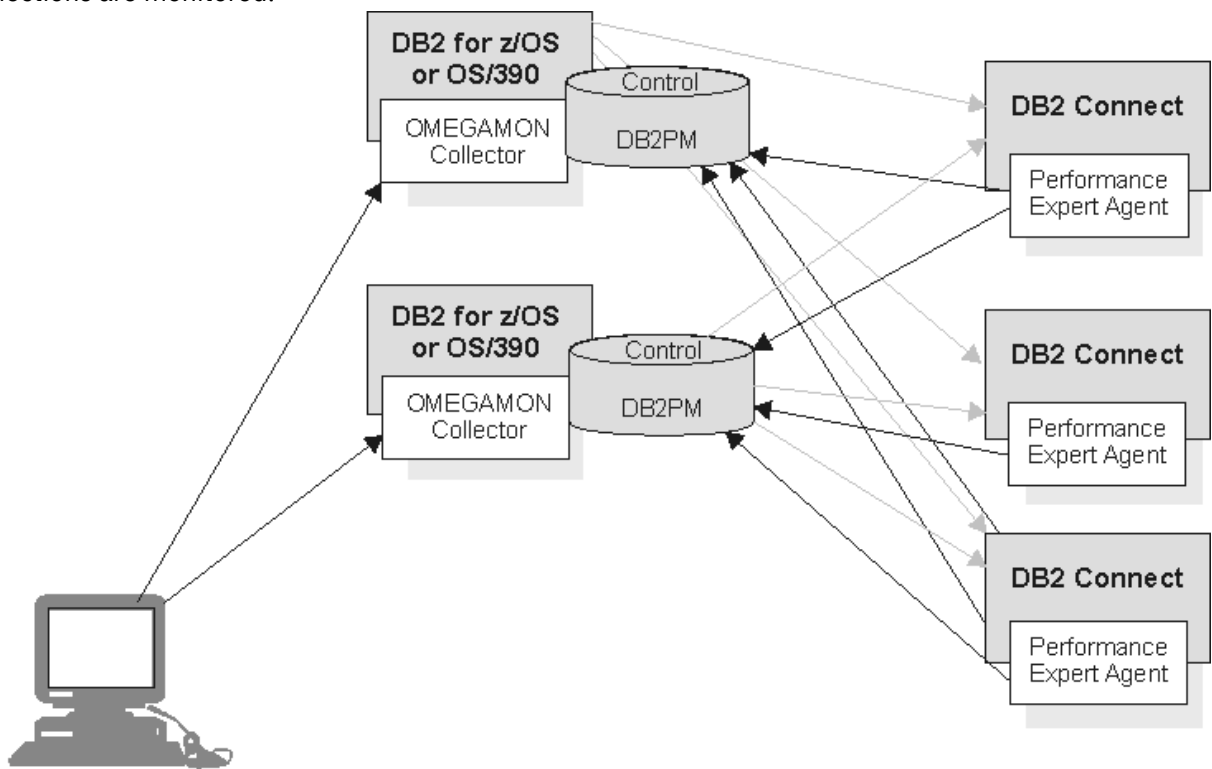
To monitor the activities of a Db2 Connect gateway and the connections of DCS applications, your system must meet the following criteria:

- Performance Expert Agent for DB2 Connect Monitoring is installed on the system, on which the Db2 Connect gateway runs, and is configured.
- At least one OMEGAMON Collector is registered for a Performance Expert Agent for DB2 Connect Monitoring.
- Performance Expert Agent for DB2 Connect Monitoring and the OMEGAMON Collector are started.

Refer to *Configuration and Customization* for information about how to install, configure, and start Performance Expert Agent for DB2 Connect Monitoring and to register the OMEGAMON Collector.

Performance Expert Agent for DB2 Connect Monitoring periodically collects snapshots of performance-related information about the Db2 Connect gateway and the connections of DCS applications and sends them to the OMEGAMON Collector. One Performance Expert Agent for DB2 Connect Monitoring can send data to one or more OMEGAMON Collectors, and one OMEGAMON Collector can receive data from one or more Performance Expert Agents for DB2 Connect Monitoring. The collected data is stored in the DB2PM database of the OMEGAMON Collector.

Figure 12 on page 41 provides an overview of how the activities of Db2 Connect gateways and the DCS connections are monitored.



Performance Expert Client

Figure 12. Monitoring the activities of DB2 Connect gateways and DCS connections

Performance Expert Client enables you to view the collected data from various perspectives. You can view:

- A list of all Db2 Connect gateways that are connected to a specific DB2 subsystem. You can also view Statistics details on them.
- Statistics details on a specific Db2 Connect gateway, independent of any DB2 subsystem
- A list of all DCS applications that are connected to any DB2 subsystem and view details on them
- Information about the DCS application that is related to a specific thread


Related information:

- Some windows or panes show the key performance counters in a table. One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns” on page 54](#).

Viewing all Db2 Connect gateways connected to a DB2 subsystem

You can display a list of all Db2 Connect gateways that are connected to a DB2 subsystem and view Statistics details on them.

To view Statistics details:

1. In the System Overview window, select a DB2 subsystem in the folders pane and then click the **Statistics Details** toolbar button ().

If you are not connected to this DB2 subsystem, you are prompted to log on.

2. In the folders pane of the displayed Statistics Details window, click **DB2 Connect Server**.

The DB2 Connect Server pane opens, which lists all the Db2 Connect gateways that are connected to this DB2 subsystem.


3. To get more Statistics details on a specific Db2 Connect gateway, double-click it.

The DB2 Connect Server *ip* page opens, where *ip* stands for the IP address of the selected Db2 Connect gateway. On the various panes, you find:

- Statistics about the Db2 Connect gateway, including details on the number of agents and pooled agents, the connections waiting for the host to reply, and the connections waiting for the client to send a request
- Statistics about the processes at the selected Db2 Connect gateway, such as the CPU and the working set
- The amount of time that is required to process a sample SQL statement
- Statistical details on the packages received at the Db2 Connect gateway

Viewing Statistics details on a specific Db2 Connect gateway

You can view information about a specific Db2 Connect gateway. In the System Overview window, select a

Db2 Connect gateway in the folders pane and then click the **Statistics** toolbar button ().

If you are not connected to this Db2 Connect gateway, you are prompted to log on to the OMEGAMON Collector that is receiving information from the Performance Expert Agent for DB2 Connect Monitoring that is monitoring the Db2 Connect gateway.

The Statistics window opens. The information shown in this window is identical to the information on the DB2 Connect Server *ip* page of the Statistics Details window. See [“Viewing all Db2 Connect gateways connected to a DB2 subsystem” on page 42](#).

Viewing information about all DCS applications connected to any DB2 subsystem

You can display a list of all the DCS applications that are connected to any DB2 subsystem and view details on them.

To view details:

1. In the System Overview window, select a Db2 Connect gateway in the folders pane and then double-click **Application Summary** in the contents pane.

If you are not connected to this Db2 Connect gateway, you are prompted to log on.

The Thread Summary window opens, which shows the key data for the DCS application connected to any DB2 subsystem.

2. To get more details on a specific DCS application, double-click it.

The Thread Details window opens. On its various panes, you find:

- An overview of the DCS application, such as cumulative counters and status information.
- Information about the SQL statement that was last executed by the DCS application, including network Statistics, the host response time, and the CPU time distribution.
- Outbound and inbound data and the network time taken to transmit the data.

Viewing information about the DCS application related to a thread

You can view information about the DCS application that is related to a specific thread.

To view information about the DCS application that is related to a specific thread:

1. In the System Overview window, click a DB2 subsystem in the folders pane and then double-click **Thread Summary** in the contents pane.

If you are not connected to this DB2 subsystem, you are prompted to log on.

The Thread Summary window opens, which shows the key performance counters for all connected threads.

2. Double-click a thread.

The Thread Details window opens.

3. In the folders pane of the Thread Details window, click **DB2 Connect Server**.

The DB2 Connect Server pane is displayed, which lists the related DCS application.

4. To get more details on this DCS application, double-click it.

The DB2 Connect Server *appl_id* page is displayed. *appl_id* stands for the identifier that is generated when the application connects to the database at the database manager or when Distributed Database Connection Services (DDCS) receives a request to connect to a DRDA database. The information shown on the various panes of this page is identical to the information on the Thread Details window. See [“Viewing information about all DCS applications connected to any DB2 subsystem” on page 42.](#)

The gateway snapshot time

The Gateway Snapshot Time counter on the DB2 Connect Server pane of the Statistics Details window and on the Statistics window shows the timestamp at which the snapshot was taken at the Db2 Connect gateway. This timestamp usually differs from the snapshot timestamp of the OMEGAMON Collector.


If you are viewing current data, the OMEGAMON Collector displays the most recent data that was received from the Db2 Connect gateway. In history mode, the OMEGAMON Collector searches for a Db2 Connect gateway snapshot that is closest to the history snapshot of the OMEGAMON Collector.

If the Db2 Connect gateway, OMEGAMON Collector, Performance Expert Client, or all of them are in different times zones, the times shown in the Gateway Snapshot Time counter and for the OMEGAMON Collector snapshot are adapted to the Performance Expert Client time. For example, if the OMEGAMON Collector receives a snapshot from the Db2 Connect gateway that was taken at 11 a.m., the OMEGAMON Collector location is two hours ahead of the Db2 Connect gateway location, and the Performance Expert Client location is three hours ahead of the OMEGAMON Collector location, the Gateway Snapshot Time counter displays a time of 16:00 (4 p.m.).

Viewing DB2 system parameters

You can view the system parameter values that you specified when configuring a DB2 subsystem.

To view the database configuration do one of the following:

- In the System Overview window, select the DB2 subsystem in the folders pane and then click **System Parameters** toolbar button ().

- You can also use the menu bar or pop-up menu: click the DB2 subsystem in the folders pane, then click **Selected > Monitor > System Parameters** on the menu bar or **Monitor > System Parameters** on the pop-up menu.

If you are not connected to the selected DB2 subsystem, you are prompted to log on.

The System Parameters window shows the current values for the selected DB2 subsystem.

To view detailed information for a system parameter, click the name of the parameter in the folders pane.

Related information:

- The System Parameters window shows some of its system parameters in tables. In contents panes containing a table:
 - One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns”](#) on page 54.
 - You can choose to view details on a table entry in another window or on a notebook page. For more information, refer to [“Viewing details on a table entry in a contents pane”](#) on page 54.
- Some windows or panes show the key performance counters in a table. One or more columns might be hidden. For information about how to view hidden counters, to hide counters that you are not interested in, or to change the order of the columns, refer to [“Customizing table columns”](#) on page 54.
- Performance counter values that are too long are truncated. For information about how to view the entire value, refer to [“Viewing long performance counter values”](#) on page 54.
- If you are working in different time zones, refer to [“Working in different time zones”](#) on page 55 for more information.

Working with exception processing

Exception processing is an efficient way to monitor Db2 subsystems for performance problems. When exception processing is active, you are informed when certain events occur, or when specific counters reach levels that might produce problems.

Exceptions are recorded by the OMEGAMON Collector. To run exception processing, it must be set up during the customization of the OMEGAMON Collector.

For more information about setting up exception processing, refer to [Configuration and Customization](#).

Exception processing is started when the OMEGAMON Collector is started. It continues until you stop it, the OMEGAMON Collector or Db2 is stopped, or you log off from the Db2 subsystem that you are monitoring. When you leave Performance Expert Client you can request that the collection of exception processing data is continued.

OMEGAMON for Db2 Performance Expert supports:

- Event exception processing: This monitors a Db2 subsystem for the occurrence of particular events.
- Periodic exception processing: This periodically monitors thread-activity and Statistics fields for exception conditions.
- Ability to launch an IBM Db2zAI System Assessment from periodic and event exception detail displays for improved Root Cause Analysis.


Note: Db2zAI integration is valid only for Db2 13 and Db2 12 subsystems.

This section explains how to start and stop exception processing, to view the exceptions that occurred, and to use the exception user exit.

Starting exception processing

This topic explains how to start exception processing.

To work with exception processing, you can do one of the following:

- In the System Overview window, select the DB2 subsystem in the folders pane and then click the **Exception Processing** toolbar button ().
- In the System Overview window, click **Tools > Exceptions > Exception Processing**.

The Exception Processing window is displayed.

This window provides information about the status of both event and periodic exception processing for all DB2 subsystems to which you are logged on. The icons that are displayed when you click **Event Exceptions** in the folders pane of the window have the following meaning:



The event is supported by the OMEGAMON Collector and switched on at its startup.



The event is supported by the OMEGAMON Collector but was not switched on at its startup.

N/A

The event is not supported by the OMEGAMON Collector.

The Exception Processing window shows up to the latest 9999 exceptions recorded. If you want to change this number, click **Exception Processing > Settings**. The Exception Settings window is displayed. On its pages, you can change the number of entries that is collected in the log for event exception and periodic exception processing.

If you want to be informed whenever an exception is recorded, you can choose to be notified by a message or a beep. Click **Exception Processing > Settings**. The Exception Settings window is displayed. On its pages, you can specify the kind of notification for event exceptions and periodic exceptions.

Processing event exceptions

This topic shows how to process event exceptions.

When you monitor event exceptions, you are informed about the following events:

- Authorization failure
- Coupling facility rebuild start
- Coupling facility rebuild stop
- Coupling facility alter
- Data set extension
- Deadlock
- EDM pool full
- Global trace start
- Log space shortage
- Thread commit in doubt
- Timeout
- Unit of recovery problem

All events must be enabled on the OMEGAMON Collector.

Refer to [Configuration and Customization](#) for information about how to enable an event.

You can monitor exceptions for one or more DB2 subsystems.

Starting and stopping event exception processing

This topic show how to start and stop event exception processing.

To start monitoring for event exceptions:

1. In the Exception Processing window, click **Exception Processing > Activation - z/OS**.

The Activation window is displayed. The exact name of, the layout of, and the selections possible on, this window depend on whether you are logged on to one or to several DB2 subsystems. The following table shows the various possibilities.

If you are logged on to ...	Description of Activation window
One DB2 subsystem	The Activation – <i>subsystem</i> is displayed, where <i>subsystem</i> stands for the name of the DB2 subsystem to which you are logged on. The event list shows all events supported by z/OS. The events enabled by the OMEGAMON Collector are selected.
Several DB2 subsystems	The Activation – z/OS window is displayed. This window consists of the DB2 subsystems pane, listing all DB2 subsystems to which you are logged on, and the Activation pane. Select the DB2 subsystems you want to monitor. The event list appears and shows all events supported by z/OS. The events that are enabled on the OMEGAMON Collector of at least one DB2 subsystem are selected. The selected events that are not enabled by all OMEGAMON Collectors are followed by an asterisk (*).

