

QMF Data Service
13.1

Studio User's Guide



Note

Before using this information and the product it supports, be sure to read the general information under "Notices" at the end of this information.

May 31, 2022 edition

This edition applies to Version 13 Release 1 IBM® Db2 Query Management Facility (QMF) Classic Edition and Enterprise Edition, which are features of IBM Db2 13 for z/OS (5698-DB2). It also applies to Version 13 Release 1 of IBM Db2 QMF for z/OS (5698-QMF), which is a stand-alone IBM Db2 for z/OS tool. This information applies to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation .**

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

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

Contents

- About IBM® DB2 QMF® Data Service..... V**
- What's new in IBM® DB2 QMF® Data Service Studio User's Guide..... vii**
- Chapter 1. Data Service Studio overview..... 1**
- Chapter 2. Data Service perspective.....3**
- Chapter 3. Connecting to the Data Service server5**
 - Connecting to the Data Service server..... 5
 - Completing the configuration of DRDA access to RDBMS data sources 5
 - Locale considerations 6
- Chapter 4. Creating server metadata.....9**
 - Creating virtual source libraries 9
 - Creating virtual tables.....10
 - Creating virtual tables for Adabas data..... 11
 - Creating virtual tables for RDBMS data sources..... 13
 - Creating virtual tables for CA IDMS data..... 15
 - Creating virtual tables for IMS data..... 16
 - Creating virtual tables for IBM MQ..... 21
 - Creating virtual tables for sequential data..... 23
 - Creating virtual tables for zFS and HFS file system data..... 25
 - Creating virtual tables for VSAM, VSAM CICS, and IAM data.....27
 - Creating virtual tables for VSAM and sequential access using ADDI.....29
 - Creating virtual tables for VSAM and sequential access using RAA..... 31
 - Creating virtual views..... 33
- Chapter 5. Accessing IT Operational Analytics data..... 35**
 - System Management File sample code..... 35
- Chapter 6. Accessing DB2 unload data..... 37**
- Chapter 7. Generating and executing SQL queries.....39**
- Chapter 8. Server Trace..... 41**
 - Enabling Studio calls in the Server Trace results..... 41
 - Starting Server Trace..... 42
 - Filtering Server Trace results.....42
 - Using Server Trace Zoom.....42
 - Searching Server Trace messages.....43
 - Labeling Server Trace messages..... 43
 - Exporting Server Trace messages..... 44
 - Importing Server Trace messages.....45
- Chapter 9. Data Service preferences..... 47**
 - Data Service preferences.....47
 - Admin preferences.....47
 - Console preferences..... 48

Dictionary preferences.....	48
Driver preferences.....	48
SQL preferences.....	49
JCL preferences.....	50
Metadata Discovery preferences.....	51
SSL preferences.....	52
Notices.....	55
Trademarks.....	56
Terms and conditions for product documentation.....	57
Privacy policy considerations.....	57
Index.....	59

About IBM® DB2 QMF® Data Service

IBM® DB2 QMF® Data Service (Data Service) provides integration facilities for both mainframe data sources and other off-platform data sources. Data Service provides optimized, virtualized, and parallelized access to a wide variety of data.

Use the Data Service Studio user interface to create the virtual source libraries and virtual tables that map to and represent the data that you want to access.

What's new in IBM® DB2 QMF® Data Service Studio User's Guide

This section describes recent technical changes to IBM® DB2 QMF® Data Service.

New and changed information is marked like this paragraph. Editorial changes that have no technical significance are not marked.

Description	Related APARs
When using SMF log streams, you can use the LS_TIMESTAMP and LS_TIMESTAMP_LOCAL virtual columns to retrieve timestamp values. When used in a WHERE predicate, the timestamp is searched using the respective time zone. See “System Management File sample code” on page 35.	PH01595 PI99341
Db2 Direct is a new Data Service server access method used to access Db2 data by reading the data in the underlying Db2 VSAM linear data sets directly. The Db2 data access method is specified when creating virtual tables for access to Db2 data. See “Creating virtual tables for RDBMS data sources” on page 13.	PI95733
The process of creating maps to access VSAM and sequential data has been simplified by support of the following methods: <ul style="list-style-type: none">• Querying information in the IBM Application Discovery and Delivery Intelligence (ADDI) dictionary. See “Creating virtual tables for VSAM and sequential access using ADDI” on page 29.• Querying information in the IBM Rational Asset Analyzer (RAA) dictionary. See “Creating virtual tables for VSAM and sequential access using RAA” on page 31 and “Metadata Discovery preferences” on page 51.	PI94393
SQL query access to DB2 unload data sets is now provided. See Chapter 6, “Accessing DB2 unload data,” on page 37.	PI94374
When connecting from the Data Service Studio to the Data Service server, password phrase authentication is supported. See “Connecting to the Data Service server” on page 5.	PI94316
IMS Direct supports access to multiple IMS subsystems. The IMS ID Override field has been added to the IMS DBD Metadata wizard. See “Using the IMS DBD Metadata wizard” on page 17.	PI89834

Chapter 1. Data Service Studio overview

This topic introduces the Data Service Studio and the function it provides in IBM® DB2 QMF® Data Service.

Data Service enables users to have seamless access to mainframe data without needing to know technical details, such as how the data is formatted or where it is located. Data Service provides data integration without the complexity and cost associated with extracting, transforming, and loading the data.

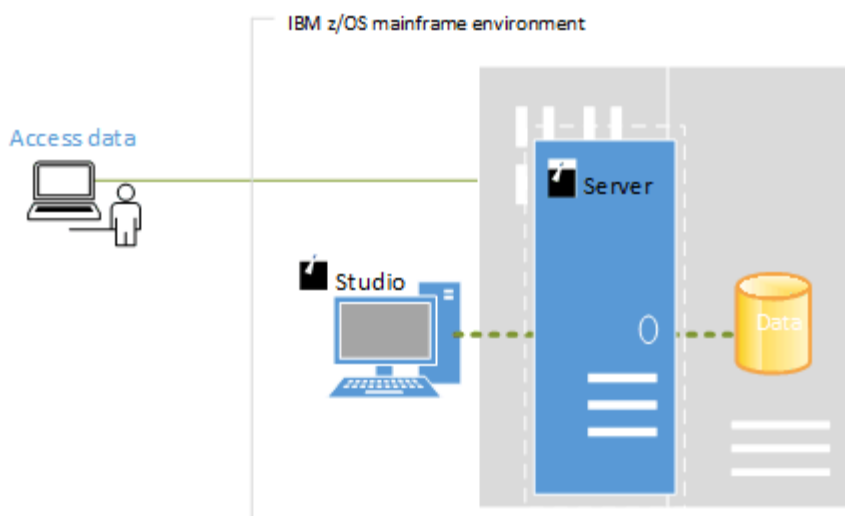
The following mainframe data sources are supported by Data Service:

- Adabas
- Db2 for z/OS
- CA IDMS
- IMS
- IBM MQ
- Sequential
- VSAM, VSAM CICS, and IAM
- zFS and HFS

Support is also provided for Oracle, Microsoft SQL Server, and Db2 for LUW data sources. In addition, you can access the systems data in System Management Files (SMF), System Log (SysLog) files, and Operational Log (OperLog) files, from which you can perform IT Operational Analytics (ITOA).

Data Service Studio is an Eclipse-based user interface that communicates with the Data Service server. You use the studio to create metadata objects, such as virtual source libraries, virtual tables and virtual views on the host Data Service server. These server metadata objects provide the information that is necessary to access your mainframe data and off-host RDBMS data sources by virtualizing your data.

The illustration that follows shows the main components within the IBM® DB2 QMF® Data Service architecture.



Getting started with the studio

The following list highlights the steps to start using Data Service Studio:

- Before you can begin using the Data Service Studio, you must first open the Data Service perspective from the **Window** menu, and then connect to the Data Service server on the mainframe that has the data that you want to access.

- To get access to the mainframe data, use the Data Service Studio to create the following components on the Data Service server:
 - *Virtual source libraries*: A virtual source library is a reference to a library that already exists on the mainframe. Virtual source libraries point to the information (metadata) that is required to virtualize the source data. You create virtual source libraries using the **Virtual Source Library Wizard**.
 - *Virtual tables*: A *virtual table* is a map to the data that you want to access from the data source. After the virtual table is created, use it to generate and execute the SQL. The resulting SQL is used to read and extract the mapped data from the mainframe. You create virtual tables using the **New Virtual Table Wizard**.
 - *Virtual views*: Optionally, you can use virtual tables to create virtual views from which you can generate SQL queries. A virtual view is the SQL SELECT statement that contains the columns from the source data that are used to read data directly from the data source. In some cases, creating a virtual view is more convenient than regenerating and editing the SQL. Virtual views are also used if your virtual table is missing columns or if you want to join different types of data. You create virtual views using the **New Virtual View Wizard**.

Chapter 2. Data Service perspective

The **Data Service** perspective provides the default views, editors, and wizards that you use to perform tasks that are associated with accessing your mainframe data.

Use this perspective to perform the following tasks:

- Explore mainframe resources and view metadata.
- Create and manage data sources.
- Generate and modify SQL queries.
- Create virtual tables from SQL.
- Create virtual views for use with complex SQL queries.
- Generate the code to use in your applications from SQL.

Views

The following views are available with this perspective:

Views	Description
Active Connections	Lists the open JDBC connections between the Studio and one or more servers. The current active connection is used by the SQL Editor to issue SQL queries over that JDBC connection. You can create new or delete existing server connections.
Server view	Lists data resources, stored procedures, and metadata. You can perform tasks on selected objects in the tree. Explorer views include the following tabs: <ul style="list-style-type: none">• Client: Lists information that is related to data sources and application development on your local machine.• Server: Lists the Data Service server to which you want to connect, view resources, or perform tasks.• Network: Lists the host and server connections within your network. You can choose to view or modify existing host and server connection settings.• Favorites: Lists shortcuts to the mainframe resources that you frequently access.
Server Trace	Applies labels to Server Trace messages for use when searching within the Server Trace view.
Lists	Use to display details for each tree node or object that is selected in an Explorer view.
Search	Use to search for a text string within the Server Trace results.
Server Trace	Use to set and gather server diagnostic information for support purposes.

Views	Description
Server Trace Import	Use to import Server Trace (. isx) files.
SQL Results	Use to display the result set returned from a SQL query in the SQL Results tab, and the resulting trace information in the SQL Messages tab.
Studio Navigator	Use to list shortcuts to key task views, wizards, and editors.
Virtualization Facility	Use to display virtual table mapping details.

Editors

The following text editors are available with this perspective:

Editors	Description
Data Source	Use to edit existing connection definitions.
SQL	Use to compose SQL statements and to invoke queries against the server.
Virtualization Facility	Use to edit metadata settings related to virtual tables and virtual views.

Wizards

This perspective includes wizards that guide you through tasks, such as:

- Setting the server connection.
- Creating virtual source libraries.
- Creating virtual tables for SQL access to data.
- Generating application code from SQL.

Chapter 3. Connecting to the Data Service server

To access data on the mainframe, connect Data Service Studio to the Data Service server that is running on an z/OS mainframe instance.

Connecting to the Data Service server

Use the Data Service Studio to connect to the Data Service server that is running on an instance of z/OS.

Before you begin

Before you can connect to the Data Service server, the server must be configured and started.

