

Precise™ for DB2

User Guide

Version 10.0.x

IDERA

Precise™ for DB2 User Guide

Copyright © 2019 Precise Software Solutions, Inc. All rights reserved.

Document release version 1.0

Precise™, Precise Software™, the Precise™ Logo, Precise i³™, Precise Indepth™, Precise Insight™, Precise Savvy™, SmarTune™, Performance Warehouse™, Application Service Dashboard™, Precise for Storage Tiering™, Precise for Storage Tiering Plus Apps™, Precise for Database & Storage™, Precise for Applications™, Precise for Storage™, Precise Insight Inquire™, Performance Management Database™, and PMDB™ are trademarks or registered trademarks of Precise Software Solutions, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, decompilation and/or reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Precise Software Solutions, Inc. and its licensors, if any.

Certain third-party software may be distributed, embedded, or bundled with this product or recommended for use in connection with its installation and use. Such third-party software is separately licensed by its copyright holder. The list that includes the names of the copyright and license agreements can be found in the Release Notes document.

THE DOCUMENTATION IS PROVIDED "AS-IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PRECISE SOFTWARE SOLUTIONS, INC. SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Global Headquarters

Brookhollow Central III
2950 North Loop Freeway West, Suite 700
Houston, Texas 77092
Phone: +1-877-693-1886
Fax: +1-650-898-1666

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Precise product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Customer Support

For technical assistance, visit our customer portal at <http://www.idera.com/support/productsupport> where you can find an extensive knowledge base, product updates and our online community forums. You can also contact our Customer Support Team using our customer portal, or go to <http://www.idera.com/support/productsupport> for a list of our support access numbers in your country.

Contents

- Introducing Precise for DB2 10
 - About Precise for DB2 10
 - Improved performance management..... 10
 - Integration with other Precise products 10
- Components of Precise for DB2 10
 - Precise for DB2 agents 10
 - Precise Listener and Precise for DB2 FocalPoint 11
 - PMDB loaders 11
 - Optional Collector extensions 11
- How Precise for DB2 works 12
- The Precise for DB2 Time Model 12
 - In DB2 states 12
 - Overall Activity Graphs 12
- What you can do with Precise for DB2..... 13
- Where to get more information..... 13
- Precise for DB2 basics 14
 - How most tabs are structured 14
 - About the Precise bar..... 15
 - About the Main area 15
 - About the Association area 16

About drilling down in context.....	16
How a drilldown affects the Tab heading.....	16
About configuring your settings.....	17
About configuring general settings	17
About configuring display settings	17
About configuring Tree view settings.....	17
About configuring Time Frame settings	17
Tasks common to most tabs	17
Switching to a different tab.....	18
Selecting a time frame	18
Selecting a predefined time frame from the Precise for DB2 toolbar.....	18
Selecting a time frame using the calendar icon.....	18
Selecting an entity.....	19
Filtering data	19
Associating entities with data that meets specific criteria	19
Focusing on information in overtime graphs	20
Determining which table columns to display	20
Copying data to the clipboard.....	21
Exporting to Precise Custom Portal.....	21
Adding, viewing, and deleting Favorites	21
Opening your Precise product from StartPoint.....	22
Getting an overview of your application	23
About the Dashboard tab.....	23
How the Dashboard tab is structured	24
About the general tab structure	25
About the Main area in the Dashboard tab	25
About the Details area in the Dashboard tab.....	26

How the Dashboard can help you identify performance problems.....	28
Identifying the most resource-consuming database partition	28
Analyzing top resource consuming entities	28
Analyzing bufferpools	29
Identifying the top resource consumers in various entities	29
Getting current application information	30
About the Current tab	30
How the Current tab is structured	32
About the Main area in the Current tab.....	32
About the Association area in the Current tab.....	32
About examining entities in the Current tab	32
About the Database Partition	32
About getting an overview of current Database activities	32
About viewing Database locks	32
About the Application entity.....	33
About the Application Holding Locks entity	34
About getting an overview of current applications holding locks.....	34
About viewing information on the lock breakdown of an application that is holding locks.....	35
About viewing the text of an application that is holding locks.....	36
About the Statement entity	36
About getting an overview of current Statement entities.....	36
About viewing current Statement text	37
How the Current tag can help you identify performance problems.....	37
Examining resource consumption of an entire database partition	37
Observing the activity of a single application.....	37
Examining a single connection	38
Analyzing currently held locks.....	38

Examining performance over time	39
About the Activity tab	39
How the Activity tab is structured.....	40
About the Tree pane in the Activity tab.....	41
About the entities you can examine	42
About All DB2 Databases	43
About Database Partitions	44
About Databases.....	46
Selected Database	47
About Programs	48
About Statements.....	50
About Objects	52
About Machines.....	52
About Modules	54
About Actions.....	56
About OS Users	58
About DB Users.....	60
About Tablespaces.....	62
About Bufferpools.....	64
Examining the relationship between database objects.....	67
About the Objects tab	67
How the Objects tab is structured.....	67
About the Main area in the Objects tab.....	67
About the Association area in the Objects tab.....	68
About Objects tab entities	68
Schemata	68
Tables.....	68

Indexes.....	69
Views.....	70
Database Partition Groups.....	71
Tablespaces.....	71
Bufferpools	72
Database Partitions.....	72
Analyzing Statement plans	73
About the SQL tab	73
How the SQL tab is structured	74
About viewing the execution plan of a statement	74
About the Execution Plan tree.....	75
About interpreting text highlighting on the Highlights tab	76
About viewing which objects are referenced by the execution plan	76
About viewing information on all operators and predicates in the Execution plan.....	78
About viewing the statement text	78
About viewing recommended indexes for a statement.....	78
About viewing additional execution plan details	79
About tuning actions	79
Creating a new statement.....	79
Opening an existing statement	80
Editing an existing statement.....	80
Re-explaining an access plan.....	81
Generating recommendations.....	81
Simulating the effect of changes	83
About the What-If tab	83

How the What-If tab is structured	83
About the Main area in the What-If tab.....	83
About the Summary area in the What-If tab.....	83
About the Association area in the What-If tab.....	83
Adding or editing an index change for evaluation.....	84
How the What-If tab can help you evaluate changes	85
Adding an index	85
Examining DB2 database statistics	86
About the Statistics tab	86
How the Statistics tab is structured	86
The Main area	86
About the entities you can examine	87
Overview	87
Prefetchers.....	88
Cleaners	88
Logging.....	89
Package Cache	89
Sorts.....	89
Locks	90
How the Statistics tab can help you identify performance problems.....	90
Examining Prefetchers	90
Examining Sorts.....	90
Examining Locks	91
Examining Logging	91

Introducing Precise for DB2

This section includes the following topics:

- [About Precise for DB2](#)
- [Components of Precise for DB2](#)
- [How Precise for DB2 works](#)
- [What you can do with Precise for DB2](#)
- [Where to get more information](#)

About Precise for DB2

Addressing the needs of DBAs, system administrators, managers, quality assurance specialists and planners, Precise for DB2 (Precise for DB2) provides a complete view of database performance by capturing, measuring, and correlating performance metrics from all critical system components.

Through its wide range of data presentations, Precise for DB2 provides the tools necessary for proactive monitoring, diagnosing, detecting, and correcting of performance problems before they become visible in your DB2 databases and related applications.

Current operational data is displayed to enable easy identification of problematic conditions. Recent operational data, pertaining to designated time periods, enables operations staff to detect problematic conditions that occurred in the past and that are likely to reoccur if not corrected. Performance data is archived and a fully detailed display of the historical data is provided for use in performance, QA, and capacity planning analysis.

Improved performance management

Precise for DB2 can help IT staff discover the answer to questions, such as, why an application ran so slowly during a particular time interval, or where an application is spending most of its time. The first step is to establish a baseline level of database resource consumption by understanding the behavior of the average application.

Since Precise for DB2 is now part of the Precise infrastructure, you can drilldown to other installed databases like J2EE, Oracle, or Microsoft SQL server. The behavior of the average application represents the average percentage of the total time that all applications spend consuming or waiting for each resource. These aggregate figures provide a broad picture of the operation of your database. They enable you to identify major bottlenecks in your system, for example, CPU and I/O limitations, DB2 database waits due to locks, remote queries, and so on.

Precise for DB2 also enables you to focus on and drill down into individual applications, SQL statements and database objects in your system that are responsible for heavy consumption of system and database resources and to quickly find the definitive root cause of performance degradation.

Integration with other Precise products

Precise for DB2 has been integrated with Insight, Alerts, and Report Manager.

Insight monitors the response times of systems and breaks these times down into the various tiers and components. Insight also provides tier-specific metrics to help understand the activities in each tier.

For more information, see the *Insight User's Guide* or *Online Help*.

Alerts provides a notification system that will alert you to application response time problems and enables you to launch Precise for DB2 in context to view more detailed data.

For more information, see the *Alerts User's Guide* or *Online Help*.

The Precise Custom Portal is a lightweight, configurable portal application that provides access to various types of data from different applications. The Precise Custom Portal features a Web-based view. It is highly customizable and extensible, allowing you to build a dashboard for each user or each function within your organization. For example, you can build a dashboard for all the information that a system administrator or director of IT operations would need to constantly follow up on. This component is automatically installed as part of the Precise framework installation.

For more information, see the *Precise Custom Portal User's Guide*.

Report Manager gathers and organizes historical information that enables the IT team to focus on infrastructure hot spots. Report Manager is a reporting tool that resides on a Web server. You can use Report Manager to compare period-to-period performance against a baseline and identify response time problems "at-a-glance", before they affect your bottom-line. It queries the PMDB and displays the results in reports that are generated on a scheduled basis or on demand.

For more information, see the *Report Manager User's Guide* or *Online Help*.

The PMDB is the repository of the Precise product suite which stores application performance data for long-term analysis and trending reports.

For more information, see the *Precise Administration Guide*.

Components of Precise for DB2

Precise for DB2 employs a client/server architecture to provide high performance database monitoring and tuning.

Precise for DB2 consists of the following components:

- One or more Precise for DB2 Collectors installed on one or more servers in your application
- Precise for DB2 FocalPoint
- Tuner
- Listener
- PMDB
- Collector extension, such as the storage collector or ERP Collector

See the relevant product's manual.

Precise for DB2 agents

The Precise for DB2 Collectors are installed on servers and communicate using TCP/IP. The Collectors collect and analyze DB2 database performance data, which is then stored in flat historical files and later in the PMDB.

The Collectors monitor and record all DB2 database activity on a 24 x 7 basis. The Collectors work using a sampling technique and thereby consume less resources than a trace agent.

The Collectors collect resource consumption data on the following Precise for DB2 data entities: databases, applications, programs, statements, batches, login names, DB users, and machines.

The following agents are available:

Application Performance agent	Precise for DB2 provides a complete view of application performance management by capturing, measuring and correlating performance metrics from all critical system components.
-------------------------------	---

SAP™	If ERP support is installed, the Collector also includes SAP. This agent provides data on the activities and resource consumption of packaged application entities. This data can help you identify ERP users, transactions, reports, and so on, that are heavy users of DB2 resources. The current version only supports SAP.
Storage agents	EMC Storage. See “Optional Collector extensions” on page 11.

Precise Listener and Precise for DB2 FocalPoint

Each Precise product has a FocalPoint that communicates with the listeners installed on the monitored servers. The Precise Listener communicates with other Precise components on the same server. For example, the PMDB FocalPoint requests to load data into the PMDB are sent to the Precise Listener. Then, the Precise Listener executes the PMDB loader to aggregate and prepare performance data to be loaded. In a similar way, InformPoints and Insight agents send requests to the Precise Listener.

PMDB loaders

PMDB Loaders periodically load data—collected by the Collectors and stored in the historical files—into the PMDB. When the PMDB Loaders run, the flat historical files—written after the last time data was loaded into the PMDB—are identified and opened. The summarized data is also loaded into the PMDB. Data on the last incomplete time slice is only loaded into the PMDB when the next scheduled load process runs.

Optional Collector extensions

Optional Collector extensions are available. They provide additional information in certain applications, namely for storage and for Enterprise Resource Planning (ERP). Precise for Storage correlates DB2 database activity with the I/O on the back-end storage. ERP extensions correlate DB2 database activity with front-end ERP users and transactions. Precise for Storage is available for the following technologies:

- EMC Symmetrix

If storage system support is installed, the Collector provides detailed information regarding the I/O resource consumption across the storage devices. By investigating this information, you can relate I/O activity to specific components of the storage device and pinpoint the cause of the I/O performance degradation.

With Precise for Storage it is possible to browse device-specific mappings and statistics information on EMC Symmetrix arrays.

Note: Support for Veritas File System (VxFS) and Veritas Volume Manager (VxVM) is available if you have installed Precise for Storage and Veritas Mapping Service (VxMS) on Sun Solaris. For more information, see the *Precise Installation Guide* or contact Customer Support or your SE.

ERP extensions are available for the following technology:

- SAP

If ERP support is installed, the Collector provides detailed information on the activities and resource consumption of “packaged application” entities. It enables you to identify ERP users, transactions, reports, and so on, that are heavy users of DB2 resources and to drill down to the statements executed by them and the objects they accessed.

See *Precise Installation Guide*

See *Precise Administration Guide*

How Precise for DB2 works

Precise for DB2 stores a complete history of your database activity in the PMDB. This provides you with a performance tracking tool that gives you the ability to assess database performance trends and analyze changes in database activities over long periods of time. This data is vital for identifying patterns in resource consumption and predicting future resource consumption in your DB2 database.

The Precise for DB2 Collector runs on your server using a unique, accurate and comprehensive sampling technology that requires low overhead. It collects performance data from an instance and then correlates and loads the data into the PMDB tables. This data is then analyzed and summarized by the FocalPoint. You can later access this data using the Precise for DB2 Web browser interface.

The Precise for DB2 Time Model

An application may wait for many events in DB2. Since the number of events is very large and the meaning of some events is unclear, Precise for DB2 groups various DB2 events into event groups or states and defines the status of sessions currently connected to DB2 by their states. Each DB2 application state is either *in DB2* or *not in DB2*.

In DB2 states

The In DB2 resource consumption graphs give a breakdown of how DB2 is spending its time executing statements and utilities. The wait states with the longest duration indicate the bottlenecks within DB2. Reducing the duration of these wait states improves the response time or alternatively increases the throughput.

The following table lists the existing In DB2 states.

Table 1-1 In DB2 states

In DB2 state	Description
Using CPU	Time spent by the application using CPU.
CPU Wait	Time spent by the application waiting for CPU.
Sync I/O Wait	Time spent by the application waiting for a synchronous I/O request to/from the bufferpool.
Direct I/O Wait	Time spent by the application waiting for a direct I/O request of LOB or long column to/from the tablespace file bypassing the bufferpool.
Prefetch Wait	Time spent by the application waiting for a prefetcher to read into the bufferpool the next group of pages.
Lock Wait	Time spent by the application waiting for the release of a lock.
I/O Error Wait	Time spent by the application waiting for all active applications using a tablespace to release it in case of I/O error. The application that encounters an I/O error in a tablespace must wait to acquiesce the entire tablespace.
Remote Query Wait	Time spent by the application waiting for pending results from a remote federated server.
Utility Wait	Time spent by the application while waiting during running of a utility (command) such as: backup, restore, load, unload, restart, acquiesce and create a db.

Overall Activity Graphs

The Overall Activity resource consumption graphs display information on total time spent in DB2 and a breakdown of activities outside of DB2. Overall Activity can help answer the question: "Is the application spending its time in DB2 or elsewhere?"

The following table lists the Overall Activity states.

Table 1-2 Overall Activity states

Overall Activity state	Description
In DB2 (Summed)	Time spent by the selected item or items in the selected category executing an SQL statement.
UOW Waiting	Time spent by the application waiting for a request or input. For example, a Unit Of Work wait takes place when (1) a user starts a Command Line Processor session, connect to the database and then do nothing, (2) a client application executes an SQL statement, displays its result set on a window and waits for next GUI request while still connected, (3) export application writes fetched rows to an output file.
Decoupled from DB agent	Time spent by the application when a database agent (process) is not assigned to the connected application. This happens when using connection concentrator (MAX_CONNECTIONS > MAXCAGENTS). It is impossible to tell whether the application is idle or the application is waiting for an available database agent.
Connect Pending	Time spent by the application waiting for DB2 to complete connect request.
Disconnect Pending	Time spent by the application waiting for DB2 to complete disconnect request.
Wait for Remote Request	Time spent by the application waiting for a remote DB2 server to complete SQL execution.
Wait for Commit Phase II	Time spent by the application waiting for the coordinator to request DB2 to start commit phase II. This state is also known as an in-doubt transaction.

What you can do with Precise for DB2

Use Precise for DB2 to detect, diagnose and resolve performance problems in your DB2 database. Precise for DB2 lets you identify real-time performance problems (for example, locking problems) and monitor a complete history of your DB2 database.

You can use Precise for DB2 to:

- See your whole application status at a glance
- Support very large installations – hundreds of databases and partitions
- Resolve performance problems by identifying their root cause and improving statement performance
- Proactively monitor, analyze, and tune Precise for DB2 data partitions
- Analyze and tune locking problems
- Determine whether an I/O problem is caused by a Precise for DB2 application
- Tune a database object rather than a specific statement that accesses the database object
- Perform I/O tuning of storage devices
- Identify the most resource-consuming SQL statement and view a detailed explanation of the execution plan of the SQL statement
- Recommend which index would provide a better execution plan and reduce resource consumption
- Monitor your system's current state as well as previous states

Where to get more information

More information on Precise, its products, technical notes, and so on, can be found in the Release Notes document for this version.

Precise for DB2 basics

This section includes the following topics:

- [How most tabs are structured](#)
- [About drilling down in context](#)
- [About configuring your settings](#)
- [Tasks common to most tabs](#)
- [Opening your Precise product from StartPoint](#)

How most tabs are structured

Though each tab is structured differently, most tabs consist of two different areas. Each area can include different control elements, such as tabs and view controls, and displays information in various formats, such as tables, graphs, or charts. The various areas are related to each other in that performing an action in one area affects the information displayed in other areas on the page. For example, the Current tab contains an upper and a lower area. The lower area (Association area) shows information in a table format. Each row in the table represents an application. The upper area (Main area) displays general information on the selected entity as well as related applications.

The entities displayed in the Association area are associated with the selected entity or related application displayed in the Main area. At times, the relationship between the entity displayed in the Main area and those displayed in the Association area is that of parent to child, and sometimes it merely represents that there is a relationship between the selected entity and the entities displayed in the Association area (such as when a statement is the selected entity displayed in the Main area and the list of programs associated with the selected statement is displayed in the Association area).

It is possible to click a row and drill down to an additional level below the selected level or use the filter feature to determine which information is displayed in the table. When you perform an action on an element in the lower area (in this case the Association area), the information displayed in the upper area (in this case the Main area) changes to reflect the action you requested.

You can control the tab by doing any of the following:

- Keeping the selected entity but changing the associated entities—by selecting a different entity from the Association controls in the Association area.
- Changing the selected entity—by selecting a row in the Association area.
- Changing the view for the selected entity—by selecting a different view in the Main area.

For example, when you want to view information on a specific application in the Current tab, choose Active Applications from the Association controls. The Association area changes to display applications-related information. The Tab heading and the Main area remain unchanged.











In the Association area, click the row of the application you want to view detailed information for. The Tab heading indicates the newly selected entity; the Main area displays information on the application you drilled down to, and the Association area shows information on statements executed by this application.

About the Precise bar

The Precise bar enables you to keep track of where you have been and provides various controls.

The following table describes the function of each of the toolbar buttons

Table 2-1 Precise bar functions

Icon	Name	Description
	Back	During a work session, Precise products keep track of where you have navigated to. The Back button enables you to navigate between previously visited views. The Back control displays your previous view.
	Forward	Enables you to navigate to the next view. This button is only enabled if you clicked Back or if you chose a history option.
	AdminPoint	Opens Precise AdminPoint.
	Home	Navigates to the highest level entity, usually the instance or Tier (all instances). The time frame settings remain the same.
	Stop	Stops a request for information from the server.
	Refresh	Updates the data currently displayed.
	Favorites	Enables you to add or remove favorites in your Favorites list.
	Send	Opens a new email message in your email program with the link to the current application in context.
	Settings	Opens the General Settings and Time Frame Settings dialog boxes.
	Help	Opens the online help in context.

About the Main area

The Main area displays general information on the selected entity. The structure of this area depends on the selected entity and tab.