2. Click the **Start** (▶) button next to **Event exception processing**.

To stop event exception processing:

1. In the Exception Processing window, click **Exception Processing > Activation**.

The Activation window is displayed.

2. Click the **Stop** (■) button next to **Event exception processing**.

Viewing event exceptions

To get more information about an event exception, double-click it in the Event Exception Log pane of the Exception Processing window or in the 100 Most Recent Exceptions pane of the System Overview window.

When an event exception occurs, you are informed in the following ways:

- An entry is added to the Event Exception Log pane of the Exception Processing window.
- An entry is added to the 100 Most Recent Exceptions pane of the System Overview window.
- A message appears, a beep is issued, or both, if you requested this kind of notification in the Exception Settings window.

Viewing deadlocks

This topic shows how to find out if a deadlock occurs.

When a deadlock occurs, you are informed in the following ways:

- An entry is added to the Event Exception Log pane of the Exception Processing window.
- In the folders pane of the System Overview window, the icon changes next to the DB2 Subsystem for which the deadlock was recorded and next to the folders that contain this DB2 Subsystem. The event exception icon (🔔) change remains, by default, for at least five minutes or for the minimum period of time that you specified on the Preferences page of the Configuration window. See [“Specifying global settings”](#) on page 14. An event exception icon overrules any warning or problem icon.
- An entry is added to the Exceptions page of the System Overview window.
- By a message, a beep, an e mail, or a mixture of them. Whether you receive these types of information, and which of them, depends on your specifications in the Exception Settings window and on the

Exception page of the OMEGAMON Collector Properties window. To receive an e mail, you must also have requested this type of information about the Activation window.

- By using a user exit. In this case, you must have enabled the user exit and specified its path on the Exception page of the OMEGAMON Collector Properties window. In addition, you must have selected the user exit in the Activation window. For information about how to write a user exit, refer to [“Working with the exception user exit”](#) on page 49.

Viewing details on a deadlock

When a deadlock occurs, you need to know what caused the deadlock and how to resolve it. OMEGAMON for Db2 Performance Expert helps you find the reasons for a deadlock by supplying information, such as which threads are involved in the deadlock, which statements led or might have led to the deadlock, and which thread holds the resource that is requested by another thread.

To view this information:

1. Double-click a deadlock in the Event Exception Log pane of the Exception Processing window or on the Exceptions page of the System Overview window. The various pages of the Deadlock Details window provide you with all the information that you need to narrow down the reasons for the deadlock.
2. To get information about the thread that was rolled back (*victim*), about the system and the DB2 Subsystem on which the deadlock occurred, and about the number of threads that are involved in the deadlock, have a look at the Overview page.

Processing periodic exceptions

You can use periodic exception processing to periodically check Db2 counters against specific thresholds. When the value of a counter falls above or below these thresholds, an exception is logged. To run periodic exception processing, you must first define thresholds for specific fields.

Establishing thresholds

Follow these steps to establish periodic exception thresholds.

About this task

Db2 provides a comprehensive set of thread activity and statistics fields. However, you usually need to define only a limited number of thresholds. Because long response times are a good indicator of a performance problem, you can start by defining exception thresholds for time fields. For example, you can set thresholds that are appropriate to your environment for class 1 and class 2 elapsed times.

- Choose the most important applications or transactions in your system.
- Always define exception thresholds for critical business applications. Applications that are frequently executed are good candidates for exception thresholds.
- In a threshold set, specify thresholds for all Db2 counters that you want to monitor together.

Procedure

1. In the Exception Processing window, click **Exception Processing > New Threshold Set**.
2. Select **z/OS** as operating system.
3. Type a unique name for the set, your name (this is usually your logon user ID), and click **OK**.
4. Add a threshold.
 - a) Click **Threshold Set > New Threshold**.
 - b) Under **Exception field**, select the counter you want to monitor, select the appropriate exception category and exception subcategory, and select the counter from the **Select the exception field** list.
 - c) Under **Warning and problem thresholds**, specify the threshold values for warnings and problems for the selected counter.

- d) Under **Qualifier**, specify filter criteria for the counter.
 - e) Click **OK** to save your settings.
5. (Optional) You can perform one or more of the following optional steps:
 - a) **Add thresholds to the threshold set:** Click **Threshold Set > New Threshold**.
 - b) **Change a threshold:** Click the threshold and select **Selected > Edit**.
 - c) **Deactivate a threshold in a threshold set:** Click the threshold and select **Selected > Active**.
 6. Click **Threshold Set > Save**.
The Exception Processing window lists the new threshold set under the **z/OS** folder in the **Threshold Sets** folder.
 7. You can also:
 - a) **Change the properties of a threshold set:** Click the threshold set in the folders pane of the Exception Processing window and select **Selected > Properties**.
 - b) **Add or remove thresholds from a threshold set:** Click the threshold set in the folders pane of the Exception Processing window and select **Selected > Edit**.

Starting and stopping periodic exception processing

This topic describes how to start and stop periodic exception processing.

To start periodic exception processing:

1. In the Exception Processing window, click, for example, **Exception Processing > Activation - z/OS** (depending on which operating system you want to start periodic exception processing).
The Activation window is displayed. For more information about the exact name of, the layout of, and the selections possible on this window, refer to [Table 2 on page 46](#).
2. If you want to start periodic exception processing for several DB2 subsystems at once, you must select the subsystems you want to monitor.
3. Ensure that your threshold set is displayed in the **Threshold set** field. In addition, type the time interval at which the thread and Statistics fields are checked for exceptions.
4. Click the **Start** (▶) button next to **Periodic exception processing**.

All active thresholds in the threshold set are used during periodic exception processing.

To stop periodic exception processing:

1. In the Exception Processing window, click **Exception Processing > Activation**.
The Activation window is displayed.
2. Click the **Stop** (■) button next to **Periodic exception processing**.

Viewing periodic exceptions

This topic describes how to view periodic exceptions.

When a periodic exception occurs that is outside of the specified warning or problem threshold, you are informed in the following ways:

- An entry is added to the Event Exception Log pane of the Exception Processing window.
- A warning icon (▲) or a problem icon (●) appears:
 - Next to the affected counter in the window that contains this counter.
 - In the folders pane of the System Overview window, next to the DB2 subsystem that contains the affected counter and next to the folders that include this DB2 subsystem.

A problem icon overrules a warning icon. No icon is shown for counters that are listed in a table.

When you move your mouse pointer over the counter, hover help shows the timestamp of the exception. If this counter is on a KPIs page of the System Overview window, hover help also contains the defined warning and problem thresholds.

The icon disappears after the counter value falls below or above the threshold or you stop periodic exception processing and after you refresh the data. In the System Overview window, you need not refresh the data.

- An entry is added to the 100 Most Recent Exceptions pane of the System Overview window.
- A message is displayed, a sound (beep) is played, or both, if you specify this kind of notification in the Exception Settings window.

You can get more information about a periodic exception from the Periodic Exception Log pane of the Exception Processing window and the 100 Most Recent Exceptions pane of the System Overview window. In the Periodic Exception Log pane, click the exception and then:

- Click **Selected > Current Statistics Details** to see the corresponding Statistics values.
- Click **Selected > Field Exception History** to see a log of the exception values for the counter.
- Click **Selected > Thread Identity** to see information about the application causing the exception.
- Click **Selected > Current Thread Details** to see thread information. (This is available only if your OMEGAMON Collector supports this function.)

In the 100 Most Recent Exceptions pane, right-click the exception and select the information you want to see.

Working with the exception user exit

OMEGAMON for Db2 Performance Expert supports a user exit to handle event and periodic exceptions. This user exit is a predefined exit point where a user-written routine can be given control.

The name of the exception processing user exit is DGOMUPXT. A sample of this user exit is provided in TKO2SAMP.

To activate the user exit, do the following in the Exception Processing window:

1. Click **Exception Processing > Activation**.

The Activation window is displayed.

2. Select the **User exit** check box.
3. Click the **Start** (▶) button next to **Periodic exception processing**.

Periodic exception processing must be running in order to use the user exit—even if you are only interested in event exceptions.

Refer to [Configuration and Customization](#) for information about customizing this user exit.

Launching IBM Db2zAI System Assessment Tool

Throughout event and periodic exceptions details, you can launch an IBM Db2zAI System Assessment by clicking the **Analyze root Cause** button.

About this task

Note: Db2zAI integration is valid only for Db2 13 and Db2 12 subsystems.

Monitoring data sharing groups

If you are working with a member of a data sharing group, you might want to view, from time to time, the monitoring information collected for another member in the data sharing group. Or you might want to compare the data collected for all members in the group. Most of the monitoring functions allow you to

switch among the individual members and between a member and the group. You find a **Show Data for** list in the toolbar.

From this list, you can select another member or the group. Depending on your selection, the window shows the data for the selected member or the group. If a member is not active, the member name in the list is followed by QUIESCED or FAILED. This depends on why the member became inactive.

Sometimes, while you are working with a member, another member is added to the group, a member is deleted, or the status of a member is changed. To refresh the displayed list, click **View > Refresh Data Sharing Group Information**.

To get monitoring information from all members of a data sharing group, ensure that the monitor traces are started on all members.

For information about how to ensure this, refer to the description of the DATASHARINGGROUP and DB2COMMAND startup parameters in *Configuration and Customization*.

This section describes the windows displayed when you monitor the thread of a data sharing group and view Statistics information about it.

Monitoring the threads of a data sharing group

When you request the display of all active threads connected to a member of a data sharing group, the Thread Summary window shows the member name in the **Show Data for** list and lists all active threads running on this member.

When you select the group from the **Show Data for** list, all active threads running on the entire group are displayed.

The **Member** column in the window shows, for each thread, the member on which it is running. In the case of CP parallelism or sysplex query parallelism, the parent thread is displayed and the values of the children are accumulated.

Viewing Statistics on a data sharing group

When you request Statistics information about a member of a data sharing group, the Statistics Details window shows the member name in the **Show Data for** list and displays the Statistics information for this member.

When you select the group from the **Show Data for** list, the entries in the folders pane are reduced to those available in group view. The performance counters are displayed in tabular format, with a column for each active member of the data sharing group.

In member view, part of the Statistics information is already displayed in tabular format, such as the information about the Buffer Management pane. When you view this pane for the entire group, the Buffer Management pane shows the values for the most important buffer-pool performance counter, which is the hit ratio, for each buffer pool and member in the data sharing group. The counter name is displayed above the buffer-pool table.

Thus, when Statistics information is shown in tabular format in member view and you want to view this information in group view, OMEGAMON for Db2 Performance Expert selects the most important performance counter in the table and displays its values for the entire group.

Viewing the system health of a data sharing group

When you want to view the system health of a member of a data sharing group, the System Health window shows the member name in the **Show Data for** list and displays the system health for this member.

When you select another member from the **Show Data for** list, the displayed data views are cleared and the snapshots of the selected member appear.

When you select the group from the **Show Data for** list, the displayed data views show the following:

- A chart for each member in the data sharing group.

- A list of all members in the data sharing group, instead of a report.

To hide the chart of a member, clear the appropriate check box in the member list. To display the charts in a different order, right-click the member you want to move and then click the appropriate command. To view the report on a member, right-click its chart.

Viewing history data

You can view history data for threads, locking conflicts, Statistics, and system parameters. This data is stored as snapshots.

History data is useful, for example, if you want to examine events that lead to or follow an exception without recreating the situation.

The OMEGAMON Collector must be enabled for collecting history data. The data is then collected at intervals that you defined during installation.

For information about setting up the OMEGAMON Collector to collect history data, see [Configuration and Customization](#).

You can display history data in the Thread Summary, Statistics Details, System Health, and System Parameters window, and in all windows dealing with threads involved in a locking situation.

To view history data:

1. Click **View > History**.

The history slider is activated.

The slider is divided into units where each unit stands for the date and time at which a snapshot was collected. The timestamp above the slider indicates the date and time of the snapshot at the current slider position. The timestamps below the slider show the time range through which data is available in history mode. This is useful, for example, if you zoom the time interval displayed to you.

2. To view the data of a specific snapshot, do one of the following:

- Move the slider to the date and time you want.
- Click **View > History Settings**. A window is displayed on which you can specify the date and time of the snapshot data you want. Click **OK** to get the snapshot data.
- To view the next snapshot, click the right arrow (→) or **View > Forward**. To view the previous snapshot, click the left arrow (←) or **View > Back**. You can repeat these commands until you have reached the snapshot you want.
- If the slider contains too many units for you to move the slider to an exact date and time, you can decrease the time interval shown by clicking the plus (+) sign above the slider as often as you need. To increase it again, click the minus (-) sign.

To display the current data again, click **View > History** again to deactivate the history slider.

Restriction: When you view SQL statements in the dynamic SQL cache in history mode, you cannot define a filter. For more information about SQL statements data, refer to [“Viewing the SQL statements in the dynamic SQL cache”](#) on page 28.

Using Explain products to analyze and improve SQL statements

This topic explains how to use Explain products to analyze and improve SQL statements.

To determine and improve the access path of dynamic and static SQL statements that contain a SELECT, UPDATE, INSERT, DELETE, or VALUES clause, you can use one of the following Explain products:

- IBM Db2 Query Workload Tuner for z/OS Version 5.1
- IBM InfoSphere® Optim™ Query Workload Tuner Version 4.1 or later (including Data Studio)
- IBM Data Server Manager V2.1.5
- IBM Db2 Administration Foundation for z/OS

You can use Explain products to view the access path of a dynamic or static SQL statement when:

- Monitoring an application or thread
- Viewing Statistics on the SQL statements stored in the dynamic SQL statement cache

If you want to use Query Workload Tuner, you must first define the local database alias of the currently monitored DB2 subsystem:

1. Open the following wizard or window:
 - The New DB2 System wizard as described in [“Defining a Db2 subsystem” on page 10](#)
 - The DB2 System Properties window as described in [“Changing the properties of a Db2 subsystem” on page 13](#)
2. Open the DB2 Subsystem page.
3. In the **Local DB Alias** field, type the local database alias of the DB2 subsystem.

Displaying the access path while monitoring an application or thread

This topic shows how to view the access path while monitoring an application or thread.

To view the access path:

1. Open the Thread Details window.
2. In the folders pane, click **SQL Statement**.

The SQL Statement pane shows the SQL statement that is currently run by the thread. The **Explain** button is enabled.