Procedure

1. From the Data Service Studio menu, click **Window > Open Perspective > Data Service**.
2. On the **Server** tab, click **Set Server**.
3. In the **Set Server** dialog box, complete the following:
 - **Host:** Select or enter the TCP/IP host name or IP address of the mainframe system on which the Data Service server is deployed.
 - **Port:** Enter the port number that the Data Service server uses. The default is 1200.
 - **Userid:** Enter your mainframe user ID.
 - **User Password:** Enter your password or password phrase for the mainframe user ID.
4. Click **OK**.

Completing the configuration of DRDA access to RDBMS data sources

To complete the configuration of DRDA access to RDBMS data sources, you must bind packages on the Data Service server, and grant users the authority to use those packages.

Before you begin

You must know the host name and the port number of the Data Service server and your log on credentials. Your log on credentials must have the authority to bind packages and grant privileges.

About this task

Perform the following task for each RDBMS data source that you want to access.

Procedure

1. From the Studio, click **Window > Open Perspective > Data Service**.
2. On the **Server** tab, click **Set Server**.
3. In the **Set Current Server** dialog box, complete the following fields:

Option	Description
Host	Enter the TCP/IP host name or IP address of the mainframe system.
Port	Enter the port number that is used to communicate with the Data Service server. The default is 1200.
Userid	Enter the mainframe user ID.

Option	Description
User Password	Enter the password for the mainframe user ID.

- Click **OK**.
- On the **Server** tab, expand **SQL > Data > Other Subsystems**.
- Right-click the subsystem and select **BIND/GRANT Packages**.
- On the **BIND/GRANT Packages** page, complete the following fields:

Field	Action
Package Prefix	Enter the two character prefix to assign to the package. The package prefix must match the prefix that is defined on the mainframe server. If you change the default prefix (DS), you must also change it in the <i>hlq</i> .SCQDEXEC (CQDSIN00) file.
Number of Cursors	Enter the number of cursors to use to process results. The default is 200.
Collection	Enter the value to use to bind packages. The default is NULLID . This value is normally determined by the DB2 Administrator.
Table Qualifier	Enter the value to use to qualify unqualified SQL. This value is normally determined by the DB2 Administrator.
Owner UserId	Enter the user ID of the package owner. This value is normally determined by the DB2 Administrator.
Grant to	Set only when granting authority for the target DB2 server. The default is PUBLIC .
Bind Package	Binds the product packages. This is the default setting.
Grant Execute	Grants execute permissions on the package to the user ID that is specified in the Grant to field.
Replace Packages	Replaces an existing package for the specified subsystem. Select this option only if the package already exists.

- Depending on the options that you select, additional dialog boxes and messages might be displayed.
- Review the results in the **Results** text box and click **BIND/GRANT**.

Locale considerations

You can modify the data source connection definitions to use different local code pages.

Before you begin

You have the option to change the default code page (US/English IBM 1047) that the Studio uses to perform character data translations between the native Java character encoding (UTF-8) and the mainframe.

Procedure

- To configure the data source connection definition, in the **Active Connections** view, close all open connections.
- On the **Client** tab, expand **Data Service > Data Sources > JDBC > Default Config File**.
- Right-click the data source that you want to modify and click **Edit**.
- In the **Data Source Editor**, click the **Connection String** tab.
- Add or modify the Charset setting to use the appropriate EBCDIC code page. For example, Charset=IBM037.
- If LGID=ENC exists in the connection string, delete it to avoid a conflict with the Charset setting.

7. Close the **Data Source Editor**.
8. When prompted, click **Yes** to save the data source definition.
9. To change the default Charset that the Studio uses when creating connection definitions, from the **Window** menu select **Preferences**, expand **Data Service > Driver**.
10. In **Connection Overrides**, enter the new Charset setting and click **OK**.
11. On the **Server** tab, expand **SQL > Data**.
12. Right-click the data source to which you want to connect and select **Create Connection Definition (DSN)**.
13. Accept the default name that is displayed or enter a new DSN name and click **OK**.
14. In the **Data Source Editor**, click the **Connection String** tab and confirm that the new Charset setting displays in the connection string.

Results

When running queries using the new data source definition, the character data (including language specific glyphs) that you chose is displayed in the **SQL Results** view.

Chapter 4. Creating server metadata

Using the Data Service Studio, you create the server metadata that provides the information necessary to virtualize your data. Server metadata includes virtual source libraries, virtual tables, and virtual views.

Creating virtual source libraries

Virtual source libraries point to the information that IBM® DB2 QMF® Data Service needs in order to access some types of mainframe data.

Before you begin

A virtual source library is a server metadata object that references a source library that exists on the Data Service (host) server. The members of the source library contain layout information specific to a type of data, for example a COBOL or PL/I copybook (copybook), Adabas Data Definition Module (DDM) views, IMS Database Definition (DBD) files, or IMS Program Specification Block (PSB) files. Virtual source libraries provide a reusable catalog of the host's data source libraries.

Note: When creating a virtual source library, the current user must have read access to the host data source library.

About this task

Virtual source libraries are a prerequisite to creating virtual tables for the following types of data sources:

- Adabas
- IMS
- IBM MQ
- Sequential
- VSAM, VSAM CICS and IAM
- zFS and HFS

When creating the virtual source libraries you specify the following data set (PDS/PDSE) names based on the type of data that you want to access:

- To access Adabas data, you specify the name of the PDS/PDSE that contains the Data Definition Module (DDM) views that have been set up for the Adabas data in your environment.
- To access IMS data, you may need to create multiple virtual source libraries that reference multiple types of source libraries. You may create a separate virtual source library that references the IMS DBD files, the IMS PSB files, and the copybooks that describe the layout of each IMS segment. In each case, you specify the PDS/PDSE that is specific to the source library.
- To access IBM MQ data, you specify the name of the PDS/PDSE that contains the copybook that describes the data written to the queue.
- To access sequential data, you specify the name of the PDS/PDSE that contains the copybook that describes the structure of the sequential data records.
- To access VSAM, VSAM CICS, and IAM data, you specify the name of the PDS/PDSE that contains the copybook that describes the structure of the VSAM, VSAM CICS, and IAM data records.
- To access z/FS and HFS data, you specify the name of the PDS/PDSE that contains the copybook that describes the structure of the records in the data file.

Procedure

1. On the **Server** tab, expand **Admin > Source Libraries**.

- Right-click **Create Virtual Source Library** and select **Create Virtual Source Library**.
- On the **Virtual Source Library** page, complete the following fields:

Field	Action
Name	Enter a unique, meaningful name for the virtual source library you are creating.
Description	Enter an optional description for the virtual source library.
Library Name	Enter the name of the PDS/PDSE that contains the layout information for the data you want to access.

- Click **Finish**.

Results

The new virtual source library is displayed in the **Source Libraries** folder.

Creating virtual tables

To access your data, create a virtual table or virtual view that maps to your source data and that matches the definition of the source data structure on the mainframe.

From the virtual table or virtual view, you generate the SQL that is used to read and access the mapped data from the mainframe. You create virtual tables using the **New Virtual Table Wizard** that is specific to the type of data that you want to access. Some virtual tables, including SMF virtual tables, are made available during the product installation.

The following high-level procedures must be completed prior to creating a virtual table:

- Start the Data Service Studio.
- Open the Data Service perspective.
- Connect to the Data Service server.
- Run the **Create Source Library** wizard to create a virtual source library to map to your mainframe data. This procedure is not required to create virtual tables for RDBMS data.

Virtual table tasks

When a virtual table is selected on the **Server** tab, you can perform the following tasks:

- Edit:** Edit the virtual table properties in the editor.
- Copy** and **Paste:** Copy the virtual table and paste the copy under the **Virtual Tables** node.
- Disable:** Disable the virtual table on this server.
- Delete:** Delete the virtual table from the server.
- Create Virtual View:** Create a virtual view from the virtual table.

Key and index information

To view a summary of key and index information for an existing virtual table, select the virtual table on the **Server** tab and from the **Window** menu, select **Show View > Properties**. The properties for the selected table are displayed in the **Properties** view.

If a virtual table includes columns that have a primary key or an index, the column is notated using the following symbols:

- Key symbol – This column is associated with a primary key.
- Superscript numeral 1 – This column is associated with a unique index, but does not have an associated primary key.
- Superscript asterisk – This column is associated with a non-unique index.

This primary key and index information is also highlighted when you browse RDBMS tables under the **Other Subsystems** tree.

You can control the identification of primary keys and indexes using settings in [SQL preferences](#).

Creating virtual tables for Adabas data

Create a virtual table that maps to the Adabas data that you want to access, and from which the SQL used to access the data is generated and executed.

Before you begin

Have the Adabas database ID and password, the file number, and the subsystem name available.

Procedure

1. Expand the **SQL > Data > SSID** node, where *SSID* is the name of your server.
2. Right-click **Virtual Tables** and select **Create Virtual Table(s)**.
3. Under **Wizards**, select the **ADABAS** wizard and click **Next**.
4. On the **New Virtual Table Wizard** page, complete the following fields and click **Next**:

Field	Action
Name	Enter a unique name. The name can contain a maximum of 50 characters. The name must consist of an uppercase letter followed by zero or more characters, each of which is an uppercase letter, a digit, or the underscore character.
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Arrays Handling	Enable one of the following array management options: <ul style="list-style-type: none"> • Flatten arrays into a single fixed table at runtime: This relates to multiple occurring (MU) fields and periodic (PE) groups. • Return arrays into separate tables at runtime: This relates to multiple occurring (MU) fields and periodic (PE) groups. A subtable is generated for each array. Subtables only support read access.

5. On the **ADABAS Details** page, complete the following fields and click **Next**:

Field	Action
DB ID	Enter the Adabas database ID.
File Number	Enter the number of the file to use.
Adabas Password	If the file is password-protected, enter the password. This password is stored and encrypted in the virtual table so that future queries use the same password to access the data.
SubSystem	Enter the name of the Adabas subsystem.
Max MU Count	Enter the maximum number of times to repeat the MU field. The default is 10.
Max PE Count	Enter the maximum number of times to repeat the PE field. The default is 10.
Create Count Field	Select this check box to index every MU or PE field so that the index (count) field created precedes the repeating field. This index field tells the caller how many repeating fields are being used.
Secure	Select this check box to choose the Adabas file ID number to be used for file name security.

Field	Action
DE Search only	Select this check box if you want the utility to generate control definitions that allow the client to only use WHERE columns that are Adabas descriptors (such as superde, subde, and hyperde).
Search by PE index	Select this check box to allow the client to target rows that match a particular occurrence of the PE field when searching rows using the WHERE clause. If this parameter is not specified, all rows where any occurrence of that PE field match the value specified will be targeted.
Unpacked to Packed	Select this check box to convert all unpacked format fields to packed format.
Binary to Integer	Select this check box to convert all 2-byte and 4-byte binary fields to short integer and integer formats, respectively.
Binary to Packed	Select this check box to map the binary fields in the Adabas file to SQL decimal columns (numeric packed decimal format) in the generated virtual table. Note the following points: <ul style="list-style-type: none"> • If the precision of the Adabas binary field allows for the possibility of a numeric value that would cause data overflow when converted to SQL decimal, the column in the virtual table will be mapped to SQL binary instead. This means that Adabas fields with precision greater than 12 will continue to be mapped to SQL binary. • If you select the Binary to Integer check box and the Binary to Packed check box, the precision of the Adabas binary field will determine if it gets mapped to an SQL integer (that is, 2-byte or 4-byte fields) or a decimal type.
Advanced	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

6. Optional: On the **Data Definition Module** page, if you have a Natural Data Definition Module (DDM) listing of the file, you can complete the following to get additional metadata information:

Field	Action
Available Source Libraries	From the list of Available Source Libraries , select the virtual source library that contains the data structure definition that you want the virtual table to use.
Source Library Members	Select the names of each virtual source library member that represents the data structure that you want to include. The green arrow next to a DDM indicates that it is a suggested member, not that it is selected.