Information can be displayed as follows:

- As two over-time graphs that are displayed side-by-side and describe two data series.
- As a pie chart and a table that are displayed side-by-side. The table displays more details regarding a specific slice in the pie chart.

The View Control drop-down menu allows you to choose different view options. You can access this menu by clicking the arrow icon.

In some tabs, you can filter the contents of the view by clicking the Filter button.

See [“Filtering data”](#) on page 19.

About the Association area

The Association area lists entities that have similar attributes and that are related to the selected entity named in the Tab heading. Besides the Tab bar, the Association area is the main navigation tool. It enables you to drill down from one entity to related entities (usually of a different type).

The entities displayed in this area depend on the selected tab and the selected entity. The same entity may be displayed in several tabs, with different data displayed for each one (for example, a DB2 table will be displayed in the Objects tab with data from the DB2 data dictionary, and in the Activity tab with data sampled by the Collector).





The Association control, accessed by the arrow icon, allows you to associate with different entities. The last option, More..., opens a dialog box that lets you carry out additional operations, such as to view additional associations, change the sort order, or control the number of returned rows.

The information displayed in the Association area is arranged in a tabular format. Each row in the table represents an entity, and each column describes an attribute of the entity, such as text or a graph. It is possible to change the sort order of the entities in the table by clicking a column title.

The Association tabs located above the table let you display different sets of attributes for the entities displayed in the table. The Association tabs are entity-dependent. For some entities, no tabs are displayed.

A numeric column can optionally display a graph, and at times, it can display a stacked bar graph as well. Switching to a graph view in these columns offers some level of breakdown to the information displayed in the column.

The following icons are used to indicate which format information is displayed.

-  Displays the information in numerical format.
-  Displays the information in a bar graph.
-  Displays the information as a stacked bar graph.
-  Displays the information in percentage format

The Association area lets you drill down to another entity by clicking a row. The information displayed in both the Main area and the Association area will change to reflect your selection.

See [“Filtering data”](#) on page 19.

See [“Associating entities with data that meets specific criteria”](#) on page 19.

About drilling down in context

The term 'in-context' means that you can display additional information on a selected item by drilling down to another tab or view. The filter settings you defined (for example, the selected time frame you chose) and the entity you selected are carried over to the other view or tab, to allow you to continue analyzing your subject from a different perspective. This concept takes on slightly different meanings depending upon where you are attempting to drill down in context from.

For example, when viewing information on a database partition in the Dashboard tab, you can click a link in the Details area (right pane) to view additional information on the related tab, in context of your original selection.

Or, when viewing a list of alerts for your product in Alerts, you can open your product in context and continue investigating the factors that led the system to issue that alert.

In short, the information displayed when drilling down in context is always related to your original selection's settings.

How a drilldown affects the Tab heading

The Tab heading displays the name of the currently selected entity on this screen. When you drill down to a new entity in the Association area, the Tab heading changes to reflect the name of the newly selected entity.

About configuring your settings

The following settings allow you to control the appearance and behavior of the user interface. They can be configured from the Settings menu on the Precise bar:

- General settings - Includes the General, Display, and Tree view tabs.
- Time Frame settings.

About configuring general settings

On the General tab of the General Settings dialog box, you can adjust the following:

- Number of items that will be displayed for each association
From the list, select the maximum number of items that will be displayed in a table. The default is 50 items. You can control the number of items returned for a specific association by using the More... option in the Association controls. This way, you can display up to 500 items.
- Maximum number of characters to display in a text ToolTip
When you move your cursor over text columns in a table listing Statements or PL/SQLs, a ToolTip appears, displaying more text than can normally fit into a table cell. This field controls the maximum number of characters that can be displayed. The default is 500 characters.
- How often the information in the Current tab is refreshed
If you select the Auto-refresh Current tab every check box, the Current tab automatically refreshes at the specified interval. By default, the tab is not refreshed automatically. If you select this option, the default refresh interval is every five minutes.

About configuring display settings

When you drill down to or choose a different entity, selecting the option on the Display tab of the General Settings dialog box allows you to maintain the selected view or tab when changing from one tab to another. This is sometimes referred to as sticky tabs.

About configuring Tree view settings

(Only applies to Precise products that include a Tree view)

On the Tree View tab of the General Settings dialog box, you can adjust the following:

- Number of categories to display when expanding an item.
- Number of items to display when expanding a category.

About configuring Time Frame settings

You can determine the resolution of the data that is displayed in the overtime graphs using the Time Frame Settings dialog box. By using this dialog box you can define the default time frame to display.

Tasks common to most tabs

The following tasks are commonly performed in most tabs:

- [Switching to a different tab](#)
- [Selecting a time frame](#)
- [Selecting an entity](#)
- [Filtering data](#)
- [Associating entities with data that meets specific criteria](#)
- [Focusing on information in overtime graphs](#)

- [Sending an email message](#)
- [Determining which table columns to display](#)
- [Copying data to the clipboard](#)
- [Exporting to Precise Custom Portal](#)
- [Exporting to Precise Custom Portal](#)
- [Adding, viewing, and deleting Favorites](#)

Switching to a different tab

You can easily switch between the different tabs using the Tab Selection bar. When you start your Precise product, the Dashboard tab opens by default. The tab of the selected tab is displayed in blue.

To select a tab

- 1 Click a tab on the Tab Selection bar to display information on the selected entity in a different tab.

Selecting a time frame

You can configure Precise for DB2 to display transaction performance data for a specific time frame using the predefined time frame options or calendar icons.

Selecting a predefined time frame from the toolbar displays transaction performance data for the selected time period up to the current time.

See [“Selecting a predefined time frame from the Precise for DB2 toolbar”](#) on page 18.

Selecting the time frame using the calendar icon, you can choose to define a time range independent of the current time, or to define a time range up to the current time.

See [“Selecting a time frame using the calendar icon”](#) on page 18.

The predefined time frame options are:

- Last 8 hours (8h) (default)
- Last 1 day (1d)
- Last 2 days (2d)
- Last 2 weeks(2w)
- Last 5 weeks (5w)

The time frame selected affects all information displayed in Precise for DB2. Only data that falls within the selected time frame is shown in these areas.

Selecting a predefined time frame from the Precise for DB2 toolbar

To select a predefined time frame

- 1 From the Precise for DB2 toolbar, select one of the predefined time frames.

Selecting a time frame using the calendar icon

To select a time frame

- 1 Click the calendar icon. In the dialog box that is displayed perform one of the following:
 - a To define a time frame independent from the current time, select the ‘Time Range’ option and select the Start and End dates and times.
 - b To define a time frame up to the current time, select the ‘Last’ option and enter the desired time frame.
 - c To use one of the three previously used time frames, select the ‘Recently used’ option and from the drop-down menu select the desired time frame.

- d To use a previously saved time frame, select **Use a previously saved time frame** and from the drop-down menu select the desired time frame.
- 2 To save your settings for future access, select **Save these definitions for future use as:** and enter a name in the corresponding field.
- 3 Click **OK**.

Selecting an entity

You can select the entity that you want to examine in more detail.

To select an entity

- 1 On the Main area of the tab, select the entity you want to examine.

Filtering data

You can focus on specific contents of the Activity tab by showing a subset of the information. This shows the contribution of entities such as programs and users.

You can define flexible selection criteria in the Filter dialog box. To open the dialog box, you need to click the Filter Off icon next to the Time Frame selector. Filter On or Filter Off indicates the current state of the filtering mechanism. Filter Off means all information is shown; Filter On means a subset of information is shown.

When you apply your selections in the Filter dialog box, the information displayed in both the Main area and the Association area is modified to reflect your selections. Also, the filtering continues to apply when you drill down to associated entities.

Note: You can enter multiple criteria, in which case each criterion is applied using the logical AND operator.

To filter data

- 1 Click **Filter is Off/Filter is On**.
- 2 In the Filter dialog box, do the following for each entity you want to filter:

From the left drop-down list Select an entity.

From the middle drop-down list Select an operator, such as, Like, <>, Not Like, In, Not In.

In the text box Type the criteria (case-sensitive) for the selected entity.
If you select the operator Like or Not Like, you can use the % wildcard character to represent 0 or more characters, and the _ wildcard character to represent exactly one character. If you select the operator In or Not In, type a comma to separate values.

- 3 Click **OK**.

Associating entities with data that meets specific criteria

You can associate the displayed entity with specific data to focus your analysis.

Note: The criteria no longer apply when you drill down to another entity.

To associate entities with data that meets specific criteria

- 1 Click the arrow located to the left of the Association controls and select **More...** (applies to Current and Objects tabs only)

- 2 In the Associate With dialog box, on the Entries tab, select the entity you want to associate data with from the Populate table with list.
- 3 In the Sort entries by list, determine according to which criteria you want the information to be sorted and in which order.
- 4 From the Display top list, select the number of rows to display.
- 5 On the Criteria tab, do the following for each entity you want to associate data with:

From the left drop-down list Select an entity.

From the middle drop-down list Select an operator, such as, Like, <>, Not Like, In, Not In.

In the text box Type the criteria (case-sensitive) for the selected entity.
If you select the operator Like or Not Like, you can use the % wildcard character to represent 0 or more characters, and the _ wildcard character to represent exactly one character. If you select the operator In or Not In, type a comma to separate values.

- 6 Click **OK**.

Focusing on information in overtime graphs

Some entities display an overtime graph. The overtime graph displays entity statistics over a specified time period. Depending on the number of points displayed in the graph, you may need to zoom in or out. The text displayed on the x-axis varies according to the time frame. If there is a year or day change, x-axis labels will display accordingly.

To focus on information in overtime graphs, select the desired time frame on the overtime graph, click and drag the left or right handle to increase or decrease the time range. The small zoom (spyglass) icon will display on the upper right of the selected time range, and above the overtime graph legend. Click the zoom icon to zoom in according to the selected time frame.

Sending an email message

You can send an email message to one or more recipients from the Precise toolbar.

The default subject for the message will be "Link to a Precise application".

The email will include a link to the Precise product in the current context (time frame and selected entries).

To send an email message

- 1 Click the email icon on the Precise toolbar. The default email program opens.
- 2 Fill in the required fields and click **Send**.

Determining which table columns to display

Tables are used to display information about a set of related entities in the Main and Association areas. It is possible to determine which columns to display in the Association area tables.

To determine which columns to display in the Association area

- 1 Click the Table icon on the upper right-hand side of a table and select **Column Chooser**.
- 2 In the Table columns dialog box, click the arrows to move the names of the columns that you want to display to the Visible box and the ones that you do not want to display to the Invisible box.
- 3 Click **OK**.

Copying data to the clipboard

At times you may want to save data displayed in the table area in a Microsoft Excel spreadsheet for further analysis or save an image of a graph to the clipboard.

To copy data displayed in the Association area to the clipboard

- 1 Click the Table icon on the upper right-hand side of a table and select **Copy to clipboard**.
The table can be pasted into Microsoft Excel or as an HTML file.

To copy a graph to the clipboard

- 1 Right-click a graph and choose **Copy to clipboard**.
You can now paste the image into any application that works with the clipboard.

Exporting to Precise Custom Portal

The Export to Precise Custom Portal Portlet feature enables you to export the view of the chosen table or graph and generate a portlet with that view in Precise Custom Portal, so that it will provide you with another way of monitoring your application.

Prerequisites

To be able to use this feature, you need to have the following rights in Precise:

- View permissions to all tiers in the application.

If you do not have sufficient rights, you will get an error message when trying to execute this feature.

Exporting the information

You can either export a table view or a graph view.

Note: The name field has the following restrictions: maximum 100 characters.

To export a table view

- 1 Click the Column Chooser icon.
- 2 Select **Export to Precise Custom Portal Portlet**.
- 3 Insert a name in the name field that clearly describes the table view.
- 4 Click **OK**.

To export a graph view

- 1 Right-click the graph.
- 2 Select **Export to Precise Custom Portal Portlet**.
- 3 Insert a name in the name field that clearly describes the graph view.
- 4 Click **OK**.

Adding, viewing, and deleting Favorites

The Favorites feature allows you to save a specific location in your application and to retrieve the same location later without having to navigate to it.

About the Favorites feature

The new Favorites feature includes the following options:

- Relative Time Frame – Saving relative time frame instead of static date. For example, saving the last 7 days will always display the last 7 days, depending on the day entered.

- One click to specific location – Once you open Precise by opening a saved Favorite item, you will not have to enter a login credential nor click the login button.
- IE Favorites support – Adding a new Favorite item in Precise will also add it to the IE Favorites menu.

UI description

A Favorites menu has been added to the Precise bar in each product including StartPoint.

An Add/Delete Favorites option under the Favorites menu allows you to save the current location or delete an existing one.

To add a new Favorite location

- 1 On the Add/Delete Favorites dialog box, enter the name of the new Favorites entry.
- 2 Click **Add**. The dialog box is closed and the new Favorite is added to the list.

To view a Favorites location

- 1 On the Precise bar, click **Favorites**.
- 2 Select the Favorites location you want to view.

To delete an existing Favorite location

- 1 On the Add/Delete Favorites dialog box, select the Favorite location to be deleted.
- 2 Click **Delete**. The dialog box closes and the selected Favorite is deleted from the list.

Note: The favorite address is displayed in the Address field and cannot be edited.

Opening your Precise product from StartPoint

Precise is a Web-based application. You can access the Precise user interface using the Internet Explorer browser, version 6.0 or later. The syntax of the Precise URL address is: `http://<server>:<port>`, where `<server>` refers to the Precise FocalPoint server and `<port>` refers to the port number used by the GUI Web server. By default, the port number is 20790 (20760 for J2EE). For example: `http://beanotown:20700`.

This URL provides secure access to the StartPoint using authorized roles. From here, you can open all Precise products. It gives you a quick overview of the status of your applications and access to the AdminPoint, where you can perform various management tasks (see the *Precise Administration Guide* for details).

Note: You must have local administrator privileges on the server where the StartPoint is running.

To open your product using Precise StartPoint

- 1 Type the address of the StartPoint user interface into the Address bar of your browser and click **Enter**. The Precise login page opens. The login page provides secure access to Precise and to your specific product.
- 2 Specify your authorized role name and password. By default, both role name and password are `admin`. For more information about role names, see the *Precise Administration Guide*.
- 3 Click **Login**. The StartPoint page opens. This is the Precise home page.
- 4 On the Product Selection bar, from the drop-down list, select the product you want to open.

Getting an overview of your application

This section includes the following topics:

- [About the Dashboard tab](#)
- [How the Dashboard tab is structured](#)
- [How the Dashboard can help you identify performance problems](#)

About the Dashboard tab

The Dashboard tab (displayed in provides an overview of the information displayed in the other Precise for DB2 tabs. It lets you quickly identify the most active database partitions which are most likely to benefit from a tuning operation.

The Dashboard tab lets you identify the top resource consumers of your DB2 database partitions within the selected time frame. This will guide you in your search for the root cause of a problem. You can also select a database partition and watch its activity over the last hours, days, or months.

The Dashboard tab provides answers to questions like: "Are all my DB2 database partitions up and running?"; "Which database partition is consuming the most resources?"; "Which application is consuming the most resources for the specified database partition and across all database partitions?"; "Is the problem related to a slow I/O, a low hit-ratio, a memory shortage, high CPU consumption, or another specific problem?"

The All row in the All DB2 Databases table located in the Main area (left pane) displays the average availability of all database partitions. When you select the All row, a summary of the activities of all the database partitions and a comparison of application identifiers is displayed in the tab views in the Database Details area (right pane).

If you are working with SAP and ERP support is installed, the Dashboard tab displays information on DB2 activity in relation to the ERP application and users, and specific information on executed views and transactions.

When an extension for ERP is installed, the Precise for DB2 Collector divides DB2 activities into several groups, called work types.

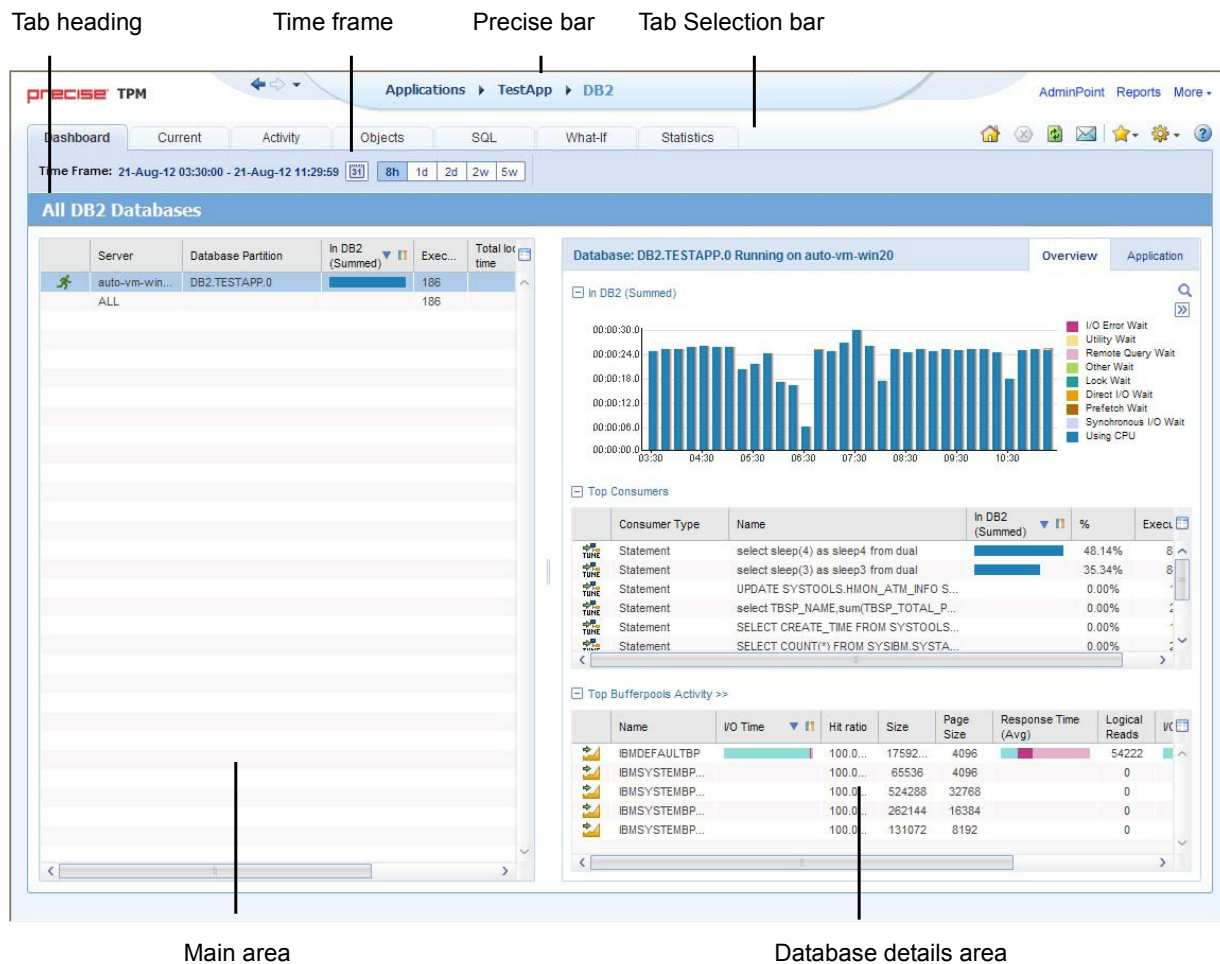
Technology	Work Type (can be any of the following:)
SAP Dialog	Activities generated from the SAP Dialog work process.
SAP Update	Activities generated from the SAP Update work process.
SAP Spool	Activities generated from the SAP Spool work process.
SAP Background	Activities generated from the SAP Background work process.

Note: All databases, including those with only one partition, are referred to as database partitions. In the Multipartition column, databases with only one partition will have a value = 'No'.

- See “[Precise for DB2 basics](#)” on page 14.
- See “[Getting current application information](#)” on page 30.
- See “[Examining performance over time](#)” on page 39.
- See “[Examining the relationship between database objects](#)” on page 67.
- See “[Analyzing Statement plans](#)” on page 73.
- See “[Simulating the effect of changes](#)” on page 82.
- See “[Examining DB2 database statistics](#)” on page 86.

How the Dashboard tab is structured

Figure 3-1 Precise for DB2 Dashboard tab



The Dashboard tab is divided into two areas—the Main area and the Details area. The Main area lists all the databases or database partitions that are monitored by Precise for DB2. The Details area provides comparative information regarding the selected database partition.

You can filter the information shown in the Details area so that only the contribution of certain entities is included.

