

IBM Business Intelligence Pattern with BLU Acceleration  
Version 1.0.0

*Dynamic Query Analyzer User Guide*



**Note**

Before using this information and the product it supports, read the information in "Notices" on page 17.

**Product Information**

This document applies to IBM Business Intelligence Pattern with BLU Acceleration 1.0.0 and may also apply to subsequent releases.

Licensed Materials - Property of IBM

© **Copyright IBM Corporation 2013.**

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

---

# Contents

<b>Introduction</b> . . . . .	<b>v</b>
<b>Chapter 1. Getting started with Cognos Dynamic Query Analyzer</b> . . . . .	<b>1</b>
Cognos Dynamic Query Analyzer . . . . .	1
Optimize dynamic cube performance with Aggregate Advisor . . . . .	1
Running Aggregate Advisor . . . . .	3
Processing the Aggregate Advisor results . . . . .	3
Applying in-database recommendations . . . . .	4
Stored procedures . . . . .	6
Changing the event schedule for loading aggregate tables . . . . .	7
Analyzing query log files . . . . .	7
Running a report that uses dynamic query mode . . . . .	7
Opening query log files . . . . .	8
Exploring query log files . . . . .	9
DQM server logs . . . . .	11
<b>Chapter 2. Troubleshooting Cognos Dynamic Query Analyzer</b> . . . . .	<b>13</b>
<b>Chapter 3. Accessibility features in Cognos Dynamic Query Analyzer</b> . . . . .	<b>15</b>
IBM and accessibility . . . . .	15
<b>Notices</b> . . . . .	<b>17</b>
<b>Index</b> . . . . .	<b>21</b>



---

## Introduction

IBM® Cognos® Dynamic Query Analyzer is a query visualization tool. With it you can visualize and troubleshoot the query logs generated when running dynamic query reports in IBM Cognos Business Intelligence. In addition, IBM Cognos Dynamic Query Analyzer includes Aggregate Advisor, a tool that analyzes the performance and structure of dynamic cubes and provides suggestions for the creation of aggregates to improve cube performance.

### Audience

This document is intended to help report administrators use IBM Cognos Dynamic Query Analyzer to analyze reports and find performance issues in dynamic query reports.

In addition, modelers and administrators can use the Aggregate Advisor to optimize dynamic cubes.

### Finding information

To find BI Pattern with BLU Acceleration product documentation on the web, access the IBM Business Intelligence Pattern with BLU Acceleration Information Center (<http://pic.dhe.ibm.com/infocenter/cbiima/v1r0m0/index.jsp>).

### Forward-looking statements

This documentation describes the current functionality of the product. References to items that are not currently available may be included. No implication of any future availability should be inferred. Any such references are not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of features or functionality remain at the sole discretion of IBM.

### Samples disclaimer

The Sample Outdoors Company, Great Outdoors Company, GO Sales, any variation of the Sample Outdoors or Great Outdoors names, and Planning Sample depict fictitious business operations with sample data used to develop sample applications for IBM and IBM customers. These fictitious records include sample data for sales transactions, product distribution, finance, and human resources. Any resemblance to actual names, addresses, contact numbers, or transaction values is coincidental. Other sample files may contain fictional data manually or machine generated, factual data compiled from academic or public sources, or data used with permission of the copyright holder, for use as sample data to develop sample applications. Product names referenced may be the trademarks of their respective owners. Unauthorized duplication is prohibited.

### Accessibility Features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products. IBM Cognos Dynamic Query Analyzer has accessibility features. For information about these

features, see Chapter 3, “Accessibility features in Cognos Dynamic Query Analyzer,” on page 15.

---

## Chapter 1. Getting started with Cognos Dynamic Query Analyzer

IBM Cognos Dynamic Query Analyzer is a tool that provides graphical representations for the query logs produced by dynamic query mode queries. In addition, Cognos Dynamic Query Analyzer has Aggregate Advisor, a tool that analyzes dynamic cubes and suggests aggregates that, when implemented, improve cube performance.

**Important:** You must configure IBM Cognos Dynamic Query Analyzer before you can access the content store and query logs of your IBM Cognos Business Intelligence server. For configuration instructions, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Cubes Installation and Configuration Guide*.

