



Simba Google BigQuery JDBC Data Connector

Installation and Configuration Guide

Version 1.8

September 2026

Copyright

This document was released in September 2026.

Copyright ©2014-2026 insightsoftware. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission from insightsoftware.

The information in this document is subject to change without notice. insightsoftware strives to keep this information accurate but does not warrant that this document is error-free.

Any insightsoftware product described herein is licensed exclusively subject to the conditions set forth in your insightsoftware license agreement.

Simba, the Simba logo, SimbaEngine, and Simba Technologies are registered trademarks of Simba Technologies Inc. in Canada, the United States and/or other countries. All other trademarks and/or servicemarks are the property of their respective owners.

All other company and product names mentioned herein are used for identification purposes only and may be trademarks or registered trademarks of their respective owners.

Information about the third-party products is contained in a third-party-licenses.txt file that is packaged with the software.

Contact Us

www.insightsoftware.com

About This Guide

The *Simba Google BigQuery JDBC Data Connector Installation and Configuration Guide* explains how to install and configure the Simba Google BigQuery JDBC Data Connector on all supported platforms. It also provides details about the connector's features.

The guide is intended for end users of the Simba Google BigQuery JDBC Connector.

To use the Simba Google BigQuery JDBC Connector, the following knowledge is helpful:

- Familiarity with the platform on which you are using the Simba Google BigQuery JDBC Connector
- Ability to use the data store to which the Simba Google BigQuery JDBC Connector is connecting
- An understanding of the role of JDBC technologies in connecting to a data store
- Experience creating and configuring JDBC connections
- Exposure to SQL

Document Conventions

The following conventions are used throughout this guide to emphasize important concepts:

Italics are used when referring to book and document titles.

Bold is used in procedures for graphical user interface elements that a user clicks and text that a user types.

Monospace font indicates commands, source code or contents of text files.



Note:

A text box with a blue exclamation mark indicates a short note appended to a paragraph.



Important:

A text box with a yellow exclamation mark indicates an important comment related to the preceding paragraph.

Contents

Copyright	2
About This Guide	3
Document Conventions	3
Contents	4
About the Simba Google BigQuery JDBC Connector	8
System Requirements	9
Simba Google BigQuery JDBC Connector Files	10
Installing and Using the Simba Google BigQuery JDBC Connector	11
Referencing the JDBC Connector Libraries	11
Registering the Connector Class	11
Building the Connection URL	13
Configuring Authentication	15
Using a Google User Account	15
Using a Google Service Account	16
Using Pre-Generated Access and Refresh Tokens	16
Using Application Default Credentials	17
Using an External Account	18
Configuring Logging	20
Features	22
SQL Connector	22
Data Types	22
Nested and Repeated Records	23

Security and Authentication	27
Catalog and Schema Support	27
Large Result Set Support	27
Dataset Locations	28
Write-Back	28
Positional Parameters	28
BigQuery Storage API	29
Google Cloud Dedicated (GCD)	29
Connector Configuration Options	31
AdditionalProjects	31
AllowLargeResults	31
BYOID_AudienceUri	32
BYOID_CredentialSource	32
BYOID_PoolUserProject	32
BYOID_SA_Impersonation_Uri	33
BYOID_SubjectTokenType	33
BYOID_TokenUri	33
DefaultDataset	33
EnableHighThroughputAPI	34
EnableSession	34
EnableTimestampPicos	35
EnableWriteAPI	35
FilterTablesOnDefaultDataset	35

HighThroughputActivationRatio	36
HighThroughputMinTableSize	36
JobCreationMode	37
KMSKeyName	37
LargeResultDataset	37
LargeResultsDatasetExpirationTime	38
LargeResultTable	38
Location	39
LogLevel	39
LogPath	40
MaxResults	40
MetaDataFetchThreadCount	40
OAuthAccessToken	41
OAuthClientId	41
OAuthClientSecret	41
OAuthPvtKey	41
OAuthPvtKeyPath	42
OAuthRefreshToken	42
OAuthServiceAcctEmail	42
OAuthType	42
PrivateServiceConnectUri	43
ProjectId	44
ProxyHost	44

ProxyPort	44
ProxyPwd	45
ProxyUid	45
QueryDialect	45
QueryProperties	45
RequestGoogleDriveScope	46
ServiceAccountImpersonationEmail	46
ServiceAccountImpersonationScopes	46
ServiceAccountImpersonationTokenLifetime	46
SSLTrustStore	47
SSLTrustStorePwd	47
StringColumnLength	47
Timeout	47
TimestampFallback	48
Universe Domain	48
UnsupportedHTAPIFallback	49
useQueryCache	49
Third-Party Trademarks	50

About the Simba Google BigQuery JDBC Connector

The Simba Google BigQuery JDBC Connector enables Business Intelligence (BI), analytics, and reporting on data that is stored in BigQuery. The connector complies with the JDBC 4.2 data standard.

JDBC is one of the most established and widely supported APIs for connecting to and working with databases. At the heart of the technology is the JDBC connector, which connects an application to the database. For more information about JDBC, see *Data Access Standards* on the Simba Technologies website: <https://www.simba.com/resources/data-access-standards-glossary>.

This guide is suitable for users who want to access data residing within BigQuery from their desktop environment. Application developers might also find the information helpful. Refer to your application for details on connecting via JDBC.

System Requirements

For details on the specific versions of the data source and Java runtimes supported, please refer to the connector's release notes.

Simba Google BigQuery JDBC Connector Files

The Simba Google BigQuery JDBC Connector is delivered in a ZIP archive named `SimbaBigQueryJDBC42-[Version].zip`, where *[Version]* is the version number of the connector.

The archive contains the connector supporting the JDBC API version indicated in the archive name, as well as release notes and third-party license information.

Installing and Using the Simba Google BigQuery JDBC Connector

To install the Simba Google BigQuery JDBC Connector on your machine, extract the files from the ZIP archive to the directory of your choice.

**Important:**

If you received a license file through email, then you must copy the file into the same directory as the connector JAR file before you can use the Simba Google BigQuery JDBC Connector.

To access a BigQuery data store using the Simba Google BigQuery JDBC Connector, you need to configure the following:

- The list of connector library files (see [Referencing the JDBC Connector Libraries](#))
- The `Driver` or `DataSource` class (see [Registering the Connector Class](#))
- The connection URL for the connector (see [Building the Connection URL](#))

Referencing the JDBC Connector Libraries

Before you use the Simba Google BigQuery JDBC Connector, the JDBC application or Java code that you are using to connect to your data must be able to access the connector JAR files. In the application or code, specify all the JAR files that you extracted from the ZIP archive.