See “[Filtering data](#)” on page 19.

About the general tab structure

The following table displays information on the general structure of the Dashboard tab.

Table 3-1 Elements of the Dashboard tab

Element	Description
Time Frame	Allows selecting the time frame we want to work with. The time frame influences all counters displayed (e.g., executions, I/O time).
Filter On/Off	Allows filtering the displayed information according to user-defined conditions.
Tab heading	Displays All DB2 databases.
Main area	Each row displays a single database or a single database partition. It displays all databases/partitions which belong to the current tab. Also displayed is a summary row (All).
Details area heading	Displays either a database partition and the name of the server it is running on, or all databases running on all servers. Note: <number> is displayed only if the partition belongs to a multi-partitioned database.
Details area	Displays graphs and tables describing the row selected in the Main area.
Details area tabs	Allows switching between the different views (Overview, Application)

About the Main area in the Dashboard tab

The database table displayed in the Main area lists all the databases or partitions that are monitored by Precise for DB2. Each row corresponds to a database or database partition. The All row displays the average activity and the average availability (%) of all database partitions.

The following table displays information on the different database partitions in the application.

Table 3-2 Database partitions table

Column name	Description	Sortable
Status	Icon representing the status of the collector (running, stopped, starting, etc.)	No
Server	Name of the server on which the database partition is running.	Yes
Database partition	<instance>.<database>[.number]	Yes
In DB2 (summed)	Stacked bar of 'In DB2' data with regular 3 display modes. Information limited to time frame.	Yes
Executions	# of statements executed in time frame. Note: this value is taken from db-stats information.	Yes
Total lock time	Time period this database partition spent in lock wait.	Yes
Multi-partitioned	Icon representing if this is part of a multi-partitioned installation or not	Yes
Availability	% of availability (0-100)	Yes
Catalog partition	Name of the database partition (see format above) which contains the database catalog. For single partition, this column will be left blank.	Yes

Note: When you open the Dashboard tab, only the first six columns are displayed. To view additional columns (such as: Executions and Last Update columns), use the left and right arrows to the right of the column titles.

About the Details area in the Dashboard tab

The Details area displays information on the database or partition selected in the Main area. It has two tabs: an Overview tab and an Application tab.

Note: If you select the All row, the information displayed reflects the average availability of all Database Partitions.

The Overview tab

The Overview tab displays information in three views:

- In DB2 (summed) – displays database activity over time. Each time period is divided according to the ‘In DB2’ groups (using CPU, locks, etc.).
- Top consumers – displays the top statements and objects (tables/indexes) using database resources.
- Bufferpool activity – displays the overall activity of each bufferpool.

The In DB2 (summed) view

This graph displays the database partition In DB2 activity over time, split according to normal In DB2 division.

The Top consumers view

The following table describes the different elements of the Top consumers table.

Table 3-3 Top consumers table

Column name	Description	Sortable
Launch	Launch to the activity's statement or object. You will see activity over time. From a statement, launch to the SQL tab. From an Index or Table, Launch to the Activity tab.	No
Consumer Type	Table, Index or Statement	Yes
Text	Table name, index name or statement unformatted text	Yes
In DB2 (summed)	In DB2 time consumed by the object during the time period.	Yes
%	% (0-100) representing the percentage of time (of the total In DB2 time of the database partition) spent on that statement or object	Yes
Executions	Total statement executions within the selected time frame. Available for statements, tables, and indexes.	Yes

The Bufferpool activity view

The Bufferpool activity view displays the activity of each bufferpool in the system.

When clicking the view's heading (Bufferpool activity), we launch to activity and view the comparison screen of the bufferpools, at the database partition level.

The following table describes the different elements of the Bufferpool activity table.

Table 3-4 Bufferpool activity table

Column name	Description	Sortable
Launch	Launch to activity's bufferpool. You will be able to see the bufferpool's activity over time.	No
Name	The bufferpool name.	Yes
Total I/O time	A stacked bar of I/O time break into the following: Synchronic read time, A-synchronic read time, direct read time, Synchronic write time, A-synchronic write time, direct write time.	
Hit ratio	Bufferpool hit ratio	Yes
Size	Bufferpool size (in MB)	Yes
Bufferpool size (in MB)	Bufferpool page size	Yes
Response Time (Avg)	Bufferpool's average response time	No
Logical Read	Logical reads	Yes
Total I/O Count	Bufferpool's total I/O count, broken down into the following: Synchronic reads, A-synchronic reads, direct reads, Synchronic writes, A-synchronic writes, directs writes	Yes
Physical Reads	Total physical reads	Yes
Pages	Number of pages	Yes

The Application tab

The Application tab displays information in three views:

- Programs
- Modules
- DB Users

All views share the same columns headings. The following table describes the different elements of these tables.

Table 3-5 Application tab tables

Column name	Description	Sortable
Program	Name of the program.	Yes
In DB2 (summed)	The In DB2 view displays a stacked bar graph on the In DB2 resource consumption over time, for the selected time frame, broken down into the In DB2 application states.	Yes
%	Percentage of time the program took out of the total In DB2 time for the selected database partition.	Yes
Executions	Number of program executions.	Yes
Applications	Number of applications opened by the program.	Yes

Note: If you are working with SAP and ERP support is installed, the Application tab is named SAP instead, and the entities displayed are Work Types, Transactions, and Modules.

Launching from the Dashboard to other tabs

When you open the Dashboard tab and select a row on the Main area (left pane) and then click on another tab, the target tab is opened in the context of the selected database partition.

How the Dashboard can help you identify performance problems

You can identify a performance problem by doing one or more of the following:

- [Identifying the most resource-consuming database partition](#)
- [Analyzing top resource consuming entities](#)
- [Analyzing bufferpools](#)
- [Identifying the top resource consumers in various entities](#)

Identifying the most resource-consuming database partition

Precise for DB2 enables you to determine which are the most resource-consuming database partitions, indicating a potential performance problem.

The following example shows how to determine which database partition consumes the most resources. This example can show the top resource-consuming database partition, as measured by the total time spent executing In DB2.

To identify the most resource-consuming database partition

- 1 In the Database Partition list, choose the Database Partition you want to analyze.
- 2 In the Time Frame list, choose the period of time you want to analyze.
- 3 In the Database Partition Details area view the In DB2 (Summed) overtime resource consumption graph. You can then launch to the Activity tab to check what the application was doing during the specified time period.
- 4 In the Database Partition Details area view the Top Consumers grid to see which were the top statements and objects (tables and indexes) that consumed most of the resources. Then launch in context to the SQL tab to further investigate the heavy statement, or directly launch in context to the Activity tab to see which statements were using the heavy object.
- 5 In the Database Partition Details area view the Bufferpool Activity grid to see which bufferpools were most active, and what are the statistics regarding these bufferpools. Then launch in context to the Activity tab to view overtime and extended information on these bufferpools.
- 6 Use any of the following options to investigate further:
 - Launch to the Statistics tab for the selected database partition.
 - Launch to the Activity tab to check what the application was doing at the specified time.

Analyzing top resource consuming entities

Precise for DB2 enables you to determine which are the most resource-consuming statements and objects (tables and indexes), as measured by the total time spent executing In DB2, indicating a potential performance problem.

The following example shows how to identify the top consumers that are the most resource-consuming statements and objects.

To identify the top resource consumers (statements and objects)

- 1 In the Database Partition list, choose the Database Partition you want to analyze.
- 2 In the Time Frame list, choose the period of time you want to analyze.
- 3 In the Database Partition Details area view the Top Consumers grid to see which were the top statements and objects (tables and indexes) that consumed most of the resources.

- 4 In case of a heavy statement directly launch in context to the SQL tab to further investigate the heavy statement. In case of a heavy table or index directly launch in context to the Activity tab to see which statements were using the heavy object.

Analyzing bufferpools

Precise for DB2 enables you to determine which are the most activated bufferpools. It also presents additional statistics-related information on each bufferpool, such as hit ratio, average response time etc.

The following example shows how to identify the top active bufferpools.

To identify the top active bufferpools

- 1 In the Database Partition list, choose the Database Partition you want to analyze.
- 2 In the Time Frame list, choose the period of time you want to analyze.
- 3 In the Database Partition Details area view the Bufferpool Activity grid to see which were the most active bufferpools. Examine additional statistics information on the active bufferpool.
- 4 In case you are interested in other statistics (e.g. the hit ratio), order the grid according to the required column.
- 5 Launch directly to the Activity tab (using the launch icon) to view overtime and additional information on the specified bufferpool. Then drill from there to tablespaces and objects to see which were the most activated tablespaces and objects within the specified bufferpool.

Identifying the top resource consumers in various entities

A look at the Applications tab quickly helps you identify the top resource-consuming entities.

The following example shows how to identify the top resource-consuming entities.

To identify the top resource consumers in various entities

- 1 In the Database Partition list, choose the Database Partition you want to analyze.
- 2 In the Time Frame list, choose the period of time you want to analyze.
- 3 In the Database Partition Details area, click Applications. For example, a SAP database, PS8 on hp8, is monitored by a SAP extension.

The Application tab shows the following SAP-related top entities:

- Background work type
- SAP user SUPER
- Transaction SAP_CCMS_MONI_BATCH_DP

Most likely, these entities are related.

- 4 To investigate further, click an entity to launch the Activity tab. In the Activity tab you can confirm the connection between the problematic entities. For this example, you can identify a background transaction run by the SAP CCMS administration tool.

Getting current application information

This section includes the following topics:

- [About the Current tab](#)
- [How the Current tab is structured](#)
- [About examining entities in the Current tab](#)
- [How the Current tag can help you identify performance problems](#)

About the Current tab

The Current tab displays near real-time information on all applications that were connected to your DB2 database during the last minute or, in the Application context, from the beginning of the time slice. If you have discerned an immediate performance issue and are looking for an indication as to where to begin your analysis, start with the Current tab. It is also worthwhile to start with the Current tab if your system appears to be 'hung' and there is no throughput.

You can use this tab to monitor your system's current activities and resource consumption and to identify performance problems, such as current locking problems and applications involved in locks, as they occur.

The Current tab provides answers to the following types of questions: "How many applications are currently running?" or "Why doesn't my application work when I click OK?" The Current tab also shows you which applications are currently active, and which machines are connected to which users.

Using the Current tab, you can monitor current connections to your DB2 database partition, check the state of each connection, and identify the resources currently utilized by each application. Further investigation of each application provides a thorough understanding of the SQL statement being executed, the login name executing the statement, the program, and the resource utilization for the application. In addition, you can view the statements executed by the selected application within the last time slice.

In the Current tab, data on a current application is continuously updated to provide information on what happened during the last minute or time slice. Because the displayed application information is derived from a dynamic application, it is recommended to regularly refresh the screen, either manually or automatically, to ensure that the information displayed is up-to-date.

Note: The Current tab displays information on the last minute of activity or, in the Application context, from the beginning of the time slice. For example, if you've worked at your screen for ten minutes and then refresh the information displayed on your screen, you will be able to view new information that occurred during the last minute of activity or from the beginning of the time slice in the Application context.

See "[Analyzing currently held locks](#)" on page 38.

See "[Getting an overview of your application](#)" on page 23.

How the Current tab is structured

The Current tab displays information on a selected entity and its associated entities. For example, if the parent entity is an application, its associated entities would be Statements.

When you open the Current tab from the Dashboard tab, the selected entity is by default a database partition, meaning that information is displayed on the database partition level. If you open the Current tab from another tab, the historical settings (meaning those settings which were selected when you left the tab, such as a drill-down to a Locked Objects entity) are taken into account and the information displayed the last time you viewed this tab is displayed (similar to clicking the History button and returning to a previous tab).

The selected entity is always reflected in the Tab heading, which serves as a point of orientation. The highest-level entity you can view information for in the Current tab is database partition. You can view another database partition by selecting it from the database partition list.

About the Main area in the Current tab

The Main area shows comprehensive information on the selected entity. You can choose from several views to examine the entity from different angles. For example, if you are analyzing a database partition entity, you can focus exclusively on the data displayed in the overview, such as, In DB2 data (default).

About the Association area in the Current tab

The Association area provides corresponding information on the entities associated with the selected entity (displayed in the Main area). The selection you make is reflected in the Association area only; the Main area remains unchanged. For some entities, tabs above the Association area enable you to view additional information. For example, the tabs displayed for a database entity are: Application, In DB2, and Connection. Clicking a tab displays different table columns with the same type of child entities.

From the Association area, you can also drill down to another entity by clicking a table row. A drilldown affects the entire tab. When you drill down to another entity, the Tab heading displays the new selection, the Main area displays information on the newly selected entity, and the Association area displays the entities associated with the selected entity.

For example, when you want to view information on a specific application in the Current tab, choose Applications from the Association controls. The Association area changes to display application statement information. Note that the Tab heading remains unchanged and the Main area changes to application context.

Locks can be viewed in both an application context and a data partition context. In the application context you can see which application is locking you. In the data partition context a lock tree is displayed showing details of all locks holding the application.

If you want to view detailed information on a specific application holding lock, click the row of the applications holding locks you want to view detailed information for. The Tab heading indicates the newly selected entity; the Main area displays locking and In DB2 information on the application you drilled down to, and the Association area shows information on the statements associated with this application.

See [“Precise for DB2 basics”](#) on page 14.

See [“Getting an overview of your application”](#) on page 23.

About examining entities in the Current tab

The Current tab displays information on different entities. This section provides an overview of all entities, their meaning, and their views.

The following entities can be examined in the Current tab:

- [About the Database Partition](#)
- [About the Application entity](#)
- [About the Application Holding Locks entity](#)
- [About the Statement entity](#)

About the Database Partition

Displays information on the activity and resource consumption of the database, including the number of applications connected to the database during the last sample taken and the database resource consumption breakdown for the last minute.

About getting an overview of current Database activities

Displays information on the activity and resource consumption of the database, including the number of applications connected to the database during the last sample taken and the database resource consumption breakdown for the last minute.

The following table describes the information displayed in the Database Overview view.

Table 4-1 Database Overview view

View area	Description
Application	<p>Displays the following application-related information:</p> <ul style="list-style-type: none"> ■ Applications—number of applications (different applications) that were connected to the database during the last sample. ■ Active Applications—number of applications (different applications) that were in one of the In DB2 states during the last sample. ■ Locked Applications—number of blocked applications connected to the database during the last sample.
In DB2	<p>Displays the resource consumption breakdown of the database during the last minute or, in the Application context, from the beginning of the time slice. In addition to the resource distribution, the view displays the number of applications in each state during the last sample.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> ■ State—displays the In DB2 state. ■ In DB2—graphical representation of the Time column. ■ Time—amount of time the selected entity was in this state. ■ %—percentage of time the selected entity was in this state. <p>Note: All the information displayed in this table represents information collected during the last minute, with the exception of the Applications column, which represents the number of applications, in the last sample, that were in a particular state.</p>

About viewing Database locks

Displays lock wait duration for the selected Database during the last minute. The applications that are currently involved in a locking or blocking situation are displayed in tree format.

Note: Since the Lock tree represents the locks that are currently occurring, and the lock breakdown represents the duration of locks occurring during the last minute, it is possible that a situation may occur whereby a lock was freed, but since the last minute has not yet passed, the lock tree will display the message 'No locks detected' while the Lock Breakdown grid still shows lock duration.

The following table provides a detailed explanation of the Lock Tree

Table 4-2 Database Locks view

View area	Description
Lock tree	<p>It is important to determine the following information for each blocking chain:</p> <ul style="list-style-type: none"> ■ Which applications are involved in the lock? ■ Which application is blocking the others? <p>The Lock Tree provides the necessary information to get to the root of the problem by providing information on applications currently involved in the blocking situation.</p> <p>A yellow lock indicates that the application is currently locking other applications but is not itself blocked.</p> <p>A red lock indicates the application is currently blocked by another application. Note that an application may also block other applications while being blocked by itself.</p> <p>To gain a better understanding of the lock tree take a closer look at its structure. Notice that the application located in the root of each chain is the application that blocks all others. Each application directly blocks all applications positioned one level beneath it.</p> <p>The Lock Tree displays the application handle, program name, the object being blocked and the lock type.</p>

About the Application entity

Displays information on a single connection with Microsoft SQL Server. Precise for SQL Server identifies an application by the following attributes:

- Application ID—a unique integer assigned to each user connection when the connection is made.
- Application identifiers such as program or database user.

When one of these attributes changes, a new application is recorded and displayed.

The following table describes the information displayed in the Application overview.

Table 4-3 Application Overview

View area	Description
Application tab	<ul style="list-style-type: none"> ■ Displays the following information: ■ Application Handle ■ Application ID ■ Program ■ State ■ In DB2 ■ Text

Table 4-3 Application Overview

View area	Description
In DB2 tab (last minute)	<p>Displays the resource consumption breakdown of the application during the last minute (that is, 60 samples of the Collector). In addition to the resource distribution, the view displays the number of applications in each state during the last sample.</p> <p>It is important to remember that since this view only displays the last minute's activities, you will always be viewing the activities that occurred during the last minute and not the activities that occurred since the last refresh action.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> ■ Application Handle ■ Application ID ■ In DB2 ■ Using CPU ■ CPU Wait ■ Sync I/O Wait ■ Prefetch Wait ■ Direct I/O Wait ■ Utility Wait ■ Lock Wait ■ I/O Error Wait ■ Remote Query Wait
Connection tab	<p>A connection includes application identifiers such as User, Login, and Machine.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> ■ Application Handle ■ Application ID ■ Program ■ Connect Time ■ User ■ Login ■ Machine ■ Module ■ Action

The following table describes the information displayed in the Application Text view.

Table 4-4 Application Text view

View area	Description
Text	<p>Displays the text of the current statement. If the application is not active and executing a statement, a message notifying you that the application is not active and that no text is available is displayed.</p>

About the Application Holding Locks entity

The fact that a currently active application is holding a resource does not necessarily mean that there is a contention between resources. The Application Holding Locks entity displays information that lets you differentiate between applications that are locking resources and applications that are holding resources.

About getting an overview of current applications holding locks

Displays information on the application identifier, including connection details, application-related statistics and a breakdown of application resource consumption, during the last minute.

The following table describes the information displayed by the Application Holding Locks overview.

Table 4-5 Applications Holding Locks overview

View area	Description
Application tab	<ul style="list-style-type: none"> ■ Displays the following information: ■ Application Handle ■ Application ID ■ Program ■ State ■ In DB2 ■ Locked Tablespace ■ Locked Schema ■ Locked Table ■ Text
In DB2 tab (last minute)	<p>Displays the resource consumption breakdown of the application during the last minute (that is, 60 samples of the Collector). In addition to the resource distribution, the view displays the number of applications in each state during the last sample.</p> <p>It is important to remember that since this view only displays the last minute's activities, you will always be viewing the activities that occurred during the last minute and not the activities that occurred since the last refresh action.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> ■ Application Handle ■ Application ID ■ In DB2 ■ Using CPU ■ CPU Wait ■ Sync I/O Wait ■ Prefetch Wait ■ Direct I/O Wait ■ Utility Wait ■ Lock Wait ■ I/O Error Wait ■ Remote Query Wait
Connection tab	<p>A connection includes application identifiers such as User, Login, and Machine.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> ■ Application Handle ■ Application ID ■ Program ■ Connect Time ■ User ■ Login ■ Machine ■ Module ■ Action

About viewing information on the lock breakdown of an application that is holding locks

The Locks view displays lock wait duration broken down into lock levels such as Table, Page and Key, for the selected application, during the last minute.

The applications that are currently involved in a locking or blocking situation with the selected application are displayed in tree format. Since the Lock tree represents the locks that are currently occurring, and the lock breakdown represents the duration of locks occurring during the last minute, it is possible that a situation may occur whereby a lock was freed, but since the last minute has not yet passed, the lock tree will display the message 'No locks detected' while the Lock Breakdown grid still shows lock duration.

To gain a better understanding of the lock tree take a closer look at its structure. Notice that the application located in the root of each chain is the application that blocks all others. Each application directly blocks all applications positioned one level beneath it.

The Lock Tree displays the SPID, program name, the object being blocked and the lock type.

The following table describes the information displayed in this view.

Table 4-6 Locks view for the Application Holding Locks entity

View area	Description
Lock Tree	<p>The Lock tree provides the necessary information to determine the following information for each blocking chain:</p> <ul style="list-style-type: none"> ■ A yellow lock indicates that the application is currently locking other applications but is not itself blocked. ■ A red lock indicates the application is currently blocked by another application. Note that an application may also block other applications while being blocked by itself.