3. Under **SQL statement text**, click **Explain**. If the statement contains a SELECT, UPDATE, INSERT, DELETE, or VALUES clause, Query Workload Tuner is opened to view the access path for this statement. Alternatively, if Administration Foundation is configured, a popup will be shown, returning Administration Foundation Tuning Job ID and clickable URL to Administration Foundation Tuning Job summary panel.

Note: Usually only one of the Explain products is started at the same time. (This depends on the data bridge server connection that is used.) You can also cut and paste the SQL statement text from Performance Expert Client to manually create a project.

Displaying the access path while viewing Statistics on SQL statements


This topic shows how to display the access path while viewing Statistics on SQL statements stored in the dynamic SQL statement cache.

To display the access path:

1. Open the Statistics Details window.
2. In the folders pane, click **SQL Statements**.

The SQL Statements pane is displayed without any entries.

3. To display the SQL statements that are in the cache:

- a. Click the **Filter**  toolbar button to check the filter values and set the appropriate values for your system. For more information about defining filter criteria, see step [“5” on page 28](#).
- b. Select the **Receive statement cache information** check box.

The list is refreshed whenever you click **View > Refresh** or when an automatic refresh occurs.

4. Double-click a statement.

The SQL Statements *text* page opens, where *text* stands for the first couple of characters of the statement text.

5. Under **SQL statement**, click **Explain**.

6. If the statement contains a SELECT, UPDATE, INSERT, DELETE, or VALUES clause, the following window opens depending on the Explain products installed:
 - If Query Workload Tuner is installed and launched, a new Project is defined and the Statement text is shown. Use the corresponding function of Query Workload Tuner to generate the Access Plan.
 - If Administration Foundation is configured, a popup will be shown, returning Administration Foundation Tuning Job ID and clickable URL to Administration Foundation Tuning Job summary panel.

Exporting snapshot data for saving and printing

You can export snapshot data to HTML format. You can then print the data or save it as an HTML file.

To export snapshot data:

1. Do one of the following in the window containing the snapshot data:

- Click the **Print/Export** (🖨️) toolbar button.
- Click the leftmost menu and then click **Print/Export**.

The Print/Export window is displayed.

2. You can change the settings in this window as follows:

- If the window consists of several pages and you want to export them all, click **All** under **Pages**.
- If the window contains a large table, you can:

- Specify how many columns and rows you want to export.

Example: To export only the first five columns, select **Only first** for **Columns** and then specify 5 in the field.

- Specify after which column the table is split. Select the **Split after** check box and then select the column number.

3. Click **OK**.

A Web browser window is displayed showing the window contents in HTML format. You can use your Web browser to save or print the contents of the window.

If your window contains a table with many columns and you specified to split the table after a specific number of columns, for example three, your HTML file looks similar to the example shown in [Figure 13](#) on page 53.

Buffer Manager

Number	Buffer Pool	Hit Ratio	Getpage Requests
1	BPO	75.0	20
2	TOTAL	75.0	20

Number	Synch Read I/O	Buffer Updates
1	5	0
2	5	0

Figure 13. Example of a split wide table

To make the table easier to read, the rows are numbered in the **Number** column.

Working with windows displaying snapshots

This section provides you with the information that applies to all windows displaying snapshots.

Customizing table columns

Some contents panes of the Thread Details, Statistics Details, and System Parameters window show the key performance counters in a table. One or more columns might be hidden.

To view hidden counters, to hide counters that you are not interested in, or to change the order of the counters, perform the following steps:

1. Click **View > Customize Columns**.

The Customize Columns window is displayed.

2. To make a column visible, click the column in the Hidden columns list and then click Add (►). The column is removed from the Hidden columns list and displayed at the bottom of the Visible columns list.

To make all hidden columns visible, click ►►.

3. To hide a column, click the column in the Visible columns list and then click Remove (◄). The column is removed from the Visible columns list and displayed in the Hidden columns list.

To hide all columns, click ◄◄.

4. To move a column up in the hierarchy, click the column in the Visible columns box and then click Move up. To move this column to the top of all columns, click Move to top.

To move a column down in the hierarchy, click the column in the Visible columns box and then click Move down. To move this column to the bottom of all columns, click Move to bottom.

5. Click OK.

The previous window is displayed with the columns rearranged, hidden, or made visible according to your specifications.

Viewing details on a table entry in a contents pane

Some windows show part of the performance information and system parameters in tables. In the contents panes containing a table, you can get further details on a table entry by double-clicking the entry. When you work with Performance Expert Client for the first time, the window becomes a notebook where the initial counter categories or system parameters are shown on the Main page and where further details on the additional notebook pages are displayed.

If you want to compare the data on the individual pages, you can change to window view where each notebook page becomes a window within the main window that you can move, cascade, and tile according to your needs. Click **View > Switch to window view** in the main window. To return to page view, click **View > Switch to page view** in the main window. Your selection remains valid throughout Performance Expert Client until you change it again.

Viewing long performance counter values

Some performance counter values exceed the length of the field in which they are displayed and are therefore truncated.

To view the entire value, move your mouse pointer over the field containing the value. Hover help is displayed showing the value in its full length. If the value is longer than 40 characters, a line break is inserted after every 40 characters.

To view the invisible parts of the value, click in the field containing the value and then use the left or right arrow, or press the Home or End key to go to the beginning or end of the value.

You can also copy the value by marking it and using copy and paste.

Working in different time zones

The windows displaying snapshots contain timestamps to inform you when the snapshots were taken. If the OMEGAMON Collector and Performance Expert Client are in different time zones, the snapshot times are adapted to the Performance Expert Client time.

For example, if a snapshot is taken at 4 p.m. and the time at the OMEGAMON Collector location is six hours ahead of your Performance Expert Client, the window displaying this snapshot shows a time of 10:00 (10 a.m.).

Chapter 5. Working with Performance Warehouse

This topic first describes the tasks you need to perform before you can start working with Performance Warehouse and then how to use Performance Warehouse to perform analysis tasks on DB2 performance data using processes, rules of thumb, and queries.

Performance Warehouse provides a process-oriented view of performance-analysis tasks. It allows you to automate tasks that previously required user interaction, such as loading DB2 data into the Performance Warehouse database or generating reports.

A process can consist of steps that link together these manual tasks. For example, you might want to transfer DB2 trace records on your MVS system to the Performance Warehouse database for further analysis. Your corresponding Performance Warehouse process might consist of a:

- *CONVERT* step that converts a save data set to a Save-File data set.
- *LOAD STATISTICS* step that loads the data set into the Performance Warehouse database.

Processes represent recurring performance-analysis tasks, they can be scheduled for automatic and repeated execution. You can trace these process executions from the workstation.

Processes produce output data that provides DB2 performance-related information. You can load this data into the Performance Warehouse database and then use rules of thumb and SQL queries to analyze it. Rules of thumb represent a category of expert rules. These rules help you identify more complex performance problems and provide you with tuning recommendations. For example, you can combine several performance counters in a value expression and relate this expression to problem and warning thresholds. You can also create, store, and run SQL queries on the Performance Warehouse database.

Sample processes, rules of thumb, and queries are provided to help you get started.

Running Performance Warehouse in a data-sharing environment

The Performance Warehouse database is maintained by the OMEGAMON Collector that runs with the startup parameter `PERFORMANCEWAREHOUSE=YES`. In a data-sharing group, the OMEGAMON Collectors might be installed on more than one member.

Start the Performance Warehouse server task only for one OMEGAMON Collector in the data-sharing group to avoid conflicts.

[Figure 14 on page 58](#) shows a typical data-sharing environment running Performance Warehouse.

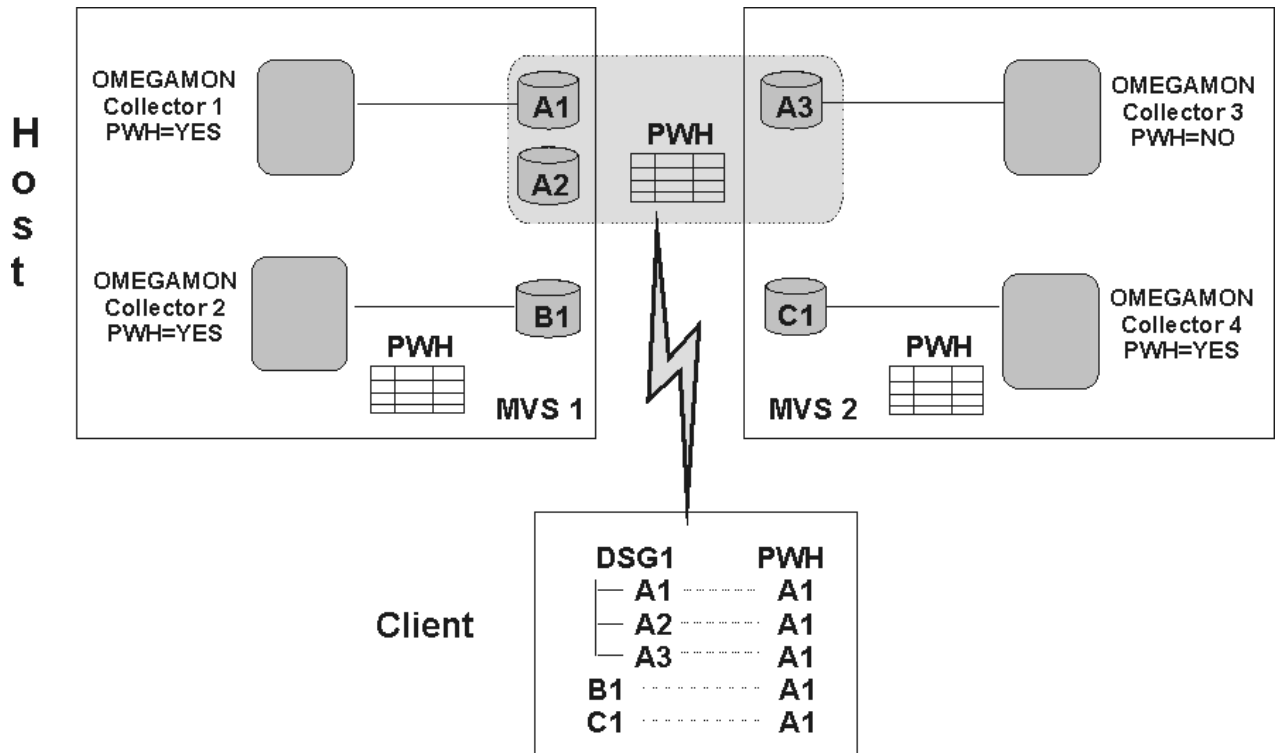


Figure 14. Performance Warehouse in a data-sharing environment

The environment includes:

- Two MVS systems in a Parallel Sysplex® environment with one data-sharing group, DSG 1, that comprises the members A1, A2, and A3, and additional Db2 subsystems B1 and C1.
- A Performance Warehouse that contains the performance data of the Db2 subsystems A1, A2, and A3, and a Performance Warehouse each for B1 and C1.
- OMEGAMON Collector 1 that is dedicated to data-sharing group DSG 1 on MVS 1 maintains the Performance Warehouse database in a data-sharing group.

The OMEGAMON Collectors 1, 2, and 4 run the processes to produce reports and load data into the Performance Warehouse database. If SMF data is to be loaded, the OMEGAMON Collector requires read access to the corresponding SMF data set.

Performance Expert Clients that want to use Performance Warehouse connect to the Db2 subsystem or the Db2 data sharing group that contains Performance Warehouse. If you wanted to connect to the data-sharing group and access Performance Warehouse, you would specify the location of DSG 1 as the database name and MVS 1 or MVS 2 as the host name when cataloging the DB2PM database of the OMEGAMON Collector regardless of which member of DSG 1 you want to connect to.

Logging on to Performance Warehouse

You log on to a Performance Warehouse by connecting to its Performance Warehouse database.

You can log on to Performance Warehouse automatically or manually:

- You can automatically log on to a Performance Warehouse associated with a Db2 subsystem. You can associate a Performance Warehouse with a Db2 subsystem when defining the Db2 subsystem for monitoring from Performance Expert Client (see step “7” on page 11) or when changing the properties of the Db2 subsystem (see “Changing the properties of a Db2 subsystem” on page 13).
- You can manually log on to an associated Performance Warehouse or to any Performance Warehouse that was configured. When you log on manually the Db2 subsystem that it is associated with, the database need not exist any more. If you removed the associated Db2 subsystem from the System Overview window but kept the Db2 catalog entries, you can still work with Performance Warehouse.

If you log on to a Performance Warehouse database that is maintained by an OMEGAMON Collector that is not running, you can define processes, rules of thumb, and queries but you cannot run them.


Logging on to a Performance Warehouse automatically

To log on to a Performance Warehouse automatically, select a DB2 subsystem ID in the folder pane of the System Overview window, then click **Selected > Monitor > Performance Warehouse - Expert**.

The Expert page of the Performance Warehouse window is displayed. The folders pane lists, under the z/OS folder, all Performance Warehouse systems to which you previously logged on or that you added manually (see *Managing Performance Warehouse systems*). The names that appear are the DB2 connection aliases, which associate a Performance Warehouse with a DB2 subsystem.

In addition, you are logged on to the Performance Warehouse that is associated with the selected DB2 subsystem. You can now work with the items that are listed on the displayed page of the Performance Warehouse window, or you can move to another page to perform a different set of tasks.

Logging on to a Performance Warehouse manually

To log on to a Performance Warehouse manually, click the **Performance Warehouse** toolbar button () in the System Overview window.

The Expert page of the Performance Warehouse window is displayed. The folders pane lists, under the **z/OS** folder, the DB2 connection alias of all Performance Warehouse systems to which you previously logged on or that you added manually (see [“Managing Performance Warehouse systems” on page 59](#)).

If no Performance Warehouse systems are listed because you are working with Performance Warehouse for the first time or the Performance Warehouse of your choice is not listed, add the Performance Warehouse systems you need. For more information refer to [“Managing Performance Warehouse systems” on page 59](#).

Click the Performance Warehouse to which you want to log on and click **Selected > Connect** to open the Connect to Performance Warehouse window.

Specify your DB2 user ID and your password and then click **OK**.

You are logged on to the specified Performance Warehouse. You can now work with the items listed on the displayed page of the Performance Warehouse window or move to another page to perform a different set of tasks.

Managing Performance Warehouse systems

You can add and delete Performance Warehouse systems.

Adding a Performance Warehouse system

Add a Performance Warehouse. You can add a Performance Warehouse system automatically or manually.

About this task

A Performance Warehouse is automatically added to the folders pane of the Performance Warehouse window the first time you perform an automatic logon to a Performance Warehouse. For information on how to do this, refer to [“Logging on to a Performance Warehouse automatically” on page 59](#). The name under which a Performance Warehouse appears in the Performance Warehouse window is the database alias, which associates a Performance Warehouse with a Db2 subsystem.