Using the Connector in a JDBC Application

Most JDBC applications provide a set of configuration options for adding a list of connector library files. Use the provided options to include all the JAR files from the ZIP archive as part of the connector configuration in the application. For more information, see the documentation for your JDBC application.

Using the Connector in Java Code

You must include all the connector library files in the class path. This is the path that the Java Runtime Environment searches for classes and other resource files. For more information, see "Setting the Class Path" in the appropriate Java SE Documentation.

- For Windows: <http://docs.oracle.com/javase/8/docs/technotes/tools/windows/classpath.html>
- For Linux and Solaris:
<http://docs.oracle.com/javase/8/docs/technotes/tools/unix/classpath.html>

Registering the Connector Class

Before connecting to your data, you must register the appropriate class for your application.

The following classes are used to connect the Simba Google BigQuery JDBC Connector to BigQuery data stores:

- The `Driver` classes extend `java.sql.Driver`.
- The `DataSource` classes extend `javax.sql.DataSource` and `javax.sql.ConnectionPoolDataSource`.

The connector supports the following fully-qualified class names (FQCNs) that are independent of the JDBC version:

- `com.simba.googlebigquery.jdbc.Driver`
- `com.simba.googlebigquery.jdbc.DataSource`

The following sample code shows how to use the `DriverManager` class to establish a connection for JDBC 4.2:

```
private static Connection connectViaDM() throws Exception
{
    Connection connection = null;
    connection = DriverManager.getConnection(CONNECTION_URL);
    return connection;
}
```

The following sample code shows how to use the `DataSource` class to establish a connection:

```
private static Connection connectViaDS() throws Exception
{
    Connection connection = null;
    DataSource ds = new com.simba.googlebigquery.jdbc.DataSource();
    ds.setURL(CONNECTION_URL);
    connection = ds.getConnection();
    return connection;
}
```

The following sample code shows how to use service authentication to establish a connection:

```
private static Connection connectViaDS() throws Exception
{
    Connection connection = null;
    DataSource ds = new
```

```
com.simba.bigquery.jdbc.DataSource());  
ds.setURL(CONNECTION_URL);  
ds.setProjectId(PROJECT);  
ds.setOAuthType(0); // Service Authentication  
ds.setOAuthServiceAcctEmail(EMAIL);  
ds.setOAuthPvtKeyFile(KEYFILE);  
connection = ds.getConnection();  
return connection;  
}
```

**Note:**

When using the `DataSource` class to establish a connection, all the required properties of `CONNECTION_URL` should be configured. For service authentication, the properties can be configured separately.

Building the Connection URL

Use the connection URL to supply connection information to the data store that you are accessing. The following is the format of the connection URL for the Simba Google BigQuery JDBC Connector:

```
jdbc:bigquery://[Host]:[Port];ProjectId=[Project];OAuthType=[AuthValue];  
[Property1]=[Value1];[Property2]=[Value2];...
```

The variables are defined as follows:

- `[Host]` is the DNS or IP address of the server.
- `[Port]` is the number of the TCP port to connect to. Specifying the port number is optional if you are connecting to port 443.
- `[Project]` is the name of your BigQuery project.
- `[AuthValue]` is a number that specifies the type of authentication used by the connector. For more information, see [OAuthType](#) and [Configuring Authentication](#).
- `[Property1..N]` and `[Value1..N]` are additional connection properties supported by the connector. For a list of the properties available in the connector, see [Connector Configuration Options](#).

**Important:**

- Properties are case-insensitive.
- Do not duplicate properties in the connection URL.

Configuring Authentication

The Simba Google BigQuery JDBC Connector uses the OAuth 2.0 protocol for authentication and authorization. It authenticates your connection through Google OAuth APIs. You can configure the connector to provide your credentials and authenticate the connection to the database using one of the following methods:

- [Using a Google User Account](#)
- [Using a Google Service Account](#)
- [Using Pre-Generated Access and Refresh Tokens](#)
- [Using Application Default Credentials](#)
- [Using an External Account](#)

Using a Google User Account

You can configure the connector to authenticate the connection with a Google user account.

You must provide your Google user account credentials to connect to the server. For more information about authenticating through OAuth 2.0 with a Google user account, see "Using OAuth 2.0 to Access Google APIs" in the Google Identity Platform documentation:
<https://developers.google.com/identity/protocols/OAuth2>.

To configure user account authentication:

1. Connect to the server using a connection URL written in the following format:

```
jdbc:bigquery://[Host]:[Port];ProjectId=[Project];  
OAuthType=1;
```

The variables are defined as follows:

- *[Host]* is the DNS or IP address of the server.
- *[Port]* is the number of the TCP port to connect to. Specifying the port number is optional if you are connecting to port 443.
- *[Project]* is the name of your BigQuery project.

For example:

```
jdbc:bigquery://https://www.googleapis.com/bigquery/v2:443;ProjectId=My  
BigQueryProject;OAuthType=1;
```

2. The connector opens a browser to the Google sign-in page, and triggers a listener in the background. If the sign-in in the browser is successful, the connector gets the refresh token and stores it.

For more information about connection URL syntax, see [Building the Connection URL](#).

Using a Google Service Account

You can configure the connector to authenticate the connection with a Google service account. The service account can handle the authentication process so that no user input is required.

You must provide the full path to a private key file for the service account. The service account email address is optional, if not provided, it will be automatically derived from the key file. You can download the private key file from the Google API console web page. For more information about OAuth authentication using a service account, see "Using OAuth 2.0 for Server to Server Applications" in the Google Identity Platform documentation: <https://developers.google.com/identity/protocols/OAuth2ServiceAccount>.

To configure service account authentication:

1. Set the `OAuthType` property to 0.
2. Set the `ProjectID` property to the name of your BigQuery project.
3. Optionally, Set the `OAuthServiceAcctEmail` property to your Google service account email address. If not specified, the email will be read from the key file.
4. Set the `OAuthPvtKeyPath` property to the full path to the key file that is used to authenticate the service account email address. This parameter supports keys in `.p12` or `.json` format. Alternatively, you can set `OAuthPvtKey` to the full path of a service key file, or a raw JSON object.

For example, the following connection URL authenticates the connection using a service account:

```
jdbc:bigquery://https://www.googleapis.com/bigquery/v2:443;
ProjectId=MyBigQueryProject;OAuthType=0;
OAuthServiceAcctEmail=bqtest1@data-driver-testing.iam.gserviceaccount.com;
OAuthPvtKeyPath=C:\SecureFiles\ServiceKeyFile.p12;
```

For more information about connection URL syntax, see [Building the Connection URL](#).