About viewing the text of an application that is holding locks

Displays the text of the current statement. If the application is not active and executing a statement, a message notifying you that the application is not active and that no text is available is displayed.

About the Statement entity

The Statement entities display information on an SQL statement that is executed by the selected application during the last time slice.

About getting an overview of current Statement entities

Displays general information on the statement, including overall properties, its breakdown in DB2 during the last time slice, and the breakdown of the locks during the last time slice.

The following table describes the information displayed in the Statement Overview view.

Table 4-7 Statement Overview view

View area	Description
General tab	<ul style="list-style-type: none"> ■ Statement ID—Displays the internal statement hash value (SHV) calculated by the Collector, or, in the case of a stored procedure, displays the name of the stored procedure. ■ In DB2 ■ Text

Table 4-7 Statement Overview view

View area	Description
In DB2 Last minute)	<p>The following information is displayed:</p> <ul style="list-style-type: none"> ■ Statement ID—Displays the internal statement hash value (SHV) calculated by the Collector, or, in the case of a stored procedure, displays the name of the stored procedure. ■ In DB2 ■ Using CPU ■ CPU Wait ■ Sync I/O Wait ■ Prefetch Wait ■ Direct I/O Wait ■ Utility Wait ■ Lock Wait ■ I/O Error Wait

About viewing current Statement text

Displays text of the selected statement.

How the Current tag can help you identify performance problems

You can identify a performance problem by doing one or more of the following:

- [Examining resource consumption of an entire database partition](#)
- [Observing the activity of a single application](#)
- [Examining a single connection](#)
- [Analyzing currently held locks](#)

Examining resource consumption of an entire database partition

The Current tag provides an almost real-time picture of the performance and behavior of your DB2 database partition. The Current tag displays information on the last time slice of activity, unless otherwise configured as for the last minute. Usually you will examine the activities of the entire database partition level to determine what is happening in your system. By examining the current activity at a database partition level of display, you can see how it is currently behaving and determine if there is an overall problem that affects all applications or if there is a problem with specific applications.

To examine resource consumption of an entire instance and database

- 1 In the database partition list, choose the database partition you want to analyze. Do the following:
 - a On the View controls in the Main area, click **Overview**. This will show you a breakdown of resource consumption for the last minute.

Observing the activity of a single application

The Association area of the Current tag lets you view which applications are currently connected to the database partition. You can observe the current state of the application and determine if it is active or waiting for a resource. You can also identify the application by different identifiers such as its login information, the machine it is running on, or the programs it is running. The currently active TSQL code is displayed for active applications.

To observe current application activity

- 1 In the database partition list, choose the database partition you want to analyze.
- 2 On the Association controls in the Association area, select one of the following to focus your investigation on the applications that were active in SQL Server during the last minute:
 - Active Applications
 - Applications

Examining a single connection

You may need to explore a single running application, to obtain more information on it. In the Association area you can drill down to the specific application in question, and continue your analysis.

To examine a single connection

- 1 In the database partition list, choose the database partition you want to analyze.
- 2 In the Association area, drill down to the application whose connection you want to analyze.
- 3 Continue your analysis by examining its resource breakdown and its recent statements, during the last few minutes.

Analyzing currently held locks

To be able to avoid a blocking situation, it is important to have an accurate picture of all the locks held in the system. This will help you understand if the lock system is tuned and if it is possible to avoid a potential blocking situation. Using Precise for DB2, you can now observe the locks currently held in the database partition and group them according to lock level and lock mode.

Locks can be viewed in two levels:

- Application context - To see which application is locking you.
- DB partition context - To see the Lock tree with details of all locks holding the application.

To analyze locks currently held

- 1 In the database partition list, choose the item you want to analyze.
- 2 On the View controls in the Main area, click **Locks** and examine the locking contentions. Move the cursor over the locked application in the tree to display the text of the locked application.

Examining performance over time

This section includes the following topics:

- [About the Activity tab](#)
- [How the Activity tab is structured](#)
- [About the entities you can examine](#)

About the Activity tab

The Activity tab allows you to analyze performance behavior over time or within a given time period. When investigating a performance issue, the Activity tab is an excellent place to start. You can easily identify the performance profile of your DB2 database partition and determine who is executing which application, what they are waiting for and how many resources they are consuming. This is the critical information you need to identify the top consumers of your DB2 database partition and its major latencies. By focusing on the top consumers and latencies (wait states) you can maximize the tuning impact on your system while minimizing the time it takes to obtain results.

The Activity tab is also very useful for a scale-up exercise. Too frequently companies add CPUs to an existing server in their attempt to improve performance. Although occasionally successful, step one of any scale-up activity should be to identify the major latencies in the DB2 database partition and resolve them. Only once the major wait states are resolved, will a scale-up actually be capable of improving performance. Otherwise, you may be simply wasting money and resources.

The information displayed in the Activity tab is derived from historical data and enables you to access information for different time intervals at varying levels of detail. You can analyze what happened at a specific time, pinpoint problematic time periods, and drill down to analyze a bottleneck that caused a performance problem.

The Activity tab was designed to answer questions such as, "What happened yesterday, as well as a month ago and a year ago?" and "What has been the performance level of my applications over time".

Data in the Activity tab is available up to the last time slice that ended (information on the last minute is available in the Current tab). The Activity tab is usually the place to identify and investigate tuning problems.

When working in the Activity tab, you can choose to work on a specific entity or to view information on all the databases in the application.

Note: You usually switch to the Activity tab, after focusing on a database partition in the Dashboard tab, or by launching Precise for DB2 from another product (Insight, Alerts, or Report Manager). In these cases, the product automatically focuses on the relevant database partition.

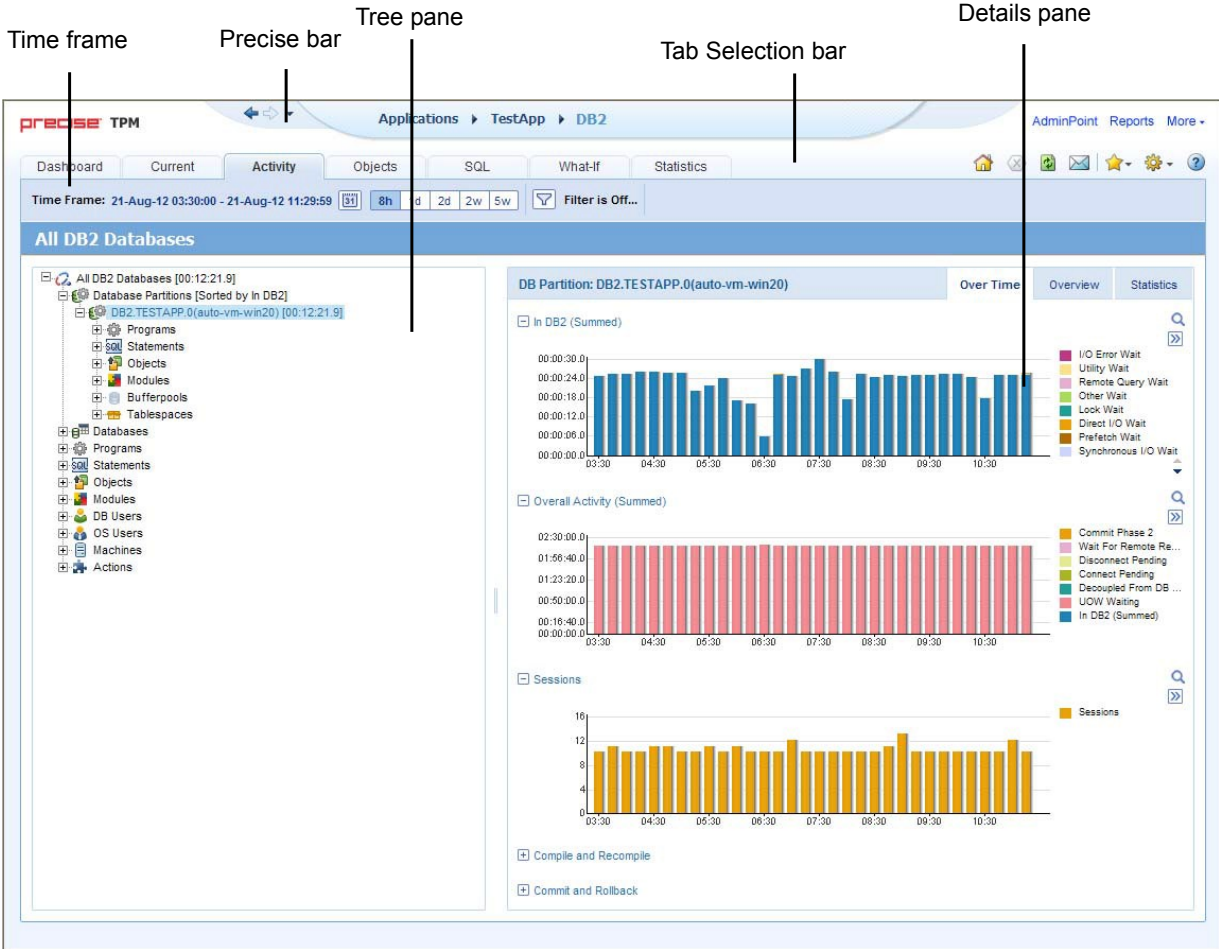
When you do not select a specific database partition to view in-context, a summary of the performance information on All database partitions in the application is displayed. You can analyze the average behavior of the entire Precise for DB2 Tier. You can then identify the database partition which consumes the most resources in your system and focus on it to further your investigation.

When you focus on a specific database partition (either as a result of launching or from the Activity tree) you can also identify which is the dominant resource consumed, in terms of CPU, I/O lock, etc. By identifying your system's wait type you can formulate an idea as to what kind of solutions will help you improve your system.

Note: If you are working with SAP, PeopleSoft, or Siebel, and ERP support is installed, the Activity tab displays information on the historical Precise for DB2 activity in relation to the ERP application and users, and specific information on executed views and transactions.

How the Activity tab is structured

Figure 5-1 Precise for DB2 Activity tab



When you open the Activity tab from the Dashboard tab, information on the database partition selected in the Dashboard is displayed, or on All DB2 Databases if the All row was selected. When you open the Activity tab from another tab, the settings from your last visit to the Activity tab are retained, such as a drill down to a Statement entity. This is similar to clicking the History button and returning to a previous tab.

The highest-level entity you can view information for in the Activity tab is All DB2 Databases. You can view a particular database partition by selecting it on the tree.

You can use the Filter dialog box to filter the Details pane so that only the contribution of certain entities is included.

See "Filtering data" on page 19.

The Activity tab is divided into two areas, the Tree pane and the Details pane. A breakdown of the entities that are associated with the selected entity is displayed in the Tree pane. Information on the entity you selected is displayed in the Details pane.

About the Tree pane in the Activity tab

Which entries are displayed, and how many items are displayed can be changed in the Settings menu.

The duration value displayed next to each entry reflects the duration of each instance in the selected entity, according to the sort parameter you selected. The sort parameter is displayed next to each entity. The default entity for the top entity in the tree is In DB2 time.

See [“About configuring your settings”](#) on page 17.

Determining what information is displayed in the Tree

You can also control the number of entries displayed and the order in which they are displayed in the tree using the sort and filter feature.

To determine what information to display in the tree

- 1 Right-click on an entity in the tree, and then click **Sort and Filter**.
- 2 In the Associate With dialog box, on the Entries tab, select the entity you want to associate data with from the Populate table with list.
- 3 In the Sort entries by list, determine according to which criteria you want the information to be sorted and in which order.
- 4 From the Display top list, select the number of rows to display.
- 5 On the Criteria tab, do the following for each entity you want to associate data with:

From the left drop-down list Select an entity.

From the middle drop-down list Select an operator, such as, Like, <>, Not Like, In, Not In.

In the text box Type the criteria (case-sensitive) for the selected entity.

If you select the operator Like or Not Like, you can use the % wildcard character. If you select the operator In or Not In, type a comma to separate values.

- 6 Click **OK**.

About viewing information for entities associated with a selected entity

You can drill down to the entities that are associated with the selected entity by clicking on the plus sign (+) located next to each entity. When you drill down on an entity, associated entities are displayed on the tree and entity details are displayed in the Details pane. When you drill down on an instance that is associated with the selected entity, you display the instances that are associated with the selected entity and you can view additional information for the selected instance in the Details pane.

Viewing information from a different time frame

You can switch to a different time frame and view additional information on the selected entity that was collected for that time frame. Only the information in the subnodes of the entity you have selected is refreshed.

To select a different time frame

- 1 Click on the entity in the tree.
- 2 In the Time Frame list, select a different time frame than the one already being displayed.
Notice that only the information for the selected entity is refreshed, and it is moved up to the top of the tree.

Refreshing a view

The Refresh option refreshes the display from the PMDB.

To refresh the information displayed

- 1 Right-click on an entity in the tree.
- 2 Click **Refresh**.

Printing the Tree display

You can print an image of what is displayed on the Tree.

To print the Tree display

- 1 Right-click on an entity in the tree.
- 2 Click **Print**.
- 3 Use the print functions of your browser to print out an image of the information displayed on the Tree.

About the Details pane in the Activity tab

The Details pane displays information corresponding to the entity (category) or instance selected in the Tree pane. If an entity (category) is selected, the Details pane displays a Comparison view of all the instances displayed under the selected category in the Tree pane. If an item is selected, the Details pane displays detailed information regarding the selected item.

The Details pane displays information corresponding only to the selected tree branch (displayed in the Tree pane). For example, if you drill down on the Programs category associated to a selected Database Partition item, the Details pane displays only information for the Programs under the selected Database Partition item. The selection you make is reflected in the Details pane only; the Tree display remains unchanged.

Note: When drilling down in the tree, we filter according to the parent entity. There are two exceptions to this rule: drilling down from Tablespaces to Objects and drilling down from Objects to other entities. For example, drilling down from Objects to Statements does not filter Statements in DB2.

For some entities, tabs at the top of the Details pane enable you to view additional information. For example, the tabs displayed for a Statement entity are: General and In DB2. Clicking a tab displays different table columns with the same type of associated entities.

When database files are stored on an EMC Storage unit and the Precise for Storage agent is installed, a tab named Storage is displayed. The tab presents a list of storage device used by the DB Partition entity. For each storage device it displays statistics regarding Cache Hit Ratio, Containers and Table Spaces, Requests amounts, Throughput, and I/O rate. See [“Storage tab”](#) on page 46.

When information on a selected entity is displayed in table and graph format in the Details pane, there is a direct correlation between the selected entry in the table and the information displayed in the graph below the table. Selecting a different row changes the information displayed in the graph displayed below the table.

You can also launch a different tab altogether. For example, clicking the Tune icon next to a statement will take you to the SQL tab, in context, to allow detailed investigation and tuning of the SQL statement.

About the entities you can examine

The Activity tab displays information on different entities. This section provides an overview of all entities, their meaning, and their views.

See [“The Precise for DB2 Time Model”](#) on page 12.

The following entities can be examined in the Activity tab:

- [About All DB2 Databases](#)
- [About Database Partitions](#)

- [About Databases](#)
- [About Programs](#)
- [About Statements](#)
- [About Objects](#)
- [About Machines](#)
- [About Modules](#)
- [About Actions](#)
- [About OS Users](#)
- [About DB Users](#)
- [About Tablespaces](#)
- [About Bufferpools](#)

About All DB2 Databases

DB2 instances aggregate information from all the databases.

The following table describes the information displayed in the Over Time tab.

Table 5-1 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	Application overall activity over time
Sessions	Number of sessions over time
Compile and Recompile	The time the application spent for compiles and recompiles over time
Commits and Rollbacks	The time the application spent for commits and rollbacks over time.

The following table describes the information displayed in the Overview tab.

Table 5-2 Overview tab

Item	Description
Details	Summed information for the chosen time frame of Total In db2, Sessions, etc.
In DB2 (Summed)	The summed In DB2 breakdowns.
Statistics	Summed statistics information such as locks escalations, locks timeouts, etc

The following table describes the information displayed in the Statistics tab.

Table 5-3 Statistics tab

Item	Description
Rows Read/Written	The number of rows read and written over time.
Rows Selected/Affected	The number of rows selected and rows affected over time
Bufferpool Physical I/O	Pages read over time
Locks	Locks escalations and locks timeouts over time.

About Database Partitions

Database partitions list all the information taking place in the database partitions.

Database Partition Comparison

The following table describes the information displayed in the General tab.

Table 5-4 General tab

Column	Description
Database Partition	The name of the database: <instance>.<database>.<partition#>
Catalog Server	The server of the catalog partition.
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement in DB2 (Avg)	The average In DB2 time for the statement
Statement Execution	Total statement executions within the selected time
Sessions	Number of sessions

Note: As pointed out in the [About the Details pane in the Activity tab](#) section, the graph at the bottom of the screen displays detailed information for the Data Partition selected in the table above.

The following table describes the information displayed in the In DB2 tab.

Table 5-5 In DB2 tab

Column	Description
Database Partition	The name of the database: <instance>.<database>.<partition#>
Catalog Server	The server of the catalog partition.
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-6 Statistics tab

Column	Description
Database Partition	The name of the database: <instance>.<database>.<partition#>
Catalog Server	The server of the catalog partition.
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected Database Partition

The following table describes the information displayed in the Over Time tab.

Table 5-7 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in Overview tab.

Table 5-8 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-9 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written

Table 5-9 Statistics tab

Item	Description
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

The following table describes the information displayed in the Storage tab.

Table 5-10 Storage tab

Item	Description
Devices in Use	List of devices used by DB2 database partition. Storage Statistics information is shown for each device. Selecting one of devices will present detailed information about the selected device in other components below.
Device <unit name> <device number>'s Cache Hits	Historical overtime graph of read, write, and combined cache hit ratio (%) of the selected device.
Containers and Tablespaces on Device <unit name> <device number>	List of Containers and Table spaces stored on the selected device.
Device <unit name> <device number>'s Requests	Historical overtime bars graph for read and write requests for the selected device.
Device <unit name> <device number>'s Throughput	Historical overtime graph for Throughput (Mb/Sec) of the selected device.
Device <unit name> <device number>'s I/O Rate	Historical overtime graph for I/O rate (Sec/Req) on the selected device.

About Databases

Databases list activity by database. All database partition information is aggregated to the database level.

Database Comparison

The following table describes the information displayed in the General tab.

Table 5-11 General tab

Column	Description
Database	Name of the database where the statement is running
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-12 In DB2 tab

Column	Description
Database	Name of the database where the statement is running
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-13 Statistics tab

Column	Description
Database	Name of the database where the statement is running
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected Database

The following table describes the information displayed in the Over Time tab.

Table 5-14 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions

Table 5-14 Over Time tab

Item	Description
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-15 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-16 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

The following table describes the information displayed in the Storage tab.

Table 5-17 Storage tab

Item	Description
Devices in Use	Information about devices that are in use
Device <unit name> <device number>'s Cache Hits	The amount of Cache Hits for the selected device
Containers and Tablespaces on Device <unit name> <device number>	Information on the Containers and Tablespaces that are defined for this specific device.
Device <unit name> <device number>'s Requests	The amount of Requests for the selected device
Device <unit name> <device number>'s Throughput	The Throughput for the selected device
Device <unit name> <device number>'s I/O Rate	The I/O Rate for the selected device

About Programs

These are the names of programs that execute Precise for DB2 activity.

Program Comparison

The following table describes the information displayed in the General tab.

Table 5-18 General tab

Column	Description
Program	Program name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-19 In DB2 tab

Column	Description
Program	Program name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-20 Statistics tab

Column	Description
Program	Program name
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written

Table 5-20 Statistics tab

Column	Description
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected Program

The following table describes the information displayed in the Over Time tab.

Table 5-21 Over Time tab

Column	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-22 Overview tab

Column	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-23 Statistics tab

Column	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts overtime

About Statements

These are the statements Precise for DB2 is executing. Statements are collapsed.

Statement Comparison

The following table describes the information displayed in the General tab.

Table 5-24 General tab

Column	Description
Icon	Launch to SQL tab for statement tuning
Statement	Statement name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
In DB2 (Avg)	The average In DB2 time per statement execution
Executions	Total statement executions within the selected time frame. Available for statements, tables, and indexes.
Text	Statement text

The following table describes the information displayed in the In DB2 tab.

Table 5-25 In DB2 tab

Column	Description
Icon	Launch to SQL tab for statement tuning
Statement	Statement name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Wait for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

Selected Statement

The following table describes the information displayed in the Over Time tab.