7. On the **Virtual Table Layout** page, complete the following fields and click **Next**:

Field	Action
Source	Expand the source file to verify that it displays the expected data layout.
Start Field	This field is not supported for Adabas because the entire data layout is used.
End Field	This field is not supported for Adabas because the entire data layout is used.

8. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7](#), “Generating and executing SQL queries,” on page 39.

Important: Use caution when using the `BASE_KEY` in `WHERE` predicates, (for example, `[PARENT TABLE].BASE_KEY = [CHILD TABLE].PARENT_KEY`) when joining the parent table with a child subtable, since this will result in a table scan of the entire Adabas file. It is recommended instead to use the `CHILD_KEY` (for example, `[PARENT TABLE].CHILD_KEY = [CHILD TABLE].PARENT_KEY`).

Creating virtual tables for RDBMS data sources

Create virtual tables that map to RDBMS data sources, such as Db2 for z/OS, Db2 LUW (Linux, UNIX, and Windows), Oracle, and Microsoft SQL Server.

About this task

It is recommended that you create a virtual table for each RDBMS table from which you want to access data. Creating a virtual table for each RDBMS table allows you to perform joins across data that may originate from different DRDA accessible RDBMS subsystems or to perform joins between your RDBMS data and other types of virtualized data, such as IMS or VSAM data.

This wizard allows you to create multiple virtual tables at a time if the selected source tables belong to the same RDBMS subsystem. In this wizard, a view is treated the same as a table; each table or view is mapped to a virtual table.

Db2 data access method:

When virtual tables are created for access to Db2 for z/OS data, an option is available to select the access method. *Db2 Direct* is a Data Service server access method that reads the data in the Db2 VSAM linear data sets directly instead of accessing the data through traditional Db2 APIs. For more information, see “Db2 for z/OS data access methods” in the *Solution Configuration Guide*.

Note: The data access method options are not displayed if the Data Service server does not support Db2 Direct.

Procedure

1. On the **Server** tab, explore the RDBMS metadata information by expanding the **SQL > Data > Other Subsystems** node, and then navigating down the appropriate subtree. The hierarchy begins with the subsystem, followed by the schema, and then the tables and views.
2. Select a single table or view from the tree, or use the following techniques to select multiple tables or views:
 - To select more than one individual node, hold down the Ctrl key and click each node to be included.
 - To select a range of tables (or views), click the first table in the range, and then hold the Shift key and select the last table in the range. All tables within the range will be included.
 - To select a group of nodes, click the parent node. All of the children under the parent node will be included. For example, select the **Tables** node to include all tables belonging to that schema. Or, select the schema node to include all tables and views under that schema.

You can use a combination of these techniques. For example, you can select two schema nodes to create virtual tables for all tables and views belonging to those two schemas.

3. Right-click the selected items and select **Create Virtual Table(s)**. The **New Virtual Tables Wizard** launches.
4. On the **New Virtual Tables for DBMS access** page, complete the following fields:

Field	Action
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.

Field	Action
Description	Enter an optional description.
Naming Pattern	Specify the format to use for the generated virtual table names. Use the following variables to create naming patterns that are derived from the RDBMS metadata: <ul style="list-style-type: none"> • {Subsystem}: Subsystem name • {Schema}: Source schema name • {Table}: Source table name
Virtual Target System	Select a virtual target system from the drop-down list. A virtual target system points to the RDBMS subsystem that contains the data that you want to access using the current virtual table. If there are no virtual target systems in the drop-down list, click Create Target System to create one. <p>By using virtual target systems, you can easily change the name of the RDBMS subsystem that is referenced in the virtual tables. For example, you create a virtual target system called TSDSN1, and specify that it will access the RDBMS subsystem DSN1. Then, you create 50 virtual tables that access data in the RDBMS source TSDSN1 (that is, pointing to DSN1). If it becomes necessary to change the name of the RDBMS source DSN1, you only have to change it in a single place by editing the virtual target system. These target systems can be located under the SQL > Target Systems > DBMS node in the server view tree.</p>
<ul style="list-style-type: none"> • Use traditional DB2 access (read/write, transactional integrity) • Use DB2-Direct access (read-only, high performance bulk data access) 	Select the access method to use when accessing Db2 for z/OS data. <p>Choose Use traditional DB2 access (read/write, transactional integrity) to use Db2 APIs such as DRDA, CAF, and RRSAF. This is the default selection.</p> <p>Choose Use DB2-Direct access (read-only, high performance bulk data access) to use Db2 Direct.</p> <p>Note: These options are available only when creating virtual tables for access to Db2 for z/OS data and if the Data Service server supports Db2 Direct.</p>
Advanced	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

5. In the results table, review the list of selected entries. Modify the selections as needed.

Tip: Use the check box in the header row of the table to control the selection of all entries.

6. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for CA IDMS data

Create virtual tables that map to the CA IDMS data that you want to access and from which the SQL used to access the data is generated and executed.

Before you begin

The Data Service server must be configured for CA IDMS access, and the CA IDMS central version referenced by the data server SYSQL DD statement must be active.

About this task

CA IDMS schema records are mapped using the CA IDMS data dictionary. Each record is mapped as a separate virtual table using the COBOL names to derive the SQL column names. In addition to records, schema sets can be mapped as well. Virtual tables created for CA IDMS sets serve as correlation tables between CA IDMS records so SQL joins can navigate the CA IDMS schema.

Procedure

1. On the **Server** tab, explore the CA IDMS metadata information by expanding the **Discovery > IDMS** node, and then navigating down the appropriate subtree. The hierarchy begins with the data dictionary, followed by the CA IDMS schema, the CA IDMS subschema, and then the associated records and sets.
2. Select one or more records, as follows:
 - To select individual records, hold down the Ctrl key and click each record to include.
 - To select a range of records, click the first record in the range, and then hold the Shift key and select the last record in the range. All records within the range will be included.
 - To select all child records under a parent, click the parent record.
3. Right-click the selected records and select **Create Virtual Table(s)**. The **New Virtual Tables Wizard** launches.

Note: You can map the relevant CA IDMS sets in the wizard.

4. On the **Create IDMS virtual tables** page, complete the following **Common Virtual Table Settings**:

Field	Description
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Arrays Handling	Select one of the following options: <ul style="list-style-type: none">• Flatten arrays into a single fixed table at runtime (Y): This option supports both OCCURS and OCCURS DEPENDING ON statements.• Return arrays into separate tables at runtime (N): This option supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables support SQL read access only.
Virtual Table Naming Patterns	Specify the format to use for the generated virtual table names. You can specify different patterns for records and sets. Use the following variables to create naming patterns that are derived from the IDMS metadata: <ul style="list-style-type: none">• {SubSchema}: Subschema name• {Record}: Record name• {Set}: Set name

Field	Description
Prune IDMS record field suffix from column names	Select this option to remove the IDMS record field suffix from the column names.

5. In the table that lists the IDMS records, review the list of selected entries. Modify the selections as needed.

Tip: Use the check box in the header row of the table to control the selection of all entries.

6. To map the sets, click **Fetch Related IDMS Sets**. The Studio collects additional metadata from the server and displays the relevant items in the table that lists the IDMS sets.

7. In the table that lists the IDMS sets, review the list of selected entries. Modify the selections as needed.

8. To disable MapReduce, click **Advanced** and select **Disable MapReduce**.

9. Click **Finish**.

Results

The Studio creates the virtual tables (the metadata maps) on the server.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for IMS data

Create a virtual table that maps to the IMS data that you want to access, and from which the SQL used to access the data is generated and executed.

Before you begin

The Program Specification Block (PSB) and Database Definition (DBD) source members, and the copybooks for each segment must exist in the virtual source libraries defined to the server. For details, see [“Creating virtual source libraries ”](#) on page 9.

To use the IMS Direct feature, the IMSDIRECTENABLED parameter must be enabled in the Data Service server IN00 file.

About this task

When an IMS SQL query is run, the SQL Engine for the server will determine if the request is best executed using IMS Direct (native file support) or if IMS APIs are required. The determination is based on the database and file types supported as well as the size of the database.

Procedure

1. Expand the **SQL > Data > SSID** node, where *SSID* is the name of your server.
2. Right-click **Virtual Tables** and select **Create Virtual Table(s)**.
3. Under **Wizards**, select the **IMS** wizard and click **Next**.
4. On the **New IMS virtual Table(s)** page, create metadata for an IMS virtual table by completing the following steps:
 - a) Choose a DBD by doing one of the following steps:
 - Select a **DBD** from the drop-down list.
 - If your DBD does not appear in the drop-down list, click **Extract DBD** to create the requisite metadata. The **New IMS DBD Metadata Wizard** launches. See [“Using the IMS DBD Metadata wizard”](#) on page 17.

b) Choose a PSB by doing one of the following steps:

- Select a **PSB** from the drop-down list.
- If your PSB does not appear in the drop-down list, click **Extract PSB** to create the requisite metadata. The **New IMS PSB Metadata Wizard** launches. See [“Using the IMS PSB Metadata wizard”](#) on page 18.

c) Click **Create Virtual Table** to create a virtual table for an IMS segment in the selected DBD and PSB. The **New Virtual Table Wizard** launches. See [“Using the IMS Virtual Table wizard”](#) on page 19.

Note: Both the DBD and PSB must be defined for this button to be enabled.

5. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Using the IMS DBD Metadata wizard

Use the **New IMS DBD Metadata Wizard** to create DBD server metadata.

About this task

This wizard is used to create server metadata containing information extracted from the selected DBD source. This DBD metadata is a prerequisite for creating IMS virtual tables. The name of each DBD map will be determined from the contents of the DBD source.

Procedure

1. On the **New DBD Metadata** page, complete the following fields and click **Next**:

Field	Action
Metadata Library	From the drop-down list, select the target library where the DBD metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.

2. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Available Source Libraries	From the list of Available Source Libraries , select the virtual source library that contains the DBD source member.
Source Library Members	Select the DBD that you want to use and click Download to copy the member from the mainframe to your desktop. Use Filter patterns to filter the list.
Downloaded Source Files	Review the list of downloaded members and ensure that the check box for the DBD that you want to use has been selected.

3. On the **Data Layout** page, complete the following fields and click **Next**:

Field	Action
Source	Expand the source file to verify that it displays the expected database definition (DBD).
Start Field	Accept the default root start field, or if multiple DBD nodes are present in the source tree, you can click on one of the DBD nodes to indicate that you only want to map that one DBD.
End Field	End Field selection is disabled when extracting DBD source.