Using Pre-Generated Access and Refresh Tokens

You can configure the connector to authenticate the connection using access or refresh tokens that have already been generated from the Google Authorization Server. When using this method, you can authenticate your connection by providing an access token, or by providing a refresh token along with a client ID and client secret.

For information about obtaining access and refresh tokens, see "Using OAuth 2.0 to Access Google APIs" in the Google Identity Platform documentation: <https://developers.google.com/identity/protocols/OAuth2>.

**Important:**

When generating the tokens to access BigQuery, you must specify the `https://www.googleapis.com/auth/bigquery` scope. If you are working with federated tables, you should also specify the `https://www.googleapis.com/auth/cloud-platform` scope.

To configure authentication using an access or refresh token:

1. Set the `OAuthType` property to 2.
2. Set the `ProjectID` property to the name of your BigQuery project.
3. Do one or both of the following:
 - Set `OAuthAccessToken` to your access token.
 - Or, set `OAuthRefreshToken` to your refresh token.
4. If you are using a refresh token, set the `OAuthClientId` property to your client ID and set the `OAuthClientSecret` property to your client secret.

For example, the following connection URL authenticates the connection using a refresh token:

```
jdbc:bigquery://https://www.googleapis.com/bigquery/v2:443;
OAuthType=2;ProjectId=MyBigQueryProject;
OAuthAccessToken=a25c7cfd36214f94a79d;OAuthRefreshToken=1jt9Pcyq8pr3lvu143pf
14r86;OAuthClientId=11b5516f132211e6;OAuthClientSecret=bCD+E1f2Gxhi3J4klmN;
```

For more information about connection URL syntax, see [Building the Connection URL](#).

Using Application Default Credentials

You can configure the connector to authenticate the connection using credentials obtained through Application Default Credentials on the environment, if they are available. For information about how to configure Application Default Credentials, see "Google Application Default Credentials" in the Google Identity Platform documentation: <https://developers.google.com/identity/protocols/application-default-credentials>.

For more information about authenticating through OAuth 2.0, see "Using OAuth 2.0 to Access Google APIs" in the Google Identity Platform documentation: <https://developers.google.com/identity/protocols/OAuth2>.

To configure authentication using Application Default Credentials:

1. Set the `OAuthType` property to 3.
2. Set the `ProjectID` property to the name of your BigQuery project.

For example:

```
jdbc:bigquery://https://www.googleapis.com/bigquery/v2:443;
OAuthType=3;ProjectId=MyBigQueryProject;
```

For more information about connection URL syntax, see [Building the Connection URL](#).

Using an External Account

You can configure the connector to authenticate the connection with an external account (workload or workforce identity federation).

You must provide a configuration file for the external account. The configuration file can be downloaded from the [Google API Console](#).

To configure External account authentication:

1. Set the `OAuthType` property to 4.
2. Set the `ProjectID` property to the desired project.
3. Set the `OAuthPvtKey` or `OAuthPvtKeyPath` property to the full path of the external account configuration file, or the `OAuthPvtKey` property to a raw JSON object containing the configuration file contents. Alternatively you can use `BYOID_` properties instead of configuration file.

The following is the format of a workforce configuration file:

```
{
  "type": "external_account",
  "audience": "//iam.googleapis.com/locations/[LOCATION]/workforcePools/[WORKFORCE_POOL_ID]/providers/[PROVIDER_ID]",
  "subject_token_type": "urn:ietf:params:oauth:token_type:id_token",
  "token_url": "https://sts.googleapis.com/v1/token",
  "workforce_pool_user_project": "[WORKFORCE_POOL_USER_PROJECT]",
  "credential_source": {"file"/"url": "[PATH_TO_OIDC_CREDENTIALS]"
  ...
}
```

For example, the following connection URL authenticates the connection using an external account:

```
jdbc:bigquery://https://www.googleapis.com/bigquery/v2:443;
ProjectId=MyBigQueryProject;OAuthType=4;OAuthPvtKeyPath=
C:\Path-to-Files\ConfigurationFile.json;

jdbc:bigquery://https://www.googleapis.com/
bigquery/v2:433;OAuthType4;ProjectId=
MyBigQueryProject;BYOID_
AudienceUri="//iam.googleapis.com/locations/global/workforcePools/pool-
id/providers/provider-id;BYOID_
PoolUserProject=workforceProjectNumber;BYOID_CredentialSource={"file":
"C:\\Token.txt"};
```

**Note:**

For more information about using external accounts, see "Workforce Identity Federation" and "Workload Identity Federation" in the Google Cloud documentation:

<https://cloud.google.com/iam/docs/workforce-identity-federation>

<https://cloud.google.com/iam/docs/workload-identity-federation>

**Note:**

For more information about generating configuration files, see "Generate a Configuration File" in the Google Cloud documentation:

- https://cloud.google.com/iam/docs/workforce-obtaining-short-livedcredentials#generate_a_configuration_file
- <https://cloud.google.com/iam/docs/using-workload-identity-federation>

Configuring Logging

To help troubleshoot issues, you can enable logging in the connector.



Important:

Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

The settings for logging apply to every connection that uses the Simba Google BigQuery JDBC Connector, so make sure to disable the feature after you are done using it.

In the connection URL, set the `LogLevel` key to enable logging at the desired level of detail. The following table lists the logging levels provided by the Simba Google BigQuery JDBC Connector, in order from least verbose to most verbose.

LogLevel Value	Description
0	Disable all logging.
1	Log severe error events that lead the connector to abort.
2	Log error events that might allow the connector to continue running.
3	Log events that might result in an error if action is not taken.
4	Log general information that describes the progress of the connector.
5	Log detailed information that is useful for debugging the connector.
6	Log all connector activity.

To enable logging:

1. Set the `LogLevel` property to the desired level of information to include in log files.
2. Set the `LogPath` property to the full path to the folder where you want to save log files. To make sure that the connection URL is compatible with all JDBC applications, escape the backslashes (`\`) in your file path by typing another backslash.

For example, the following connection URL enables logging level 3 and saves the log files in the `C:\temp` folder:

```
jdbc:bigquery://localhost;LogLevel=3;LogPath=C:\\temp
```

3. To make sure that the new settings take effect, restart your JDBC application and reconnect to the server.

The Simba Google BigQuery JDBC Connector produces the following log files in the location specified in the `LogPath` property:

- A `BigQuery_driver.log` file that logs connector activity that is not specific to a connection.

- A `BigQuery_connection_[Number].log` file for each connection made to the database, where `[Number]` is a number that identifies each log file. This file logs connector activity that is specific to the connection.

If the `LogPath` value is invalid, then the connector sends the logged information to the standard output stream (`System.out`).