Table 5-26 Over Time tab

Column	Description
In DB2 (Summed)	The In DB2 breakdowns over time
In DB2 (Avg)	The average In DB2 time per statement execution
Statement Executions	Total statement executions within the selected time

The following table describes the information displayed in the Overview tab.

Table 5-27 Overview tab

Column	Description
Details	Detailed activity information and last explain details
In DB2 (Summed)	The In DB2 breakdowns
Text	Statement text

About Objects

Lists of tables and indexes in use.

Object Comparison

The following table describes the information displayed in the General tab.

Table 5-28 General tab

Column	Description
Icon	Launch to Objects tab
Type	The object type
Object	The object name
Schema	The object schema name
Base Table Name	The name of the table
In DB2 (Summed)	The In DB2 time
%	The In DB2 percentage as a % of total rows
Sequential Prefetch I/O	I/O time used for sequential prefetch
List Prefetch I/O	I/O time used for list prefetch
No Prefetch I/O	I/O time used with no prefetch

Selected Object

The following table describes the information displayed in the Overview tab.

Table 5-29 Overview tab

Column	Description
Details	Object details
In DB2 (Summed)	The In DB2 time
Prefetched I/O (Summed)	Over time information or I/O separated by Sequential prefetch, List prefetch, and No prefetch

About Machines

This is a list of the machines from which the users are logging in.

Machine Comparison

The following table describes the information displayed in the General tab.

Table 5-30 General tab

Column	Description
Machine	Machine name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-31 In DB2 tab

Column	Description
Machine	Machine name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-32 Statistics tab

Column	Description
Machine	Machine name
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written

Table 5-32 Statistics tab

Column	Description
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected Machine

The following table describes the information displayed in the Over Time tab.

Table 5-33 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-34 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-35 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

About Modules

Modules represent the name of the package that was activated.

Module Comparison

The following table describes the information displayed in the General tab.

Table 5-36 General tab

Column	Description
Module	Module name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-37 In DB2 tab

Column	Description
Module	Module name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-38 Statistics tab

Column	Description
Module	Module name
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written

Table 5-38 Statistics tab

Column	Description
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected Module

The following table describes the information displayed in the Over Time tab.

Table 5-39 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-40 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-41 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

About Actions

These are the account strings that were activated from the client information.

Action Comparison

The following table describes the information displayed in the General tab.

Table 5-42 General tab

Column	Description
Action	Action name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-43 In DB2 tab

Column	Description
Action	Action name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-44 Statistics tab

Column	Description
Action	Action name
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written

Table 5-44 Statistics tab

Column	Description
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected Action

The following table describes the information displayed in the Over Time tab.

Table 5-45 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-46 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-47 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

About OS Users

These are the names of the users that logged into the machine and executed activity in DB2.

OS User Comparison

The following table describes the information displayed in the General tab.

Table 5-48 General tab

Column	Description
OS User	OS User name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-49 In DB2 tab

Column	Description
OS User	OS User name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-50 Statistics tab

Column	Description
OS User	OS User name
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written

Table 5-50 Statistics tab

Column	Description
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected OS User

The following table describes the information displayed in the Over Time tab.

Table 5-51 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-52 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following

Table 5-53 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

About DB Users

These are the names of the users that logged into the DB2 database.

DB User Comparison

The following table describes the information displayed in the General tab.

Table 5-54 General tab

Column	Description
DB User	Database User name
In DB2 (Summed)	The In DB2 breakdowns
%	The In DB2 percentage as a % of total rows
Statement In DB2 (Avg)	The average In DB2 time for the statement
Statement Executions	Total statement executions within the selected time
Sessions	Number of sessions

The following table describes the information displayed in the In DB2 tab.

Table 5-55 In DB2 tab

Column	Description
DB User	Database User name
In DB2 (Summed)	The In DB2 breakdowns overtime
Using CPU	In DB2 breakdown description
Waiting for CPU	In DB2 breakdown description
Sync I/O Wait	In DB2 breakdown description
Direct I/O Wait	In DB2 breakdown description
Prefetch Wait	In DB2 breakdown description
Lock Wait	In DB2 breakdown description
I/O Error Wait	In DB2 breakdown description
Remote Query Wait	In DB2 breakdown description
Utility Active	In DB2 breakdown description

The following table describes the information displayed in the Statistics tab.

Table 5-56 Statistics tab

Column	Description
DB User	Database User name
Lock Escalations	Total number of lock escalations
Lock Timeouts	Total number of lock timeouts
BP Data Read I/O	Total data reads (in pages)
BP Index Read I/O	Total index reads (in pages)
Direct Read/Write I/O	Total of direct reads and writes (in pages)
Rows Read	Number of rows read
Rows Written	Number of rows written

Table 5-56 Statistics tab

Column	Description
Rows Selected	Number of rows selected
Rows Affected	Number of rows affected
Internal Row Affected	Number of internal rows affected

Selected DB User

The following table describes the information displayed in the Over Time tab.

Table 5-57 Over Time tab

Item	Description
In DB2 (Summed)	The In DB2 breakdowns over time
Overall Activity (Summed)	The application's overall activity including In DB2 over time
Sessions	Number of sessions
Compile and Recompile	The time the application spent for commits & recompiles over time
Commits and Rollbacks	The time the application spent for commits & rollbacks over time

The following table describes the information displayed in the Overview tab.

Table 5-58 Overview tab

Item	Description
Details	Detailed activity information
In DB2 (Summed)	The In DB2 breakdowns
Statistics	Statistic counters overview

The following table describes the information displayed in the Statistics tab.

Table 5-59 Statistics tab

Item	Description
Rows Read/Written	Number of rows read/written
Rows Selected/Affected	Number of rows selected/affected
Bufferpool Physical I/O	Physical I/O over time (in pages)
Locks	Number of lock escalations and lock timeouts over time

About Tablespaces

Tablespaces list activity by Tablespace.

This entity is only available for the Database and Database Partition entities. To display activity by tablespace, drill from the Database or Database Partition to the tablespaces entity.

The following table describes the information displayed in the General tab.

Tablespace Comparison

Table 5-60 General tab

Column	Description
Tablespace	The name of the tablespace
I/O Time (Summed)	Total I/O time breakdown
I/O Count	Total I/O count breakdown (in pages)
Page Size	The tablespace's page size
Pages	The number of pages
No Victim Available	The number of times an agent did not have a pre-selected victim buffer available

The following table describes the information displayed in the I/O Time tab.

Table 5-61 I/O Time tab

Column	Description
Tablespace	The name of the tablespace
I/O Time (Summed)	Total I/O time breakdown
Synchronous Read Time	The total Synchronous Read Time
Asynchronous Read Time	The total Asynchronous Read Time
Direct Read Time	The total Direct Read Time
Synchronous Write Time	The total Synchronous Write Time
Asynchronous Write Time	The total Asynchronous Write Time
Direct Write Time	The total Direct Write Time

The following table describes the information displayed in the I/O Count tab.

Table 5-62 I/O Count tab

Column	Description
Tablespace	The name of the tablespace
I/O Count	The total I/O Count
Synchronous Data Reads	The total Synchronous Data Reads
Synchronous Index Reads	The total Synchronous Index Reads
Asynchronous Data Reads	The total Asynchronous Data Reads
Asynchronous Index Reads	The total Asynchronous Index Reads
Direct Reads	The total Direct Reads
Synchronous Data Writes	The total Synchronous Data Writes
Synchronous Index Writes	The total Synchronous Index Writes
Asynchronous Data Writes	The total Asynchronous Data Writes
Asynchronous Index Writes	The total Asynchronous Index Writes
Direct Writes	The total Direct Writes

The following table describes the information displayed in the Response Time tab.

Table 5-63 Response Time tab

Column	Description
Tablespace	The name of the tablespace
Synchronous Read Response Time	The average Synchronous Read Response Time
Asynchronous Read Response	The average Asynchronous Read Response
Direct Read Response Time	The average Direct Read Response Time
Synchronous Write Response Time	The average Synchronous Write Response Time
Asynchronous Write Response	The average Asynchronous Write Response
Direct Write Response Time	The average Direct Write Response Time

Selected Tablespace

The following table describes the information displayed in the Over Time tab.

Table 5-64 Over Time tab

Item	Description
Hit Ratio	The tablespace Hit Ratio over time
I/O Time	The tablespace I/O time breakdown over time
Response Time (Avg)	The tablespace average Response Time over time
I/O Count	The tablespace I/O Count breakdown over time (in pages)

The following table describes the information displayed in the Overview tab.

Table 5-65 Overview tab

Item	Description
Details	Details about the selected tablespace
I/O Time	Total I/O time breakdowns (in Pie chart)
I/O Count	Total I/O Count breakdowns (in Pie chart)

About Bufferpools

Bufferpools list activity by Bufferpool.

This entity is only available for the Database and Database Partition entities. To display activity by bufferpool, drill from the Database or Database Partition to the Bufferpools entity.

Bufferpool Comparison

The following table describes the information displayed in the General tab.

Table 5-66 General tab

Column	Description
Bufferpool	The name of the bufferpool
I/O Time (Summed)	Total I/O time breakdown
I/O Count	Total I/O count breakdown (in pages)
Page Size	The bufferpool's page size
Pages	The number of pages
No Victim Available	The number of times an agent did not have a pre-selected victim buffer available

The following table describes the information displayed in the I/O Time tab.

Table 5-67 I/O Time tab

Column	Description
Bufferpool	The name of the bufferpool
I/O Time (Summed)	Total I/O time breakdown
Synchronous Read Time	The total Synchronous Read Time
Asynchronous Read Time	The total Asynchronous Read Time
Direct Read Time	The total Direct Read Time
Synchronous Write Time	The total Synchronous Write Time
Asynchronous Write Time	The total Asynchronous Write Time
Direct Write Time	The total Direct Write Time

The following table describes the information displayed in the I/O Count tab.

Table 5-68 I/O Count tab

Column	Description
Bufferpool	The name of the bufferpool
I/O Count	The total I/O Count
Synchronous Data Reads	The total Synchronous Data Reads
Synchronous Index Reads	The total Synchronous Index Reads
Asynchronous Data Reads	The total Asynchronous Data Reads
Asynchronous Index Reads	The total Asynchronous Index Reads
Direct Reads	The total Direct Reads
Synchronous Data Writes	The total Synchronous Data Writes
Synchronous Index Writes	The total Synchronous Index Writes
Asynchronous Data Writes	The total Asynchronous Data Writes
Asynchronous Index Writes	The total Asynchronous Index Writes

Table 5-68 I/O Count tab

Column	Description
Direct Writes	The total Direct Writes

The following table describes the information displayed in the Response Time tab.

Table 5-69 Response Time tab

Column	Description
Bufferpool	The bufferpool name
Synchronous Read Response Time	The average Synchronous Read Response Time
Asynchronous Read Response	The average Asynchronous Read Response
Direct Read Response Time	The average Direct Read Response Time
Synchronous Write Response Time	The average Synchronous Write Response Time
Asynchronous Write Response	The average Asynchronous Write Response
Direct Write Response Time	The average Direct Write Response Time

Selected Bufferpool

The following table describes the information displayed in the Over Time tab.

Table 5-70 Over Time tab

Item	Description
Hit Ratio	The bufferpool Hit Ratio over time
I/O Time	The bufferpool I/O time breakdown over time
Response Time (Avg)	The bufferpool average Response Time over time
I/O Count	The bufferpool I/O Count breakdown over time (in pages)

The following table describes the information displayed in the Overview tab.

Table 5-71 Overview tab

Item	Description
Details	Details about the selected bufferpool
I/O Time	Total I/O time breakdowns
I/O Count	Total I/O Count breakdowns

Examining the relationship between database objects

This section includes the following topics:

- [About the Objects tab](#)
- [How the Objects tab is structured](#)
- [About Objects tab entities](#)

About the Objects tab

This tab enables you to see objects definitions, I/O activity, and space usage over time. This section provides a general description of how to use the Objects tab in your analysis of DB2 objects. For further information on a specific DB2 object, see DB2's relevant documentation.

How the Objects tab is structured

The Objects tab displays information on a selected entity and its associated entities. When you open the Objects tab the selected entity is by default Database Partition, meaning that information is displayed on the Database Partition level. When you launch the Objects tab from the Activity tab with a Table or Index entity, using the launch icon, the tab is launched in context with the entity you selected in the Activity tab.

If you open the Objects tab from another tab, the historical settings (those settings which were selected when you left the tab, such as the last entity you drilled down to) are taken into account and the information displayed the last time you viewed the tab is displayed (similar to clicking the History button and returning to a previous tab).

The selected entity is always reflected in the Tab heading, which serves as a point of orientation. The highest level entity you can view information for in the Objects tab is the Data Partition level. You can select a Data Partition from the Data Partition list.

The entities displayed in the Association area are associated with the selected entity displayed in the Main area. At times, the relationship between the entity displayed in the Main area and those displayed in the Association area is that of parent to child and sometimes it merely represents that there is a relationship between the selected entity and the entities displayed in the Association area.

About the Main area in the Objects tab

The Main area displays overview information for the entity selected in the Association area as well as its related entity. By default, the Objects tab opens displaying information on the Database entity and its last selected related entity. You can change the database to display by selecting the desired database in the Database Partition field at the top of the Main area and its related entity by clicking the Association Control button at the top of the Association area. You can also display information on a different entity by selecting the desired entity from the Association Control button list and the related entity by selecting it from the table displayed in the Association area.

About the Association area in the Objects tab

The Association area is comprised of the Association Control button for selecting the entity whose information you want to display and a table displaying the selected entity's related information. Two entities, Tables and Indexes, display information in two views: the Dictionary tab and the Storage & Statistics tab. The Dictionary tab displays dictionary information for this entity and the Storage & Statistics tab displays statistics based on the last runstat information for this entity.

About Objects tab entities

The Objects tab displays information on different entities. This section provides an overview of all entities, their meaning, and their views.

The following entities can be examined in the Objects tab:

- [Schemata](#)
- [Tables](#)
- [Indexes](#)
- [Views](#)
- [Database Partition Groups](#)
- [Tablespaces](#)
- [Bufferpools](#)
- [Database Partitions](#)

Schemata

The following table describes the information displayed for the Schemata entity.

Table 6-1 Schemata

Column	Description
Schema	Schema name
Owner	Owner name
Definer	The user who defined the schema
Creation Time	Time of schema creation

Tables

Two association selector options are available:

- Dictionary tab- shows dictionary information for the table.
- Storage & statistics tab- shows table storage information and RUNSTATS information. This is the default view.

The following table describes the information displayed in the Dictionary tab.

Table 6-2 Dictionary tab

Column	Description
Name	Table name
Schema	Table's schema
Type	Temporary, MDC, MQT or normal

Table 6-2 Dictionary tab

Column	Description
Tablespace	Name of primary tablespace
Statistics Run	Date last analyzed
Range Partitioned	Yes – for Range Partitioned table
DB Partitions	The number of DB partitions this table exists in
Indexes	The number of indexes defined for this table
Logged	No or blank
Lock Mode	The Lock Mode type
Index Tablespace	Name of the tablespace which contains the indexes
LOB Tablespace	Name of the tablespace which contains the LOBs

The following table describes the information displayed in the Storage and Statistics tab.

Table 6-3 Storage and Statistics tab

Column	Description
Name	Table name
Schema	Table's schema
Type	Temporary, MDC, MQT or normal
Tablespace	Name of primary tablespace
Statistics Run	Date last analyzed
Pages Allocated	How many pages are allocated
Pages Used	How many pages contain actual data
Rows	How many rows
Ave. Row Len	Average row length (bytes)
Clustering Factor	The clustering factor of the clustering index (if one exists) or blank if no clustering index was defined
Extent size	Size of extent (from tablespace definitions)
% Free Pages	% Free of pages
% Overflows	% of records overflowing
Rows per Page	Rows / (# of used pages)

Indexes

The following table describes the information displayed in the Dictionary tab.

Table 6-4 Dictionary tab

Column	Description
Name	Index name
Schema	Index schema

Table 6-4 Dictionary tab

Column	Description
Table Name	Base table name
Key Columns	The number of participating columns in the index
Includes	The number of columns included in the index
Type	Clustering, Dimensions block, Block or blank for regular
Unique	Yes, No, or Primary
Columns	The list of participating columns in the index
Statistics Run	Time last analyzed
Tablespace	Name of the tablespace which contains the index
Table Schema	The table's schema

The following table describes the information displayed in the Storage and Statistics tab.

Table 6-5 Storage and Statistics tab

Column	Description
Name	Index name
Schema	Index schema
Tablespace	Name of the tablespace which contains the index
Statistics Run	Time last analyzed
Clustering Factor	Clustering factor or Cluster ratio if factor is not available
Levels	How many levels exist
First Column Distincts	The number of distinct first-key values
Second Column Distincts	The number of distinct keys using the first two columns
Distinct Values	The number of distinct values
Leafs	How many leaf pages are in use
Key Length	The key length
% Deleted Rows	The percentage of rows marked 'Deleted'
Leaf Distribution	Avg. length between leaf pages
% Empty Pages	Percentage of empty pages
% Free	Percentage of each index page to be reserved during initial index building
Extent Size	Taken from tablespace definitions

Views

The following table describes the information displayed for the Views entity.

Table 6-6 Views

Column	Description
View Name	View name

Table 6-6 Views

Column	Description
Owner	View owner
Definer	Name of view creator
Valid	Is the view valid
Text	Unformatted first N characters of view's definition

Database Partition Groups

The following table describes the information displayed for the Database Partition Groups entity.

Table 6-7 Database Partition Groups

Column	Description
Name	Name of the partition group
Partitions	Number of partitions in this group
Definer	Name of the user who defined the group
Create Time	Time group was created

Tablespaces

The following table describes the information displayed for the Tablespaces entity.

Table 6-8 Tablespaces

Column	Description
Name	Tablespace name
Type	DMS or SMS
Data Type	All, LOB, System temp or Declared temp
Bufferpool	Name of the bufferpool associated with this tablespace
Bufferpool Size	Size of bufferpool on all partitions
Page Size	Page size
Pages Allocated	Total tablespace size
Pages Used	Total tablespace used (only for DMS)
Extent Size	Number of pages per extent
Prefetch Size	Defined prefetch size or Automatic
Defined in Group	The name of the partitioned group
Overhead	DB2 internal parameter describing the overhead associated with the tablespace's disks
Transfer Rate	DB2 internal parameter, describing the time it takes to read one page
Definer	Name of the user who defined the tablespace
Remarks	User comments

Bufferpools

The following table describes the information displayed for the Bufferpools entity.

Table 6-9 Bufferpools

Column	Description
Name	Bufferpool name
Size (Pages)	Bufferpool size
Page Size	Page size of the bufferpool
Defined in Groups	Comma separated list of the DB partition groups to which it belongs
All Tablespaces Size	Size, in MB, of all tablespaces associated with this bufferpool
Tablespaces	The number of tablespaces defined to use this bufferpool

Database Partitions

The following table describes the information displayed for the Database Partitions entity.

Table 6-10 Database Partitions

Column	Description
Server	Server name
ID	ID of the partition
Default on Server?	Yes if this is the default partition when connecting to the database on that machine
Total Memory	Total shared memory dedicated for this database partition
Total Bufferpools Size	Total shared memory dedicated for this partition's BPs
Defined in Groups	Comma separated list of DB partition groups in which this DB partition participates

Analyzing Statement plans

This section includes the following topics:

- [About the SQL tab](#)
- [How the SQL tab is structured](#)
- [About tuning actions](#)

About the SQL tab

The SQL tab lets you analyze execution plans and explain results so that you can tune statements and achieve optimal results.

The process of explaining statements is a prerequisite for tuning. The explain process is designed to clarify the access path chosen for a statement and translate it into a visual medium. The SQL tab helps you identify bottlenecks and guides you through the steps required to tune SQL statements.

After a statement is explained, the explain results are stored in the PMDB. This information includes the objects referenced by the statement and the operations performed on these objects. The top statements are automatically explained every day. You can control how many statements to explain using a setting for the Explain Statements PMDB process in AdminPoint. See the *Precise Agent Installer*/*Precise Administration Guide* for details.

Understanding the execution plan chosen by the DB2 Optimizer is extremely important when tuning your application. You can ensure optimal system performance by ensuring that the best plans are used for your queries.

Precise for DB2 provides you with a special tab only for this purpose, the SQL tab.

To analyze an execution plan, Precise for DB2 provides you with a full picture of the objects (table, indexes, and so on) participating in the plan. You can change a plan by modifying your query or changing the schema, for example by adding an index. To assist in this, Precise for DB2 can provide you with index recommendations for your statement.