Procedure

1. In the Performance Warehouse window, click **Performance Warehouse > New Performance Warehouse**.

2. In **New Performance Warehouse**, specify the following values:

Database Alias name for cataloged Db2 PM database

Type the alias of the Performance Warehouse maintained by the OMEGAMON Collector that monitors this Db2 subsystem. Use the database alias that you specified when cataloging the DB2PM database of the OMEGAMON Collector. JDBC T2 connectivity will be used in this case.

Database URL in T4 connection style

Type the Database connection URL of the Performance Warehouse that is maintained by the OMEGAMON Collector that monitors this Db2 subsystem. This URL is equal to the database T4 connection URL string which points to the DB2PM database of the OMEGAMON Collector.

For example: **jdbc:db2://rs22.yourcompany.com:4567/RS22DC1A**

You must use the T4 URL in case you want to use specific JDBC connection attributes not supported with T2 connectivity (such as **clientApp1Compat**).

For further information about cataloging the DB2PM database, see [Configuration and Customization](#).

3. If you want to log on to the new Performance Warehouse after leaving this window, select the **Connect immediately after configuration** check box.
4. Click **OK**.

The new Performance Warehouse is added to the **z/OS** folder and remains part of this folder until you delete it again. If you requested to log on immediately after defining the new Performance Warehouse, the Connect to Performance Warehouse window is displayed.

The name under which Performance Warehouse appears in the Performance Warehouse window is the database alias that you specified in the **Database alias** field of the New Performance Warehouse window and which associates a Performance Warehouse with a Db2 subsystem.

Deleting a Performance Warehouse system

Delete a Performance Warehouse.

Procedure

1. Select the Performance Warehouse to be deleted.
2. If you are still logged on to this Performance Warehouse, click **Selected > Disconnect**.
3. Click **Selected > Delete**.

Defining and running processes

Processes consist of one or more steps that perform tasks, such as loading existing DB2 performance data into the Performance Warehouse database or generating reports.

This section describes processes and reports in more detail, uses scenarios to show how to define and run processes, and gives hints and tips on certain tasks.

A process can consist of one or more of the following steps or step combinations:

- The *CONVERT step* reads an existing save data set and produces a save-file data set that can be used as input for a LOAD step.
- The *LOAD step* loads an existing file or save-file data set into the Performance Warehouse database.
- The *REPORT step* generates Statistics and Accounting reports from existing DB2 performance data. You can use various commands and subcommands to produce different report types. This step can include:
 - Creation, conversion, and loading of save data sets into the Performance Warehouse database.
 - Creation and loading of file data sets into the Performance Warehouse database.
- The *CRD (collect report data) step* collects report data and stores it in MVS data sets.

Report types

To create different report types, you add various commands and subcommands to the REPORT step.

The GLOBAL command automatically appears as the first command in the REPORT step. You use this command to filter input data, set default values for subcommand options, and define global processing options. The following options are, for example, supported by the GLOBAL command:

- FROM
- TO
- INTERVAL
- BOUNDARY
- INCLUDE/EXCLUDE
- TIMEZONE

Table 3 on page 61 shows the command options for Accounting commands and subcommands. Table 4 on page 61 shows the command options for Statistics commands and subcommands. For more information about the individual command options, refer to *Report Command Reference*.

Restriction: After a process has run, all output data sets on the host are stored in the Performance Warehouse database and removed from the host. If you want to use the output data sets created by a REPORT step as input for a LOAD step or a CONVERT and LOAD step, you must include a LOAD FILE or LOAD SAVE subcommand as part of your REPORT step.

<i>Table 3. Supported Accounting command options.</i>				
Note: The Accounting subcommands FILE and RESTORE are not supported.				
Command options	ACCOUNTING REDUCE	ACCOUNTING REPORT	ACCOUNTING TRACE	ACCOUNTING SAVE
FROM	*	*	*	
TO	*	*	*	
INTERVAL	*			
BOUNDARY	*			
EXCEPTION/ NOEXCEPTION		*	*	
SCOPE		*		
ORDER		*		
LAYOUT		*	*	
INCLUDE/EXCLUDE	*	*	*	
TIMEZONE				
TOP		*	*	

<i>Table 4. Supported Statistics command options</i>					
Command options	STATISTICS REDUCE	STATISTICS REPORT	STATISTICS TRACE	STATISTICS SAVE	STATISTICS FILE
FROM	*	*	*		*

Table 4. Supported Statistics command options (continued)

Command options	STATISTICS REDUCE	STATISTICS REPORT	STATISTICS TRACE	STATISTICS SAVE	STATISTICS FILE
TO	*	*	*		*
INTERVAL	*				
BOUNDARY	*				
EXCEPTION/ NOEXCEPTION		*	*		*
SCOPE		*			
ORDER					
LAYOUT		*	*		
INCLUDE/ EXCLUDE	*	*	*		*
DSETSTAT		*	*	*	*
TIMEZONE					

Scenario 1: Transferring performance data into the Performance Warehouse database

The following scenario shows how you can make use of existing data sets that were not generated by a Performance Warehouse REPORT step. It describes how to transfer save or save-file data sets on an MVS system to the Performance Warehouse database.

You can perform this scenario on the following pages of the Performance Warehouse window:

- Expert
- Report and Load

Creating a process group

A process needs a process group as a container.

To create a process group:

1. In the folders pane, select **Process Groups** and then click **Selected > Create**.

The Process Group Properties window is displayed.

The **Author**, **Creation date and time**, and **Modification date and time** fields are filled in when you click **OK**. The **Author** field will contain your DB2 user ID.

2. In the **Name** field, type a unique name for the process group.

Example: Type Legacy.

3. In the **Description** field, you can type a description for the process group.

Example: Type Transfer "legacy" performance data sets into Performance Warehouse.

4. Click **OK** to save the process group and return to the Performance Warehouse window.

The new group appears in the folders pane as a subitem of **Process Groups**.

Creating a process using a process template

Process templates are supplied with the product. You can find these templates in the Public process group. Here, you use the template DB2PM.Templates.Load Warehouses 2 to create a process.

This template contains a CONVERT step and a LOAD STATISTICS step.

1. In the folders pane, under **Public**, select DB2PM.Templates.Load Warehouses 2 and then click **Selected > Copy**.

The Copy Process window is displayed.

2. In the **New process group** list, select the Legacy process group.
3. In the **New process name** field, type a unique name for the copied process.

Example: Transfer Statistics Save.

4. Click **OK**.

When you return to the Performance Warehouse window, the process appears under **Processes** of the Legacy process group. The template's steps are also copied.

Viewing the process properties

To view the process properties, select the Transfer Statistics Save process in the folders pane, and then click **Selected > Properties**.

The Process Properties window is displayed:

1. Ensure that the following is set:
 - **Status** is set to `in definition` so that you can make changes to the process.
 - **Scope** is set to `private`. When you copy a process template, the scope changes from `public` to `private` so that this process is not available to other users.
 - **Schedule of process** is set to `immediately`.
2. Change the description in the **Description** field to match the process.
3. Click **OK** to save the changes to the process and return to the Performance Warehouse window.

Configuring the process steps

The copied template contains default values for the process steps.

You must configure the steps to add:

- An input data set to the CONVERT step
- Target tables to the LOAD STATISTICS step

To configure the CONVERT step:

1. In the folders pane, click **Steps**. In the contents pane, select the CONVERT step and then click **Selected > Properties**.

The Convert Properties window is displayed.

2. Click the **Options** tab.
3. Change the name of the input data set to `my.input.data.set` and click **OK** to save the changes.

To configure the LOAD STATISTICS step:

1. In the folders pane, click **Steps**. In the contents pane, select the LOAD STATISTICS step and then click **Selected > Properties**.

The Load Statistics Step Properties window is displayed.

2. Click the **Options** tab. Specify the types of target tables that you want to load into the Performance Warehouse database.

3. Click **OK** to save the changes.

The process definition is now complete so you can run the process.

Running the process

To run the process immediately, select the process in the folders pane, and then click **Selected > Execute**. The process is scheduled to start immediately and is activated.

The following information is recorded about processes:

- The processes (with their start and stop timestamps and their status) that are currently running and have run since you started working with Performance Warehouse.
- The individual steps of each process, with their start and stop timestamps and their status.
- The output data sets of each step in a process.

While the process is running, you can view its progress. In the folders pane, click **Process Executions** to view the progress of the process. The Process Executions pane is displayed, listing the process executions for all the processes. To view only the Transfer Statistics Save process, you must define filter criteria.

To define filter criteria:

1. Click **View > Filter**.

The Filter Process Execution window is displayed.

2. Specify the filter conditions.

You can use wildcard characters in the filter criteria: `_` (underscore) for one character and `%` for more than one character.

Example:

- In the **Name** field, type `Transfer Sta%`
- In the **Group** field, type `Legacy`
- As **Comparison** operator next to the **Name** field, select `Like` because you are using a wildcard character in the name

3. Click **OK**.

The Process Executions pane then lists only the process-execution entries for the Transfer Statistics Save process.

Viewing the output data sets

When the process finishes (status FINISHED on the Process Execution pane), you can view its output.

To view the output:

1. On the Process Executions pane, select the process execution and then click **Selected > Details**.

The Process Execution Details window is displayed. It contains information about all the steps of the process and the output data sets generated by these steps.

2. To view the list of output data sets generated by the LOAD STATISTICS step, click the **Output Data Set** folder under **LOAD**.

3. To view the contents of a data set, double-click the data set.

The contents of the data set are shown in a Web browser. You can use your browser to save and print the data set. For information about how to specify your Web browser of choice, see [“Specifying global settings” on page 14](#).

To delete the data set from the Performance Warehouse database, select the corresponding data set, and click **Delete**.

Scenario 2: Defining Accounting and Statistics reports using an existing input data set

Suppose you are a database administrator for several OS/390® locations hosted on an MVS system. Users tell you that they have DB2 performance problems during online transaction processing on certain DB2 locations every day between 11 a.m. and 12 noon.

You decide to generate Accounting reports to determine if there is a DB2 problem. You also decide to generate a Statistics report to get an overview of system activity. You think that you might also need to analyze the data later using queries and rules of thumb, so you make sure that the performance data is transferred to the Performance Warehouse database.

To create the reports and transfer the data, you must:

1. Create a REPORT step.
2. Define the Accounting and Statistics reports.
3. Run the process to generate reports and transfer data to the Performance Warehouse database.
4. View and analyze the generated reports.

“Scenario 1: Transferring performance data into the Performance Warehouse database” on page 62 describes how to work with process groups, processes and steps, how to run the process, and how to view the results. This scenario describes how to use Performance Warehouse to define the reports the DBA needs to analyze the problem. It assumes that you have an existing input data set containing the Accounting and Statistics performance records for the timeframe and DB2 locations in question. If this is not the case, proceed with “Scenario 3: Collecting performance data for the generation of Accounting and Statistics reports” on page 68.

To define the reports you must:

1. Specify an input data set.
2. Filter the performance records to produce meaningful results.
3. Set up an Accounting report.
4. Transfer data to the Performance Warehouse database.
5. Set up a Statistics report.

You can perform these tasks on the following pages of the Performance Warehouse window:

- Expert.
- Report and Load.

Specifying an input data set

The input data set on the MVS system contains the Accounting and Statistics performance records for the timeframe and the DB2 locations in question.

To specify the input data set:

1. In the folders pane, click **Steps**. In the contents pane, select the REPORT step and then click **Selected > Properties**.

The Report Step Properties window is displayed.

2. Click the **Options** tab.
3. In the **Current data set specification** pane, click **Add**.

The Add Input Data Set window is displayed.

4. Type the name of the data set and click **OK**.

Specifying filter criteria for the performance records

You know that the performance problems occur during online processing for certain locations in a particular timeframe. You can specify filter criteria to account for these symptoms and apply them to all reports.

You use the GLOBAL command to specify these global performance-record filter criteria:

1. Click **GLOBAL** in the folders pane.

The GLOBAL command options are shown in the Current data set specification pane.

2. To restrict the records to the timeframe, type 11:00:00:00 in the **From Time** field and 12:00:00:00 in the **To Time** field.
3. To include performance records from online transaction-processing applications only, click the **Include** filter button.

The Include/Exclude Filter window is displayed.

- a. Select the row containing the DB2 PE identifier CONNTYPE (connection type) and click **Add**.

The Add Include/Exclude Value window is displayed.

- b. Associate the following values with CONNTYPE:

- CICS®
- IMS-MPP
- IMS-TBMP

The symbol in the **Usage** column shows whether the DB2 PE identifier is included or excluded.

- c. Click **OK** to include the filter in the GLOBAL command options.

The performance problems occur only at certain “slow” locations but you do not know exactly which locations are affected. However, you do know that the location names starting with BOE and DVI perform well. You use a filter with wildcard characters to ensure that these “good performers” are excluded from the report.

4. To exclude performance records from a particular location, click the **Exclude** filter button.

The Include/Exclude Filter window is displayed.

- a. Select the row containing the DB2 PE identifier LOCATION and click **Add**.

The Add Include/Exclude Value window is displayed.

- b. Associate the following values with LOCATION:

- BOE*
- DVI*

- c. Click **OK** to include the filter in the GLOBAL command options.

Setting up an Accounting report

To set up an Accounting report, right-click **REPORT** in the folders pane and click **Add**.

On the Add Command window:

1. Select ACCOUNTING and click **OK**.

The ACCOUNTING folder is added to the REPORT step.

2. Right-click the **ACCOUNTING** folder and click **Add**.

The Add Command window is displayed

3. Select REPORT and click **OK**.

The REPORT subcommand is added to the ACCOUNTING REPORT set and the current ACCOUNTING REPORT command options are shown in the Current ACCOUNTING REPORT command options pane.

4. To get as much detailed information in the report as possible, select **LONG** in the **Layout** field.

You need to produce two reports with different aggregation levels:

- Aggregated information for all online transaction application executions
- Aggregated information for each online transaction application execution

5. To filter the information needed for all online transaction application executions:

a. Click the **Order** filter button.

The Order Filter window is displayed. Set 1 contains the default order criteria **PRMAUTH** and **PLANNAME**.

b. Use the appropriate buttons to remove **PRMAUTH** and **PLANNAME** and add the DB2 PE identifier **INTERVAL**.

INTERVAL represents the information needed for all online transaction application executions.

6. To filter the information needed for each online transaction application execution:

a. Select Set 2 in the **Selected identifier for** field.

b. Add the DB2 PE identifier **CORRNAME** using the appropriate button.