To disable logging:

1. Set the `LogLevel` property to 0.
2. To make sure that the new setting takes effect, restart your JDBC application and reconnect to the server.

Features

More information is provided on the following features of the Simba Google BigQuery JDBC Connector:

- [SQL Connector](#)
- [Data Types](#)
- [Nested and Repeated Records](#)
- [Security and Authentication](#)
- [Catalog and Schema Support](#)
- [Large Result Set Support](#)
- [Dataset Locations](#)
- [Write-Back](#)
- [Positional Parameters](#)
- [BigQuery Storage API](#)
- [Google Cloud Dedicated \(GCD\)](#)

SQL Connector

The SQL Connector feature of the connector enables applications to execute standard SQL queries or legacy BigQuery SQL queries against the database.

The connector does not support query prefixes, and instead determines which dialect to use based on the `QueryDialect` connection setting. By default, `QueryDialect` is set to `SQL` so that the connector executes queries using standard SQL syntax. For more information, see [QueryDialect](#).

Data Types

The Simba Google BigQuery JDBC Connector supports many common data formats, converting between BigQuery, SQL, and Java data types.

The following table lists the supported data type mappings.

BigQuery Type	SQL Type	Java Type
ARRAY	SQL_VARCHAR	STRING
BIGNUMERIC For BIGNUMERIC data, the connector always returns 77 for the precision and 38	SQL_NUMERIC	BIGDECIMAL

BigQuery Type	SQL Type	Java Type
for the scale.		
BOOL	SQL_BOOLEAN	BOOLEAN
BYTES	SQL_VARBINARY	BYTE[]
DATE	SQL_DATE	DATE
DATETIME	SQL_TIMESTAMP	STRING
FLOAT64	SQL_DOUBLE	DOUBLE
GEOGRAPHY (See note below)	SQL_VARCHAR	STRING
INT64	SQL_BIGINT	BIGINTEGER
INTERVAL	SQL_VARCHAR	STRING
JSON	SQL_VARCHAR	STRING
NUMERIC For NUMERIC data, the connector always returns 38 for the precision and 9 for the scale.	SQL_NUMERIC	BIGDECIMAL
RANGE	SQL_VARCHAR	STRING
STRING	SQL_VARCHAR	STRING
STRUCT	SQL_VARCHAR	STRING
TIME	SQL_TIME	TIME
TIMESTAMP	SQL_TIMESTAMP	TIMESTAMP


Note:

GEOGRAPHY data cannot be used for a GEOGRAPHY column parameter. To insert or filter on GEOGRAPHY data, you must use the generating function described in "Geography Functions in Standard SQL" in the Google BigQuery documentation: https://cloud.google.com/bigquery/docs/reference/standard-sql/geography_functions.

Nested and Repeated Records

The Simba Google BigQuery JDBC Connector fully supports nested and repeated records. The connector returns the base type as a text representation of the JSON object.

Querying STRUCT Data

Standard SQL syntax represents the sub-components of record data as nested sub-types. The dot operator (.) is used to select sub-components. In the examples below, `city` and `years` belong to the base record type of `address`.

If the record column is specified in a query projection list, the connector returns the base record as a text representation of the JSON record object, and no flattening occurs.

See the following examples to see how to retrieve base records and sub-components from STRUCT data.

Selecting Base Record

Sample query to retrieve a base record from a STRUCT column:

```
select (STRUCT("Vancouver" as city, 5 as years)) as address
```

The connector returns the results as a text reinterpretation of the JSON object, as shown here:

```
{
  "v": {
    "f": [
      {
        "v": "Vancouver"
      },
      {
        "v": "5"
      }
    ]
  }
}
```

The BigQuery console would represent the query results as a table, as shown here:

Row	f0_city	f0_years
1	Vancouver	5

Select Sub-Components

Sample query to retrieve a sub-component from a STRUCT column:

```
select address.city from (select (STRUCT("Vancouver" as city, 5 as years)) as address)
```

The connector returns the results as a text reinterpretation of the JSON object, as shown here:

```
[
  {
    "city": "Vancouver"
  }
]
```

The BigQuery console would represent the query results as a table, as shown here:

Row	city
1	Vancouver

Querying Arrays

The Simba Google BigQuery JDBC Connector fully supports array data types. The connector returns the base array type as a text representation of the JSON array object.

Selecting Arrays of Primitive Type

Sample query to select a primitive array:

```
SELECT [1,2,3]
```

The connector returns the results as a text reinterpretation of the JSON object, as shown here:

```
{
  "v": [
    {
      "v": "1"
    },
    {
      "v": "2"
    },
    {
      "v": "3"
    }
  ]
}
```

The BigQuery console would represent the query results as a table, as shown here:

Row	f0_
1	1
	2
	3

Selecting Arrays of STRUCT Data

Sample query to select multiple objects from a STRUCT array:

```
SELECT [STRUCT("Vancouver" as city, 5 as years), STRUCT("Boston" as city, 10 as years)]
```

The connector returns the results as a text reinterpretation of the JSON object, as shown here:

```
{
  "v": [
    {
      "v": {
        "f": [
          {
            "v": "Vancouver"
          },
          {
            "v": "5"
          }
        ]
      }
    },
    {
      "v": {
        "f": [
          {
            "v": "Boston"
          },
          {
            "v": "10"
          }
        ]
      }
    }
  ]
}
```

The BigQuery console would represent the query results as a table, as shown here:

Row	f0_city	f0_years
1	Vancouver	5
	Boston	10

Security and Authentication

To protect data from unauthorized access, BigQuery data stores require all connections to be authenticated using the OAuth 2.0 protocol. The Simba Google BigQuery JDBC Connector provides mechanisms that allow you to complete an OAuth 2.0 authentication flow using a personal Google account, a Google service account, or Application Default Credentials. You can also specify an access token or refresh token that you have already generated from the Google Authorization Server, and use those credentials to connect to Google BigQuery.

When you connect to BigQuery using a personal Google account, a Google service account, or Application Default Credentials, the connector automatically initiates an OAuth 2.0 authentication flow. The connector retrieves an access token based on the credentials specified in the connection URL, and then uses the token to authenticate the connection to the database. When you connect using an access token or refresh token, the connector authenticates the connection to BigQuery without going through an OAuth 2.0 authentication flow.

For detailed connector configuration instructions, see [Configuring Authentication](#).

For more information about OAuth 2.0, see "Using OAuth 2.0 to Access Google APIs" in the Google Identity Platform documentation: <https://developers.google.com/identity/protocols/OAuth2>.