4. On the **IMS Server configuration** page, complete the following fields:

Field	Action
<ul style="list-style-type: none"> • Use IMS/DBCTL (read/write, transactional integrity) • Use IMS-Direct (read-only, high performance bulk data access) 	<p>Select the IMS protocol to use.</p> <p>Choose Use IMS/DBCTL (read/write, transactional integrity) to use IMS API calls.</p> <p>Choose the default option Use IMS-Direct (read-only, high performance bulk data access) to enable IMS Direct for the DBD. To use this feature, IMS Direct must also be enabled in the Data Service server IN00 file. You must select this option for the DBD to be able to enable IMS Direct for a virtual table.</p>
IMS ID Override (used with IMS-Direct only)	Specify the IMS ID of the IMS subsystem to use when multiple IMS subsystems are defined for use with IMS Direct. This value will override the default IMS ID in the DBD map.
Advanced	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

5. Click **Finish**.

What to do next

Return to the **New IMS Virtual Table(s)** page and define the IMS PSB. See [“Creating virtual tables for IMS data”](#) on page 16.

Using the IMS PSB Metadata wizard

Use the **New IMS PSB Metadata Wizard** to create PSB server metadata.

About this task

This wizard is used to create server metadata containing information extracted from the selected PSB source. This PSB metadata is a prerequisite for creating IMS virtual tables. The name of each PSB map will be determined from the contents of the PSB source.

Procedure

1. On the **New PSB Metadata** page, complete the following fields and click **Next**:

Field	Action
Metadata Library	From the drop-down list, select the target library where the PSB metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.

2. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Available Source Libraries	From the list of Available Source Libraries , select the virtual source library that contains the PSB source member.

Field	Action
Source Library Members	Select the PSB that you want to use and click Download to copy the member from the mainframe to your desktop. Use Filter patterns to filter the list.
Downloaded Source Files	Review the list of downloaded members and ensure that the check box for the PSB that you want to use has been selected.

3. On the **Data Layout** page, complete the following fields and click **Next**:

Field	Action
Source	Expand the source file to verify that it displays the expected program specification block (PSB).
Start Field	Accept the default root start field, or if multiple PSB nodes are present in the source tree, you can click on one of the PSB nodes to indicate that you only want to map that one PSB.
End Field	End Field selection is disabled when extracting DBD source.

4. Click **Finish**.

What to do next

Return to the **New IMS Virtual Table(s)** page and create the virtual table. See [“Creating virtual tables for IMS data”](#) on page 16.

Using the IMS Virtual Table wizard

Use the **New Virtual Table Wizard** to create a new IMS virtual table.

About this task

This wizard is used to map an IMS segment using a copybook representation to produce a new IMS virtual table.

Procedure

1. On the **New IMS Virtual Table** page, complete the following fields and click **Next**:

Field	Action
Name	Enter a unique name. The name can contain a maximum of 50 characters. The name must consist of an uppercase letter followed by zero or more characters, each of which is an uppercase letter, a digit, or the underscore character.
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Convert VAR* fields to True VAR* fields	This is a deprecated field and should not be selected.
Arrays Handling	Select one of the following options: <ul style="list-style-type: none"> • Flatten arrays into a single fixed table at runtime (Y): This option supports both OCCURS and OCCURS DEPENDING ON statements. • Return arrays into separate tables at runtime (N): This option supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables support SQL read access only.

2. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Available Source Libraries	From the list of Available Source Libraries , select the virtual source library that contains the data structure definition that you want the virtual table to use.
Source Library Members	Select the PDS members that represent the data structures to include and click Download to copy the members from the mainframe to your desktop.
Downloaded Source Files	Select one or more previously downloaded members.

3. On the **Virtual Table Layout** page, complete the following fields and click **Next**:

Field	Action
Source	<p>Browse the source tree to verify that it displays the expected data layout. By default, all of the fields in the tree will be included in the mapping. To include only a subset of the fields for the mapping, modify the start field value and, optionally, the end field value, as follows:</p> <ul style="list-style-type: none"> • For the start field, accept the default root start field, or expand the tree and select a different start field. When selecting a different start field, Enable End Field Selection must not be selected. • For the end field, accept the default end field, or expand the tree and select a different end field. When selecting a different end field, Enable End Field Selection must be selected.
Start Field	Identifies the first field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is not selected, and select a different start field in the Source tree.
Enable End Field Selection	Use this field to control selection of the start field and end field values in the Source tree. When this option is not selected (default), you can select the start field. When this option is selected, you can select the end field.
End Field	Identifies the last field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is selected, and select a different end field in the Source tree.

4. On the **IMS Information** page, complete the following fields:

Field	Action
Segment Name	From the drop-down list, select the segment name.
<ul style="list-style-type: none"> • Use IMS/DBCTL (read/write, transactional integrity) • Use IMS-Direct (read-only, high performance bulk data access) 	<p>Select the IMS protocol to use.</p> <p>Choose the default option Use IMS/DBCTL (read/write, transactional integrity) to use IMS API calls.</p> <p>Choose Use IMS-Direct (read-only, high performance bulk data access) to enable IMS Direct on the virtual table. To use this feature, IMS Direct must also be enabled for the selected DBD and enabled in the Data Service server IN00 file.</p>
Advanced	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor;

Field	Action
	resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

5. Click **Finish**.

What to do next

Return to the **New IMS Virtual Table(s)** page and if necessary create the next virtual table. See [“Creating virtual tables for IMS data”](#) on page 16.

Creating virtual tables for IBM MQ

Create a virtual table that maps to the IBM MQ data that you want to access, and from which the SQL used to access the data is generated and executed.

Before you begin

Before creating the virtual table, verify that the MQ queue exists and that the copybook exists in the source library. If you use delimited data, configure support for delimited data processing. See "Configuring delimited data support" in the *Customization Guide*.

About this task

Data in MQ queues is described using COBOL or PLI data descriptions taken from copybooks or programs.

Procedure

1. Expand the **SQL > Data > SSID** node, where *SSID* is the name of your server.
2. Right-click **Virtual Tables** and select **Create Virtual Table(s)**.
3. Under **Wizards**, select the **MQ** wizard and click **Next**.
4. On the **New Virtual Table Wizard** page, complete the following fields and click **Next**:

Field	Action
Name	Enter a unique name. The name can contain a maximum of 50 characters. The name must consist of an uppercase letter followed by zero or more characters, each of which is an uppercase letter, a digit, or the underscore character.
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Arrays Handling	Enable one of the following array management options: <ul style="list-style-type: none"> • Flatten arrays into a single fixed table at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements. • Return arrays into separate tables at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables only support SQL read access.

5. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Available Source Libraries	Select the source library that contains the data structure to use.

Field	Action
Source Library Members	Select the PDS members that represent the data structures to include and click Download to copy the members from the mainframe to your desktop. Use Filter patterns to filter the list.
Downloaded Source Files	Select one or more previously downloaded members.

6. On the **Virtual Table Layout** page, complete the following fields and click **Next**:

Field	Action
Source	<p>Browse the source tree to verify that it displays the expected data layout. By default, all of the fields in the tree will be included in the mapping. To include only a subset of the fields for the mapping, modify the start field value and, optionally, the end field value, as follows:</p> <ul style="list-style-type: none"> • For the start field, accept the default root start field, or expand the tree and select a different start field. When selecting a different start field, Enable End Field Selection must not be selected. • For the end field, accept the default end field, or expand the tree and select a different end field. When selecting a different end field, Enable End Field Selection must be selected.
Start Field	Identifies the first field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is not selected, and select a different start field in the Source tree.
Enable End Field Selection	Use this field to control selection of the start field and end field values in the Source tree. When this option is not selected (default), you can select the start field. When this option is selected, you can select the end field.
End Field	Identifies the last field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is selected, and select a different end field in the Source tree.

7. On the **MQ Details** page, complete the following fields:

Field	Action
Queue Manager Name	Enter the IBM MQ queue manager name. The name is a four-character subsystem name.
Queue Name	Enter the IBM MQ queue name. The name can contain a maximum of 48 characters and must comply with z/OS data set naming standards.
Post-Read Exit Name	To manipulate the data after reading it from the queue, enter the name of the post-read exit to use. This is the custom exit routine that is installed on the server and is used to perform additional processing after a record is read from the data source.

8. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for sequential data

Create a virtual table that maps to the sequential data that you want to access, and from which the SQL used to access the data is generated and executed.

Before you begin

Before creating the virtual table, verify that the data set name exists and that the copybook exists in the source library.

Procedure

1. Expand the **SQL > Data > SSID** node, where *SSID* is the name of your server.
2. Right-click **Virtual Tables** and select **Create Virtual Table(s)**.
3. Under **Wizards**, select the **Sequential** wizard and click **Next**.
4. On the **New Virtual Table Wizard** page, complete the following fields and click **Next**:

Field	Action
Name	Enter a unique name. The name can contain a maximum of 50 characters. The name must consist of an uppercase letter followed by zero or more characters, each of which is an uppercase letter, a digit, or the underscore character.
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Convert VAR* fields to True VAR* fields	This is a deprecated field and should not be selected.
Arrays Handling	Enable one of the following array management options: <ul style="list-style-type: none">• Flatten arrays into a single fixed table at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements.• Return arrays into separate tables at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables only support SQL read access.• Flatten arrays now: If you select this option, you cannot change array-handling after you save the virtual table.

5. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Available Source Libraries	Select the source library that contains the data structure to use.
Source Library Members	Select the PDS members that represent the data structures to include and click Download to copy the members from the mainframe to your desktop.
Download Source Files	Select one or more previously downloaded members.

6. On the **Virtual Table Layout** page, complete the following fields and click **Next**:

Field	Action
Source	Browse the source tree to verify that it displays the expected data layout. By default, all of the fields in the tree will be included in the mapping. To include only a subset

Field	Action
	<p>of the fields for the mapping, modify the start field value and, optionally, the end field value, as follows:</p> <ul style="list-style-type: none"> For the start field, accept the default root start field, or expand the tree and select a different start field. When selecting a different start field, Enable End Field Selection must not be selected. For the end field, accept the default end field, or expand the tree and select a different end field. When selecting a different end field, Enable End Field Selection must be selected.
Start Field	Identifies the first field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is not selected, and select a different start field in the Source tree.
Enable End Field Selection	Use this field to control selection of the start field and end field values in the Source tree. When this option is not selected (default), you can select the start field. When this option is selected, you can select the end field.
End Field	Identifies the last field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is selected, and select a different end field in the Source tree.

7. Optional: On the **Virtual Table Redefines** page, accept the default table redefines or expand **Redefine** to modify your selection, and click **Next**.

8. On the **Data Source Details** page, complete the following data source fields and click **Next**:

Field	Action
Data Set Name	<p>Enter the data set name you want to use. The following data set types are supported:</p> <ul style="list-style-type: none"> PDS or PDSE: Specify the partitioned data set name. This requires that you also enter a Member name prior to validating that the member name exists on the host. Physical sequential: Specify the sequential data set name and click Validate to verify that the data set name exists on the host. Generation Data Groups (GDG): Specify the GDG data set using the GDG syntax. For example: <i>hlq.DATA.SEQ(-1)</i>. You can also specify a base GDG name so that all generations of the GDG will potentially be accessed. Click Validate to verify that the data set name exists on the host.
Member	If you selected a PDS or PDSE for the Data Set Name , you must also enter the member name to use. Click Validate to verify that the member name exists on the host.
Post-Read Exit Name	To manipulate the data after reading it from the source file, enter the name of the post-read exit to use. This is the custom exit routine that is installed on the server and is used to perform additional processing after a record is read from the data source.
Advanced	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

9. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for zFS and HFS file system data

Create a virtual table that maps to file data that you want to access on a zFS or HFS file system and from which the SQL used to access the data is generated and executed.