CORRNAME represents the information needed for each online transaction application execution.

c. Click **OK** to add the filter criteria to the **ACCOUNTING REPORT** command options.

Transferring Accounting performance data to the Performance Warehouse database

To analyze Accounting performance data produced by the **REPORT** step at a later date, you must transfer the data to the Performance Warehouse database. You need to include a **LOAD** command in the **REPORT** step to do this.

To transfer the data:

1. Right-click the **ACCOUNTING** folder and click **Add**.

The Add Command window is displayed.

2. Select **SAVE** and click **OK**.

The **SAVE** folder is added to the Accounting report.

3. Right-click the **SAVE** folder and click **Add**.

The Add Command window is displayed.

4. Select **LOAD SAVE** and click **OK**.

The **LOAD SAVE** subcommand is added to the **ACCOUNTING SAVE** set and the current **LOAD** command options are shown in the Current DB2 **LOAD** command options pane.

5. Select the target tables to be loaded with the performance records and click **Add**.

Setting up a Statistics report

To set up a Statistics report, right-click **REPORT** in the folders pane and click **Add**.

On the Add Command window:

1.

2. Select **STATISTICS** and click **OK**.

The **STATISTICS** folder is added to the **REPORT** step.

3. Right-click the **STATISTICS** folder and click **Add**.

The Add Command window is displayed.

4. Select **REPORT** and click **OK**.

The REPORT subcommand is added to the STATISTICS REPORT set and the current STATISTICS REPORT command options are shown in the Current STATISTICS REPORT command options pane.

5. To get as much detailed information in the report as possible, select LONG in the **Layout** field and click **OK**.

The REPORT step is now complete and you can run the process. For information about how to run a process and view the output, see [“Running the process” on page 64](#) and [“Viewing the output data sets” on page 64](#).

Scenario 3: Collecting performance data for the generation of Accounting and Statistics reports

If an input data set does not exist, you must collect the relevant performance data before you can define your Accounting and Statistics reports.

In [“Scenario 2: Defining Accounting and Statistics reports using an existing input data set” on page 65](#), you used an input data set that contained the Accounting and Statistics performance records for the timeframe and DB2 locations in question.

The list of tasks necessary to define the reports would then look as follows:

1. Create a CRD (collect report data) step.
2. Filter the performance records to produce meaningful results.
3. Set up an Accounting report.
4. Transfer data to the Performance Warehouse database.
5. Set up a Statistics report.
6. Schedule the process containing the CRD and REPORT steps.

You can perform this scenario on the following pages of the Performance Warehouse window:

- Expert
- Report and Load

Creating a CRD (collect report data) step

To create a CRD step, click **Steps** in the folders pane and click **Selected > Create**.

On the Create Step window:

1. Select **CRD**.
The CRD Step Properties window is displayed.
2. Click the **Options** tab.
3. To configure the data collection, select **Data** under **Options categories**.
The Current data option pane is displayed.
4. Under **Data categories**, select **Accounting** and **Statistics**.
5. Specify which IFCIDs you want to collect.
6. To specify the criteria for the data collection to stop, select **Stop** under **Options categories**.
The Current stop conditions pane is displayed.
7. In scenario 2, the database administrator wants to collect data between 11 a.m. and 12 noon. To specify this, select the **Elapsed time** check box and type a duration of one hour.
8. Click **OK** to save the configuration of the CRD step.

Steps [“2” on page 68](#) to [“5” on page 68](#) are described in [“Scenario 2: Defining Accounting and Statistics reports using an existing input data set” on page 65](#).

Scheduling the process containing the CRD (collect report data) and REPORT steps

To start data collection at 11 a.m., you must schedule the process containing the CRD and REPORT steps. In the folders pane, select the appropriate process and then click **Selected > Properties**.

In the Process Properties window:

1. Click **Schedule**.

The Process Schedule Properties window is displayed.

2. Click **At schedule time** and then click the button next to the schedule string.

The Schedule Process window is displayed.

3. Set the schedule to 11 a.m.:

- Under **Minutes**, select **Specify minutes** and type 0.
- Under **Hours**, select **Specify hours** and type 11.

4. Click **OK** until you return to the Performance Warehouse window.

The Public process group contains the template Ad-Hoc Report, which consists of a CRD step and a REPORT step. This template collects data for five minutes and generates an Accounting and Statistics report in LONG format. You can use this template to get started.

Scenario 4: Collecting performance data regularly

You might want to collect performance data on a regular basis, for example, to collect Statistics data between 12 noon and 4 p.m. every day.

To perform this task:

1. Define a process consisting of a single CRD (collect report data) step.
2. Schedule the process to run periodically and activate it.

You can perform this scenario on the following pages of the Performance Warehouse window:

- Trace
- Report and Load
- Expert

It is recommended that you use the Trace page because it only contains the items you need to configure and activate, and for which you want to view the progress, such as the CRD process.

To define a CRD process:

1. Copy the process template DB2PM.Templates.Collect Report Data from the Public process group to a process group of your own. This template collects Accounting and Statistics data for the duration of 5 minutes.
2. Display the properties for the Collect Report Data step and change the following option categories on the Options page:
 - For **Output data set**, type an output data set name and select **Append**.
 - For **Data**, clear the **Accounting** check box.
 - For **Stop**, specify an elapsed time of 4 hours.

To schedule the process, proceed as described in [“Scheduling the process containing the CRD \(collect report data\) and REPORT steps”](#) on page 69, but type 12 under **Hours**.

To activate the process, display the process properties and change the status from **in definition** to **active**.

For more details on how to define, schedule, activate, and view the progress of a process, refer to [“Scenario 1: Transferring performance data into the Performance Warehouse database”](#) on page 62.

Hints and tips

The following sections give hints and tips on process-related Performance Warehouse tasks.

Understanding the schedule string

When you specify a schedule for a process, the specifications you make in the Schedule Process window are reflected as a string in the **Schedule of process** field of the Process Properties window and in the **At scheduled time** field of the Process Schedule Properties window. This string can look as follows:
2007/1/16/*/16/58

A string consists of the following parts, in this sequence: year/month/day/day-of-the-week/hour/minute. A 0 for the day of the week stands for Saturday. An asterisk (*) represents all allowed values for a field.

Important: When you define a schedule in the Schedule Process window, you can specify the dates and weekdays when you want to run a process. However, you cannot specify that a process should run at a certain date only if this date is a specific day of the week.

For example, if you type 1 under **Days** and 5 under **Months** and select the **Monday** check box under **Day of week**, the process will run on May 1st and on every Monday. In this example, the string would look similar to this one: */5/1/2/*/*

Deleting process executions of processes scheduled ‘immediately’

If you run a process that is scheduled ‘immediately’ and delete the process execution entry within five minutes after the process has finished, the OMEGAMON Collector runs the process again.

This is because the OMEGAMON Collector replaces ‘immediately’ with the time at which the process was started and interprets this time as a schedule when no process execution is available.

Reusing output data sets produced by a REPORT step

After a process has run, all output data sets on the host are stored in the Performance Warehouse database and removed from the host.

If you want to use the output data sets created by a REPORT step as input for a LOAD or a CONVERT and a LOAD, you must define a LOAD FILE or LOAD SAVE subcommand as part of your REPORT step.

Deleting data permanently

To permanently delete data sets and files from the Performance Warehouse database, you must delete them in the Process Execution Details window or delete the appropriate process execution entry from the Process Executions pane of the Performance Warehouse window.

For information about how to delete data from the tables in the Performance Warehouse database, refer to [Chapter 6, “Overview of the Performance Warehouse database,” on page 83.](#)

Using rules of thumb to analyze performance data

Rules of thumb apply a few simple rules and ratios to key performance indicators.

The sample rules of thumb delivered with the product have been assembled by DB2 experts in the field over many years and are valuable criteria by which to measure the performance of applications in a DB2 system. They can be executed as they are.

- For more information about the sample rules of thumb, see [Reporting User's Guide](#).
- The samples are described in detail in the Online Help:
 - Click **Help** in the following windows:
 - Definition page of the Rules-of-Thumb Cluster Properties window
 - Definition page of the Rule-of-Thumb Properties window

- Rules-of-Thumb Analysis window.
- Use the links under the related topic `Predefined rule-of-thumb cluster definitions for z/OS`.

This section describes how to define rules of thumb and how to use them to analyze performance data. You always start in the Performance Warehouse window.

In the System Overview window, click the **Performance Warehouse** toolbar button to start Performance Warehouse. The Performance Warehouse window is displayed.

Managing rule-of-thumb groups

A rule-of-thumb group is a container for rule-of-thumb clusters and rules of thumb.

The Public rule-of-thumb group contains the sample rules of thumb delivered with the product and those rules of thumb that other users have made generally available.

Creating a rule-of-thumb group

To create a rule-of-thumb group, select **Rule-of-Thumb Groups** in the folders pane of the Performance Warehouse window and then click **Selected > Create**.

The Rule-of-Thumb Group Properties window is displayed. The **Author**, **Creation date and time**, and **Modification date and time** fields are filled in when you click **OK**. The **Author** field will contain your DB2 user ID. Then:

1. Specify a unique name for the rule-of-thumb group in the **Name** field.
2. You can type a description for the rule-of-thumb group in the **Description** field. This is optional.
3. Click **OK** to save this rule-of-thumb group and return to the Performance Warehouse window.

The new group appears in the folders pane as a subitem of **Rule-of-Thumb Groups**.

Changing a rule-of-thumb group

You can change the description of the rule-of-thumb group, or rename or delete the rule-of-thumb group.

You can change a rule-of-thumb group in one of the following ways:

- Change the description of the rule-of-thumb group.

To change the description, click the rule-of-thumb group and click **Selected > Properties**.

- Rename a rule-of-thumb group.

To rename a group, click the rule-of-thumb group and click **Selected > Rename**.

- Delete a rule-of-thumb group.

To delete a group, click the rule-of-thumb group and click **Selected > Delete**.

Copying a rule-of-thumb group

If you want to create several rule-of-thumb groups that only differ in a rule-of-thumb cluster property or a rule of thumb, you might prefer to copy and change a rule-of-thumb group.

When you copy a rule-of-thumb group, any rule-of-thumb clusters and rules of thumb that it contains are copied with it.

To copy a rule-of-thumb group:

1. In the folders pane of the Performance Warehouse window, under **Rule-of-Thumb Groups**, click the rule-of-thumb group and then click **Selected > Copy**.

The Copy Rule-of-Thumb Group window is displayed.

2. In the **New rule-of-thumb group name** field, specify a unique name for the copied group.
3. Click **OK**.

The copied group appears in the folders pane.

Managing rule-of-thumb clusters and rules of thumb

A rule-of-thumb cluster represents a counter category that you want to analyze and for which you want to define rules of thumb. Examples of counter categories are Accounting and Statistics.

The first rule of thumb that you define for a cluster or copy or move into the cluster determines the cluster category. Any additional rules of thumb that you create for a cluster, or copy or move into the cluster must match this cluster category.

Creating a rule-of-thumb cluster

You can create a rule-of-thumb cluster using a template. Templates are displayed under the Public rule-of-thumb group. To create a cluster using a template, copy the template, and then change the copy.

To create a rule-of-thumb cluster without a template:

1. In the folders pane of the Performance Warehouse window, click **Rule-of-Thumb Clusters** and then click **Selected > Create**.

The Rule-of-Thumb Cluster Properties window is displayed.

The **Author**, **Creation date and time**, and **Modification date and time** fields are filled in when you click **OK**. The **Author** field will contain your DB2 user ID.

2. Specify a unique name for the rule-of-thumb group in the **Name** field.
3. In the **Scope** list, click one of the following:
 - private if you do not want to make this rule-of-thumb cluster available to other users.
 - public if you want other users to view, copy, and use this rule-of-thumb cluster as a template.
4. You can type a description for the rule-of-thumb cluster in the **Description** field.
5. Click **OK** to save the new rule-of-thumb cluster and return to the Performance Warehouse window.

The new rule-of-thumb cluster appears as a subitem of **Rule-of-Thumb Clusters** and contains the subitem **Rules of Thumb**.

To ensure name uniqueness throughout the rules-of-thumb-group folder, the rule-of-thumb cluster is displayed under the following name: *user.rotgroup.rotclustername*, where *user* stands for the DB2 user ID of the rule-of-thumb cluster creator, *rotgroup* for the rule-of-thumb group that this rule-of-thumb cluster belongs to, and *rotclustername* for the name of the rule-of-thumb cluster.

Creating a rule of thumb

In a rule of thumb, you specify the tables, report blocks, and columns that you want to analyze for a specific counter category.

You also define the analysis rules, set the thresholds, and provide tuning recommendations.

To create a rule of thumb:

1. Click **Rules of Thumb** for the cluster in the folders pane of the Performance Warehouse window and then click **Selected > Create**.

The General page of the Rule-of-Thumb Properties window is displayed.

The **Author**, **Creation date and time**, and **Modification date and time** fields are filled when you click **OK**. The **Author** field will contain your DB2 user ID.

2. Specify a unique name for the rule of thumb in the **Name** field.
3. If this is the first rule of thumb in the cluster, select the counter category for the rules of thumb that you want to add in the **Type** list.

You can define rules of thumb for Accounting and Statistics.

Important: The counter category you select here determines the category for the entire cluster. Therefore, this selection is not available for the second and subsequent rules of thumb in the cluster.

4. You can type a description for the rule of thumb in the **Description** field. This is optional.
5. Click the **Definition** tab. On the Definition page you can select the:
 - Tables
 - Report blocks
 - Columns that you want to analyze and for which you want to set rules, threshold values, and tuning recommendations.

You can do the following:

- Select a report block or a table. Your selection determines the content of the **Columns** box.
 - Fill the **Filter expression** and **Additional columns** boxes by clicking a column in the **Columns** box and then clicking the appropriate **Add** (▶) button. The entry appearing in these boxes consists of the table and column names.
 - If you add several columns to the **VALUE expression** box, you must use operators to combine the columns into expressions. Simple arithmetic expressions are allowed that contain the operators (,), +, -, /, *. The result of the arithmetic expression is run against the values that you specify under **WARNING and PROBLEM thresholds**.
 - If you want to include the values of columns that are not part of the arithmetic expression in the analysis result, add these columns to the **Additional columns** box.
 - The SQL query generated from the specifications for this and all other rules of thumb in the cluster and for the rules-of-thumb cluster itself must not exceed 32 KB.
6. Make your selections and fill in the various fields.
 7. Click **OK** to save the new rule of thumb and to return to the Performance Warehouse window.