Catalog and Schema Support

The Simba Google BigQuery JDBC Connector supports both catalogs and schemas to make it easy for the connector to work with various JDBC applications. Projects are mapped to catalogs, and table datasets are mapped to schemas. For more information, see [ProjectId](#). The connector provides access to all of the schemas/databases that are listed under this catalog, ensuring compatibility with standard BI tools.

Large Result Set Support

BigQuery imposes a maximum response size on all requests. If you do not enable large result set support, when executing queries, you might encounter an error message such as "Response too large to return".

If you expect your query to return a large result set, do the following to make sure that the query results can be returned as expected:

- If you are using standard SQL, specify a destination dataset and table for storing the query results. To do this, set the `LargeResultDataset` and `LargeResultTable` properties, respectively.
- If you are using legacy SQL, enable support for large result sets, and then specify a destination dataset and table for storing the query results. To do this, enable the `AllowLargeResults` property, and then set the `LargeResultDataset` and `LargeResultTable` properties to the destination dataset and table, respectively.

**Important:**

- When the `LargeResultDataset` and `LargeResultTable` properties are set, all query results are written to and read from the those tables regardless of the query and its result size. Because of this, the result cache is not available for subsequent queries, and you are billed for every query that you make.
- Working with large data sets may cause you to reach the query limits defined in Google BigQuery. For information about query limits, see "Quotas & Limits" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/quotas>.

In both standard and legacy SQL, if you do not specify a destination dataset or table, the connector stores large result sets in a temporary, default location. The default dataset is hidden, and is named "`_simba_jdbc`". The default table has a name consisting of the prefix "`temp_table`" followed by the time of table creation and a randomized ID. These datasets and tables are deleted after 24 hours.

Dataset Locations

The Simba Google BigQuery JDBC Connector supports auto-routing for regional dataset locations. If multiple datasets are available in different geographic regions, the connector automatically queries a dataset in the correct region.

For more information about dataset locations, see "Dataset Locations" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/locations>.

**Important:**

If you are using a large result set, make sure that the `LargeResultDataset` option is specified, and that the specified dataset is in the same region as the queried table.

Write-Back

The Simba Google BigQuery JDBC Connector supports Data Manipulation Language (DML) statements such as INSERT, UPDATE, and DELETE.

For example, the following INSERT statement is supported:

```
INSERT INTO MyTable (Col1, Col2) VALUES ("Key", "Value");
```

The connector also supports Data Definition Language (DDL) statements. Be aware that BigQuery supports specific syntax for DDL statements, and your statements must be written in that syntax. For more information, see "Using Data Definition Language Statements" in Google BigQuery's *Standard SQL Query Reference*: <https://cloud.google.com/bigquery/docs/data-definition-language>.

Positional Parameters

A parameterized query contains placeholders that are used for parameters. The values of those parameters are supplied at execution time.

Query parameters can be used as substitutes for arbitrary expressions. Parameters cannot be used as substitutes for identifiers, column names, table names, or other parts of the query.

The Simba Google BigQuery JDBC Connector supports SQL positional parameters. Parameters are specified in queries with a question mark (?).

For example, the following parameterized query is supported:

```
SELECT * FROM MyTable WHERE Col1=?
```

BigQuery Storage API

The connector can leverage the BigQuery Storage API, which allows higher throughput than the standard API. This enables the connector to handle large result sets more efficiently. For more information about the API, see "BigQuery Storage API Overview" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/reference/storage/>.

If this feature is enabled, the connector checks the number of rows in an incoming result set table and the number of pages needed to retrieve all the results. If the number of rows and pages exceeds the defined threshold, the connector switches to using the BigQuery Storage API. If the connector encounters any issues initializing the storage API for retrieval, it falls back to using the standard API, unless this is a permissions issue. To ensure best performance, do not use this feature with a named destination dataset or table.

You can customize the thresholds for using the BigQuery Storage API. For information about the configuration options used to determine when the API is used, see the following:

- [EnableHighThroughputAPI](#)
- [HighThroughputActivationRatio](#)
- [HighThroughputMinTableSize](#)
- [UnsupportedHTAPIFallback](#)

The `EnableHighThroughputAPI` setting only specifies whether Storage API can be used when large results are enabled, and can disable Storage API for large sets when set to 0. The actual activation of Storage API is controlled by the `HighThroughputActivationRatio` and `HighThroughputMinTableSize` settings.

In order to take advantage of this feature, the BigQuery project you are querying must have the BigQuery Storage API enabled. For more information, see "Enabling the API" in the Google BigQuery documentation: https://cloud.google.com/bigquery/docs/reference/storage/#enabling_the_api.



Important:

Pricing for the BigQuery Storage API is different than pricing for the standard API. For more information, see "BigQuery Storage API Pricing" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/pricing#storage-api>.

Using the large result set support feature can reduce some of the performance gains from the BigQuery Storage API. If you enable the BigQuery Storage API, we recommend setting `AllowLargeResults` to 0

Google Cloud Dedicated (GCD)

The Google Cloud Dedicated (GCD) is Google's solution to data sovereignty. You can configure the connector to use specific universe domains for Google services like BigQuery and Secure Token

Service by setting either PrivateServiceConnectUri or UniverseDomain.

PrivateServiceConnectUri : A list of service URIs with the desired GCD universe domain. This updates only the specified services.

UniverseDomain: A single domain that applies to all services the driver uses. These two settings are mutually exclusive. If both are set, only PrivateServiceConnectUri is used.

For security, any mismatch between universe domains (e.g., between the connection string and key file) will trigger a DomainMismatchException, preventing cross-domain credential sharing.

For example, if you set UniverseDomain=test-domain.smb in the connection string and connect with KeyFile=my_keyfile.json, where the key file contains a "universe_domain": "test-domain.xyz" field, the connector returns an exception and fail to connect. This is to prevent inter-universe credential sharing, which is a critical security feature of GCD.

Examples:

1. No UniverseDomain & No PrivateServiceConnectUri is specified: Connection string: ...; OAuthType=0; Keyfile=/path/to/my_keyfile.json;..

Behavior: The target URI have no changes.

2. Only PrivateServiceConnectUri is specified: Connection string: ...; OAuthType=0; Keyfile=/path/to/my_keyfile.json; PrivateServiceConnectUri=BIGQUERY=https://my-private-service-bigquery.p.googleapis.com, IAM=https://my-private-service-iam.p.googleapis.com, ...

Behavior: The target URI domain is replaced with the PrivateServiceConnectUri domain provided in the connection string, like: https://my-private-service-bigquery.p.googleapis.com.

3. Only UniverseDomain is specified: Connection string: ...; OAuthType=0; Keyfile=/path/to/my_keyfile.json; UniverseDomain=regional.googleapis.ca;...