One technique to identify the source of the problem is to view the historical information of statements, showing performance degradation. A change in the schema, volume, or execution plan may explain the impact on performance.

The following table shows from which tabs and entities it is possible to launch to the SQL tab, in context

Table 7-1 Launch in-context to following entities

Tab	Entities
Dashboard	Statement
Current	Statement, Active or Current Session that is currently executing a statement.
Activity	Statement
What-If	Statement that is affected by one or more index evaluations.

Note: If the SQL tab is opened with no statement in-context, the following error message is displayed:

No statement in context. Please open a statement.

In the Current tab, you can either launch the SQL tab by clicking the SQL tab button when the selected entity in the Main area is a statement, or by selecting a statement in the Association area and clicking the Tune icon that appears before the SQL text.

In the Activity tab, you can either launch to the SQL tab by clicking the Tune icon in the first column of the table on the Details pane, or by clicking the Tune hyperlink at the end of a statement row on the Tree pane.

If the SQL tab is opened with no statement in context, a message prompts you to open a statement.

If Precise has already collected the steps of a real plan, it appears highlighted in blue; otherwise, it appears highlighted in gray. If you click on a gray statement, Precise for DB2 will access the monitored database partition and try to retrieve the selected plan's steps.

There can be statements with many execution plans.

See [“About the Dashboard tab”](#) on page 23.

See [“About the Current tab”](#) on page 30.

See [“About the Activity tab”](#) on page 39.

See [“About the Objects tab”](#) on page 67.

How the SQL tab is structured

The following sections explain how to view information and navigate through the SQL tab:

- [About viewing the execution plan of a statement](#)
- [About interpreting text highlighting on the Highlights tab](#)
- [About viewing which objects are referenced by the execution plan](#)
- [About viewing information on all operators and predicates in the Execution plan](#)
- [About viewing the statement text](#)
- [About viewing recommended indexes for a statement](#)
- [About viewing additional execution plan details](#)

About viewing the execution plan of a statement

The Plan Tree area displays the execution plan of a statement and related information, such as statistics, referenced objects, and operations performed in the execution plan. It lets you assess and tune statements based on real information.

Two types of execution plans are available when launching to the SQL tab:

Virtual execution plan	<p>The plan which was collected from DB2 V\$ tables. Additional real-time statistics are available for these plans. For example, you can view a breakdown of In DB2 time for a plan as compared to other plans, or a breakdown of its I/O when accessing DB2 objects.</p> <p>When launching to the SQL tab with a specific SQL statement in context, Precise for DB2 presents actual execution plan information, including actual plan steps, and information and statistics for multiple plans. For example, if a statement has several different execution plans, all are displayed.</p>
Estimated execution plan	<p>Depending upon the version of DB2 running on your system, you may be only able to generate and run an estimated execution plan. At times, real execution plan information is lost (as when they are removed from the system before Precise for DB2 could access them). If this happens, all real-time statistics (such as, I/O, and In DB2) are considered to belong to other plans.</p>

The Plan view is divided into three panes. The Plan Tree area and the Text View area are displayed in the left pane. The Details area is displayed in the right pane. The information displayed in the Details area is controlled by the information tabs (Highlights, Objects, Operators, Text, and Recommendations) that are located at the top of this area.

The following table describes the information displayed in each tab of the Plan view.

Table 7-2 Information tabs

Information Tabs	Description
Highlights	The default tab, it displays the statement's Optimized text. When opening a new plan tree or when clicking on the root of the tree, the original text is displayed (Text tab). When an operator is clicked, the text will be switched to the optimized text of the statement.
Objects	Displays information on all referenced objects in the execution plan, including their tables, indexes, and columns. Statistical and general details are displayed for each object.
Operators	Displays information on all operators in the Execution plan tree. When you sort a table, you can locate operators by their total cost or by any other column in the table.
Text	Displays the Original text of the statement.
Recommendations	Displays new index recommendations for the statement. These are additional indexes that can improve the statement. Note: This tab is only displayed after you generate Recommendations (by clicking the Recommend button or clicking Actions>Recommend).

About the Execution Plan tree

The Execution Plan tree displays both the Virtual execution plan and the Estimated execution plan, of the specified statement.

The execution plan's text appears both in the Text View area (Text tab) at the bottom of the execution plan tree and at the top of the Details area (also a Text tab). The text that relates to the selected step in the Plan tree is highlighted. This lets you view the text of the statement, the execution plan, and additional information, such as the objects referenced by the statement, all at once. In the Text View area you can also view two additional tabs displaying Optimized Text and Remote Query Text.

Actions that can be performed on the tree

You can perform the following actions on the execution plan tree:

- The Playback controls, located at the top of the tree, enable you to freely move within the execution plan of an explained statement (step by step).
- By moving the pointer over the execution plan steps, you can view a ToolTip that contains statistical information, such as Estimated Cost, Estimated Rows, and Estimated I/O Cost, and a description of the specified step.
- Selecting a specific step (operator) highlights nested steps and affects the information displayed in the Detail tabs: the relevant Optimized text is shown on the Highlights tab; the referenced tables, indexes, and columns are shown on the Objects tab; the operators and predicates are shown on the Operators tab; and the Original text is shown on the Text tab. The Recommendations tab, only displayed when you generate recommendations, shows index details and the DDL for creating the indexes.
- Clicking the hatchet mark (☒), located at the bottom of the tree, displays the Tree View area with three tabs allowing you to select which statement text to view:
 - Optimized Text
 - Original Text
 - Remote Query Text

Note: When launching to the SQL tab, the Highlights tab in the Details area and the Text tab in the Text View area are selected by default.

About interpreting text highlighting on the Highlights tab

The Highlights tab displays the statement's Optimized text as generated by the DB2 optimizer. Depending upon the operator selected on the Tree, the text will have different colors highlighting relevant tables, columns, and predicates.

When an SQL statement is explained, parts of the statement's text are highlighted for each branch of the Explain tree. Highlighting specific parts of the statement's text can help you understand the access path chosen by the DB2 optimizer. For example, color highlighting references to a table and its columns enables you to easily identify tables and distinguish between the various tables accessed by the statement.

Predicate text is highlighted and underlined. Referenced columns in the predicate text are highlighted in blue or red. Non-column predicate text (e.g., constants, parameter markers, logical operators, functions, etc.) is highlighted in green (i.e., range delimiting predicate) or in pink (i.e., non-range delimiting predicate).

Only applied predicates are highlighted. In the following example, the predicate on the bid column is range delimiting

The screenshot shows the DB2 Explain interface. The title bar indicates the session ID '00789.55258.35576.55499 (Running on DB2.TEST(db-vm-db2))'. The main window is titled 'Virtual Execution Plan as of Nov 25, 2008 11:12'. On the left, a tree view shows the execution plan steps, with the final step being an index scan on 'I4SQL_PL_B_BRANCHES_BID'. On the right, the 'Highlights' tab is active, displaying the 'Optimized Text' of the query. The text shows an update statement followed by a select statement with a range predicate on the 'bid' column: 'select (q2.bbbalance + +7.088050000000000e+005) from i4sql.b_branches as q2 where (q2.bid = 0)'. The predicate text is highlighted in pink, and the column name 'bid' is highlighted in blue.

Note: Not all operators include color fonts. Color fonts are used only when relevant for performance tuning.

About viewing which objects are referenced by the execution plan

The Objects tab displays three tables (Tables used in Plan, Indexes on Table, Columns in Table) that list all referenced objects in the execution plan, including their indexes and columns. Statistical details and general details are displayed for each object and its sub-entities.

The following table describes the information displayed for the referenced tables.

Table 7-3 Tables used in Plan



Column	Description
	Click to launch to Objects tab—with the selected index in context—to get additional catalog information about the object.
	Click to locate and highlight in Bold font all the operators in the execution plan that access the specified index.
Table Name	Name of the table.
Contribution Ratio	The Contribution ratio of the table, computed by I/O Cost.
Prefetch	The Prefetch type (LIST, SEQUENTIAL, or NONE).
Statistics Time	The last statistics update.
Rows	The number of rows based on last statistics update.
Total Pages	The total pages based on last statistics update.

Table 7-3 Tables used in Plan

Column	Description
Data Pages	The data pages based on last statistics update.
Overflow Records	Overflow records based on last statistics update.
Preferred Locksize	Preferred locksize based on last statistics update.



The following table describes the information displayed in the indexes defined on a selected table.

Table 7-4 Indexes on Table.

Column	Description
	Click to launch to Objects tab—with the selected index in context—to get additional catalog information about the object.
	Click to locate and highlight in Bold font all the operators in the execution plan that access the specified index.
Index Name	Name of the index.
Contribution Ratio	The Contribution ratio of the table, computed by I/O Cost.
Prefetch	The Prefetch type (LIST, SEQUENTIAL, or NONE).
Statistics Time	The last statistics update.
Unique Rule	Unique rule based on last statistics update.
Cluster Ratio	Cluster ratio based on last statistics update.
Cluster Factor	Cluster factor based on last statistics update.
Index Leafs	Index leafs based on last statistics update.
Index Levels	Index levels based on last statistics update.
Distinct First Keys	Distinct first key based on last statistics update.
Distinct Full Keys	Distinct full keys based on last statistics update.
Sequential Pages	Sequential pages based on last statistics update.
Density	Density based on last statistics update.

The following table describes the information displayed in Columns in Table.

Table 7-5 Columns in table


Column	Description
	Click to launch to Objects tab—with the selected table in context—to get additional catalog information about the object.
	The column definition in the chosen index (Ascending, Descending, or Included).
Column	Column name.
Type	Data type of the column.
Length	The length of the column.
Key Number	The position within the chosen index.

About viewing information on all operators and predicates in the Execution plan

The Operators tab displays information on all operators in the Execution plan tree. When you sort a table, you can locate operators by their total cost or by any other column in the table.

The following table describes the information displayed in the Operators table.

Table 7-6 Operators table

Column	Description
	Click to locate and highlight in Bold font the same operator in the execution plan.
Operator ID	Operator ID issued by the DB2 Optimizer.
Execution Order	The execution order of operators in the plan.
Operator	Operator type.
Subtree Total Cost	The total cost of this operator subtree.
Cost(%)	The percentage cost of this operator in the plan.
First Row Cost	First row cost of the operator subtree.
I/O Cost	I/O Cost.
CPU Cost	CPU Cost.
Prefetch	Prefetch type.
Estimated Rows	The estimated rows of the operator.
Estimated Pages	The estimated pages of the operator.
Potential Filter Factor	Potential filter factor calculated by the predicate's filter factor of the operator.
Table Partitions	The estimated number of table partitions to be accessed.

The following table describes the information displayed in the Predicates table.

Table 7-7 Predicates table

Column	Description
Text	Predicate text.
Filter Factor	The estimated filter factor of the predicate.
How Applied	How predicate is being used by the specified operator.
Sub Select	Indicates when the subquery used in this predicate is evaluated.

About viewing the statement text

The Text tab displays the statement's original text.

About viewing recommended indexes for a statement

The Recommendations tab displays the output of the Recommend algorithm. The Recommendations tab is divided into three areas:

- Recommendations table
- Details area
- DDL area

The following table describes the information displayed in the Recommendations table.

Table 7-8 Recommendations table

Column	Description
Object	The name of the index.
Actions	Create index.
Columns	Columns participating in the index or Include Columns in the index.

The Details area displays statistical information of the recommended index.

The DDL area displays the create index statement.

About viewing additional execution plan details

You can view additional execution plan details by hovering over a specific execution plan row with the mouse. A ToolTip is displayed showing some of the following details, depending upon which row you are hovering over:

- Operator Type
- Description
- Estimated Total Cost
- First row cost
- Estimated I/O cost
- Estimated CPU cost
- Estimated communication cost
- Contribution Ratio
- Estimated total cost
- Operator Arguments such as:
 - MAXPAGES=ALL
 - PREFETCH=NONE
 - ROWLOCK=EXCLUSIVE
 - SCANDIR=FORWARD
 - TABLOCK=INTENT EXCLUSIVE

About tuning actions

Use the Actions menu to perform the following operations:

- [Creating a new statement](#)
- [Opening an existing statement](#)
- [Editing an existing statement](#)
- [Re-explaining an access plan](#)
- [Generating recommendations](#)

Creating a new statement

You can create a new statement and save it in the PMDB in a logical cabinet and folder hierarchy. You can also rewrite statements and view their access plans.

To create a new statement

- 1 On the Actions menu, click **New**.
- 2 In the New Statement dialog box, choose the relevant properties for the new statement from the drop-down lists and enter a statement name, as follows:

Instance	The name of the instance that the statement belongs to.
User	The statement's parsing user.
Cabinet	The name of the cabinet that the statement is saved in (creates a cabinet if it does not already exist).
Folder	The name of the folder that the statement is saved in (creates a folder if it does not already exist).
Statement	The name of the statement.
Comment	User free text.

- 3 Type the SQL text for the statement in the Text frame.
- 4 Click **OK**.

The new statement is saved in the PMDB in a logical cabinet and folder hierarchy.

Opening an existing statement

You can load a previously saved statement.

To open an existing statement

- 1 On the Actions menu, click **Open**. The Open Statement dialog box is displayed.
- 2 In the Open Statement dialog box, select the statement you want to open by choosing its relevant properties from the drop-down lists, as follows:

Database	The name of the database where the statement is running.
Cabinet	The name of the cabinet that the statement is saved in.
Folder	The name of the folder that the statement is saved in.
Statement	The name of the statement.
Comment	Displays free text entered by the user.

The statement's text is displayed in the Text frame.

- 3 Click **OK**.

Note: If the SQL tab is opened with no statement in-context, the following error message is displayed:

No statement in context. Please open a statement.

To open a statement you can close the error message dialog and click the **Open** button displayed on the upper right-hand side of the Details area or click **Actions>Open**. To open a new statement click **Actions>New**.

Editing an existing statement

You can edit a statement that was saved in the PMDB. The statement that was collected by the Collector is saved as a new statement in a logical cabinet and folder hierarchy.

To edit an existing statement

- 1 On the Actions menu, click **Edit**.
- 2 In the Edit Statement dialog box, select the statement you want to edit by choosing its relevant properties from the drop-down lists, as follows:

Database	The name of the database where the statement is running.
User	The statement's parsing user.
Cabinet	The name of the cabinet that the statement is saved in.
Folder	The name of the folder that the statement is saved in.
Statement	The name of the statement.
Comment	Displays free text entered by the user.

- 3 Enter a new statement name in the Statement field.
- 4 Edit the SQL text for the statement in the Text frame.
- 5 Click **OK**.

The edited statement is saved in a logical cabinet and folder hierarchy.

Re-explaining an access plan

You can refresh a statement plan by generating a new estimated access plan and saving it, if it is different from the most recent access plan. You can then display the most recent access plan.

To re-explain a statement

- 1 Go to **Actions>Re-Explain** or, alternatively, click the **Re-Explain** button.

Generating recommendations

You can generate new index recommendations for the statement. These are additional indexes that can improve the statement.

To generate a recommendation

- 1 Go to **Actions>Recommend** or, alternatively, click the **Recommend** button.

Simulating the effect of changes

This section includes the following topics:

- [About the What-If tab](#)
- [How the What-If tab is structured](#)
- [How the What-If tab can help you evaluate changes](#)

About the What-If tab

The What-If tab allows you to simulate the effect of creating an index on the performance of all statements that Precise for DB2 has collected, without physically making the change.

The What-If tab is intended to be used in conjunction with the SQL tab. The normal sequence is to examine any index recommendations made on an individual SQL statement on the Recommendation tab in the SQL tab and click the Quick Link to launch the What-If tab, with the index recommendation as the suggested change. The What-If tab may also be used as a standalone tab where you can simulate the effect of your own index changes.

What-If scans statements in the PMDB. Precise for DB2 can accurately calculate the revised cost of a statement by generating a new explain plan. Statements that have a measurable change in cost are marked as definitely affected. However, even if the revised cost is the same as before, there may still be an impact on response time. For example, Precise for DB2 does not change the cost of a DML statement to cater for the overhead of maintaining extra indexes. Such statements are marked as possibly affected because the effect is immeasurable.

As may be expected, adding an index may improve the response times of select statements that can use the new index but decrease the response time of insert, update, and delete statements that have to maintain the extra index. However, the effect will be immeasurable because usually the cost of these affected statements will not be changed. To help you make an informed choice about the desirability of the change, the What-If tab also displays the previous In DB2 resource consumption and execution figures for the statement.

Having selected an impacted statement, you can launch the SQL tab in context, to compare the statement's execution plan before and after the simulation.

You should bear in mind the following things when using the What-If tab:

- It is important that statistics are up to date so that the simulation is as accurate as possible.
- What-If is only as good as the data it has to work on. It is dependent on statements being loaded into the PMDB and explained by the "Explain Statements" PMDB process.
- Until all functions within an application are exercised, Precise for DB2 cannot give a true picture of all affected statements.
- The index recommendation feature in the SQL tab only examines a single statement. It may be the case that the What-If tab suggests that this recommendation should not be implemented because of its impact on other statements. It is always a good idea to use both tabs together.

See ["Getting current application information"](#) on page 30.

See ["Examining performance over time"](#) on page 39.

See ["Analyzing Statement plans"](#) on page 73.

How the What-If tab is structured

The What-If tab is slightly different in appearance from other tabs. It is split into the following areas: Main area, Summary area, and Association area.

About the Main area in the What-If tab

The main area lists index changes that you would like to make or that were recommended in the SQL tab.

The following table describes the different index simulation parameters

Table 8-1 Index simulation parameters

Parameter	Description
Change type (Simulation)	Indicates the type of the simulation, which can be one of the following: <i>Index Creation</i> , <i>Index Modifications</i> , or <i>Index Deletion</i> .
Owner	DB2 owner of the index.
Table	Name of the table that the index belongs to.
Index	Name of the index.
Unique	Indicates whether the index is unique or not.
Allow Reverse Scan	Indicates the value of the reverse scan option (Yes/No).
Columns	List of columns in the index.
Include Columns	List of include columns in the index.
Cluster Ratio	The estimated cluster ratio of the index.

About the Summary area in the What-If tab

The following table describes the information that is shown in the Summary area after a simulation has been completed..

Table 8-2 Simulation summary

Category	Item	Description
Improved statements	Statements whose cost has improved	The number of statements that are positively impacted
Deteriorated statements	Statements whose cost has deteriorated	The number statements that are negatively impacted
Unaffected statements	Statements whose cost has not been affected	The number of statements that are not affected by the change.
Affected statements by object type	Statements that were changed categorized by object	The number of statements that were affected for each object category (a row for each index change)

About the Association area in the What-If tab

The Association controls in the Association area allow you to choose statements that will definitely be affected and statements that will possibly be affected. You can use the More... option to filter the results.


For statements that will definitely be affected, Precise for DB2 was able to accurately calculate the impact, for example when a change in the explain plan occurred and the cost figures are available. For statements that will possibly be affected, Precise for DB2 did not have enough information to make an accurate assessment, for example because cost information is missing, the rule-based Optimizer is in use, the explain plan is the same but extra indexes must be maintained, or simply because an older version of DB2 is in use.

Selecting either of these options displays a list of impacted statements. They both contain the following tab.

Effect tab

The effect tab displays the following table columns

Table 8-3 Effect tab

Column	Description
	The Tune icon launches the SQL Tab on the Plan tab.
Statement	Name of the impacted statement.
In DB2	Resource consumption of the impacted statement during the selected time frame. The format is: hh:mm:ss.s. Use this column to assess the heavy resource consuming statements that are either positively or negatively impacted.
Cost	Current estimated cost of the DB2 Optimizer, as taken from the latest explain plan.
Predicted Cost	Predicted cost of the DB2 Optimizer for the new plan.
Change Ratio	The ratio between the predicted cost and the current cost, expressed as a percentage. Therefore, 100% represents no change; a value greater than 100% is a degradation, and a value less than 100% is an improvement. If cost information is not available, the cost, predicted cost, and change ratio are all N/A.
Executions	Number of times the statement was executed during the selected time frame.
Comment	Precise for DB2 may include its own comment on the simulated change, or why, for example, costs are not available.

Evaluating index changes

The buttons below the Main area let you add and edit index changes for evaluation, simulate an index change, or delete a suggested change from the evaluation table.

Adding or editing an index change for evaluation

If you want to test the consequences of an index change that is not yet listed in the evaluation table in the Main area, you can add a new entry. If you want to modify an index change that has already been suggested, you can edit it.

