

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 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 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.

 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, “About the Current tab” on page 30, “About the Activity tab” on page 39, and “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. 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. 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.

Estimated execution plan. 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.

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 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. <div>  This tab is only displayed after you generate Recommendations (by clicking the Recommend button or clicking Actions>Recommend). </div>

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



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.

Figure 1 Highlights tab content



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 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.
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 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 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 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 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 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**.



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, 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, go to **Actions>Recommend** or, alternatively, click the **Recommend** button.