Behavior: The target URI domain will be replaced with the domain provided in above connection string, like: https://bigquery.regional.googleapis.ca.

4. Both UniverseDomain & PrivateServiceConnectUri is specified: Connection string: ...; OAuthType=0; Keyfile=/path/to/my_keyfile.json; PrivateServiceConnectUri=BIGQUERY=https://my-private-service-bigquery.p.googleapis.com, IAM=https://my-private-service-iam.p.googleapis.com, ...

UniverseDomain=regional.googleapis.ca; ...

Behavior: The target URI domain will be replaced with the domain provided in above connection string, like: https://my-private-service-bigquery.p.googleapis.com.

Connector Configuration Options

Connector Configuration Options lists and describes the properties that you can use to configure the behavior of the Simba Google BigQuery JDBC Connector.

You can set configuration properties using the connection URL. For more information, see [Building the Connection URL](#).



Note:

Property names and values are case-insensitive.

AdditionalProjects

A comma-separated list of BigQuery projects that the connector can access and use as catalogs. These projects are available as catalogs in metadata functions.

Default Value	Data Type	Required
None	String	No

AllowLargeResults

This option specifies whether the connector supports query results larger than 128MB when working in legacy SQL (the `QueryDialect` property is set to `BIG_QUERY`).

- 1: The connector allows query results that are larger than 128MB in size.
- 0: The connector returns an error when query results are larger than 128MB in size.

When working in standard SQL (the `QueryDialect` property is set to `SQL`), this option is always considered to be enabled. For more information about the supported SQL dialects, see [QueryDialect](#).



Important:

When this property is enabled for legacy SQL, all query results are written to and read from the destination tables, regardless of the query and its result size. Because of this, the result cache is not available to subsequent queries, and you are billed for every query that you make.

For detailed information about how the connector stores large result sets, see [Large Result Set Support](#).

Default Value	Data Type	Required
0	Integer	No

BYOID_AudienceUri

When configuring External Account Authentication, set this option to the full audience URI. This option corresponds to the audience property in the configuration file.

For example:

```
BYOID_AudienceUri=//iam.googleapis.com/locations/${LOCATION}/workforcePools/${PROVIDER_POOL_ID}/providers/${PROVIDER_ID}
```

Default Value	Data Type	Required
None	String	Yes, if OAuthType is set to External Account Authentication (OAuthType=4) and OAuthPvtKey or OAuthPvtKeyPath are not set.

BYOID_CredentialSource

When configuring External Account Authentication, set this option to a JSON object that has credential source information. This option corresponds to the credential_source property in the configuration file.

For example:

```
BYOID_CredentialSource={"file": "C:\\Token.txt"};
```

Default Value	Data Type	Required
None	String	Yes, if OAuthType is set to External Account Authentication (OAuthType=4) and OAuthPvtKey or OAuthPvtKeyPath are not set.

BYOID_PoolUserProject

When configuring External Account Authentication, set this option to the project number or ID associated with the workforce pools user project.

This option corresponds to the WORKFORCE_POOL_USER_PROJECT property in the configuration file.

Default Value	Data Type	Required
None	String	Yes, if OAuthType is set to External Account Authentication (OAuthType=4) and OAuthPvtKey or OAuthPvtKeyPath are not set. Only work force need this property.

BYOID_SA_Impersonation_Uri

When configuring External Account Authentication, set this option to service account impersonation url for workload Identity federation.

This option corresponds to the `service_account_impersonation_url` property in the configuration file.

Default Value	Data Type	Required
None	String	Yes, if OAuthType is set to External Account Authentication (OAuthType=4) and OAuthPvtKey or OAuthPvtKeyPath are not set. Only work load needs this property.

BYOID_SubjectTokenType

When configuring External Account Authentication, set this option to subject token type. the default value is work force toke type. This option corresponds to the `subject_token_type` property in the configuration file.

For example:

`BYOID_SubjectTokenType=urn:ietf:params:oauth:token-type:jwt;`

Default Value	Data Type	Required
<code>urn:ietf:params:oauth:token-type:id_token</code>	String	No

BYOID_TokenUri

When configuring External Account Authentication, set this option to token URI. This option corresponds to the `token_url` property in the configuration file.

For example:

`BYOID_TokenUri=https://sts.googleapis.com/v1/token;`

Default Value	Data Type	Required
<code>https://sts.googleapis.com/v1/token</code>	String	No

DefaultDataset

The name of a dataset that the connector queries by default.

Specifying a default dataset enables you to use unqualified table names in SQL statements. The connector treats unqualified tables as part of the default dataset. It can be a dataset ID in which the dataset is in projectID or it can include both project ID and dataset ID in the format of 'myProject:myDataset'. For information about specifying the project to bill, see [ProjectId](#).

Default Value	Data Type	Required
None	String	No

EnableHighThroughputAPI

This option specifies whether the connector uses the BigQuery Storage API for large result sets.

- 1: The connector uses the Storage API for result sets that exceed the activation ratio.
- 0: The connector does not use the Storage API.

Be aware that the storage API must be enabled for the BigQuery project you are querying.



Important:

Pricing for the storage API is different than pricing for the standard API. For more information, see "BigQuery Storage API Pricing" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/pricing#storage-api>.

For information on the BigQuery Storage API, see [BigQuery Storage API](#).

For information on customizing the activation ratio, see [HighThroughputMinTableSize](#) and [HighThroughputActivationRatio](#).

Default Value	Data Type	Required
0	Integer	No

EnableSession

This option specifies whether the connector creates a session ID from the first executed query on that connection. The session ID is passed to all subsequent executed queries as a BigQuery connection property.

- 1: The connector creates a session ID from the first executed query on that connection.
- 0: The connector does not create a session ID.



Note:

A session ID passed through the `QueryProperties` list is treated as a direct pass-through and is not specially handled by the connector.

For information on BigQuery connection properties, see the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/reference/rest/v2/ConnectionProperty>.

Default Value	Data Type	Required
0	Integer	No

EnableTimestampPicos

This connection property when enabled retrieves timestamp picoseconds precision.

- 0: The connector does not use picoseconds and follow default behavior.
- 1: The connector enables picoseconds timestamp retrieval.

Default Value	Data Type	Required
0	Boolean	No

EnableWriteAPI

The storage write API property can now be used by the connector for bulk inserts.

Default Value	Data Type	Required
1	Integer	No

FilterTablesOnDefaultDataset

This option determines whether the connector filters tables in the `DatabaseMetaData.getTables` call and columns in the `DatabaseMetaData.getColumns` call to return only tables and columns that belong to the default dataset.