Before you begin

Before creating the virtual table, verify that the PDS members that represent the data structures for the data you want to virtualize already exist in the source library.

Procedure

1. Expand the **SQL > Data > SSID** node, where *SSID* is the name of your server.
2. Right-click **Virtual Tables** and select **Create Virtual Table(s)**.
3. Under **Wizards**, select the **zFS** wizard and click **Next**.
4. On the **New Virtual Table Wizard** page, complete the following fields and click **Next**:

Field	Action
Name	Enter a unique name. The name can contain a maximum of 50 characters. The name must consist of an uppercase letter followed by zero or more characters, each of which is an uppercase letter, a digit, or the underscore character.
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Convert VAR* fields to True VAR* fields	This is a deprecated field and should not be selected.
Arrays Handling	Enable one of the following array management options: <ul style="list-style-type: none">• Flatten arrays into a single fixed table at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements.• Return arrays into separate tables at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables only support SQL read access.

5. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Download Folder	Verify that the appropriate download folder is displayed.
Available Source Libraries	Select the source library that contains the data structure to use.
Source Library Members	Select the PDS members that represent the data structures to include and click Download to copy the members from the mainframe to your desktop.
Downloaded Source Files	Select one or more previously downloaded members. Selecting previously downloaded members is optional.

6. On the **Virtual Table Layout** page, complete the following fields and click **Next**:

Field	Action
Source	<p>Browse the source tree to verify that it displays the expected data layout. By default, all of the fields in the tree will be included in the mapping. To include only a subset of the fields for the mapping, modify the start field value and, optionally, the end field value, as follows:</p> <ul style="list-style-type: none"> • For the start field, accept the default root start field, or expand the tree and select a different start field. When selecting a different start field, Enable End Field Selection must not be selected. • For the end field, accept the default end field, or expand the tree and select a different end field. When selecting a different end field, Enable End Field Selection must be selected.
Start Field	Identifies the first field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is not selected, and select a different start field in the Source tree.
Enable End Field Selection	Use this field to control selection of the start field and end field values in the Source tree. When this option is not selected (default), you can select the start field. When this option is selected, you can select the end field.
End Field	Identifies the last field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is selected, and select a different end field in the Source tree.

7. On the **zFS Virtual Table Details** page, complete the following fields:

Field	Action
Pathname	<p>Enter the path name of the zFS file.</p> <p>If the absolute path name of the zFS file is less than 255 characters in length, you must include the root slash "/" in the path name. For example, /u/tsado/data/stuff.txt.</p> <p>If the absolute path name of the zFS file is greater than 255 characters in length, you must enter the relative path name. The relative path name starts with the name of the target system to indicate the top-level directory and does not include the leading root slash. For example, data/stuff.txt, where "data" is the name of the target system.</p>
Target System	<p>If you plan to map several zFS files under the same zFS directory location, specify a target system to use.</p> <p>You can click Create to add a new path name to use, or if a relative path name is already specified in the Pathname field, you must select an existing target system from the drop-down list.</p> <p>If you choose to create a new target system, complete the following fields and click Finish:</p> <p>Name – Enter the name for the new target system.</p> <p>CCSID – Enter the CCSID of the character set in which the zFS file data is encoded. The default setting is EBCDIC 1047.</p> <p>Base Pathname – Enter the absolute path name under which the zFS file resides. Typically, this is the path name of the zFS subdirectory that contains your zFS file. At runtime, the server will determine the location of the zFS file by concatenating the path name with the value specified in the virtual table Pathname field. The server does not insert additional slash (/) separators when concatenating the target system path name and the virtual table path name. If the target system path name represents a complete directory name, include the trailing slash (/tmp/).</p>

Field	Action
Advanced	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

8. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for VSAM, VSAM CICS, and IAM data

Create a virtual table that maps to the VSAM, VSAM CICS, and IAM data that you want to access, and from which the SQL used to access the data is generated and executed.

Before you begin

You must have the VSAM or VSAMCICS cluster name available (*sourcelibrary.copybook.filename*).

Procedure

1. Expand the **SQL > Data > SSID** node, where *SSID* is the name of your server.
2. Right-click **Virtual Tables** and select **Create Virtual Table(s)**.
3. Under **Wizards**, select the **VSAM** wizard and click **Next**.
4. On the **New Virtual Table Wizard** page, complete the following fields and click **Next**:

Field	Action
Name	Enter a unique name. The name can contain a maximum of 50 characters. The name must consist of an uppercase letter followed by zero or more characters, each of which is an uppercase letter, a digit, or the underscore character.
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Convert VAR* fields to True VAR* fields	This is a deprecated field and should not be selected.
Arrays Handling	Enable one of the following array management options: <ul style="list-style-type: none"> • Flatten arrays into a single fixed table at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements. • Return arrays into separate tables at runtime: This supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables only support SQL read access. • Flatten arrays now: If you select this option, you cannot change array-handling after you save the virtual table.

5. On the **Source Download** page, complete the following fields and click **Next**:

Field	Action
Available Source Libraries	From the list of Available Source Libraries , select the virtual source library that contains the data structure definition that you want the virtual table to use.
Source Library Members	Select the PDS members that represent the data structures to include and click Download to copy the members from the mainframe to your desktop.
Download Source Files	Select one or more previously downloaded members.

6. On the **Virtual Table Layout** page, complete the following fields and click **Next**:

Field	Action
Source	Browse the source tree to verify that it displays the expected data layout. By default, all of the fields in the tree will be included in the mapping. To include only a subset of the fields for the mapping, modify the start field value and, optionally, the end field value, as follows: <ul style="list-style-type: none"> For the start field, accept the default root start field, or expand the tree and select a different start field. When selecting a different start field, Enable End Field Selection must not be selected. For the end field, accept the default end field, or expand the tree and select a different end field. When selecting a different end field, Enable End Field Selection must be selected.
Start Field	Identifies the first field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is not selected, and select a different start field in the Source tree.
Enable End Field Selection	Use this field to control selection of the start field and end field values in the Source tree. When this option is not selected (default), you can select the start field. When this option is selected, you can select the end field.
End Field	Identifies the last field within the data layout that will be mapped. To change this value, make sure Enable End Field Selection is selected, and select a different end field in the Source tree.

7. Optional: On the **Virtual Table Redefines** page, accept the default table redefines or expand **Redefine** to modify your selection, and click **Next**.

8. Complete the following VSAM related fields:

Field	Action
Cluster Name	Enter the cluster name for the VSAM data set, and click Validate . The server searches the catalog on the mainframe to confirm that the data set exists. If the data set exists, a dialog displays the data set type.
Post-Read Exit Name	To manipulate the data after reading it from the source file, enter the name of the post-read exit to use. This is the custom exit routine that is installed on the server and is used to perform additional processing after a record is read from the data source.
Pre-Write Exit Name	To manipulate the data before writing it to the source file, enter the name of the pre-exit to use. This is the custom exit routine that is installed on the server and is used to perform additional processing before a record is read from the data source.
Alternate Indexes	If the VSAM file has been defined to include alternate indexes, you can click Get to add index information to the virtual table, or you can click Delete to remove the information. Alternate indexes are used to improve query performance when the search criteria includes columns that are not part of the primary index. Alternate

Field	Action
	indexes have an indirect relationship to the cluster name, but they must be defined separately. If you are using a KSDS VSAM or ESDS cluster, you can specify alternative indexes that are associated with the cluster.
Advanced (VSAM only)	When reading large volumes of data from tables, click Advanced to display and configure the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

9. Click **Finish**.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for VSAM and sequential access using ADDI

Create virtual tables that map VSAM and sequential data for COBOL applications by using information made available through IBM Application Discovery and Delivery Intelligence (ADDI).

Before you begin

The Data Service server must be configured to access one or more ADDI projects hosted on Microsoft SQL Server. The studio recognizes ADDI when virtual views and target system maps are installed. Map recognition is based on target systems starting with the string TSIAD and virtual views starting with the name IADV_. For more information on configuring the server, see the *Solution Configuration Guide*.

About this task

To create the virtual tables that are used to access VSAM and sequential data for COBOL applications, information is queried in the ADDI project. Information is retrieved about the z/OS data sets and the COBOL copybooks used to access the z/OS data sets.

The following restrictions and considerations apply:

- Virtual table creation is restricted to data sets in the ADDI project that are processed by COBOL programs using JCL. Data sets accessed using CICS as well as other databases (such as IMS, CA IDMS, or Adabas) are not supported.
- When retrieving data sets from the ADDI project, the studio provides a list of all data sets discovered in the ADDI project that correspond to copybook information. If the data set does not have a corresponding copybook, the data set will not be presented in the studio.
- When creating virtual tables in the studio, duplicate records may appear in the generated list. (Duplicate records have the same project and copybook record names but different ID values.) This is due to multiple copies of the same copybook existing in the ADDI project. The studio provides a feature that compares the definitions of the records and allows you to remove any duplicates.
- When mapping COBOL copybooks containing REDEFINES clauses, default mapping rules related to REDEFINES will be applied which will result in disabled columns in the maps. Editing of virtual maps may be required after generation to enable or disable generated columns.
- ADDI project names are limited to 13 characters due to location name restrictions in the z/OS server.

Procedure

1. On the **Server** tab, explore the ADDI metadata information by expanding the **Discovery > IBM Application Discovery** node, and then navigating down the appropriate subtree. The hierarchy begins with the project, followed by the data sets, and then the associated records.
2. Optional: Right-click a record and select **Display Data Layout** to show the copybook for the record.
3. Select one or more data sets or records to map, as follows:
 - To select individual data sets or records, hold down the Ctrl key and click each data set or record to include.
 - To select a range of data sets or records, click the first data set or record in the range, and then hold the Shift key and select the last data set or record in the range. All data sets or records within the range will be included.
 - To select all records under a data set, click the data set.
4. Right-click the selected data sets or records and select **Create Virtual Table(s)**.
The **New Virtual Tables Wizard** launches, presenting a list of proposed virtual table names and the COBOL structure names that will be used as a basis to create columns for the virtual tables.
5. On the **Create virtual tables using IBM Application Discovery** page, complete the following fields:

Field	Description
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Naming Pattern	Specify the format to use for the generated virtual table names. You can specify different patterns for the project name and records. Use the following variables to create naming patterns that are derived from the ADDI metadata: <ul style="list-style-type: none"> • {Project}: ADDI project name • {Record}: Record name
Arrays Handling	Select one of the following options: <ul style="list-style-type: none"> • Flatten arrays into a single fixed table at runtime (Y): This option supports both OCCURS and OCCURS DEPENDING ON statements. • Return arrays into separate tables at runtime (N): This option supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables support SQL read access only. • Flatten arrays now (C): If you select this option, you cannot change array-handling after you save the virtual table.

6. In the table that lists the records, review the list of selected entries and perform the following steps:
 - a) Optional: If duplicate target virtual table names appear, which are identified with a description in the **Errors** column, click **Remove Duplicates**.
The studio compares the definitions of the records and removes any duplicates.
 - b) Click **Validate** to validate each data set and determine the data set type.
The studio populates the **Type** column with the correct data set type.
 - c) Modify the selections to map as needed.
Tip: Use the check box in the header row of the table to control the selection of all entries.
7. Optional: Click **Advanced** to display and complete the following fields:

Field	Description
MapReduce (Server Parallelism Overrides)	When reading large volumes of data from tables, you can use the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

8. Click **Finish**.

Results

The virtual tables are created on the server and are visible under the **SQL > Data > SSID > Virtual Tables** tree node, where *SSID* is the name of your server.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual tables for VSAM and sequential access using RAA

Create virtual tables that map VSAM and sequential data for COBOL applications by using information made available through IBM Rational Asset Analyzer (RAA).