---

### Cognos Dynamic Query Analyzer

IBM Cognos Dynamic Query Analyzer includes an Eclipse-based client user interface. Data is presented in a series of visual components called views. You can perform operations on the data contained in the active view. The results of these operations affect the content of other views, or cause new views to be opened automatically.

Views are contained within tab groups.

A tab group contains **Minimize** and **Maximize** buttons and can also contain buttons required by the currently active view. A tab group can be resized by dragging the side next to another tab group.

When a view is opened, it is displayed in the same location where it was last opened, unless it is being opened for the first time, in which case it is opened in a default location. A view can be moved from one tab group to another, or detached as a separate window.

To move a view to a different tab group, drag the title bar to the new tab group, and release it. To detach a view as a new window, right-click the title bar, and click **Detached**. To drag an entire tab group to a new location, right-click the title bar, and click **Move, Tab Group**.

---

### Optimize dynamic cube performance with Aggregate Advisor

Aggregate Advisor can analyze dynamic cubes and suggest aggregates that, when implemented, will improve cube performance. Aggregate Advisor can also analyze previously run reports and suggest aggregates that correspond directly to these reports.

Aggregate Advisor makes two types of recommendations.

- **In-Memory Recommendations** are aggregate recommendations that can be applied by the IBM Cognos Business Intelligence server the next time the cube is started. These aggregates are stored in the Content Store.
- **In-Database Recommendations** are aggregate recommendations that Aggregate Advisor can apply automatically to the database and the model. For more

information on how to apply the in-database recommendations, see “Applying in-database recommendations” on page 4.

Aggregate Advisor will create in-memory aggregate recommendations estimated to fit within the selected size limit for in-memory aggregates. It can also create in-database aggregate recommendations estimated to fit within the selected size limit for in-database aggregates.

When you run Aggregate Advisor, the following options are available.

*Table 1. Aggregate Advisor general options*

Option	Description
<b>Query Workload Information</b>	<p>Select <b>Cube structure and query workload information</b> to request Aggregate Advisor to consider information from the cube structure and workload logs when making recommendations.</p> <p>Select <b>Cube structure only</b> to request Aggregate Advisor to consider information from the cube structure only when making recommendations.</p> <p>Select <b>Query workload information only</b> to request Aggregate Advisor to consider information from workload logs only when making recommendations.</p> <p>If you want to query workload information, you must enable the <b>Enable workload logging</b> property for the dynamic cube before you run Aggregate Advisor. For information about how to do this, see the <i>IBM Cognos Dynamic Cubes User Guide</i>.</p>
<b>In-memory aggregates</b>	<p>If selected, Aggregate Advisor makes aggregate recommendations that can be applied by the Cognos BI server when the dynamic cube is restarted. A size limit for in-memory aggregates can also be supplied.</p>
<b>In-database aggregates</b>	<p>If selected, Aggregate Advisor makes aggregate recommendations that can be applied to automatically create aggregate tables in the database and aggregate cubes in your model.</p> <p>A size limit for in-database aggregates can also be supplied. The size limit for in-database aggregate recommendations refers to the uncompressed size of the database.</p>
<b>Advisor run time limit</b>	<p>If selected, Aggregate Advisor will end after the specified time and return the aggregate recommendations made at that time. If no recommendations are yet available, Aggregate Advisor continues to run until it has an initial set of recommendations and finishes.</p>

*Table 2. Aggregate Advisor filter workload options*

Option	Description
<b>Only consider the workload information that relates to any of the following</b>	<p>If any conditions are selected, workload information is used only from reports that satisfy the selected conditions. If more than one condition is selected, workload information from reports that satisfy any of the selected conditions is used.</p>



Table 2. Aggregate Advisor filter workload options (continued)

Option	Description
<b>And only consider workload information for the following periods</b>	If any condition other than <b>All</b> is selected, only workload information from reports run in the specified time periods is used.

For more information about aggregates and how their use can improve the performance of dynamic cubes, see the *IBM Cognos Dynamic Cubes User Guide*.

## Running Aggregate Advisor

You can run Aggregate Advisor in IBM Cognos Dynamic Query Analyzer to generate aggregate recommendations for a dynamic cube.