The new rule of thumb appears in the contents pane of the Performance Warehouse window.

Example of a rule-of-thumb definition

This example shows how to define your rule of thumb on the Definition page of the Rule-of-Thumb Properties window.

Assume the following:

You want to analyze specific data of the Buffer Pool Read report block of the DB2PM_STAT_BUFFER table. If the result of dividing the value in the NO_PREFETCH_NO_BUF column by the value in the SEQ_PREFETCH_REQ column is greater than 0.01, you want to get a warning. If the result exceeds 0.05, you want to be informed that a problem exists. In both cases, you want to give the recommendation that the EDM pool must be increased. In addition, you want the analysis result to include the value of the SEQ_PREFETCH_REQ column.

Your selections and specifications would then look as follows:

1. In the **Report block** list, click BUFFER POOL READ.
2. In the **Columns** box, click NO_PREFETCH_NO_BUF and then click **Add** (▶) next to the **VALUE expression** box.

The selected column is added to the **VALUE expression** box, preceded by the table name.

3. In the **Columns** box, click SEQ_PREFETCH_REQ and then click **Add** (▶) next to the **VALUE expression** box.

The selected column is added to the **VALUE expression** box, preceded by the table name. Performance Expert leaves a space between this and the previously added column.

4. In the **VALUE expression** box, type the operator / between the two columns to specify the division.
5. In the **Columns** box, click SEQ_PREFETCH_REQ and then click **Add** (▶) next to the **Additional columns** box.

6. In the **VALUE** list (under **WARNING and PROBLEM thresholds**), click the operator **>** because you want to be notified when the result of the previously selected columns exceeds a certain value.
7. In the **WARNING threshold** field, type 0.01.
8. In the **PROBLEM threshold** field, type 0.05.
9. In both **Recommendation** fields, type Increase BP size as needed.

Defining filter criteria and additional columns for the rule-of-thumb cluster

You can specify additional columns and define filter criteria to reduce the analysis result. Your specifications apply to all rules of thumb in the cluster.

To define the additional columns and filter criteria:

1. Click the rule-of-thumb cluster, and then click **Selected > Properties**.

The Rule-of-Thumb Cluster Properties window is displayed.

2. Click the **Definition** tab to do the following:

- Select a report block or a table. Your selection determines the content of the **Columns** box.
- Fill the **Filter expression** and **Additional columns** boxes by clicking a column in the **Columns** box and then clicking the appropriate **Add (▶)** button. The entry appearing in these boxes consists of the table and column names.
- If you want the values of additional columns included in the analysis result of all rules of thumb, add these columns to the **Additional columns** box.
- If you add several columns to the **Filter expression** box, you must use operators to combine the columns into an SQL query.
- You can use variables in the filter expression that you can resolve during the rules-of-thumb analysis. See [“Resolving variables” on page 76](#) for more information.
- The SQL query generated from the specifications on this page and for all rules of thumb in the cluster must not exceed 32 KB.

3. Make your selections and fill in the various fields.

4. Click **OK** to save the changes.

Changing a rule-of-thumb cluster or a rule of thumb

You can change the properties of a rule-of-thumb cluster or a rule of thumb, or rename or delete a rule-of-thumb cluster or a rule of thumb.

You can change a rule-of-thumb cluster and its rules of thumb in one of the following ways:

- Change the properties of a rule-of-thumb cluster or a rule of thumb.

To change the properties, click the cluster or the rule of thumb and click **Selected > Properties**.

- Rename a rule-of-thumb cluster or a rule of thumb.

To rename a cluster or a rule of thumb, click the cluster or the rule of thumb and click **Selected > Rename**.

- Delete a rule-of-thumb cluster or a rule of thumb.

To delete a cluster or a rule of thumb, click the cluster or the rule of thumb and click **Selected > Delete**.

Copying or moving a rule-of-thumb cluster to another rule-of-thumb group

The tasks of copying and moving rule-of-thumb clusters are similar. You move a rule-of-thumb cluster to another rule-of-thumb group by copying it to the other rule-of-thumb group and then deleting it from the original rule-of-thumb group.

You can copy a rule-of-thumb cluster within the same rule-of-thumb group.

To copy a rule-of-thumb cluster:

1. In the folders pane of the Performance Warehouse window, under **Rule-of-Thumb Clusters**, click the rule-of-thumb cluster and then click **Selected > Copy**.

The Copy Rule-of-Thumb Cluster window is displayed.

2. In the **Rule-of-thumb group** list, click an existing rule-of-thumb group to which you want to copy the rule-of-thumb cluster. You can select the same rule-of-thumb group that already contains the rule-of-thumb cluster.
3. In the **New rule-of-thumb cluster name** field, specify a name for the copy that is unique throughout the destination process group.
4. Click **OK** to return to the Performance Warehouse window.

The copied rule-of-thumb cluster appears under **Rule-of-Thumb Clusters** of the selected rule-of-thumb group.

If you copied a public rule-of-thumb cluster, its scope has changed from `public` to `private`.

Copying or moving a rule of thumb to another rule-of-thumb cluster

The tasks of copying and moving rules of thumb are similar. You move a rule of thumb to another rule-of-thumb cluster by copying it to the other rule-of-thumb cluster and then deleting it from the original rule-of-thumb cluster. You can copy a rule of thumb within the same rule-of-thumb cluster.

To copy a rule of thumb, the following must apply:

- The destination rule-of-thumb group must contain at least one rule-of-thumb cluster.
- The rule of thumb that you want copy must be the first rule of thumb in the rule-of-thumb cluster or fit the cluster category. The cluster category is determined by the counter category of the first rule of thumb that you copy or move into, or create for, this cluster.

To copy a rule of thumb:

1. In the folders pane of the Performance Warehouse window, click **Rules of Thumb**. In the contents pane, click the rule of thumb and then click **Selected > Copy**.

The Copy Rule of Thumb window is displayed.

2. In the **Rule-of-thumb group** list, click a rule-of-thumb group to which you want to copy the rule of thumb. You can select the same rule-of-thumb group that already contains this rule of thumb.
3. In the **Rule-of-thumb cluster** field, click a rule-of-thumb cluster to which you want to copy the template. You can select the same rule-of-thumb cluster that already contains this rule of thumb.
4. In the **New rule-of-thumb name** field, specify a name for the copy that is unique throughout the destination rule-of-thumb cluster.
5. Click **OK** to return to the Performance Warehouse window.

The copied rule of thumb appears under the selected rule-of-thumb group and rule-of-thumb cluster.

Analyzing the performance data

You can use rules of thumb to analyze data in the Performance Warehouse database.

For information about how to load data, see [“Scenario 1: Transferring performance data into the Performance Warehouse database” on page 62](#). You can analyze the data using all rules of thumb in a cluster or using individual rules of thumb. If you analyze an entire rule-of-thumb cluster, the analysis result will show a column for each of the rules of thumb in the cluster.

Several users can analyze the same data at the same time. To avoid a locking conflict, make sure that you are not analyzing the data while another user is loading data into the Performance Warehouse database.

1. To analyze the data, do one of the following:
 - If you want to use a single rule of thumb:
 - a. In the folders pane of the Performance Warehouse window, select the appropriate rule-of-thumb cluster and then click its subitem **Rules of Thumb**.

The contents pane lists all rules of thumb that you defined for this rule-of-thumb cluster.

b. Select the rule of thumb you want and then click **Selected > Analyze**.

- If you want to use all rules of thumb defined for a rule-of-thumb cluster:

In the folders pane of the Performance Warehouse window, select the appropriate rule-of-thumb cluster and then click **Selected > Analyze**.

The Options page of the Rules-of-Thumb Analysis window is displayed, which shows the filter expression specified for the rule-of-thumb cluster.

2. Specify filter criteria. For more information refer to [“Specifying filter criteria” on page 76](#).
3. Resolve variables. For more information refer to [“Resolving variables” on page 76](#).
4. Sort the analysis result by the result of the arithmetic expression defined for a specific rule of thumb. For more information refer to [“Sorting the analysis result” on page 77](#).
5. Click **Analyze** to run the analysis.

The Result matrix page of the Rules-of-Thumb Analysis window is displayed. For more information about this and the other pages of this window, see [“Viewing and understanding the analysis result” on page 77](#).

Specifying filter criteria

The LOADLOG table holds report and trace interval data for loaded data. You can use the trace start and end time, and the ID of the job step as criteria for filtering rules of thumb. This can be useful if you want to analyze only those trace records loaded by a specific user for a particular step at a certain time.

To specify filter criteria:

1. On the Options page of the Rules-of-Thumb Analysis window, click **Time** in the folders pane.

The **Time setting** area is displayed in the lower-right corner of the contents pane.

This area shows information from the LOADLOG table. It contains the user ID of the user that loaded the data, the trace interval data, and the load-job step ID associated with the load job. For more information about this table, refer to [“The LOADLOG table” on page 84](#).

2. Select a row that holds the appropriate trace-start time, trace-end time, or step-log ID and click **Start time**, **End time**, or **LOAD step**.

If more than one table is referenced, the Select Table window is displayed. Select a table from the list and click **OK**.

The corresponding filter expression is displayed in the text field. You can combine expressions using the AND, OR, and NOT operators.

For example, if you clicked **Start time**, the filter expression comprises the table name, the trace-start time column name, and the value of the trace-start time for the selected row in the time-setting table.

If you clicked **LOAD step**, the filter expression comprises the table name, the step-log ID column name, and the value of the step-log ID for the selected row in the time-setting table.

3. Click **Apply**. The filter expression is inserted in the **Variables context** field directly after the WHERE clause.
4. To remove the filter expression from the WHERE clause, click **Reset**.

Resolving variables

Unresolved variables in the filter expression are prefixed by a colon.

To resolve these variables:

- On the Options page of the Rules-of-Thumb Analysis window, click **Variables** in the folders pane.

The **Variables setting** area is displayed in the contents pane.

This area shows a list of variables found in the filter expression with the last-saved values for these variables. Values are saved each time you analyze the trace data or execute a query using the same variables. You can enter a value for each of the variables shown. These values are used when you analyze the trace data.

- Click **Reset** to change the values of the variables back to the saved values.

Sorting the analysis result

On the Result matrix and the Column details pages of the Rules-of-Thumb Analysis window, the analysis result is sorted by timestamp or as shown in the **Variables context** field on the Options page. You can change the order by selecting the result of the arithmetic expression specified for a rule of thumb as sort criterion. In this case, the Column details page shows the results of the arithmetic expression instead of the attention values.

To sort the analysis result:

1. On the Options page of the Rules-of-Thumb Analysis window, click **Order by expression value** in the folders pane.

If you selected a single rule of thumb, the **Select rule of thumb** list contains this rule of thumb. If you selected a rule-of-thumb cluster, all rules of thumb in the rules-of-thumb cluster are listed.

2. Select the rule of thumb for the arithmetic expression that you want to use to sort the analysis result.

You specified the arithmetic expression in the **VALUE expression** box on the Definition page of the Rule-of-Thumb Properties window for a rule of thumb. The sort direction (ascending or descending) depends on whether you set lower or upper limits for the **WARNING and PROBLEM thresholds**.

Viewing and understanding the analysis result

The first page that you see after the analysis finishes is the Result matrix page.

The Result matrix page shows:

- A list of the entire trace data that was created between the trace-start and trace-end times that you specified on the Time pane of the Options page. The data is sorted as follows:
 - By timestamp
 - As shown in the **Variables context** field on the Options page
 - By the result of the arithmetic expression of the rule of thumb selected on the Order by value expression pane on the Options page
- Additional columns you specified for the rule-of-thumb cluster (cluster definition columns).
- A column for each rule of thumb that you used for the analysis.
- The attention value (OK, warning, problem) of the analysis for each timestamp and rule of thumb. The result depends on your specifications on the Definition page of the Rule-of-Thumb Properties window.

A dash (–) means that a value is not available.

To reduce the list displayed on the Result matrix page, you can display only rows with at least one warning or a higher attention level or only rows with at least one problem:

- Click **View > Warning** to view all rows that contain at least one warning or at least one problem.
- Click **View > Problem** to view only rows that contain at least one problem.

Your selection remains valid until you click **View > All** again.

If there are a lot of entries in the list, they are not displayed all at one time. You can move to the next or previous set of entries by clicking **View > Next** or **View > Previous**. If you previously reduced the list, for example, to rows that show only problems, this selection remains valid when you move to the next or previous set of entries.

Viewing the results for a specific timestamp

In the Result matrix page, click a row in the **Timestamp** column to view the result of all rules of thumb for a specific timestamp. The Row details page of the Rules-of-Thumb Analysis window is displayed.

This page lists all the rules of thumb, the attention values, and the cluster-definition column values for a specific date and time.

The **Cluster definition columns** field shows the values for the columns that you added as additional columns to the rules-of-thumb cluster that contains the analyzed rules of thumb (see the Definition page of the Rules-of-Thumb Cluster Properties window).

The **Rule of thumb details** shows the result of the first rule of thumb in the list for which the **Attention** column contains a value.

To view the result of another rule of thumb, click it in the **ROT Name** column. For more information about the individual fields, see [Explanation of the individual fields in the Rules of thumb details](#).

Viewing the results for all timestamps

In the Result matrix page, click a column other than the **Timestamp** column or the **Cluster definition columns** to view the result of a rule of thumb for all timestamps. The Column details page of the Rules-of-Thumb Analysis window is displayed.

This page lists the timestamps, the attention values, and the cluster-definition column values for the selected rule of thumb. If the analysis result is sorted by the result of the arithmetic expression of a specific rule of thumb, the results of the arithmetic expression are shown instead of the attention values and you are informed of the sort.

The **Cluster definition columns** field shows the values for the columns that you added as additional columns to the rules-of-thumb cluster that contains the selected rule of thumb (see the Definition page of the Rules-of-Thumb Cluster Properties window).

The **Rules of thumb details** shows the result of the first timestamp in the list for which the **Attention** column contains a value. To view the result of another timestamp, click it in the **Timestamp** column.

Explanation of the individual fields under **Rules of thumb details** on the Row Details and Column Details pages:

ROT description

The description that you specified for the rule of thumb on the General page of the Rule-of-Thumb Properties window.

VALUE expression

The arithmetic expression that you specified in the **VALUE expression** box on the Definition page of the Rule-of-Thumb Properties window.

VALUE

The result of the arithmetic expression and its relation to the appropriate threshold value, as defined on the Definition page of the Rule-of-Thumb Properties window.

Recommendation

The recommendation that you specified on the Definition page of the Rule-of-Thumb Properties window in the case of a problem or a warning.