Before you begin

The Data Service server must be configured to access one or more RAA database schemas hosted on DB2 for z/OS. The studio recognizes RAA when RAA virtual views and target system maps are installed. Map recognition is based on target systems starting with the string TSRAA and virtual views starting with the name RAAV_. For more information on configuring the server, see the *Solution Configuration Guide*.

The preferred method to collect COBOL information is to retrieve record layouts directly from the WebSphere Application Server that hosts RAA. The WebSphere Application Server must be configured using the Metadata Discovery preferences. For more information, see [“Metadata Discovery preferences”](#) on page 51.

About this task

To create the virtual tables that are used to access VSAM and sequential data for COBOL applications, information is queried in the RAA database and from the host. Information is retrieved about the z/OS data sets and the COBOL copybooks used to access the z/OS data sets. If the WebSphere Application Server has been configured, all access to the host for record layout information will first be attempted using the WebSphere Application Server hosting RAA. If access to the RAA host fails and the record layout is stored in a PDS, layout retrieval will be attempted using the current Data Service server.

The following restrictions and considerations apply:

- Virtual table creation is restricted to data sets in the RAA database that are processed by COBOL programs using JCL. Data sets accessed using CICS as well as other databases (such as IMS, CA IDMS, or Adabas) are not supported.
- When retrieving data sets from the RAA database, the studio provides a list of all data sets discovered in the RAA database that correspond to copybook information. If the data set does not have a corresponding copybook, the data set will not be presented in the studio.
- When creating virtual tables in the studio, duplicate records may appear in the generated list. (Duplicate records have the same database and copybook record names but different ID values.) This is due to multiple copies of the same copybook existing in the RAA database. The studio provides a feature that compares the definitions of the records and allows you to remove any duplicates.

- When mapping COBOL copybooks containing REDEFINES clauses, default mapping rules related to REDEFINES will be applied which will result in disabled columns in the maps. Editing of virtual maps may be required after generation to enable or disable generated columns.

Procedure

1. On the **Server** tab, explore the RAA metadata information by expanding the **Discovery > IBM Rational Asset Analyzer** node, and then navigating down the appropriate subtree. The hierarchy begins with the database, followed by the data sets, and then the associated records.
2. Optional: Right-click a record and select **Display Data Layout** to show the copybook for the record.
3. Select one or more data sets or records to map, as follows:
 - To select individual data sets or records, hold down the Ctrl key and click each data set or record to include.
 - To select a range of data sets or records, click the first data set or record in the range, and then hold the Shift key and select the last data set or record in the range. All data sets or records within the range will be included.
 - To select all records under a data set, click the data set.
4. Right-click the selected data sets or records and select **Create Virtual Table(s)**.
The **New Virtual Tables Wizard** launches, presenting a list of proposed virtual table names and the COBOL structure names that will be used as a basis to create columns for the virtual tables.
5. On the **Create virtual tables using IBM Rational Asset Analyzer** page, complete the following fields:

Field	Description
Metadata Library	From the drop-down list, select the target library where the virtual table metadata will be stored (for example, <i>hlq.USER.MAP</i>). The target libraries are specified in the server's started task JCL.
Description	Enter an optional description.
Naming Pattern	Specify the format to use for the generated virtual table names. You can specify different patterns for the database name and records. Use the following variables to create naming patterns that are derived from the RAA metadata: <ul style="list-style-type: none"> • {Database}: RAA database name • {Record}: Record name
Arrays Handling	Select one of the following options: <ul style="list-style-type: none"> • Flatten arrays into a single fixed table at runtime (Y): This option supports both OCCURS and OCCURS DEPENDING ON statements. • Return arrays into separate tables at runtime (N): This option supports both OCCURS and OCCURS DEPENDING ON statements. A subtable is generated for each array. Subtables support SQL read access only. • Flatten arrays now (C): If you select this option, you cannot change array-handling after you save the virtual table.

6. In the table that lists the records, review the list of selected entries and perform the following steps:
 - a) Optional: If duplicate target virtual table names appear, which are identified with a description in the **Errors** column, click **Remove Duplicates**.
The studio compares the definitions of the records and removes any duplicates.
 - b) Click **Validate** to validate the data set and determine the data set type.
The studio populates the **Type** column with the correct data set type.
 - c) Modify the selections to map as needed.

Tip: Use the check box in the header row of the table to control the selection of all entries.

7. Optional: Click **Advanced** to display and complete the following fields:

Field	Description
MapReduce (Server Parallelism Overrides)	When reading large volumes of data from tables, you can use the MapReduce feature. The MapReduce feature enables you to divide the data into logical partitions and process those partitions in parallel using the Thread Count value. At runtime, the number of zIIP processors is verified and one thread is used for each zIIP processor; resulting in improved performance. The Thread Count value you specify overrides the default value (2) and the discovered value. To disable MapReduce , select the Disable MapReduce check box.

8. Click **Finish**.

Results

The virtual tables are created on the server and are visible under the **SQL > Data > SSID > Virtual Tables** tree node, where *SSID* is the name of your server.

What to do next

Use the studio to easily compose and execute SQL queries using your new virtual tables. See [Chapter 7, “Generating and executing SQL queries,”](#) on page 39.

Creating virtual views

Consider creating a virtual view if columns in your virtual table are missing or if you want to join columns from different virtual tables.

Before you begin

The virtual tables representing the data that you want to access or join must already exist.

About this task

A virtual view comprises the SELECT statement that contains the columns from the source data that are used to read data directly from the data source. For example, `SELECT * FROM HLS_JOIN_VSAM LIMIT 1000 ;`. In some cases, creating virtual views is more convenient than regenerating and editing SQL each time you want to access the same data.

Procedure

1. In the **Server View**, expand **SQL > Data > Data Service server > Virtual Tables**.
2. Right-click the virtual table that represents the data that you want to access, and select **Create Virtual View**.
3. In the **Name** field, enter a name for the virtual view.
4. From the **Target** drop-down list, select the target to use for this virtual view.
5. Optional: In the **Description** field, enter a description.
6. Click **Next**.
7. In the **Table Browser**, expand the **Virtual Tables** folder, and select an existing virtual table to use to compose the SQL statement.
8. Click **Next**.
9. Optional: Review the resulting SQL statement and make any necessary modifications.
10. Click **Validate** to validate the SQL.
11. If valid, on the **SQL Validation** message that displays, click **OK**.
12. Click **Finish**.

Results

In the **Server** view, locate the new virtual view by expanding **SQL > Data > Data Service server > Virtual Views**.

Chapter 5. Accessing IT Operational Analytics data

To access, analyze, and report IT Operational Analytics (ITOA) data, generate the SQL from ITOA virtual tables.

When you configure the Data Service server, you have the option to include pre-defined data maps that administrators can use to access the following types of ITOA data:

- IBM System Management Facilities files (SMF)
- Operations Log files (OPERLOG_SYSLOG)
- System Log files (SYSLOG)

After you have configured the Data Service server to use ITOA pre-defined data maps, you can generate the SQL that is used to access ITOA data from the ITOA virtual tables.

System Management File sample code

Use SMF virtual tables to get SQL access to data in System Management Files (SMF).

About this task

When accessing data in SMF files, you use predefined virtual columns that are defined in the SMF virtual table map.

When using SMF log streams, you can use the following virtual columns to retrieve timestamp values:

LS_TIMESTAMP

Timestamp for log stream in GMT. When used in a WHERE predicate, the timestamp is searched in GMT.

LS_TIMESTAMP_LOCAL

Timestamp for log stream in local time zone. When used in a WHERE predicate, the timestamp is searched as local time.

To get SQL access to SMF data, complete the procedure that follows.

Procedure

1. From the Server view, expand **SQL > Data > server name > Virtual Tables**.
2. Right-click the SMF virtual table or view from which you want to access the data.
3. Right-click **Generate Query**, and then review the resulting SQL statement. If necessary, you can modify the statement to meet your needs. The following shows a sample SQL statement:

```
-- ----- Name : SMF_00000
-- This statement will return all rows and all columns from the
-- following table:
-- Name : SMF_00000 : null
-- Catalog : null
-- Schema : DVSQL
-- Remarks : DATA - SMFDATA
-- Tree Location: rs28/1200/SQL/Data/VDBS/Virtual Tables/SMF_00000
-- The sql statement:
SELECT SMF_LEN, SMF_ZERO, SMF_FLAG, SMF_RTY, SMF_TIME, SMF_SID, SMF_SSI,
SMF_STY, SMF_SEQN, SMF0JWT, SMF0BUF, SMF0VST, SMF0OPT, SMF0RST, SMF0RSV,
SMF0OSL, SMF0SYN, SMF0SYP, SMF0TZ, SMF0MSWT, SMF0MTWT
FROM SMF_00000 LIMIT 1000;
```

4. Optional: Execute the SQL statement to view, test, or save the resulting data.

Chapter 6. Accessing DB2 unload data

Using existing DB2 virtual table definitions, you can issue SQL queries against your DB2 sequential unload data sets.

Before you can access your DB2 unload data using your DB2 virtual tables, you must configure access to the DB2 sequential unload data set. This access is configured using a virtual table rule. VTB rule CQDMDLDU is provided to demonstrate redirecting a DB2 virtual table to a DB2 unload data set. For information about setting up access, see "Configuring access to DB2 unload data sets" in the *Customization Guide*.

After you have performed the configuration steps, you can generate the SQL that is used to access the DB2 unload data using your existing DB2 virtual tables.

As an example, consider a virtual table named DSN*_EMPLOYEES that maps the EMPLOYEES table in DB2 subsystem DSN*. With the virtual table rule that specifies the DB2 unload data set enabled, you can query an unload sequential dataset named EMPLOYEE.UNLOAD.SEQ by issuing the following query:

```
SELECT * FROM MDLDU_DSN*_EMPLOYEES__EMPLOYEE_UNLOAD_SEQ
```

The rule performs the necessary steps to access the unload data set directly.

The following restrictions and considerations apply when using this feature:

- SQL access to DB2 unload files is limited to SQL queries only.
- The columns in the DB2 virtual table definition must exactly match the table unloaded in DB2.

Chapter 7. Generating and executing SQL queries

To test SQL access to your data, generate and execute a SQL query from an existing virtual table or virtual view.

Before you begin

To avoid fetching large result sets that are memory intensive, the Data Service Studio provides settings related to SQL generation and retrieval that can limit the amount of data that is actually retrieved for a particular query execution. For more information, see [“SQL preferences” on page 49](#).

Important: When writing SQL to access Adabas data, use caution when using the `BASE_KEY` in WHERE predicates, (for example, `[PARENT TABLE].BASE_KEY = [CHILD TABLE].PARENT_KEY`) when joining the parent table with a child subtable, since this will result in a table scan of the entire Adabas file. It is recommended instead to use the `CHILD_KEY` (for example, `[PARENT TABLE].CHILD_KEY = [CHILD TABLE].PARENT_KEY`).