### Procedure

1. Run a series of reports that are a representative workload against the dynamic cube. This step is optional. If you are designing a dynamic cube, you may wish to run the Aggregate Advisor without using workload information.
  - a. Enable workload logging for the dynamic cube. For instructions, see the topic on dynamic cube configuration settings in the *IBM Cognos Dynamic Cubes User Guide*.
  - b. Run a series of reports that are a representative workload against the dynamic cube to gather data about the report execution characteristics.
  - c. Optionally, disable workload logging after the reports are run.
2. In Cognos Dynamic Query Analyzer, click **File, Run Aggregate Advisor**.
  - a. In the **Select the cube** screen, select the cube to run in Aggregate Advisor.
  - b. In the **Specify general options** screen, specify the general options described in Table 1 on page 2.
  - c. If you selected **Query Workload Information** on the **Specify general options** screen, you can specify filtering options on the **Filter Workload Information** screen described in Table 2 on page 2. If you specify filters, only workload data that relates to the specified reports, packages, and users, and that falls within the specified time period, are considered when making aggregate recommendations.

Filters are not available if you select the **Cube structure only** option.

While Aggregate Advisor runs, a graph displays progress as a function of time. The **Details** tab displays the same information as text. Click the **Run in Background** button to dismiss the progress display and continue to run the Aggregate Advisor in the background. You can end the Aggregate Advisor run and receive the recommendations made so far by clicking the **End** button. Click the **Cancel** button to end the run without receiving any recommendations.


When the slope of the graph starts to level out, further recommendations will not have much impact and the Aggregate Advisor run can be ended.

When Aggregate Advisor finishes, the **Advisor results** view displays a summary of Aggregate Advisor recommendations. For more information on the recommendations, see "Processing the Aggregate Advisor results."

## Processing the Aggregate Advisor results

When Aggregate Advisor finishes, the **Advisor results** view displays a summary of Aggregate Advisor recommendations.

## Procedure

1. Double click an item in the **Advisor results** view for the detailed results.  
If there are previous Aggregate Advisor results, new results are shown first by default. To change the order in which Aggregate Advisor results are shown, click **View Menu** , **Sort** in the **Advisor results** view. Select the required items on which to sort, and sort direction, and then click **OK**.  
By default, the title consists of the cube name that is appended with a time stamp. To change the title of a specific result, click **View Menu, Set Description** in the **Advisor results** view. Type the new title and then click **OK**.  
To delete results, click **View Menu, Delete Aggregate Advisor Results** in the **Advisor results** view. Select the results to delete, and then click **OK**.  
A tabbed pane is displayed with more detail about the Aggregate Advisor results. The **General** tab contains a summary of the Aggregate Advisor results. The **In-database** tab describes the in-database recommendations in detail. The **Options** tab lists the options that are used for the Aggregate Advisor wizard run.
2. To clear previously applied in-memory recommendations, click **File, Clear In-Memory Recommendations**, and select the required cube.  
Previously applied in-memory recommendations can be applied again later.
3. To save in-memory recommendations, click **File, Apply Selected In-Memory Recommendations**, and select which recommendations to apply. The selected recommendations are automatically applied when the dynamic cube is restarted.  
If you apply only a subset of the recommendations, a new result is added to the **Advisor results** view with the title “Saved to Server” followed by the cube name that is appended with a time stamp. This special result consists of only the **General** tab, which lists the subset of applied recommendations.
4. To save the in-database recommendations, click **File, Save In-Database Recommendations**, and select where to save the recommendations. This file contains the information that is displayed in the **In-database** tab. The database administrator creates the aggregate tables, and the modeler models the aggregate cube for the in-database aggregate and redeploys the dynamic cube.

## What to do next

You can apply the in-database recommendations immediately. For information about how to do this, see “Applying in-database recommendations.”

## Applying in-database recommendations

You can apply the in-database recommendations to a dynamic cube after they are generated by Aggregate Advisor.

### Before you begin

Before you apply in-database recommendations, ensure that you complete the following prerequisites:

- Run Aggregate Advisor to generate the aggregate recommendations for the dynamic cube. For information about how to run Aggregate Advisor, see “Running Aggregate Advisor” on page 3.

- Create trusted credentials for the user that you are logged in as. For information about how to create trusted credentials, see the *IBM Cognos Dynamic Cubes User Guide*.