VALUE expression columns

The individual table columns that you selected for the **VALUE expression** box on the Definition page of the Rule-of-Thumb Properties window and the value produced by the analysis.

Additional columns

The table columns that you added as additional columns to the selected rule of thumb (see the Definition page of the Rule-of-Thumb Properties window) and the values produced by the analysis.

Defining and running queries

You can use the query facility to define performance-management queries to run on the Statistics and Accounting tables in the Performance Warehouse database.

For example, Accounting tables contain performance data that shows how many resources are consumed by the applications. Appropriate queries on the ACCOUNTING table can help you achieve performance improvements. A set of queries is provided. They can be run as they are but you can also define and save your own queries.

You always start on the Performance Warehouse window. In the System Overview window, click the **Performance Warehouse** toolbar button. to start Performance Warehouse. The Performance Warehouse window is displayed.

Managing query groups

A query group is a container for queries. The Public query group contains the sample queries delivered with the product and those queries that other users have made generally available.

In the sample queries, the query strings and value expressions are not formatted. However, you can copy the queries and rules, and insert blanks or new lines to get a more readable format.

Creating a query group

Before you define a query, you must create a query group to contain this query.

To create a query group:

1. In the folders pane of the Performance Warehouse window, select **Query Groups** and then click **Selected > Create**.

The Query Group Properties window is displayed.

The **Author**, **Creation date and time**, and **Modification date and time** fields are filled in when you click **OK**. The **Author** field will contain your DB2 user ID.

2. In the **Name** field, specify a unique name for the query group.
3. You can type a description for the query group in the **Description** field. This is optional.
4. Click **OK** to save this query group and return to the Performance Warehouse window.

The new group appears in the folders pane as a subitem of **Query Groups**.

Changing a query group

You can change the description of the query group, or rename or delete a query group.

You can change a query group in one of the following ways:

- Change the description of the query group.

To change the description, click the query group and click **Selected > Properties**.

- Rename a query group.

To rename a group, click the query group and click **Selected > Rename**.

- Delete a query group.

To delete a group, click the query group and click **Selected > Delete**.

Copying a query group

If you want to create several query groups that only differ in a query property, you might prefer to copy and change a query group. When you copy a query group, any queries that it contains are copied with it.

To copy a query group:

1. In the folders pane of the Performance Warehouse window, under **Query Groups**, click the query group and then click **Selected > Copy**.

The Copy Query Group window is displayed.

2. In the **New query group name** field, specify a name for the copy that is unique within the query-group folder.
3. Click **OK**.

When you return to the Performance Warehouse window, the copy appears in the folders pane.

Managing queries

After you define query groups you can create and change queries.

Creating a query

To create a query, select **Queries** in the folders pane of the Performance Warehouse window and then click **Selected > Create**.

The General page of the Query Properties window is displayed. The **Author**, **Creation date and time**, and **Modification date and time** fields are filled in when you click **OK**. The **Author** field will contain your DB2 user ID. Then:

1. Specify a unique name for the query in the **Name** field.
2. You can type a description for the query in the **Description** field. This is optional.
3. Click the **Definition** tab. Here you can select or specify the SQL string that makes up the query.
4. Select a report type.
5. Select a report block or a table. Your selection determines the content of **Columns**.
6. Select a row in **Columns** that contains the column or table name you want to include in your query and then select this name to include it in the query. For more information about the **Columns** contents, refer to [“The Performance Warehouse data tables” on page 83](#).
7. Click **Add column** or **Add table** to add the column or table name to the query.
You can use SQL query operators to combine the names.
8. Click **OK** to save the query and return to the Performance Warehouse window.

Changing a query

You can change the properties of a query, or rename or delete a query.

You can change a query in one of the following ways:

- Change the properties of a query, such as the scope, description, or query definition.

To change the properties, click the query and click **Selected > Properties**.

- Rename a query.

To rename a query, click the query and click **Selected > Rename**.

- Delete a query.

To delete a query, click the query and click **Selected > Delete**.

Copying or moving a query to another query group

The tasks of copying and moving queries are similar. You move a query to another query group by copying it to the other query group and then deleting it from the original query group.

You can copy a query within the same query group.

To copy a query:

1. In the folders pane of the Performance Warehouse window, click **Queries**. In the contents pane, click the query you want to copy and then click **Selected > Copy**.

The Copy Query window is displayed.

2. In the **Query group** list, click an existing query group to which you want to copy the query. You can select the same query group that already contains this query.
3. In the **New query name** field, specify a unique name for the copied query.
4. Click **OK** to return to the Performance Warehouse window.

The copied query appears under the selected query group.

Running a query

You can run queries on the data in the Performance Warehouse database.

For information about how to load data, see [“Scenario 1: Transferring performance data into the Performance Warehouse database”](#) on page 62. Several users can run their queries on the same data. In this case, however, to avoid a locking conflict, make sure that you are not running the query while another user is loading data into the Performance Warehouse database.

1. In the folders pane of the Performance Warehouse window, select the appropriate query group and then click its subitem **Queries**.

The contents pane lists all the queries in the query group.

2. Click the query you want to run and then click **Selected > Execute**.

The View SQL page of the Query Execution window is displayed. It shows the query definition. From here you can edit the query, specify filter criteria, and resolve any unresolved variables in the query expression. Note that variables must begin with a colon followed by at least one alphanumeric character, such as `:c`. It can also consist of alphanumeric and underscore characters, for example, `:connect_type6`.

3. Specify filter criteria. For more information refer to [“Specifying filter criteria”](#) on page 81.
4. Resolve variables. For more information refer to [“Resolving variables”](#) on page 82.
5. Click **Execute** to start the query.

The View Result page of the Query Execution window is displayed. For more information about this window, see [“Viewing the query results”](#) on page 82.

Specifying filter criteria

The LOADLOG table holds report and trace interval data for loaded data. You can use the trace start and end time, and the ID of the job step as criteria for filtering queries. This can be useful if you want to analyze only those trace records loaded by a specific user for a particular step at a certain time.

To specify filter criteria:

1. On the View SQL page of the Query Execution window, click Time.

The Time setting area is displayed in the contents pane.

This area shows information from the LOADLOG table. It contains the user ID of the user that loaded the data, the trace-interval data, and the load-job step ID associated with the load job. For more information about this table, refer to [“The LOADLOG table”](#) on page 84.

2. Select a column that holds the appropriate trace-start time, trace-end time, or step-log ID and click Start time, End time, or LOAD step.

The Select Table window is displayed.

3. Select a table from the list and click OK.

The corresponding filter expression is displayed in the text field. You can combine expressions using the AND, OR, and NOT operators.

For example, if you clicked Start time, the filter expression comprises the table name, the trace start time column name, and the value of the trace start time of the selected row in the time-setting table. If you clicked LOAD step, the filter expression comprises the table name, the step-log ID column name, and the value of the step-log ID of the selected row in the time-setting table.

4. Click Apply. The filter expression is inserted in the query at the current cursor position.

Resolving variables

Unresolved variables in the query expression are prefixed by a colon.

To resolve these variables:

1. On the View SQL page of the Query Execution window, click **Variables**.

The **Variables setting** area is displayed in the contents pane.

This area shows a list of variables found in the query expression with the last-saved values for these variables. Values are saved each time you analyze the trace data or execute a query using the same variables. You can enter a value for each of the variables shown. These values are used when you run the query.

2. Click **Reset** to change the values of the variables back to the saved values.

Viewing the query results

After the query has run, the View SQL page of the Query Execution window is displayed.

The query expression is shown at the top of the page. The table shows the result of the query. The column headings are defined by the query. The number of rows read from the result set is shown at the bottom of the table. To sort the contents of the table, click a column heading.

- To save the results in a format that can be used by standard spreadsheet programs, click Save.

The result matrix, the query properties, and the query expression are saved.

- To view the query results in a Web browser, click Browse.

You can use the Web browser's standard functions to print the query results.

Chapter 6. Overview of the Performance Warehouse database

To analyze your performance data, you can run SQL queries against the tables in the Performance Warehouse database. The sample queries delivered with Performance Warehouse help you understand the contents of, and relationships between, the tables and how data can be joined to evaluate performance data and identify performance problems.

However, to use your queries efficiently, you must understand the data model designed for the performance-relevant tables. In addition, you must be familiar with the LOADLOG table, which you need to access and manage the data in the Performance Warehouse database, for example, if you want to delete or archive it.

The Performance Warehouse database consists of:

- Performance Warehouse data tables that store the DB2 performance data collected from DB2.
- Control tables that establish the infrastructure of Performance Warehouse.

The Performance Warehouse data tables

The Performance Warehouse database contains the following Performance Warehouse data tables for table space STAT.

Table name	Table description
DB2PM_STAT_BUFFER	One row for each buffer pool that is active at the start of: <ul style="list-style-type: none">• The corresponding delta record for file data.• The interval record for save-file data.
DB2PM_STAT_DATASET	One row for each open data set that has an I/O event rate of at least one event per second during the reporting interval.
DB2PM_STAT_DDF	For each delta record (file data) or interval record (save-file data), one row per remote location participating in distributed activity using the system-directed access method and one row for all remote locations using the application-directed access method.
DB2PM_STAT_GBUFFER	One row per group buffer pool that is active at the start of the corresponding delta record for file data or of the interval record for save-file data.
DB2PM_STAT_GENERAL	For file data: One row for each Statistics delta record, containing data from IFCID 1 and 2. A delta record is a set of counters describing the DB2 activity between two consecutive DB2 Statistics records pairs. For save-file data: One row for each Statistics interval record, containing data from IFCID 1 and 2. A Statistics interval record is a set of counters describing the DB2 activity within the interval specified by the user.

The Performance Warehouse database contains the following Performance Warehouse data tables for table space ACCS:

Table name	Table description
DB2PMSACCT_BUFFER	One row per buffer pool used.
DB2PMSACCT_DDF	One row per remote location participating in distributed activity.
DB2PMSACCT_GBUFFER	One row per group buffer pool used.
DB2PMSACCT_GENERAL	One row per DB2 Accounting record, which is written when a thread terminates, is reused, or becomes inactive.
DB2PMSACCT_PROGRAM	One row per package and DBRM executed.
DB2PMSACCT_RLF	One row per resource limit type encountered.

For more information about these tables, refer to the part on the Performance Warehouse database in *Reporting User's Guide*.

The Performance Warehouse database described in the *Reporting User's Guide* contains more tables than the one described here. However, it must be created and managed manually using the sample CREATE and LOAD parts shipped with OMEGAMON for Db2 Performance Expert. The Performance Warehouse database described here is created automatically, managed by the OMEGAMON Collector, and is accessible using the Performance Warehouse client.

The Accounting and Statistics tables are associated in a 1:n relationship as described in *Reporting User's Guide*.

The general table is the parent table for all other tables. Each row in the general table is associated with zero, one, or several rows in the other tables. For Accounting SAVE tables, use the columns LOCAL_LOCATION to CLIENT_TRANSACTION for joining related data. For Statistics tables, use the columns LOCAL_LOCATION, GROUP_NAME, SUBSYSTEM_ID, MEMBER_NAME, INTERVAL_TSTAMP, and BEGIN_REC_TSTAMP for joining related data.

The Accounting and Statistics tables in the Performance Warehouse database contain the following additional columns: STEPLOGID of type INTEGER and ROW_ID of type ROWID. These columns are populated by the OMEGAMON Collector and DB2 when performance data is loaded into the Performance Warehouse data tables. The value of STEPLOGID references the LOADLOG table and can be used to correlate the performance data and the LOAD step that loaded the data. DB2 generates the ROW_ID values automatically when rows are added to the Performance Warehouse data tables. They are used in the rules-of-thumb analysis to correlate general and detailed analysis results.

The LOADLOG table

The LOADLOG table contains the information you need to access and manage the performance data.

Column name	Data type	Key	Description
LL_ID	INTEGER	P	Unique LOADLOG ID
LL_INSERTEDBY	VARCHAR(8)		User ID of the user who inserted the data
LL_STARTTS	TIMESTAMP		Date and time of the first event recorded
LL_STOPTS	TIMESTAMP		Date and time of the last event recorded
LL_LOADTYPE	CHAR(10)		Load type:ACCOUNTING, STATISTICS

Each row in the table is associated with a LOAD step executed by the OMEGAMON Collector and has a unique LL_ID. This LL_ID is referenced by an Accounting or a Statistics table with its foreign key STEPLOGID, which is defined with delete rule ON DELETE SET NULL. If you delete a row from the LOADLOG table, the column STEPLOGID of the dependent rows is set to NULL. To delete data from the DB2PM_STAT and DB2PMSACCT tables:

1. Determine the entries using the LOADLOG table and apply a search condition on the STEPLOGID column in the DELETE statements.
2. Delete the corresponding LOADLOG entries.

The LOADLOG entries related to the LOAD steps are displayed under **Time setting** on the View SQL page of the Query Execution window and on the Options page of the Rules-of-Thumb Analysis window. They help you restrict your queries and rules of thumb to the time intervals during which data was loaded.

Chapter 7. Accessibility features in Performance Expert Client

This section describes the major accessibility features in Performance Expert Client.

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, use software products successfully.

You can:

- Operate all features using the keyboard instead of the mouse.
- Access the individual help files.
- Change the system settings to high contrast for all user interface controls and all client area contents.

This section explains how to use these accessibility features.

Operating all features by using the keyboard

You can use keys or key combinations to perform operations that can also be done through mouse actions. All menu items can be accessed from the keyboard. In those cases, the keyboard equivalent appears to the right of the menu item, or the shortcut letter is underlined. Some keyboard items also have shortcuts.