- 0: The connector returns all tables in the `DatabaseMetaData.getTables` call and all columns in the `DatabaseMetaData.getColumns` call.
- 1: The connector only returns tables and columns that belong to the default dataset.



Note:

To filter tables and columns, you must define a default dataset. For more information, see [DefaultDataset](#).

When this option is set to 1, the connector behaves as described below for the functions `DatabaseMetaData.getTables` and `DatabaseMetaData.getColumns`.

For the function `DatabaseMetaData.getTables`:

Catalog	Schema	Table	Table Type	Returned List
NULL or %	NULL or %	NULL or %	NULL or %	All tables that belong to the default dataset under the default catalog
NULL or %	[schema]	NULL or %	NULL or %	All tables that belong to the specified schema under all catalogs
[catalog]	NULL or %	NULL or %	NULL or %	All tables that belong to the default dataset under the

Catalog	Schema	Table	Table Type	Returned List
	%	%	%	specified catalog
[catalog]	[schema]	NULL or %	NULL or %	All tables that belong to the specified schema under the specified catalog

For the function `DatabaseMetaData.getColumns`:

Catalog	Schema	Table	Column	Returned List
NULL or %	NULL or %	NULL or %	NULL or %	All columns of all tables that belong to the default dataset under the default catalog
NULL or %	[schema]	NULL or %	NULL or %	All columns of all tables that belong to the specified dataset under all catalogs
[catalog]	NULL or %	NULL or %	NULL or %	All columns of all tables that belong to the default dataset under the specified catalog
[catalog]	[schema]	NULL or %	NULL or %	All columns of all tables that belong to the specified dataset under the specified catalog

Default Value	Data Type	Required
0	Boolean	No

HighThroughputActivationRatio

When the number of pages in your query results exceeds this number, and the number of rows in the results exceeds the `HighThroughPutMinTableSize` value, the connector switches to using the BigQuery Storage API instead of the standard API.

If you define this, you must also set the `EnableHighThroughPutAPI` option to 1. For more information, see [EnableHighThroughputAPI](#).

Default Value	Data Type	Required
3	Integer	No

HighThroughputMinTableSize

When the number of table rows in your query results exceeds this number, and the number of pages in the results exceeds the `HighThroughPutRatio` value, the connector switches to using the BigQuery Storage API instead of the standard API.

If you define this, you must also set the `EnableHighThroughPutAPI` option to 1. For more information, see [EnableHighThroughputAPI](#).

Default Value	Data Type	Required
100	Integer	No

JobCreationMode

To improve performance, the connector can now run a query without creating a job.

- 1: The queries will always be executed as jobs (JOB_CREATION_REQUIRED).
- 2: The queries that do not require a job to execute will be executed jobless-ly (JOB_CREATION_OPTIONAL).

Default Value	Data Type	Required
2	Integer	No

KMSKeyName

The key name of the customer-managed encryption key (CMEK) that you want the connector to use when executing queries. When this property is not set, the connector uses the default encryption key from Google.

For information about CMEKs and Cloud KMS encryption, see "Protecting Data with Cloud KMS Keys" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/customer-managed-encryption>.



Important:

- Do not set this property unless you are certain that you are specifying the correct CMEK. If you execute an INSERT statement with an incorrect CMEK, the connector returns an error or corrupts the table.
- When this property is set, the connector uses the specified CMEK for all queries.

Default Value	Data Type	Required
None. The connector uses the default encryption key from Google.	String	No

LargeResultDataset

A persistent destination dataset for storing query results. For more information, see [Large Result Set Support](#).


Important:

- If you specify a persistent destination table and dataset, all query results are written to and read from the destination tables regardless of the query and its result size. Because of this, the result cache is not available for subsequent queries, and you are billed for every query that you make.
- If you are using different datasets for different locations, make sure that this option is specified, and that the specified dataset is in the same region as the queried table.

If `QueryDialect` is set to `BIG_QUERY`, this option is only used when the `AllowLargeResults` property is enabled.

Default Value	Data Type	Required
<code>_simba_jdbc</code> , if <code>QueryDialect</code> is set to <code>BIG_QUERY</code> , or if <code>QueryDialect</code> is set to <code>SQL</code> and <code>LargeResultTable</code> is specified. None, if <code>QueryDialect</code> is set to <code>SQL</code> and no value is provided for <code>LargeResultTable</code> .	String	No

LargeResultsDatasetExpirationTime

The length of time, in milliseconds, before the tables in a user-specified large result dataset expire.

This expiration time is applied to all tables created in this dataset.

Default Value	Data Type	Required
3600000	Long	No

LargeResultTable

A persistent destination dataset for storing query results. For more information, see [Large Result Set Support](#).


Important:

If you specify a persistent destination table and dataset, all query results are written to and read from the destination tables regardless of the query and its result size. Because of this, the result cache is not available for subsequent queries, and you are billed for every query that you make.

If `QueryDialect` is set to `BIG_QUERY`, this option is only used when the `AllowLargeResults` property is enabled.

Default Value	Data Type	Required
A value consisting of the prefix <code>temp_table_</code>	String	No

Default Value	Data Type	Required
followed by the time of table creation, if <code>QueryDialect</code> is set to <code>BIG_QUERY</code>, or if <code>QueryDialect</code> is set to <code>SQL</code> and <code>LargeResultDatasetSet</code> is specified. None, if <code>QueryDialect</code> is set to <code>SQL</code> and no value is provided for <code>LargeResultDataSet</code>.		

Location

The location where the BigQuery datasets are stored. If this property is specified, the connector can only query datasets that are in this location.

For a list of locations, see "Dataset Locations" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/locations>.

Default Value	Data Type	Required
None	String	No

LogLevel

Use this property to enable or disable logging in the connector and to specify the amount of detail included in log files.



Important:

Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

The settings for logging apply to every connection that uses the Simba Google BigQuery JDBC Connector, so make sure to disable the feature after you are done using it.

Set the property to one of the following numbers:

- 0: Disable all logging.
- 1: Enable logging on the FATAL level, which logs very severe error events that will lead the connector to abort.
- 2: Enable logging on the ERROR level, which logs error events that might still allow the connector to continue running.
- 3: Enable logging on the WARNING level, which logs events that might result in an error if action is not taken.
- 4: Enable logging on the INFO level, which logs general information that describes the progress of the connector.

- 5: Enable logging on the DEBUG level, which logs detailed information that is useful for debugging the connector.
- 6: Enable logging on the TRACE level, which logs all connector activity.

When logging is enabled, the connector produces the following log files in the location specified in the `LogPath` property:

- A `BigQuery_driver.log` file that logs connector activity that is not specific to a connection.
- A `BigQuery_connection_[Number].log` file for each connection made to the database, where `[Number]` is a number that identifies each log file. This file logs connector activity that is specific to the connection.