## About this task

Aggregate Advisor performs the following steps to apply the in-database recommendations:

1. Back up the model. The backup file is called *original\_model\_name.bck* and it is stored in the same location as the original model.
2. Create aggregate tables in the database.
3. Create stored procedures that load the data in to the aggregate tables.
4. Create aggregate cubes in the model. Aggregate Advisor uses the in-database recommendations to create the aggregate cubes.
5. Republish the dynamic cube that contains the aggregate cubes to the content store.
6. Run the stored procedures to load the data in to the aggregate tables.
7. Restart the dynamic cube.
8. Remove temporary files.

If you cancel the process at any time or the process fails during steps 1 through 4, all previously completed steps are rolled back. If the process fails after step 4, no roll back is done.

If you try to apply in-database recommendations that were previously applied, a warning message is displayed. You can reapply the recommendations or cancel the process.

When it applies the in-database recommendations, Aggregate Advisor reuses existing aggregate tables and aggregate cubes where possible. For example, if you add a measure or level to the model, or have new workload information, Aggregate Advisor adds new items and columns to the existing aggregate table and aggregate cube.

You have the following choices for loading the data to the tables:

### Load now

Aggregate Advisor loads the data immediately by using the stored procedures. Select this option only for initial load of the aggregate tables.

### Load later using scheduling in Cognos Administration

Aggregate Advisor creates an event that runs the stored procedures and restarts the dynamic cube. The event is scheduled to run later. For information about how to modify the time that the event runs, see “Changing the event schedule for loading aggregate tables” on page 7.

You can select this option for initial load or incremental load of the aggregate tables.

### Load later using external ETL tools

Use your own extract, transform, and load (ETL) tool to load the tables. In your ETL tool, you can use the stored procedures that are created by Aggregate Advisor to load the aggregate tables.

**Note:** Ensure that you update your fact data before you load the aggregate tables.

## Procedure

1. Select **Window > Preferences**.
2. In the **Preferences** window, do the following steps:
  - a. Click **Model** and select the dynamic cube to which you want to apply in-database recommendations.
  - b. Click **Incremental Load**. For each cube listed, set the **Changed Date Column** field to the column that represents the date when the rows in the atomic fact table were created or updated. A stored procedure uses the column to determine the new and changed rows for incremental load. For example, `Changed_Date`. If there are values that are defined, then verify that they are correct.
  - c. Click **OK**.
3. In the **Advisor Results** panel, select the set of recommendations that you want to apply.

**Tip:** To see the details of the in-database recommendations, double-click a set of recommendations.

4. Select **File > Apply In-Database Recommendations**.
5. In the **Apply In-Database Recommendations** window, select when to load the data.
6. Click **OK**.

## What to do next

If you chose to load data later by using scheduling, verify that the schedule of the event meets your business requirements by performing “Changing the event schedule for loading aggregate tables” on page 7.

If you chose to load data later by using your own ETL tools, you must restart the dynamic cube after the data is loaded.

## Stored procedures

A stored procedure does both initial load and incremental load for an aggregate table. If the aggregate table is empty, an initial load is done. If there is data in the aggregate table, an incremental load is done.

A stored procedure uses a change data capture field from the atomic fact to determine the new and changed transactions for an incremental load. For example, `CHANGED_DATE`. You must set this change data capture field in the Aggregate Advisor component of IBM Cognos Dynamic Query Analyzer in the following situations:

- The first time that you apply in-database recommendations.
- A new atomic cube has been added to the model.

For information about how to set the change data capture field, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Query Analyzer User Guide*.

A stored procedure writes the execution status of the stored procedure to a metadata table.

The stored procedures are stored in the IBM DB2<sup>®</sup> with BLU Acceleration database. An atomic cube can have one or more associated aggregate cubes. Each aggregate

cube is mapped to one aggregate table. The name of the aggregate cube is the same as the name of the aggregate table. There is one stored procedure for each aggregate table. A stored procedure is named *aggregate\_cube\_name\_SP*. For example, to find the code of the stored procedure for aggregate cube AGG\_1785921059\_SP, run the following query:

```
SELECT TEXT FROM SYSCAT.ROUTINES WHERE ROUTINENAME='AGG_1785921059_SP'
```

## Changing the event schedule for loading aggregate tables

The event loads all aggregate tables and aggregate cubes for the dynamic cube. You can change the schedule of the event in IBM Cognos Administration.