To navigate through a window or dialog by using the keyboard instead of the mouse, use the following keyboard shortcuts:

Navigating through a window or dialog	Shortcut
Access the window menu	Alt, then down arrow (consecutively)
Access the menu bar in a window	Alt+underlined letter or F10
Activate a menu item in the menu bar	Enter
Access controls in a dialog	Alt+underlined letter
Navigate through the menu bar	Right arrow, left arrow, down arrow, up arrow
Move to the next set of controls	Tab or F6
Move to the previous set of controls	Shift+Tab
Move within tables	Tab or right arrow, Shift+Tab or left arrow, down arrow, up arrow
Move within folders panes	Up arrow, down arrow
Expand a folder	Right arrow
Collapse a folder	Left arrow
Move within list boxes	Up arrow, down arrow
Move within combo boxes	Up arrow, down arrow
Move to a split bar in a window	F8
Resize a window by using the split bar	Arrow keys

To perform main tasks in windows by using the keyboard instead of the mouse, use the following keyboard shortcuts:

Table 9. Keyboard shortcuts in a window

Task	Shortcut
Display Help	F1
Close the active window	Ctrl+W or Alt+F4
Exit the application	Ctrl+Q
Cut	Ctrl+X
Copy	Ctrl+C
Paste	Ctrl+V
Print	Ctrl+P
Refresh displayed data	F5
Enter or leave history mode	Ctrl+H
Move back in history mode	Alt+left arrow
Move forward in history mode	Alt+right arrow
Add a new DB2 system in System Overview	Ctrl+N
Add a new configuration in Trace Configurations	Ctrl+N
Cancel thread in Thread Summary and Details	Del
Delete Trace Configuration in Trace Configurations	Del
Open Trace Configuration	Ctrl+Alt+N
Open Trace Activation	Ctrl+Alt+A
Open Exception Processing	Ctrl+Alt+E
Open System Overview	Ctrl+Alt+O
Open Buffer Pool Analysis	Ctrl+Alt+B
Open DB2 command for the active DB2 subsystem	Ctrl+Alt+D
Open Statistics Details for the active DB2 subsystem	Ctrl+Alt+S
Open System Health for the active DB2 subsystem	Ctrl+Alt+H
Open Threads in Lock Conflicts for the active DB2 subsystem	Ctrl+Alt+L
Open Locking Conflicts for the active DB2 subsystem	Ctrl+Alt+C
Open System Parameters for the active DB2 subsystem	Ctrl+Alt+Y
Open Performance Warehouse for the active DB2 subsystem	Ctrl+Alt+P
Open Thread Summary for the active DB2 subsystem	Ctrl+Alt+T

To perform main tasks in dialogs by using the keyboard instead of the mouse, use the following keyboard shortcuts:

Table 10. Keyboard shortcuts in a dialog

Task	Shortcut
Confirm a dialog	Enter

<i>Table 10. Keyboard shortcuts in a dialog (continued)</i>	
Task	Shortcut
Cancel a dialog	Esc
Activate a button that has the focus	Spacebar or Enter
Select and deselect check boxes and radio buttons	Spacebar
Navigate within check box and radio button groups	Tab
Open combination box menu	Alt+down arrows
Close combination box menu	Esc
Move up and down in combination box menu	Up arrow, down arrow
Activate combination box menu item	Spacebar or Enter
Move within a list box	Up arrow, down arrow
Activate a list box entry	Enter
Move to a tab in a window containing tabbed pages	Ctrl+Up
Move between the pages of a window containing tabs if the tab has the focus	Right arrow, left arrow

Changing the system settings for colors, fonts, and accessibility options

If you are not satisfied with the colors, fonts, and accessibility options supplied by OMEGAMON for Db2 Performance Expert you can change them within your operating system or within Performance Expert Client

You must restart OMEGAMON for Db2 Performance Expert before the new settings take effect.

To change the settings within Performance Expert Client:

1. In the System Overview window, click **Monitor > Configuration**.

The Configuration window is displayed.

2. Click the **Accessibility** tab.

3. To specify the appearance scheme for high contrast and font size, select one of the following options in the **High contrast scheme** list:

- To use the default settings, select **[None]**.
- To specify a white background containing black text in normal font size, select **Black on white**.
- To specify a white background containing black text in large font size, select **Black on white (large font)**.
- To specify a black background containing white text in normal font size, select **White on black**.
- To specify a black background containing white text in large font size, select **White on black (large font)**.

Product legal notices

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

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

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

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated

through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

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

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

Programming interface information

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

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

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

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

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

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

Trademarks

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

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

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

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

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

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

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

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

Terms and conditions for product documentation

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

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

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

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

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

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

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

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

Privacy policy considerations

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

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

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

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

Index

A

- access path of SQL statements [51](#)
- accessibility
 - overview [4](#)
- ACCOUNTING REDUCE subcommand [61](#)
- ACCOUNTING REPORT command, filter options [67](#)
- ACCOUNTING REPORT subcommand [61](#)
- Accounting report, setting up [66](#)
- ACCOUNTING SAVE subcommand [61](#)
- ACCOUNTING TRACE subcommand [61](#)
- accumulated values, definition [25](#)

B

- buffer pool analysis, introduction [6](#)
- buffer pool, refreshing Statistics [29](#)

C

- chart, changing type [33](#)
- command
 - understanding syntax diagrams [1](#)
- conventions [2](#)
- CONVERT step
 - changing [63](#)
 - definition [60](#)
 - input data set, specifying [63](#)
- cookie policy [91](#), [93](#)
- CRD (collect report data) step
 - creating [68](#)
 - definition [60](#)
 - scheduling [69](#)

D

- data group
 - configuration, reusing [40](#)
 - creating [30](#)
- data set, buffer pool Statistics refreshing [29](#)
- data sharing group
 - lock conflicts [24](#)
 - Statistics, viewing [50](#)
 - system health, viewing [50](#)
 - threads, monitoring [50](#)
- data view
 - changing [33](#)
 - chart color and pattern, changing [39](#)
 - configuration, reusing [40](#)
 - creating [30](#), [33](#)
 - definition [30](#)
 - displaying as performance graph in the System Overview window [40](#)
 - global threshold
 - changing [32](#)
 - specifying [38](#)

- data view (*continued*)
 - grid, adding [39](#)
 - percentages [33](#)
 - single threshold, specifying [38](#)
 - snapshot
 - changing the number of visible snapshots [39](#)
 - freezing their display [39](#)
- DB2 commands, issuing [16](#)
- DB2 Connect gateway
 - changing its properties [13](#)
 - defining [12](#)
 - logging on [14](#)
 - monitoring [40](#)
- Db2 subsystem
 - defining [10](#)
- DB2 subsystem
 - DB2 commands, issuing [16](#)
 - exception processing [44](#)
 - logging off [18](#)
 - logging on [14](#)
 - SQL activity tracing [23](#)
 - starting several monitoring functions [15](#)
 - system parameters, viewing [43](#)
 - threads, viewing all active [19](#)
- DB2 system parameters, viewing [43](#)
- DCS applications, monitoring connections of [40](#)
- defining the DB2 systems to be monitored [10](#)
- delta processing [26](#)
- delta values, definition [25](#)
- DGOMUPXT user exit [49](#)
- diagnostics
 - e mail settings, specifying [14](#)
- disability [87](#)
- dynamic SQL statement cache [28](#)

E

- e mail settings, specifying
 - for sending diagnostics reports [14](#)
- event exception processing [45](#)
- exception processing
 - definition [44](#)
 - event exception processing
 - event exception log [46](#)
 - events, list [45](#)
 - starting [45](#)
 - stopping [46](#)
 - exception notification [45](#)
 - interval processing [26](#)
 - periodic exception processing
 - periodic exception log [48](#)
 - starting [48](#)
 - stopping [48](#)
 - types, definition [44](#)
 - user exit, activating [49](#)
- exporting snapshot data [53](#)

F

FILE subcommand [61](#)
filtering threads [21](#)
font size, how to change [89](#)

G

gauge values
 Statistics [25](#)
GLOBAL command
 command options [61](#)
 exclude filter [66](#)
 filter criteria, specifying [66](#)
 include filter [66](#)
 timestamp option [66](#)

H

help
 how to get help [17](#)
high contrast, how to change [89](#)
high-water mark
 Statistics [25](#)
history data
 delta processing [26](#)
 enabling collection [51](#)
 interval processing [26](#)
 viewing [51](#)

I

interval processing [26](#)

K

keyboard shortcuts [87](#)

L

legal notices
 cookie policy [91, 93](#)
 notices [91](#)
 programming interface information [91, 92](#)
 trademarks [91–93](#)
LOAD command, specifying tables [67](#)
LOAD step
 changing [63](#)
 definition [60](#)
 output data set, viewing [64](#)
 target tables, specifying [63](#)
LOADLOG table [84](#)
locking conflict
 history data, viewing [51](#)
 threads associated with resource [24](#)
 threads, viewing all [24](#)
long performance counter values [54](#)

M

message DGOK3638, reactivating [14](#)

N

notices [91, 92](#)

O

OMEGAMON Collector
 history data, enabling collection [51](#)
order filter, specifying for ACCOUNTING REPORT command [67](#)
overview [1](#)

P

password, saving [14](#)
performance counter values, long [54](#)
Performance Expert Agent for DB2 Connect Monitoring [40](#)
Performance Expert Client
 components and features [5](#)
 leaving [18](#)
 starting [8](#)
Performance Warehouse
 adding systems [59, 60](#)
 architecture overview [57](#)
 deleting systems [59](#)
 introduction [5](#)
 logging on [58](#)
 OMEGAMON Collector option for [57](#)
Performance Warehouse data tables [83](#)
Performance Warehouse database
 description [83](#)
 loaded tables, specifying [63](#)
 loading performance data [62](#)
 performance data, transferring data to [67](#)
periodic exception processing
 starting [48](#)
 stopping [48](#)
process
 Accounting report, setting up [66](#)
 creating [63](#)
 definition [57](#)
 executing [64](#)
 output data set, viewing [64](#)
 properties, viewing [63](#)
 running
 filter criteria, specifying [64](#)
 limiting processes shown [64](#)
 progress, viewing [64](#)
 scheduling [69](#)
 Statistics report, setting up [67](#)
 step types [60](#)
 template, using [63](#)
 understanding the schedule string [70](#)
process executions
 limiting processes shown [64](#)
 viewing [64](#)
process group, creating [62](#)
programming interface information [91, 92](#)

Q

query
 copying [80](#)

- query (*continued*)
 - creating [80](#)
 - deleting [80](#)
 - filter criteria, specifying [81](#)
 - job-step ID filter, specifying [81](#)
 - properties, changing [80](#)
 - renaming [80](#)
 - results
 - printing [82](#)
 - spreadsheet, saving [82](#)
 - viewing [82](#)
 - running [81](#)
 - scope, changing [80](#)
 - timestamp filter, specifying [81](#)
 - variables, resolving [82](#)
- query group
 - changing description [79](#)
 - copying [79](#)
 - creating [79](#)
 - deleting [79](#)
 - renaming [79](#)
- Query Workload Tuner [51](#)

R

- REDUCE subcommand [61](#)
- regular processing [26](#)
- report
 - Accounting, setting up [66](#)
 - commands, allowed [61](#)
 - Statistics, setting up [67](#)
 - subcommands, allowed [61](#)
- REPORT step
 - definition [60](#)
 - filter criteria, specifying [66](#)
 - input data set, specifying [65](#)
 - loading data [67](#)
- REPORT subcommand [61](#)
- rule of thumb
 - analysis result
 - sorting [77](#)
 - changing [74](#)
 - copying to another rule-of-thumb cluster [75](#)
 - creating [72](#)
 - defining [73](#)
 - deleting [74](#)
 - filter criteria, specifying [76](#)
 - job-step ID filter, specifying [76](#)
 - moving to another rule-of-thumb cluster [75](#)
 - renaming [74](#)
 - timestamp filter, specifying [76](#)
 - variables, resolving [76](#)
- rule-of-thumb cluster
 - analysis result
 - sorting [77](#)
 - changing [74](#)
 - copying to another rule-of-thumb group [74](#)
 - creating [72](#)
 - definition [72](#)
 - filter criteria, specifying [74](#)
 - moving to another rule-of-thumb group [74](#)
- rule-of-thumb group
 - changing [71](#)
 - changing description [71](#)

- rule-of-thumb group (*continued*)
 - copying [71](#)
 - creating [71](#)
 - deleting [71](#), [74](#)
 - renaming [71](#), [74](#)

S

- SAVE subcommand [61](#)
- schedule string [70](#)
- screen readers and magnifiers [4](#)
- service information [1](#)
- shortcuts [87](#)
- snapshots
 - changing the number of visible snapshots [39](#)
 - definition [25](#)
 - exporting [53](#)
 - history data, viewing [51](#)
 - long performance counter values [54](#)
 - printing [53](#)
 - saving [53](#)
 - time zones [55](#)
- SQL activity traces
 - creating [23](#)
 - viewing [23](#)
- SQL statements
 - determining the access path [51](#)
 - viewing [28](#)
- starting several monitoring functions [15](#)
- Statistics
 - buffer pool, refreshing [29](#)
 - data sharing group [50](#)
 - delta processing [26](#)
 - gauge values, definition [25](#)
 - graphical format, viewing [29](#)
 - high-water mark values, definition [25](#)
 - history data, viewing [51](#)
 - interval processing [26](#)
 - processing modes [26](#)
 - regular processing [26](#)
 - snapshot values, definition [25](#)
 - viewing [25](#)
- STATISTICS FILE subcommand [61](#)
- STATISTICS REDUCE subcommand [61](#)
- STATISTICS REPORT subcommand [61](#)
- Statistics report, setting up [67](#)
- STATISTICS SAVE subcommand [61](#)
- STATISTICS TRACE subcommand [61](#)
- support information [1](#)
- syntax diagrams [1](#)
- system health
 - chart color and pattern, changing [39](#)
 - chart, changing type [33](#)
 - configuration, reusing [40](#)
 - current, viewing [31](#)
 - data group, creating [30](#)
 - data sharing group [50](#)
 - data view
 - changing [33](#)
 - creating [30](#), [33](#)
 - displaying data views as performance graphs in the System Overview window [40](#)
 - global threshold

- system health (*continued*)
 - global threshold (*continued*)
 - changing [32](#)
 - specifying [38](#)
 - grid, adding [39](#)
 - past, viewing [32](#)
 - single threshold, specifying [38](#)
 - snapshot
 - changing the number of visible snapshots [39](#)
 - freezing their display [39](#)
 - viewing [29](#)
- system parameters
 - history data, viewing [51](#)
 - viewing [43](#)
- system settings, how to change [89](#)

T

- terminology used [3](#)
- threads
 - active, viewing all [19](#)
 - canceling [22](#)
 - children threads [20](#)
 - data sharing group [50](#)
 - details, selected [19](#)
 - filtering [21](#)
 - history data, viewing [51](#)
 - locking conflict, viewing [24](#)
 - parallel threads [20](#)
 - qualifying [21](#)
 - resource, locking conflict [24](#)
 - windows, customizing columns [54](#)
- threshold
 - global threshold
 - changing [32](#)
 - specifying [38](#)
 - single threshold, specifying [38](#)
- time zones [55](#)
- timestamp data
 - query filtering [81](#)
 - rule-of-thumb filtering [76](#)
- TRACE subcommand [61](#)
- trademarks [91–93](#)

U

- user exit, exception processing [49](#)

W

- water marks
 - Statistics [25](#)
- Web browser settings, specifying [14](#)



Product Number: 5655-W37

SH12-7069