Procedure

1. On the **Server** tab, right-click a virtual table and select **Generate Query**.
2. Choose from the following options:
 - **Execute** – Generate the SQL query in the **Data Source Editor** and execute the query.
 - **Cancel** – Generate the SQL query in the **SQL Editor** without executing the query. The generated SQL SELECT statement has all columns from the selected table. If the table contains a large number of columns, to avoid enumerating the various column names you can choose all columns using the **Generate Query with *** option.
3. Optional: In the **SQL Editor** view, modify the SQL to select only the data that you want to access. Any ANSI compliant SQL is acceptable.
4. To view or test the data that the SQL statement returns, right-click the highlighted SELECT statement and click either **Execute SQL** to view results in the **SQL Results** view, or **Execute SQL and File results** to save the results in a .csv file.
5. Optional: To create a virtual view of the SQL, highlight the SELECT statement, right-click and select **Create a virtual view**.

Results

In the **SQL Results** view:

- Double-click a row to view additional details about that row.
- Select the **Export Result Set** view option to export the SQL results to a .csv file.
- Click **SQL Messages** to view query-related system messages.

By default, if a result set includes 25 or more columns, each set of 25 columns are displayed incrementally as groups. You can choose which group you want to view using the **Columns Group** field. You can set the number of columns that you want to include in each group, ranging from 25-200, in the **Columns per group** field.

To change how SQL results display in the **SQL Results** view, see [“Data Service preferences” on page 47](#).

What to do next

After the SQL statement is generated, you can perform any of the following tasks:

- Modify the SQL to meet your needs
- Execute the SQL to test and view the resulting data
- Create virtual views to join data or include missing columns

Chapter 8. Server Trace

Use the **Server Trace** view to record and view Data Service server messages.

In the **Server Trace** view, you can perform the following tasks:

To collect and view diagnostics for the client, run the **Gather Diagnostics** wizard, which saves the information to a .zip folder.

Enabling Studio calls in the Server Trace results

To include Studio trace calls in your Server Trace results, enable the Data Service **Enable Server Tracing of Studio Calls** preference.

Before you begin

You must be able to connect to the Data Service server from which you want to collect trace information.

Procedure

1. From the **Window** menu, select **Open Preferences > Data Service**.
2. To enable tracing, select the **Enable Server Tracing of Studio Calls** check box. **Enable Server Tracing of Studio Calls** is enabled by default.
3. In the Studio **HTTP Debug Option** drop-down list, select one of the following HTTP debug options:

Field	Action
Off	Do not collect HTTP messages. All trace activities are deactivated, including interactive tracing.
Normal	Commands that complete with a failure status are traced after execution, including the return codes.
All	All instructions are traced before execution.
Commands	All commands are traced before execution. Return codes are also traced for commands that complete with an error or failure status.
Error	Commands that complete with error status are traced after execution, including the return codes.
Failure	Commands that complete with a failure status are traced after execution, including the return codes.
Intermediates	All instructions are traced before execution. All terms, intermediate results, and substituted variable names are traced during expression evaluation. The final results of any expression that is evaluated also displays. Values assigned by arg , parse , or pull instructions are also traced.
Labels	Shows all labels when executed.
Results	All instructions are traced before execution. The final result of any expression that is evaluated also displays. Values assigned by arg , parse , or pull instructions are also traced.

Starting Server Trace

Start tracing Data Service server records in the **Server Trace** view.

Before you begin

Before running **Server Trace**, you must be able to connect to the Data Service server from which you want to collect the trace information.

Procedure

1. From the **Studio Navigator** view, on the **Common Tools** tab, click **Server Trace**.
2. To start tracing, click **Play** (the blue arrow).
The **Server Trace** table displays trace records.
3. Optional: To view message details, double-click the message and the details are displayed on the **Server Trace Zoom** page.
You can also choose to search for specific details within the message.

Filtering Server Trace results

Use the **Profile** option to filter the records that display in the **Server Trace** view.

Before you begin

You must be able to connect to the Data Service server from which you want to filter trace information. You can set filtering criteria before or after you run a Server Trace. Your most current filtering selections are automatically saved as your default filtering profile.

Procedure

1. On the **Server Trace** view, click **Profile**.
2. On the **Server Trace Profile** page, enable the fields that you want to include in the results.
3. For each enabled field, click **Add** to further filter your results. You can either select from the values that are displayed or enter the value when prompted.
4. Click **OK** to save changes to your profile and to apply the profile to the results in the **Server Trace** table.

What to do next

Use the **Display** option to select and sort columns that display in the filtered table. You can also choose to export the trace results.

Using Server Trace Zoom

Use **Server Trace Zoom** view to view Server Trace message details.

Before you begin

Server Trace must be running before you can open the **Server Trace Zoom** view.

Procedure

1. In the **Server Trace** view, double-click the message for which you want to view details.
2. In the **Server Trace Zoom** view, view message details and choose from the following options:

Field	Action
Previous	Click Previous to search for the previous occurrence of the text string entered.
Next	Click Next to search for the next occurrence of the text string entered.
Search	Click Search and enter a search string. To search for the next occurrence of the text string, click Search again.
Close	Click Close to close the search dialog.

Searching Server Trace messages

You can search Server Trace message results for a particular text string or message ID.

Before you begin

You must start the Server Trace before you can begin searching within the resulting Server Trace messages.

Procedure

1. On the **Server Trace** view, click the drop-down menu, and select **Search**.
2. On the **Search** page that is displayed, in the **From** section, select one of the following options to specify how to search within the results:

Field	Action
First	Search for the first occurrence of the text string.
Last	Search for the last occurrence of the text string.
ID	Search starting from the message ID you enter.

3. In the **For** field, enter the text string to use for searching within the message control blocks. Text strings cannot include spaces or special characters, and wild card searches are not supported.
4. Select **Previous** to find previous occurrences of the text string, or select **Next** to find the next occurrence of the text string.
5. Click **Search** to begin the search.

What to do next

View messages that meet the search criteria in the **Server Trace** view.

Labeling Server Trace messages

Create labels to bookmark server trace messages that you frequently access.

Before you begin

You must start the Server Trace before you can begin labeling messages.

Procedure

1. In the **Server Trace** view, right-click the message that you want to label and select **Add Label**.
2. On the **Message Label** dialog, enter text for the **Label** and click **OK**.
3. Optional: In the **Labels** view, double-click the label to locate the message in the **Server Trace** view.

Exporting Server Trace messages

Use the **Server Trace** view to export server trace messages as either ISX or CVS files.

About this task

You can limit the number of messages that you can export into a file by setting the **Server Trace export size limit** on the **Admin** preferences page.

Procedure

1. In the **Server Trace** view, from the drop-down menu, select **Export**.
2. Under **Export Type**, select one of the following message export options:

Field	Action
Summary	Exports the following minimum message information: <ul style="list-style-type: none">• Message ID• Date• Time• User ID• Message text
Full	Exports all available message information and all data about that message including: <ul style="list-style-type: none">• Message ID• Date• Time• User ID• Message text• Zoom
Comma Separated Format	Exports all table information to a CVS file. This file type cannot be imported for viewing in the Server Trace view.

3. Under **Export Content**, select one of the following message content options:

Field	Action
Message ID Range	Select a range of messages to export by entering the first message ID in From , and the last message ID to include in To .
Transaction ID	Exports only those messages with the RRS transaction ID value that you specify.
Global Transaction ID	Exports only those messages with the RRS global transaction ID that you specify.
Connection ID	Exports only those messages that are associated with a specific client that is currently connected to the server.
Message ID List	Lists message IDs. This option is only available if the Full export type option is selected.

4. Click **Next**.
5. On the **Export File** page, click **Browse** to specify a file name and export location.
6. Click **Finish**.

Importing Server Trace messages

To import and view Server Trace messages, use the **Import File Viewer** tab.

Before you begin

Server Trace must be running before you can import a file.

Procedure

1. In the **Server Trace** view, click the **Import File Viewer** tab and click **Import**.
2. Navigate to the ISX file that you want to import.
3. Double-click the ISX file. Messages and message details display on the **Import File Viewer** tab.
4. Optional: To view more details about a message, right-click on the message and select **Zoom**.
5. Optional: To change how the messages display, click **Display**.

Chapter 9. Data Service preferences

Preferences allow you to customize several IBM® DB2 QMF® Data Service settings.

To view preferences, from the **Window** menu, select **Open Preferences > Data Service**.

Data Service preferences

Use **Data Service** preferences to set preferences such as general session and SQL results settings.

General **Data Service** preferences are identified and described in the table that follows.

Field	Description
Enable Server Tracing of Studio Calls	Includes the Studio trace calls in your server trace results. This setting is disabled by default.
Studio HTTP Debug Option	The Studio type of debug option to be used. The default setting is Normal .
Studio Fixed Width Font	Determines the font, font style, and font size that displays in Studio. The default setting is Courier New-regular-9 .
Hex Encoding	Sets the Hex encoding to use. The default setting is UTF-8 .
File Encoding	Determines the file encoding setting to use. The default setting is windows-1252 .
CSV File Delimiter	Determines the type of file delimiter to use for CSV files. The default setting is a comma (,).
New Connection (DSN) Naming Pattern	Determines the naming pattern to use when new connections are made. The default setting is {SubSystem} .
Studio Connection Timeout (secs)	The number of seconds to wait before a server connection is determined to be unsuccessful. The default setting is 10 .
Studio Operation Timeout (secs)	The number of seconds to wait before determining that the Studio operation is unsuccessful. The default setting is 30 .
Studio Remote Control Port	The port number that the Studio uses for remote connections. The default setting is 31416 .
Use UPPER case logon credentials for both JDBC and HTTP connections	Select this check box to require that logon credentials use uppercase characters for JDBC driver and HTTP connections. This setting is enabled by default. For systems that have mixed-case password support, you must clear this check box and add the following statement to your <i>hlq</i> .SCQDEXEC (CQDSIN00) file: <pre>"MODIFY PARM NAME(PASSWORDCASE) VALUE(ASIS) "</pre>

Admin preferences

Use **Admin** preferences to set the maximum number of Server Trace messages that you want to export and to enable the tracing of Data Service Studio calls in the **Server Trace** view.

Admin preferences are identified and described in the table that follows.

Field	Description
Server Trace export size limit	Sets the maximum number of messages to export. The default value is 5000. Specifying a value greater than 5000 can cause a MAX CPU TIME EXCEEDED error to occur.
Enable Server Tracing of Studio Calls	Includes Studio trace calls in your Server Trace results. This setting is disabled by default.

Console preferences

Use **Console** preferences to view or modify console display settings.

Console preferences are identified and described in the table that follows.

Field	Description
Fixed width console	Enable to specify a maximum number of characters to display in the console. This setting is disabled by default. Maximum character width: If Fixed width console is enabled, enter the maximum number of characters to display in the console. The default is 80 characters.
Limit console output	Enable to limit the console buffer and entry sizes by setting the maximum number of characters permitted: <ul style="list-style-type: none"> • Console buffer size (characters). The default setting is 80000. • Console entry size limit (characters). The default setting is 500.

Dictionary preferences

Use **Dictionary** preferences to add or delete reserved words in dictionaries, and add or delete dictionaries based on the languages being used.

Dictionary preferences are identified and described in the table that follows.

Field	Description
Dictionary	Lists the default dictionaries. You can add new dictionaries to the list or delete existing dictionaries from the list.
Reserved word	Lists reserved words for each dictionary. You can add new words to the list or delete existing words from the list.