### About this task

This task changes the schedule for an event. If you want to make other changes to an event, such as adding conditions, you must do it in IBM Cognos Event Studio. For information about Cognos Event Studio, refer to the *IBM Cognos Event Studio User Guide*.

### Procedure

1. Open IBM Cognos Administration and select the **Status** tab.
2. Under **Current Activities**, click **Schedules**.
3. In the list, look for a schedule that is named *aggregate\_cube\_name\_data\_load*.
4. Click the down arrow beside the schedule and select **Modify the schedule**.
5. Change the frequency of the schedule and click **OK**.

---

## Analyzing query log files

You can view and analyze query log files in IBM Cognos Dynamic Query Analyzer. The graphical representation of the log files allows you to see performance bottlenecks.

### Procedure

1. Run a report that uses dynamic query mode.
2. In Cognos Dynamic Query Analyzer, click **Window, Analyze Logs**.
3. Open the query log files.
4. Explore the query log files to find performance bottlenecks.

## Running a report that uses dynamic query mode

You can run a report using the IBM Cognos Business Intelligence web interface, or you can run it using IBM Cognos Dynamic Query Analyzer.

### Before you begin

Before running the report, the query execution trace option of Query Service must be enabled. This option can be enabled in either of two ways.

- Your administrator can enable the query execution trace option in IBM Cognos Administration. This option is enabled for all reports run either through the Cognos BI web interface or through Cognos Dynamic Query Analyzer. For more information, see the topic on Query Service administration in the *IBM Cognos Business Intelligence Administration and Security Guide*.
- You can enable query execution trace for reports run through Cognos Dynamic Query Analyzer. To enable this option, click **Window, Preferences**. On the

**General** tab, check **query execution trace**. If logging is enabled through Cognos Dynamic Query Analyzer, the system administrator does not need to enable it on a server-wide basis.

If you are running reports through Cognos Dynamic Query Analyzer, the versions of Cognos Dynamic Query Analyzer and the Cognos BI server must be the same. Otherwise, the report may not run.

## Procedure

Run a report that uses dynamic query in one of the following ways.

- Use any of the methods available through the Cognos BI web interface. For more information, see the *IBM Cognos Report Studio User Guide*.
- In Cognos Dynamic Query Analyzer, navigate to the report in the **Content Store** view. Right-click the report, and click **Run Report**. You cannot set prompts for a report when you run it using Cognos Dynamic Query Analyzer.

## Results

You can now open and analyze the log files in Cognos Dynamic Query Analyzer.

## Opening query log files

When you run a report, query log files are created if the **query execution trace** option is enabled. Log files are XML files that capture the native MDX queries executed during a report run, along with execution and waiting time metrics for query constructs.

The query log files are contained in a folder named `<time-stamp>_<report_name>` inside the XQE logs folder. This folder contains one or more profile files named `profilingLog-0.xml`, `profilingLog-1.xml`, and so on. If the report execution requires executing subqueries, there is a subfolder named `subqueries` that contains folders with profile files for each subquery.

After a report is run, you can open the query log files in IBM Cognos Dynamic Query Analyzer. You can browse and open log files using any of three methods.

### Opening query log files from the Content Store view

You can use the **Content Store** view to browse and open query log files.

#### Procedure

1. If the **Content Store** view is not open, click **Window, Show View** and select the **Content Store** view.
2. In the **Content Store** view, navigate to the required report and click the arrow next to it. The report entry expands to display the profile log files and subquery folders for the report.
3. Double-click a profile entry to open the file, or click the arrow next to a subquery folder, and Double-click a profile entry to open the file.

#### Results

For each open query log file, a graph displays the profile file. **Summary**, **Query**, and **Properties** views displaying data related to the active graph also open.

## Opening query log files from the File menu

You can browse and open query log files using the **File** menu

### Procedure

1. Click **File, Open log**.
2. You can browse query log files in the following ways.
  - Select **Previously opened** to see a list of log files that were previously opened and locally cached by IBM Cognos Dynamic Query Analyzer.
  - Select **From directory** to see a list of log files accessible from a file location.
  - Select **From URL** to see a list of log files accessible using the `http://` or `file://` web browser protocol.
3. Click the arrow next to a report folder to display the profile log files and subquery folders for the report.
  - Select the check box next to a profile log file to open that query log.
  - Select the check box next to a report folder to open all query log files contained in the folder.
  - Select multiple check boxes to open multiple query log files.