To add or edit an index change

- 1 Select the Time Frame and Database Partition to work on.
- 2 Click **Add** or **Edit**, as required.
- 3 Select an **owner**
- 4 Select a **table**.
- 5 If indexes already exist for this table, select the name of the relevant index, if not, enter an index name.
- 6 If you want to allow a reverse scan, select the relevant check box.
- 7 If you want the index to be a unique index, select the relevant check box.
- 8 In the Cluster Ratio field, enter a number between 0 and 100. This ratio indicates the degree to which the table data is clustered in relation to this index. The higher the number, the better rows are ordered in index key

sequence. If table rows are in close to index-key sequence, rows can be read from a data page while the page is in the buffer.

- 9 To add/remove columns to or from the Index Columns table, use the left and right arrows. If the Unique check box is selected, you may also add/remove columns to or from the Include Columns table.
- 10 Click **OK** to update the table in the Main area.

Simulating an index change

After you add or edit an index change, you should simulate the change to evaluate its effects.

To simulate an index change

- 1 In the What-If... dialog box, click **Evaluate**.
The simulation takes into account all the changes listed in the Main area. When the simulation is complete, the Summary and Association areas are populated with data.

Deleting a suggested index change

If you decide that a specific index change should not be evaluated, you can remove it from the evaluation table.

Note: Removing an entry from the table only means that the impact of this index change will not be evaluated. It is not the same as dropping an index.

To delete an index change

- 1 In the What-If... dialog box, select a table row and click **Remove**.

How the What-If tab can help you evaluate changes

You can evaluate changes by doing the following:

- [“Adding an index”](#) on page 85

Adding an index

Simulating an index change is a two-step process. First you need a list of proposed index changes. Typically this is done automatically for you by following the link on the Recommendation tab in the SQL tab. Alternatively, you can manually propose the addition, modification, or deletion of indexes.

Second, you need to run the impact analysis against all the SQL statements in this database partition. To do this, click **Evaluate**.

The Summary area will show how many statements will be improved by the creation of the recommended index and how many will have their performance deteriorate. In the Association area on the right, you can view which statements will be affected.

The DB2 Optimizer is used to compare the cost of the original execution plan with the cost of the new execution plan.

To simulate removing an index, you still need to add a proposed index change to the table in the Main area. This is either done for you on the Advice tab in the SQL tab, or you can click **Add** to open the What if... dialog box.

Examining DB2 database statistics

This section includes the following topics:

- [About the Statistics tab](#)
- [How the Statistics tab is structured](#)
- [About the entities you can examine](#)
- [How the Statistics tab can help you identify performance problems](#)

About the Statistics tab

The Statistics tab displays database statistical information on your DB2 databases. All information is derived from Precise for DB2's own statistics and loaded into the PMDB to provide you with a long-term historical view. Analyzing database and I/O statistics over a long period of time is the only way to predict future resource consumption of the DB2 database and identify I/O patterns. It also provides a performance-tracking tool that gives you the ability to assess database statistics trends and analyze changes in database activity.

The Statistics tab is primarily intended for database tuning. You can also use the tab to periodically monitor the health of your database. It lets you provide answers to the following types of questions: Is DB2 using resources efficiently? Are the memory structures correctly sized? Are we performing too many logical I/Os?

Note: This section describes how to use the Statistics tab and which DB2 statistics are available in Precise for DB2. It is not within the scope of this chapter to define the meaning of each DB2 statistic or how to tune them. For additional information on DB2 statistics, see the relevant DB2 documentation.

How the Statistics tab is structured

The Statistics tab is divided into two areas—the Main area and the Details area. The Main area lists all the databases or database partitions that are monitored by Precise for DB2. The Details area provides summarized or overtime information regarding the selected database partition.

When you open the Statistics tab from another tab, you will be focused on an individual database partition. The tab is always launched at the database partition level, even if the currently selected entity is not a database partition.

The Main area

The database table displayed in the Main area lists all the databases or partitions that are monitored by Precise for DB2. Each row corresponds to a database or database partition.

The following table describes the information displayed in the Main area table.

Table 9-1 Main area table

Column	Description
Data Partition	<instance>.<database>[.number]
Reads & Writes	Total reads and total writes within the selected time frame
Logical Reads	Total logical reads within the selected time frame
Timeouts & Deadlocks	Number of timeouts and number of deadlocks occurred within the selected time frame
Hit Ratio	The buffer cache hit ratio
SQL Types	Total number of Selects, Inserts, Updates and Deletes occurred within the selected time frame
Lock Escalations	Number of lock escalations and exclusive lock escalation occurred within the selected time frame

About the Details area

The Details area displays statistics information on the database or database partition selected in the Main area. It has seven tabs: An Overview tab, a Prefetchers tab, a Cleaners tab, a Logging tab, a Package Cache tab, a Sorts tab, and a Locks tab.

About the entities you can examine

The Statistics tab displays information on different entities in the Details area. This section provides an overview of all entities (tabs), their meaning, and their views.

The following entities can be examined in the Statistics tab:

- [Overview](#)
- [Prefetchers](#)
- [Cleaners](#)
- [Logging](#)
- [Package Cache](#)
- [Sorts](#)
- [Locks](#)

Overview

The Overview tab displays information in five views (only first three views are opened by default):

Table 9-2 Overview tab

Item	Description
Overview	Displays database summed statistics like timeouts and deadlocks.
Hit Ratios	Displays an overtime graph of the hit ratios of the database, package cache, and catalog cache.
Timeouts & deadlock	Displays overtime graph of timeouts and deadlocks.

Table 9-2 Overview tab

Item	Description
I/O Activity	Displays overtime stacked-bar graph of the I/O activity of the database. The I/O activity is displayed as an over time stacked bar of synchronic reads and writes, a-synchronic reads and writes, and direct reads and writes.
Connections & Agents	Displays an overtime graph of the connections and agents activity. It shows the high water mark of connection, application executing and application connected (idle) as well as associated agents, high water mark of coordinator agents, and high water mark of associated agents.

Prefetchers

The Prefetchers tab displays information in four views (only the first three views are opened by default):

Table 9-3 Prefetchers tab

Item	Description
Bufferpools Hit Ratio	This graph displays the bufferpool's hit ratio over time. It displays separate graphs for database hit ratio, data pages hit ratio, and index pages hit ratio.
Physical Reads	This graph displays a stacked bar of the physical reads over time. The stacked bar is divided into synchronic, a-synchronic (for data and for index), and direct reads.
Read Response Time	This graph displays overtime line graphs of asynchronous read response time, synchronous read response time and direct read response time. Also displayed is a line graph for the average prefetch wait time.
Read Time	This graph displays a stacked bar of read time over time. The stacked bar is divided into synchronic, a-synchronic, and direct read time.

Cleaners

The Cleaners tab displays information in four views (only the first three views are opened by default):

Table 9-4 Cleaners tab

Item	Description
Physical Writes	This graph displays a stacked bar of the physical writes over time. The stacked bar is divided into synchronic, a-synchronic (for data and for index), and direct writes.
Cleaners Efficiency	This graph displays overtime line graphs for Victim Page Cleaners Execution [calculated as (bp victim page cleaners triggered) / (bp victim page cleaners triggered+ bp thresh cleaners triggered+ bp log space cleaners triggered)] and synchronous writes to asynchronous writes ratio.
Write Response Time	This graph displays overtime line graphs of asynchronous write response time, synchronous write response time, and direct write response time.
Write Time	This graph displays a stacked bar of write time over time. The stacked bar is divided into synchronic, a-synchronic, and direct write time.

Logging

The Logging tab displays information in five views (only the first three views are opened by default):

Table 9-5 Logging tab

Item	Description
Log Utilization	This graph displays an overtime stacked bar of Primary log used, secondary log used, and log space available.
Log Pages Time (ns)	This graph displays a stacked bar of log read time and log write time (in nanoseconds) over time.
Commits & Rollbacks	This graph displays a stacked bar of commits and rollbacks over time.
Log Buffer Reads - Hits & Misses	This graph displays the log buffer reads ratio between hit and misses.
Log Pages Count	This graph displays a stacked bar of log pages reads and log pages writes over time.

Package Cache

The Package Cache tab displays information in four views (only the first three views are opened by default):

Table 9-6 Package Cache tab

Item	Description
Package Cache Hit Ratio	This graph displays the package cache hit ratio, the package cache overflows and the application section hit ratio over time, in a graph with two Y-axis (the left Y-axis is for hit ratio, and the right Y-axis is for overflows)
SQL Types	This graph displays a stacked bar of Selects, DMLs and DDLs over time. This graph can be used for findings a correlation between package cache low hit ratio to the high number of DDL.
Compile & Recompile	This graph displays a stacked bar of compile time and recompile time over time, as taken from Precise for DB2 collector. This indicate the amount of parsing time the application is spending while running.
Dynamic vs. Static	This graph displays a stacked bar of dynamic SQL attempts and Static SQL attempts over time.

Sorts

The Sorts tab displays information in three views:

Table 9-7 Sorts tab

Item	Description
Sort Heap Overflows (To disk) Ratio	This graph displays the sort ratio $\{ (\text{Sort overflow}) / (\text{total sorts}) \}$ and hash joins ratio $\{ (\text{hash join overflows}) / (\text{total hash joins}) \}$ over time. This graph can indicate whether operations are performed on disk (sorts, hash joins), and whether it is needed to adjust the sort heap size.
Overflow Count	This graph displays a stacked bar of sort overflows, hash joins overflows, and small hash joins overflows over time.
Sort Time	This graph displays the sort time over time. This graph shows data only if the DB2 sort switch is on.

Locks

The Sorts tab displays information in three views:

Table 9-8 Locks tab

Item	Description
Timeouts & Deadlock	This graph displays a stacked bar of timeouts and deadlocks over time.
Lock Escalations	This graph displays a stacked bar of shared lock escalations and exclusive lock escalations over time.
Lock Time (Avg)	This graph displays the average lock time $\{(\text{lock wait time})/(\text{lock waits})\}$ over time. This graph will have values only in case DB2 Lock Switch is on.
Lock List in Use	This graph displays lock list in use over time.

How the Statistics tab can help you identify performance problems

To determine whether DB2 is performing optimally, it is necessary to monitor the health of your database and database partitions, and at times to tune them. The Statistics tab provides many performance counters grouped into several predefined graphs that display related counters that enable you to locate performance problems in your system.

You can identify a performance problem by doing one or more of the following:

- [Examining Prefetchers](#)
- [Examining Sorts](#)
- [Examining Locks](#)
- [Examining Logging](#)

Examining Prefetchers

You can examine if your prefetchers are working properly by viewing the Prefetchers tab.

- 1 In the Time Frame list, choose the period of time you want to analyze.
- 2 In the Database Partition list, choose the item you want to analyze.
- 3 Check the relation between the different reads (synchronic, asynchronic and direct) by viewing the Physical Reads view. Check the average response time of the different reads by viewing the Response Time (Avg) view. This may indicate whether your I/O related parameters are configured well.

Examining Sorts

You can examine if your application is sorting optimally, and if your sort-related parameters are configured adequately by viewing the Sorts tab.

- 1 In the Time Frame list, choose the period of time you want to analyze.
- 2 In the Database Partition list, choose the item you want to analyze.
- 3 Check whether sorts overflows and hash joins overflows occur in the Overflows Count view. Check the ratio between total sorts and sort overflows and the ratio between hash joins and hash joins overflows in Sort Heap Overflows (To Disk) Ratio graph. If the ratio is high then your sort-related parameters are probably not configured adequately.

Examining Locks

You can examine if your application is suffering from locks and if your lock-related parameters are configured adequately by viewing the Locks tab.

- 1 In the Time Frame list, choose the period of time you want to analyze.
- 2 In the Database Partition list, choose the item you want to analyze.
- 3 Check whether lock escalation occurs in your database by viewing the Lock Escalations view.
- 4 Check whether deadlocks and timeouts occur by viewing the Timeouts and Deadlocks view. This may indicate that one or more applications are holding locks for long durations or that lock-related parameters are not configured adequately.

Examining Logging

You can examine if your logging is performing well by viewing the Logging tab.

- 1 In the Time Frame list, choose the period of time you want to analyze.
- 2 In the Database Partition list, choose the item you want to analyze.
- 3 Check whether you allocated enough log space in the Log Space view. Identify whether secondary log space is used. If so, you probably need to allocate more primary log space.
- 4 Check whether your log read and write time are adequate by looking in the Log Time view. This may indicate whether the current disk is adequate for logging.

abandonment rate	In Web, a counter that keeps track of the percentage of users that abandon the loading of the Web page before it completes downloading.
action	An operation that Alerts FocalPoint automatically performs when detecting a warning or critical status for a specific metric. According to the defined action, Alerts FocalPoint opens a message box, sends an email or SNMP trap, or executes a program.
action item (manual)	Some installation steps cannot be executed automatically by Precise Installer. The administrator is required to execute them manually. Action Items are presented (when necessary) at end of installation or product update.
AdminPoint	The central administration console of Precise that facilitates the maintenance, configuration, and management of all installed Precise components, such as monitoring the status of all Precise agents and PMDB processes, getting license information, starting and stopping the agents, getting log data on agents and events, changing PMDB settings, and installing patches. See also PMDB . See also agent .
advice	In Oracle, an algorithm that is designed to recommend on gathering statistics, creating new indexes, change existing indexes, and add or delete hints to make Oracle's Optimizer choose a better access path and make the statement perform better. Can be activated from any DML (Data Manipulation Language) statement.
agent	A program that runs on a server machine to collect, process, or communicate performance information. The Precise installation consists of multiple agents.
alert	The state of an Alerts metric that has reached a near-critical or critical status. An alert is issued by Alerts, triggering an action and informing of a problem that has occurred or is likely to occur within the area sampled by the specific metric. See also action .
Alerts	These product provide alerts to problematic conditions before they turn into performance problems, based on predefined metrics and thresholds. Alerts can automatically perform an action, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program.
Alerts FocalPoint	An agent that receives data from the InformPoint agents, stores it, and performs any action that has been user-defined for that specific alert, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program. See also InformPoint .
alert type	In Alerts, the status of all metrics belonging to a metric group or a monitored instance, indicating the current performance level through colors. See also alert . See also metric .
application server metric	In J2EE, a metrics that is provided by the application server or by customer code. This can include metrics published by the Java Management Extension (JMX) APIs or vendor-specific APIs, such as IBM's Performance Management Interface (PMI).
AppTier	The abbreviation for an application tier in a Precise environment. An AppTier contains one or more instances of the same technology and purpose. Application tiers do not necessarily correspond to distinct physical servers: in many cases, the tiers are logical, with one server running multiple AppTiers or one AppTier spans a cluster of servers. A Precise environment may contain multiple AppTiers on the same technology. For example, you may group J2EE instances (JVMs) into a J2EE Presentation AppTier and a J2EE Business Logic AppTier. Segmenting application service time into the contribution of individual application tiers is helpful in identifying the source of performance problems. Analyzing the performance and behavior of each tier separately is crucial for isolating the root causes of performance problems.
cabinet	In Oracle, SQL Server, DB2, and Sybase, the highest logical level in the SQL workspace hierarchy. A cabinet contains folders and, within folders, statements.

call path	A subset of a stack trace including only those methods that have been instrumented.
client cache	In Web, a counter that keeps track of the percentage of requests taken from the client cache (http status 304).
client-side collection	See Web client browser-side .
cluster	A group of servers or instances that are configured to be treated as a single entity. In Oracle, all Oracle instances of the same Oracle database (RAC configuration) are treated as a cluster and presented in the Precise for Oracle UI as a Database. In J2EE, a group of JVMs that serve the same application can be defined as Cluster (even if they are not defined as a cluster in the application server configuration). All JVMs of the same cluster share the same configuration.
collapsed access plan	In SQL Server and Sybase, the access plan of a unique group of statements or batches that belong to the same collapsed statement or batch but have different access plans. This can differ due to the constants in the text of the original (not collapsed) statements. See also explained statement . See also collapsed statement .
collapsed statement	In Oracle, SQL Server, Sybase, and DB2, a statement whose constants are replaced with parameters. Each collapsed statement can have several access plans, according to the occurrences of its statements. See also collapsed access plan .
Collector	The program that runs on the monitored server to collect performance information. Some technologies allow a single Collector (agent) to serve multiple instances running on the same server. Other technologies require a dedicated Collector per monitored instance. In most cases, it must be installed on the monitored server. Collectors for SAP, SQL Server, and Sybase may reside on a remote server. See also agent .
collector slice size	In J2EE, the smallest unit of data presentation and storage. It can be configured as 30 seconds or five minutes. If SmartLink is configured, relevant Web and Microsoft .NET Collectors will report to SmartLink using the same collector slice size. Regardless of the defined collector slice size, the most granular slice size presented on both the Precise for Web UI and the Precise for Microsoft .NET UI will be 15 minute intervals.
completion rate	In Web, a counter that keeps track of the percentage of users that completed the loading of a Web page.
count	The number of occurrences observed during a measurement interval.
CPU time	The average amount of time consumed by the operating system actively processing instructions on behalf of a running activity.
critical status	In Alerts, the status represented by a red bullet indicating that the value of the sampled metric has exceeded the near-critical and critical threshold values. See also metric .
Cross-AppTiers	In Insight, Report Manager, and Alerts, a perspective that provides high-level information about the performance of all AppTiers in an environment, including operating system data. It provides a holistic view of the entire environment and helps understand how the AppTiers interact. See also AppTier .
current data	In SmartLink, the end-to-end activity, detected by Precise over the last five minutes, including activities that are still in progress. In all databases products, a workspace that presents Current Sessions and the SQL statements that were executed including statements that are still in progress.
custom	Represents invocations of a type not specifically defined. See also work time .
customized metric	In Alerts, a user-defined metric that measures site-specific parameters.
drilldown	Within Precise products and Insight SmartLink, the filtering of analyzed data by clicking a specific entity. Then, additional information about the selected entity is presented, plus a breakdown of its activity by another entity. For example, drilldown on a transaction and get an overtime graph of the selected transaction, plus a list of top users that executed the selected transaction.
entry point	Usually a top-level HTTP or EJB request. An entry point can originate when a user clicks in a browser or an E-commerce server invokes a remote EJB. J2EE makes a distinction between service request invocations and other invocations. The first HTTP or EJB invocation within a call path is designated as an entry point.