Driver preferences

Use **Driver** preferences to enable JDBC driver tracing and to specify the default location of the driver configuration files.

Driver preferences are identified and described in the table that follows.

Field	Description
Enable Tracing	Enables tracing for the JDBC driver. If you change this option, you must restart the Data Service Studio to complete the change. This setting is disabled by default. Note: You can also access data sources that are stored in other configuration files, by adding those configuration files from the Client view.
Default DSN Config File	Specifies the default location of the DSN file. This file is used to store the JDBC connection definitions that are generated for use in the Active Connections view.
Connection Overrides	To override the connection settings that the Data Service Studio uses when it creates JDBC connection definitions, specify a single name-value pair or a semicolon-delimited list to be used. The default setting is a blank field ().

SQL preferences

Use **SQL** preferences to specify settings related to SQL query generation, the SQL Results view, and SQL metadata retrieval.

SQL preferences are identified and described in the table that follows.

Field	Description
SQL Generate Query Behavior	Determines whether you are prompted to execute SQL or if SQL executes automatically. Options include: <ul style="list-style-type: none"> • Generate query and issue user prompt. This is the default setting. • Generate and execute query (no prompt) • Generate query but do not execute query (no prompt)
SQL Results Max Rows	Maximum number of rows to return in the SQL Results view. The default value is 1000.
SQL Results Max Bytes	Maximum number of data bytes to return in the SQL Results view. The default value is 1000000.
SQL Results values accessed as	Specifies how data values are returned. Options include String or Object . The default setting is String .

Field	Description
Use prepared statement to retrieve SQL column info for DB2 or DRDA tables	<p>The Data Service Studio obtains column metadata information from the server for Db2 and DRDA tables and views when you expand a table or view node under the Other Subsystems tree in the Server view, or in other situations where column information needs to be retrieved.</p> <p>The Data Service Studio supports two different ways of retrieving this column metadata information:</p> <ul style="list-style-type: none"> • Using a prepared statement. Typically, this server call will be faster; however, this option requires that the user have SELECT privileges to the table in the remote database. This method is the default and will be used when this preference is selected. • Using the JDBC getColumnns() API. This method is the more conventional approach; however, in some cases (for example, Oracle), the remote DRDA subsystem may take a long time to process the metadata query. This method will be used when this preference is cleared.
Fetch primary key and index information for virtual tables	<p>If this preference is selected, then when you expand a virtual table or view in the Server view, any primary key or indexed column nodes will be identified. This identification process requires the Data Service Studio to make additional metadata calls to the server. To disable these calls and the associated identifications, you can clear this preference and thus speed up the time taken to populate the column nodes. This preference is selected by default.</p>
Fetch primary key and index information for DB2 or DRDA tables	<p>If this preference is selected, then when you expand a table or view node under the Other Subsystems tree in the Server view, any primary key or indexed column nodes will be identified. This identification process requires the Data Service Studio to make additional metadata calls to the server (and subsequently to the remote database). In some cases, these additional calls may be rather expensive (for example, Oracle). To disable these calls and the associated identifications, you can clear this preference to speed up the time taken to populate the column nodes. This preference is cleared by default.</p>

JCL preferences

Use **JCL** preferences to specify JCL settings, such as JOB statement details and to define the trace information to include.

JCL preferences are identified and described in the tables that follows.

JCL

Field	Description
JCL Submit/Poll Timeout (secs)	<p>The number of seconds that can pass before the Data Service Studio stops polling the host for a job to complete. If the host does not complete the job within that number of seconds, the job status is checked on the mainframe. The default setting is 300.</p>

Generation

Fields	Descriptions
Job Name Suffix	Used to generate a default job name. This one-character suffix is appended to the user ID. The default setting is A .
Job Account	Optional accounting information that you can add to the JOB statement. The default setting is blank.
Execution Class	The execution class to be used in the JOB statement. The default setting is A .
Message Class	The message class to be used in the JOB statement. The default setting is X .
Region Size	The region size to be used in the JOB statement. The default setting is OM .
Temporary DASD Name	Generic unit name to be used in the job step for use in allocating temporary work files. The default setting is SYSDA .

Trace/Debug

Fields	Descriptions
Request Status Values	Determines the type of status values to include in the trace: <ul style="list-style-type: none">• ALL• TERSE• VERBOSE• NONE The default setting is ALL .
Level	Determines the trace level to use (1, 2, 3, or 4). The default setting is 1 .
Volume	Determines the trace volume to use (QUIET , SILENT , or NOISY). The default setting is QUIET .
Trace Function Stems	Enables the tracing of function stems. This setting disabled by default.
Dump REXX Variables	Enables the tracing of REXX dump variables. This setting is disabled by default.
Enable SSI Tracing	Enables SSI tracing. This setting is disabled by default.
Enable SSI SSOB dumps	Enables SSI SSOB dump tracing. This setting is disabled by default.

Metadata Discovery preferences

Use **Metadata Discovery** preferences to define settings for the WebSphere Application Server that hosts IBM Rational Asset Analyzer (RAA).

When using RAA to access VSAM or sequential data sets for COBOL applications, complete COBOL layout information that is required to map data is not available in the DB2 database. The mapping wizard uses a RESTful HTTP query to collect record layouts when data is mapped. While this query can be done directly to the Data Service server for data in PDS files, the preferred method to collect COBOL information is to retrieve record layouts directly from the WebSphere Application Server that hosts RAA.

Metadata Discovery preferences are identified and described in the table that follows.

Field	Description
RAA REST Root URL	Location of the RAA WebSphere Application Server. For example: <code>https://<host>:<port></code>
Alternate User ID	User ID for the RAA WebSphere Application Server. You can leave this field blank if the credentials are the same as those used to connect to the current Data Service server (using Set Server).
Alternate Password	Password for the RAA WebSphere Application Server user ID. Specify a value in this field only if a user ID has been specified in the Alternate User ID field.

SSL preferences

Use **SSL** preferences to secure JDBC and HTTP network communications between the Data Service Studio and the Data Service server.

SSL preferences are identified and described in the table that follows.

Field	Description
Use SSL for Studio-Server communications (JDBC and HTTP)	Enables secure JDBC and HTTP network communications between the Data Service Studio and the Data Service server. If enabled, select the Protocol version to use for communications between the Data Service Studio and the Data Service server. The default setting is TLS 1.2 .

Field	Description
<p>Server Authentication</p>	<p>Select the authentication strategy to use:</p> <ul style="list-style-type: none"> • Require Server Validation: Enable to require that all Data Service server certificates be authenticated and complete the following fields: <ul style="list-style-type: none"> – Truststore: The path name of the file on the local machine. The file must contain the Data Service server certificate authority (CA). – Password: The password for the truststore file. – Type: The truststore file type. For example, JKS, PKCS12, BKS, UBER. • Allow Self-Signed Certificates: Enable to allow the Data Service server to use self-signed certificates and complete the following fields: <ul style="list-style-type: none"> – Truststore: The path name of the file on the local machine. The file must contain the self-signed server CA (certificate authority) certificate. – Password: The password for the truststore file. – Type: The truststore file type. For example: JKS, PKCS12, BKS, UBER. • Trust All: Enable to allow all Data Service server certificates. If enabled, the Data Service Studio does not validate the server certificate. <p>The default setting is Require Server Validation.</p>
<p>Client Authentication</p>	<p>To enable client authentication by the Data Service server, select Enable Client Authentication and complete the following fields:</p> <ul style="list-style-type: none"> • Keystore: The path name of the file on the local machine. The file must contain a client certificate which has been signed by the server CA. • Password: The password for the keystore. • Type: The keystore file type. For example: JKS, PKCS12, BKS, UBER. • Alias: To confirm that the password is valid and that the alias (label) appears, click Refresh. <p>This setting is disabled by default.</p>

Notices

This information was developed for products and services offered in the US. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

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

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

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as shown below.

© *(your company name) (year)*.

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

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

Trademarks

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

Linux[®] is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Other company, product, and service names may be trademarks or service marks of others.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions:

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM®.

Commercial use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

Index

A

about [v](#)
Adabas
 accessing Adabas data [11](#)
ADDI [29](#)

B

bind and grant [5](#)

C

CA IDMS data
 accessing [15](#)
 virtual tables [15](#)
client authentication [52](#)
connecting
 Data Service server [5](#)
Console
 display settings [48](#)

D

data source connections [6](#)
DB2 subsystems
 accessing [5](#)
DB2 unload data
 virtual tables [37](#)
DBMS data
 accessing [13](#)
 virtual tables [13](#)
DNS default file [48](#)
documentation changes [vii](#)

H

HTTP
 debug [41](#)
 messages [41](#)

I

IBM Application Discovery and Delivery Intelligence [29](#)
IBM MQ
 accessing [21](#)
 virtual tables [21](#)
IBM Rational Asset Analyzer [31](#)
IBM® DB2 QMF® Data Service [v](#)
IMS
 accessing [16](#)
 DBD metadata [17](#)
 PSB metadata [18](#)
 virtual tables [16, 19](#)
internationalization [6](#)

J

JDBC driver settings [48](#)

L

links
 non-IBM Web sites
 [57](#)
locale considerations [6](#)

N

New IMS DBD Metadata Wizard [17](#)
New IMS PSB Metadata Wizard [18](#)
New Virtual Table Wizard [19](#)
notices
 legal [55](#)

O

overview
 Studio [1](#)

P

perspectives
 Data Service [3](#)
preferences
 Admin [47](#)
 Console [48](#)
 Dictionary [48](#)
 Drivers [48](#)
 JCL [50](#)
 Metadata [51](#)
 SQL [49](#)

R

RAA [31](#)

S

searching Server Trace
 messages [43](#)
sequential data
 accessing [23, 29, 31](#)
 virtual tables [23, 29, 31](#)
server authentication [52](#)
Server Trace
 enabling [41](#)
 exporting messages [44](#)
 filtering results [42](#)
 importing messages [45](#)
 labeling [43](#)
 messages [43](#)

Server Trace (*continued*)

starting [42](#)

view [41](#)

zoom [42](#)

SMF

virtual tables [35](#)

SQL

executing queries [39](#)

generating queries [39](#)

SQL data access

virtual table [10](#)

SSL [52](#)

Studio

overview [1](#)

summary of changes [vii](#)

zFS data (*continued*)

accessing [25](#)

virtual tables [25](#)

T

troubleshooting

Server Trace [41](#)

U

unload data

virtual tables [37](#)

V

virtual source libraries

creating [9](#)

virtual tables

Adabas [11](#)

CA IDMS data [15](#)

DB2 unload data [37](#)

HFS data [25](#)

IBM Application Discovery and Delivery Intelligence [29](#)

IBM MQ [21](#)

IBM Rational Asset Analyzer [31](#)

IMS

DBD [17](#)

PSB [18](#)

IMS data [16](#), [19](#)

sequential data [23](#), [29](#), [31](#)

SMF [35](#)

VSAM [27](#)

VSAM data [29](#), [31](#)

zFS data [25](#)

virtual views

creating [33](#)

VSAM

accessing data [27](#)

virtual tables [27](#)

VSAM data

accessing [29](#), [31](#)

virtual tables [29](#), [31](#)

W

what's new [vii](#)

Z

zFS data



Product Number: 5697-QMF
5698-DB2
5698-DB2
5698-QMF

SC28-3267-00