Opening many query log files at the same time can cause an out-of-memory error in Cognos Dynamic Query Analyzer.

### Results

For each open query log file, a graph displays the profile file. **Summary, Query,** and **Properties** views displaying data related to the active graph also open.

## Opening query log files from the Report Logs view

You can use the **Report Logs** view to browse and open query log files.

### Procedure

1. If the **Report Logs** view is not open, click **Window, Show View** and select the **Report Logs** view.
2. In the **Report Logs** view, click the arrow next to the required report. The report entry expands to display the profile log files and subquery folders for the report.
3. Double-click a profile entry to open the file, or click the arrow next to a subquery folder, and double-click a profile entry to open the file.

### Results

For each open query log file, a graph displays the profile file. **Summary, Query,** and **Properties** views displaying data related to the active graph also open.

## Exploring query log files

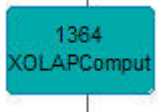


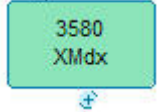

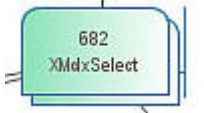
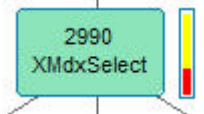
When you open query log files in IBM Cognos Dynamic Query Analyzer, you can view and analyze the data in a number of ways.

### Graph

The graph that is displayed when a query log file is opened shows a series of linked nodes. Each node represents an operation that occurred when the report

was run or represents an attribute of an operation (such as the data being processed). The meaning of the different node representations is defined in the following table.

Table 3. Node representations

Node	Description
	Represents an operation that occurred when the report was run.
	Represents an attribute of an operation, such as the data being processed. By default, most of these nodes are suppressed in the graph. To display all of them, click <b>Window, Preferences</b> , and select <b>Visualization, Node filtering</b> .
	Represents a collapsed node. To display the hidden nodes, double-click the node, or right-click it and click <b>Show Subtree</b> .
	Represents a node that contains hidden children due to the <b>Node filtering</b> settings. To display the hidden nodes, right-click the node and click <b>Expand Filtered</b> .
	Represents the node that is currently selected. The properties of this node are displayed in the <b>Properties</b> view. You can select a node by clicking it.
	Represents a node that has subqueries that can be opened in another graph. To display the subquery graphs, right-click the node and click <b>Open sub queries</b> . Some subqueries do not have an associated node in the parent trace and can be opened using this option.
	Represents timing information for the node. The color red represents the time for the report spent in the node. The color yellow represents the proportion of time spent in the children of the node. The color gray represents the time spent outside the node and its children. If the node is selected, the times are also displayed in the <b>Properties</b> view.

## Summary view

The **Summary** view provides information about the active graph.

The **Summary** section displays the name, package, and timestamp of the query. It also indicates whether the data source type is relational, OLAP, or DMR (Dimensionally Modeled Relational).

The **Timing** section displays timing information. It displays the working times and waiting times for each node that has timing information attached to it, in




descending time order. If you click any node, the corresponding node in the graph is selected. Timing information for a query is not displayed if all the timing information is contained in subqueries.


The **Analysis** section shows information if the query planning trace option is enabled.

The **Node Shapes and Colors** section provides the meaning of node shapes and colors.

## Query view

The **Query** view shows the MDX or SQL query used in this report.

You can link MDX commands in the **Query** view with nodes in the graph by using the **Link MDX to graph** icon . If you select an MDX command in the **Query** view and click the **Link MDX to graph** icon, the nodes in the graph where the MDX commands are executed are selected.


If the query is a SQL query, you can run the report again by clicking the **Execute Sql query** icon .

## Properties view

The **Properties** view displays the properties for the currently selected graph node.

## Navigation view

The **Navigation** view displays the graph as a tree structure. To open the **Navigation** view for the active graph, click **File, Show in Navigation**. You can have multiple **Navigation** views open, one for each graph.