environment	<p>The highest Precise logical group. It may contain multiple AppTiers of various technologies that serve an application together. For example, a Payroll Production environment may contain all Web servers, application servers, transaction managers, databases, and servers that serve this application. Alternatively, it may contain any set of instances that form together an administrative group. Since Precise version 8.5, an instance can belong to more than one environment. Let's say a single Oracle database serves two different applications plus the DBA wants to associate this database with a group of other databases under his responsibility regardless of served applications. In this case the Oracle database will be associated to three different Precise environments.</p> <p>See also AppTier.</p> <p>See also instance.</p>
ERP Extension	<p>In Oracle, DB2, and SQL Server, an extension to the Collector that provides detailed information on the activities and resource consumption of packaged application components. It correlates the database information, and the packaged application information and lets you see users, transactions, reports, and other elements of ERP applications, such as Oracle Applications, SAP, PeopleSoft, and Siebel.</p>
event	<p>In Alerts, the occurrence of a sampling or progress. A sampling occurrence occurs every time a metric samples. A progress occurrence occurs when a metric's progress status is changed or when the investigated status reaches the end of the given investigation time. In AdminPoint, all occurrences reported by Precise agents, including informational events, warnings, and errors related to one of the agents. All events are shown in the Events view.</p> <p>See also metric.</p> <p>See also progress.</p>
executions	<p>In Oracle, SQL Server, DB2, and Sybase, the number of times a SQL statement was executed during the selected time frame. In SmartLink and SAP, the number of times a transaction was executed during the selected time frame.</p>
explained statement	<p>In Oracle, SQL Server, DB2, and Sybase, a statement whose access path (chosen by the RDBMS Optimizer) is clarified and translated into a visual display. Explained results include information on the objects referenced by the statement and the operations performed on these objects.</p>
extended collection	<p>In Oracle, a function that proactively specifies a future period during which Oracle activity data is collected and organized for subsequent analysis. Extended collections are an easy means to view collected information, assess application resource consumption, and identify bottlenecks that are inhibiting application performance and end-user productivity.</p>
Federated Precise	<p>Federated Precise is a version of Precise (starting with version 8.5) that can manage multiple Precise installations within unified StartPoint and AdminPoint screens, displaying and managing all environments, instances, and installations.</p>
findings	<p>A ranked list of top performance problems in the selected context. They are presented in Oracle, SQL Server, .NET, and J2EE. Each finding provides: a short explanation of the problem, detailed background information, and most important: one or more links to further investigate the problem and possibly get more granular findings on the selected context.</p>
first byte time	<p>In Web, a counter that keeps track of the time that it takes from the moment a new Web page is called until the first byte arrives back from the Web server.</p>
FocalPoint	<p>An agent that communicates with the Listeners installed on the monitored servers, receives data from the Collectors, periodically processes and stores this data in the PMDB, and serves UI requests.</p>
folder	<p>In SQL Server, Oracle, DB2, and Sybase, the intermediate logical level in the SQL workspace hierarchy. Folders are grouped into cabinets and contain SQL statements.</p> <p>See also cabinet.</p>
Framework Installer	<p>The application that facilitates the installation of Precise framework components. It can be invoked from the installation DVD to install a new Precise deployment. It can also be invoked from an existing Precise deployment to install an additional framework node and attach it to the originating Precise deployment.</p> <p>See also Framework node.</p>
Framework node	<p>A set of FocalPoint agents that are installed together and manages a set of monitored instances. The performance data of these instances will be loaded into a dedicated PMDB. A single Precise deployment may contain multiple framework nodes (using a separate PMDB for each node). An environment cannot span over multiple framework nodes.</p>

garbage collection	An automatic process in the Java runtime environment that periodically reclaims memory used by objects that are no longer referenced. The process can impact an application's performance while memory is being reclaimed. Java programmers may initiate garbage collection explicitly.
group	In Web, the identifier that is used to group other identifiers, such as sites or URLs.
hint	In Oracle, an instruction directed at the Oracle Optimizer that includes considerations for an execution plan. The Oracle Optimizer will build an execution plan based on the hint, ignoring its own set of considerations.
hour group	A unit that reflects the type and level of activity within the system at different times. By defining the times of the day that are peak and off-peak, or day and night, the performance analysis can be focused on those particular times of the day. If, for example, most performance problems occur within nighttime and weekend batches, it can be useful to focus only on them.
InformPoint	In Alerts, an agent that retrieves performance data from all installed Precise products, analyzes it, and sends an alert to the Alerts FocalPoint when the predefined thresholds are exceeded. See also agent .
Insight	The Precise product family that facilitates the process of monitoring and correlating system performance. It consists of Insight. See also Report Manager . See also FocalPoint .
Insight FocalPoint	An agent that receives performance information from Insight Savvies, which monitor the environment. Insight FocalPoint then correlates, processes, and stores this information in a centralized location. The Insight performance history is stored in the PMDB. See also Savvy . See also PMDB .
instance	A monitored object of a specific technology. The following list specifies what constitutes an instance for the various supported technologies: J2EE - a Java Virtual Machine (JVM - a logical name set by the user), Microsoft .NET - a Common Language Runtime (CLR - a logical name set by the user), Oracle Applications - an Oracle Applications Form server, Oracle - an Oracle instance, SAP - a SAP system, SQL Server - an SQL instance, Sybase - a Sybase instance, Tuxedo - a Tuxedo domain, Web - a Web server, WebSphere MQ - an IBM WebSphere Queue Manager, DB2 - a DB2 non-partitioned database or a DB2 database partition. During installation, the instance is associated with one AppTier and environment. An Instance can be moved to a different environment or associated with multiple environments without re-installation nor losing historical data. See also environment .
instance statistics	See performance counter .
instance/database changes	In SQL Server and Oracle, the part of the Collect Schema Changes process that captures instance definition changes and database option changes and saves them in the PMDB.
instrumentation	The process of inserting fault-tolerant recording hooks in Java byte code, .NET MSIL, HTML pages, or other monitored components, resulting in the capture of performance metrics. In J2EE, a mechanism that enables collecting performance information when an application is executed. The process involves inserting special fault-tolerant recording hooks into application class objects. In WEB, the insertion of recording hooks into HTML pages can either be in memory (Dynamic Instrumentation) or file-based (Static Instrumentation).
instrumentation context	See invocation context .
internal invocation	The process of invoking a request from an HTTP request (Servlet or JSP) or an EJB. J2EE displays an internal invocation like any other invocation.
invocation	An execution of a J2EE entity (a Servlet, EJB, SQL Statement, method, and so on). When mentioned in plural (Invocations), means the amount of times that the entity was executed.
invocation context	The context within which a method is invoked. For example, if Method A is invoked both from Method B and from Method C, there will be two different invocation contexts for the performance metrics collected for Method A, one for when it is invoked from Method B and one for when it is invoked from Method C. See also call path and instrumentation context .
JRE (Java Runtime Environment)	As the runtime part of the Java software, the combination of the components that enable the execution of a Java program: a Java virtual machine, the core class libraries, and the files that form the Java platform.

JSP (Java Server Page)	An HTML page with special tags for Java scripting. An application server processes the tags and generates a Servlet.
JVM (Java Virtual Machine)	An instance of a JRE that executes Java programs. A server-side Java application server is itself a Java program that runs inside a JVM. Servlets, JSPs and EJBs are Java programs (applications) that run within the application server's JVM. J2EE monitors the JVM running the application server and the server-side Java applications within the application server.
JVM Heap Memory key metric	The amount of real computer memory that is allocated to the JVM for executing Java programs. In Alerts, a parameter that monitors a very important performance aspect. The status of the Oracle instance (up/down), for example, is crucial for system performance. If the instance is down, this is the first problem that needs to be solved. Marking a metric as key metric ensures that a critical alert raised for this metric receives top priority by the person that is responsible to handle alerts at any time. It can also help determine which alert to handle first in case of multiple alerts. In Alerts, key metrics are always displayed at the top of a metric table when it is sorted by alerts so that they get immediate attention.
Listener	The agent that facilitates the communication between the various Precise agents across different servers must be installed on every server where Collectors or FocalPoint agents are installed. The Listener allows communication with all other agents installed on the monitored server, while only the Precise Listener port is known by other servers.
login name	In SQL Server and Sybase, the session identifier that represents the credential used to connect to the database. When an ERP extension is installed, the user name of the packaged application's client overrides the login name. For example, when SAP extension is installed, the SAP user name overrides the login name.
machine	A session identifier. A machine as sampled, for example, by the SQL Server Collector is the identification of the machine where the client process executes. Machine is also sampled by Oracle, DB2 and Sybase. In Insight terminology Machine is called Client machine.
Main framework node	A main framework node is the single point for login and also serves as the Precise FocalPoint for the entire deployment.
major collection count	Number of estimated major garbage collection events that occurred during the last J2EE collector slice. A major garbage collection can stop the application while JVM heap memory is being reclaimed. See also garbage collection .
major garbage collection time	Percentage of time spent by the JVM executing major garbage collection events during the displayed interval. See also garbage collection .
MBean	A Java object that represents a manageable resource. In J2EE, MBeans, or Managed Beans, are used for application server metrics.
memory logger interval	The interval at which the J2EE Collector gathers JVM heap memory data. All snapshots of the memory logger's data collected according to the memory logger interval are summarized with counters in the current aggregation interval. The memory logger interval's time span is typically a small fraction of the aggregation interval.
metric	In Alerts, a query that helps measure performance in the environment. Three types of metrics are available: System metrics relate to the internal resources, operations, and objects of the monitored infrastructure; application metrics reflect the way the applications perform; user-defined metrics can be customized to specifically relate to a site. When a metric's value exceeds one of the defined thresholds, its status changes to near-critical (yellow bullet) or critical (red bullet). See also key metric .
metric set	In Alerts, a unit that groups metrics that measure related performance aspects. The following metric sets exist: Status: includes metrics that alert to functional problems related to the instance. Performance: includes metrics that alert to performance problems related to the instance. Load: includes metrics that alert to instance-related load problems that may later cause errors or crashes. Service: includes metrics that alert to instance-related SLA breaches. The metrics in this set are sampled by Insight. Performance Trending: includes metrics that alert to potential future performance problems. The metrics in this set are sampled by Report Manager.

	<p>Load Trending: includes metrics that alert to potential future load problems. The metrics in this set are sampled by Report Manager.</p> <p>Customized: includes user-defined metrics.</p> <p>Precise Status: includes metrics that alert to the near-critical or critical status of the installed Precise environment.</p>
module	A session identifier. In Oracle it contains the value of MODULE column in V\$SESSION table. In DB2 it contains the command name of non-SQL statements or the package name for SQL statements.
near-critical status	In Alerts, a status indicating that the value of the sampled metric has exceeded the defined near-critical threshold. A near-critical status is indicated by a yellow bullet.
network time	In Web, a counter that keeps track of the time spent on network activity from the server side perspective. This includes the time to read the request from the network and the network time to send the response back to the client.
node	<p>See Framework node.</p> <p>Note: when creating a new installation from DVD, it would be called "Framework." In the UI screens (columns in tables, choosing a system for a new environment etc), we would call them "Nodes." For example: choose a node for the new environment.</p> <p>When adding a new system within AdminPoint, it would be called a "Framework node".</p>
packaged application	<p>An application that is created and/or maintained by a third party and is not custom-built to one's specific needs. The following packaged applications have special treatment by Precise: SAP, Oracle Applications, PeopleSoft, and Siebel. They are harder to modify because the application code is either not available or hard to understand.</p> <p>See also ERP Extension.</p>
page size	In Web, a counter that keeps track of the amount of data loaded from the server to display the page.
page views	In Web, a counter that keeps track of the number of Web pages viewed at a specific Web site during a selected time period.
parent metric	<p>In Alerts, a joining of several child metrics. Each time a parent metric samples, it gathers data from a set of child metrics and presents it as a single metric query. The individual child metric values are displayed on the Thresholds tab of the Properties dialog box in Alerts.</p> <p>See also submetric.</p>
performance counter	In SQL Server, a Windows performance counter as reported by the operating system. In Oracle, an instance level statistic as gathered from some Oracle V\$ views. In Sybase, a sysmonitors counter collected by Sybase. In Web, an operating system or a Web server performance counter.
PMDB	The Precise data warehouse of performance and availability data. It can be hosted on an Oracle or SQL Server database.
portal server	An application server for Web-based applications that commonly provide personalization, single sign-on, and content aggregation from different sources and that host the presentation layer of Information Systems. Aggregation is the action of integrating content from different sources within a Web page. A portal may have sophisticated personalization features to provide customized content to users. Portal pages may have a different set of portlets creating content for different users.
portlet	A Java-based Web component, managed by a portlet container, that processes requests and generates dynamic content. Portlets are used by portals as pluggable user interface components that provide a presentation layer to Information Systems.
Precise deployment	An independent Precise system. It contains and manages various agent types and provides centralized monitoring and administration. A Precise deployment may contain multiple framework nodes. One of them is defined as the main framework node and it manages all other nodes of the Precise deployment.
Precise FocalPoint	See Precise FocalPoint agent .
Precise FocalPoint agent	An agent that manages all agents in a single Precise deployment. Additional product FocalPoint agents manage specific technologies and resources.
Precise Proxy agent	When multiple Framework Nodes are managed by a single Precise deployment, the Precise FocalPoint agent manages all of the agents of the main Framework node, while every other framework node has a Precise Proxy agent to manage all its agents.
Precise_root	This is the term used in a path for the Precise installation directory. The terms <i><i3_root></i> or <i>i3 root</i> can appear in text too.

program	<p>A session identifier in Oracle, DB2, SQL Server, Sybase, Insight, and Alerts. A program as sampled, for example, by the SQL Server Collector, is the name of the executable that connects to the database. Applications that do not set the application name have N/A as program. When a SAP extension is installed, the SAP transaction overrides the program, and it may change during the application's lifetime. When a Siebel extension is installed, the Siebel views override the program, and it may change during the application's lifetime. In Insight, program is the name of an executable as recognized by the operating system. If an executable is invoked from a script (a batch or a shell), the script is displayed as the command entity. In Alerts, program is part of the customized metrics definition and is the name of the executable or stored procedure executed in the database that runs when the metric is sampled. In action definitions, program is the name of the executable that will run if the metric exceeds its threshold.</p>
progress	<p>In Alerts, the management state of a metric for which an alert has been raised. The following statuses exist:</p> <p>Open: An alert is raised.</p> <p>Investigated: The alert is taken care of.</p> <p>Closed: The problem has been solved.</p>
Proxy FocalPoint	<p>See Precise Proxy agent</p>
recommended index	<p>In SQL Server, a function that uses the Microsoft Index Tuning Wizard to recommend on adding indexes or statistics for a selected statement, batch, or table. Based on the results of this function, the Optimizer will choose a better access plan and make the respective statement or batch perform better. For statements and batches, recommendations are based on the content of the statement or batch. For tables, recommendations are based on all the statements that are stored in the PMDB, are executed during the selected time period, and have an average duration time greater than the value defined in the registry (where 0 is the default).</p> <p>See also advice.</p>
related SQL	<p>In Oracle, a generated statement that uses alternative syntax to access the database in different ways and returns the same output as the original statement.</p>
relative frequency	<p>Number of contributor invocations per entry point. For example, if a service request calls three methods each time it is invoked, the Relative Frequency for the method is three invocations per service request execution. Similarly, if a service request calls one method every other time it is invoked (half of the time), the Relative Frequency for the method is 0.5 invocations per service request execution.</p>
rendering time	<p>In Web, a counter that keeps track of the time that it takes for a Web page to be loaded from the moment the first byte arrives until the Web page is fully loaded or the user interrupts or abandons the loading process.</p>
report	<p>A collection of queries, programming code, and layout settings that Report Manager executes to generate graphical results like tables and charts.</p>
Report Manager	<p>Uses historical information to identify problematic conditions, track long-term performance, volume trends and patterns, view availability problems over long periods of time and on different levels, compare the performance of similar systems, correlate between performance metrics of different products, assist in capacity planning, and generate demand-driven, user-defined reports.</p>
Report Manager FocalPoint	<p>An agent that examines the PMDB tables to produce scheduled performance reports.</p> <p>See also PMDB.</p>
report parameter	<p>In Report Manager, a keyword used in a report. Its value is set during the report execution. A parameter's value can be updated either permanently or for the current execution only.</p>
report property	<p>In Report Manager, the attributes that define a specific report, consisting of report parameters and scheduling information.</p> <p>See also report parameter.</p>
request	<p>In Web, a counter that keeps track of the number of HTTP requests sent for a viewed entity. request</p>
error	<p>In Web, a counter that keeps track of the percentage of requests completed with an HTTP error.</p>
sampling	<p>In Alerts, the process during which a metric queries a Precise product for a specific instance, retrieves values, and calculates the metric's alert level.</p>
sampling base	<p>In Alerts, the start time of a scheduled sampling process (by default Sunday, 00:00 AM).</p>
sampling period	<p>In specific Alerts metrics, the time frame during which statistical data is returned from other Precise products. Such a metric is for example Oracle's General Behavior metric, which returns database behavior for a certain period of time.</p>

sampling rate	In Alerts, the frequency of a metric's regular sampling schedule. The sampling rate is measured in minutes.
Savvy	An Insight agent that collects AppTier-specific performance data.
scalability	A system's ability to withstand load. For example, positive scalability means that the system continues to function properly even when it is called upon to service a larger number of users.
schema changes	In Oracle and SQL Server, a process that captures schema changes and saves them in the PMDB. In SQL Server, instance configuration parameters and database option changes are also captured.
Script Installer	A component of Precise Agent Installer that enables adding a small JavaScript script, also known as a Web browser-side agent, to the Web pages of your Web site.
server	Also called a host machine or server machine. The combination of a computer and associated operating system software that is accessed by multiple clients to perform requested functions and typically returns results to the clients.
server-side collection	In Web, the collecting of performance data from the Web server instance through the Web server agent.
service + network time	In Web, a counter that keeps track of the service and network time, including the total amount of time the request took to reach the server.
Service Level Agreement	See SLA (Service Level Agreement) .
service time	In Web and SmartLink, the time elapsed from when a request is received by the server to the time a response is sent back to the computer that made that request. The service time is measured on the server side. In J2EE, the time it takes an invocation to complete execution. In other words, service time is the average length of time between the start time and end time of a Java method execution. For example, the SQL service time is the time it takes the JDBC method executing the SQL statement to be completed. The service time includes CPU and wait time. The service times are reported in the interval in which they complete execution. Though a contributor's average service time may be very small, the contributor may be called many times. As a result, a contributor's overall contribution to performance may be large even though its average service time is low.
servlet	A class that is loaded only once and for which the application server uses multithreading to process requests. The servlet generates an HTML page that is sent back to the Web browser.
size	In Web, a counter that keeps track of the bytes sent and received.
SLA (Service Level Agreement)	A formal definition of an information system's performance goals. Within Precise, an SLA consists of clauses corresponding to various system activities. Once a system's SLA is defined, its SLA compliance can be analyzed, and breaches can be isolated to identify their causes.
SmartLink	In Insight, a component that provides transactive, correlated information across AppTiers from the user's perspective. Insight SmartLink functions in PeopleSoft and Web applications environments.
SmarTune	In Oracle, SQL Server, and .NET, an algorithm that analyzes performance metrics, identifies and ranks potential problem areas, and provides advice for correcting the problems.
socket	An I/O abstraction layer that enables processes to communicate with each other, regardless of whether they are on the same server. Sockets are bi-directional FIFO pipes that look like regular file I/O to the developer with the abstraction layer handling all of the low-level communication details.
StartPoint	The opening page of Precise. It provides a quick overview of the environment status and links to launch any of the Precise products.
Statistics workspace	In Oracle, DB2, SQL Server, Sybase, J2EE, and Web this workspace presents instance level statistics such as hit ratios, utilization, JMX metrics, logging, and more.
submetric	Also called child metric. In Alerts, the subquery of a parent metric. A child metric gathers its own data and combines it with the data gathered by other child metrics to form the result of the parent metric. Each child metric has its own thresholds and may be enabled or disabled individually. See also parent metric .
summary table	A container that stores the information collected by the Precise agents and loads it into the PMDB. Summary tables store the same data at different levels of granularity: time slice, hourly, daily, weekly, and monthly. By storing data in multiple summary tables, Precise can present a detailed view and progressively higher-level views of the same data. Summary tables are particularly useful for data aging. A data purging policy can be implemented for each summary table so that detailed data is retained for short-term historical analyses while more summarized data is used for long-term analyses and trending.

technology	A technology identifies the monitored object. For example, Oracle, SQL Server, Sybase, and DB2 are different technologies, while all Web servers (such as: Apache, IIS, and WebSphere) are defined as a single Web technology. A single monitored object can be monitored by two different technology's Collectors. For example, WebLogic server can be monitored by both a Web Collector and a J2EE Collector.
throughput	The number of average completions per second that are observed during an interval.
time slice	<p>A unit used to break up long sessions into smaller time periods. The length of a time slice is preset and cannot be changed. It represents the maximum time that passes before the data collected can be displayed. For example, if the length of a time slice is 15 minutes, the collection is only updated at 15-minute intervals.</p> <p>The length of a time slice is different for each technology: In Oracle, SQL Server, Sybase, DB2, Web, .NET, MQ, Tuxedo, Network, and OS, a time slice is 15 minutes. In J2EE and SmartLink, a time slice is either 30 seconds or 5 minutes. In Precise for SAP, a time slice is 5 minutes.</p> <p>See also collector slice size.</p>
URI (Uniform Resource Identifier)	The relative path to a resource after the location (network node) is found.
URL mapping	In Web, a function that defines rules that map URLs (Uniform Resource Locators) with dynamic parameters originating from a specific domain to a format that identifies the Web pages and prevents them from having different URLs.
Web client browser-side	In Web, the gathering of performance data from the browser running on the desktop of the Web application's user, through the static or dynamic instrumentation of Web pages.
Web server	A program that receives client requests for Web pages, retrieves the static pages and/or issues a request for dynamic page creation to an application server, and sends the pages to the client. In Web, a computer that delivers (serves up) Web pages. Every Web server has an IP address and possibly a domain name.
work time	The time spent in the specific entity excluding time spent by called (other) entities. In J2EE, work time of a method is calculated as the service time of that method minus the service time of the methods it invoked.
work type	In Web, a counter that keeps track of the type of the Web instance, such as Web, PeopleSoft, SAP, or Siebel.
workspace	A display unit in Precise products. All workspaces display data from different perspectives. For example, in Oracle, SQL Server, and Sybase, the Current workspace shows information on the sessions currently active in an application, and the Objects workspace displays information on Oracle or SQL Server database objects that can be used to understand relationships and associations between database schema objects.