To navigate quickly between a **Navigation** view and the corresponding graph, click the **Link with editor** icon . Clicking an item in the **Navigation** view selects the corresponding node in the graph, and vice versa.

---

## DQM server logs

The **DQM Server Logs** view displays the dynamic query mode server logs. These logs contain information about the state of the query server. If you use SAP data sources these logs will log calls to the data sources. Dynamic cube logging information is also shown when cubes are started.

You can enable dynamic query mode server logs for reports run in IBM Cognos Dynamic Query Analyzer by enabling **Dynamic query logging** in the **General** tab of **Preferences**. To view server log entries for an active report, click **File, Show in Server Log**.

Server logs are organized by server session. Click the arrow next to a server session to see detailed logging information.



---

## Chapter 2. Troubleshooting Cognos Dynamic Query Analyzer

This section provides information about potential problems you might encounter when using IBM Cognos Dynamic Query Analyzer and provides solutions and workarounds.

For troubleshooting information that is not specific to Cognos Dynamic Query Analyzer, see the Troubleshooting section of the *IBM Cognos Business Intelligence Administration and Security Guide*. You can also refer to component-specific documents.



---

## Chapter 3. Accessibility features in Cognos Dynamic Query Analyzer

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products.

The major accessibility features are accelerators and command keys that you can use to navigate through IBM Cognos Dynamic Query Analyzer. You can also use keyboard shortcuts to move among graphs and views.

- You can use command keys, or shortcut keys, to navigate through Cognos Dynamic Query Analyzer. Shortcut keys directly trigger an action and usually make use of the Ctrl keys.
- Graphs use color to display timing information. However, if you select a node, timing information is also displayed in the **Properties** view.
- The **Details** tab of the Aggregate Advisor progress indicator displays the advisor progress as text.
- Colors and fonts used by Cognos Dynamic Query Analyzer can be modified on the **Colors and Fonts** tab of the **Preferences** window.

### Keyboard shortcuts

This product uses the following keyboard shortcuts.

*Table 4. Keyboard shortcuts for Cognos Dynamic Query Analyzer*

Action	Keyboard shortcuts
Display the list of open graphs.	Ctrl+F6
Display the list of supporting views for the active graph.	Ctrl+F7
Traverse nodes in graphs and in the <b>Navigation</b> view.	Arrow keys
Cycle through a list of graphs or views.	Arrow keys
Cycle through global toolbar commands and toolbar commands for the active view.	Tab key
Open the context-sensitive menu for a view or graph.	Alt+-
Open the context-sensitive menu for a node.	Shift+F10

---

## IBM and accessibility

For more information about the commitment that IBM has to accessibility, see the IBM Accessibility Center (<http://www.ibm.com/able>).



---

## Notices

This information was developed for products and services offered worldwide.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. This document may describe products, services, or features that are not included in the Program or license entitlement that you have purchased.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group  
Attention: Licensing  
3755 Riverside Dr  
Ottawa, ON K1V 1B7  
Canada

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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



---

## Trademarks

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

The following terms are trademarks or registered trademarks of other companies:

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





---

# Index

## A

- accessibility features 15
- Advisor results view 3
- Aggregate Advisor 1, 3, 6
- Aggregate Advisor recommendations 3
- aggregate tables
  - loading 4
  - reusing 4
  - schedule loading of 7

## C

- change data capture 4, 6

## D

- DQM Server Logs view 11
- dynamic cube
  - optimizing 1
- Dynamic Query Advisor
  - navigating 1

## E

- event
  - schedule 7

## G

- graph 9

## I

- in-database recommendations 4
- incremental load 4, 7
- initial load 4

## N

- Navigation view 11

## P

- Properties view 11

## Q

- query execution trace
  - enabling in Cognos Administration 7
- query log files
  - analyzing 7
  - location 8
  - open using Content Store view 8
  - open using File menu 9
  - open using Report Logs view 9
  - viewing 9
- Query view 11

## R

- reports 7
- running in BI server 7

## S

- stored procedures 6
- Summary view 10

## T

- tab group 1
- troubleshooting 13

## V

- view 1
  - Advisor results 3
  - detaching 1
  - DQM Server Logs 11
  - moving 1
  - Navigation 11
  - Properties 11
  - Query 11
  - Summary 10