If the `LogPath` value is invalid, then the connector sends the logged information to the standard output stream (`System.out`).

Default Value	Data Type	Required
0	Integer	No

LogPath

The full path to the folder where the connector saves log files when logging is enabled.



Note:

To make sure that the connection URL is compatible with all JDBC applications, it is recommended that you escape the backslashes (`\`) in your file path by typing another backslash.

Default Value	Data Type	Required
The current working directory	String	No

MaxResults

The maximum number of results that are displayed per page.

Default Value	Data Type	Required
10000	Integer	No

MetaDataFetchThreadCount

The number of threads used to call a `DatabaseMetaData` method.

Default Value	Data Type	Required
32	Integer	No

OAuthAccessToken

The pre-generated access token you are using to authenticate into BigQuery. For details, see [Using Pre-Generated Access and Refresh Tokens](#).

Default Value	Data Type	Required
None	String	No

OAuthClientId

The Client ID you are using to authenticate into BigQuery with a pre-generated refresh token. For details, see [Using Pre-Generated Access and Refresh Tokens](#).

Note: It is recommended to use your own client ID and secret.

Default Value	Data Type	Required
Branded default Client ID	String	No

OAuthClientSecret

The client secret you are using to authenticate into BigQuery with a pre-generated refresh token. For details, see [Using Pre-Generated Access and Refresh Tokens](#).

Note: It is recommended to use your own client ID and secret.

Default Value	Data Type	Required
***	String	No

OAuthPvtKey

The keyfile used for a service account. This can be a path to the `.p12` or `.json` keyfile, or a raw JSON keyfile object.

Default Value	Data Type	Required
None	String	No

OAuthPvtKeyPath

The full path to the `.p12` or `.json` key file that is used to authenticate the service account email address, if you are authenticating your connection using a service account. For more information, see [Using a Google Service Account](#).

Default Value	Data Type	Required
None	String	Yes, if OAuthUserAuth=0.

OAuthRefreshToken

The pre-generated refresh token you are using to authenticate into BigQuery. For details, see [Using Pre-Generated Access and Refresh Tokens](#).

Default Value	Data Type	Required
None	String	No

OAuthServiceAcctEmail

The service account email ID that is used for authentication if you are using service authentication. For more information, see [Using a Google Service Account](#).

Default Value	Data Type	Required
None	String	No

OAuthType

This option specifies how the connector obtains or provides the credentials for OAuth 2.0 authentication.

- 0: The connector uses service-based OAuth authentication (see [Using a Google Service Account](#)).
- 1: The connector uses user-based OAuth authentication (see [Using a Google User Account](#)).
- 2: The connector uses pre-generated tokens for authentication (see [Using Pre-Generated Access and Refresh Tokens](#)).
- 3: The connector uses Application Default Credentials for authentication (see [Using Application Default Credentials](#)).
- 4: The connector authenticates using workload or workforce identity federation (see [Using an External Account](#)).

Default Value	Data Type	Required
0	Integer	No

PrivateServiceConnectUris

A comma-separated list of base URIs to substitute when accessing Private Service Connect URLs. The following are valid URIs:

- **ACCOUNTS:** The baseline accounts service URI, used only for interactive authentication (OAuthType value 3). The default value is: `https://accounts.google.com`.

For example:

ACCOUNTS=https://accounts-<myprivateserver>.p.googleapis.com

- **OAUTH2:** The baseline OAuth 2.0 service URI, used to retrieve access tokens for OAuth 2.0 authentication flows. The default value is: `https://oauth2.googleapis.com`.

For example:

OAUTH2=https://oauth2-<myprivateserver>.p.googleapis.com

- **READ_API:** The host and port required to access the BigQuery Storage Read API service, used to read data from tables via the Storage Read API. The default value is: `bigquerystorage.googleapis.com:443`.

For example:

READ_API=bigquerystorage-<myprivateserver>.p.googleapis.com:443



Note: The format must be `[Host] : [Port]`, with no protocol specifier or URL components.

- **STS:** The baseline security token service, used to retrieve access tokens for External Account Authentication flows. The default value is: `https://sts.googleapis.com`.

For example:

STS=https://sts-<myprivateserver>.p.googleapis.com

- **BIGQUERY:** The baseline BigQuery REST API service, used to interface with the BigQuery data source, via the REST API. The default value is: `https://bigquery.googleapis.com`.

For example:

BIGQUERY=https://bigquery-<myprivateserver>.p.googleapis.com

Default Value	Data Type	Required
None	String	No


Note:

When `OAuthType=0`, the connector prioritizes the `OAUTH2` URI from the key file specified in the `KeyFile` or `KeyFilePath` property. In order, the precedence is:

1. `KeyFile{ _Enc }/KeyFilePath{ _Enc }`
2. `PrivateServiceConnectUri=..., OAUTH2=<YOUR_OAUTH2_URL>, ...`
3. Default



Note: `PrivateServiceConnectUri` may also be used to configure Google Cloud Dedicated behaviour. For more information, refer [Google Cloud Dedicated \(GCD\)](#).

When `OAuthType=4`, the connector prioritizes the `STS` URI from either the configuration file specified in the `KeyFile` or `KeyFilePath` property. In order, the precedence is:

1. `KeyFile/KeyFilePath{ _Enc }`
2. `PrivateServiceConnectUri=..., STS=<YOUR_STS_URL>, ...`
3. Default

For more information about Private Service Connect, see "Private Service Connect" in the Google Cloud documentation: <https://cloud.google.com/vpc/docs/private-service-connect>.

ProjectId

The name of your BigQuery project. This project is the default project that the Simba Google BigQuery JDBC Connector queries against, and also the project that is billed for queries that are run using the DSN.

Default Value	Data Type	Required
None	String	Yes

ProxyHost

The IP address or host name of your proxy server.

Default Value	Data Type	Required
None	String	No

ProxyPort

The listening port of your proxy server.

Default Value	Data Type	Required
None	Integer	No

ProxyPwd

The password, if needed, for proxy server settings.

When using a proxy that requires credentials, the following JVM arguments must be used:

- `Djdk.http.auth.tunneling.disabledSchemes=`
- `Djdk.http.auth.proxying.disabledSchemes=`

Default Value	Data Type	Required
None	String	No

ProxyUid

The user name, if needed, for proxy server settings.

When using a proxy that requires credentials, the following JVM arguments must be used:

- `Djdk.http.auth.tunneling.disabledSchemes=`
- `Djdk.http.auth.proxying.disabledSchemes=`

Default Value	Data Type	Required
None	String	No

QueryDialect

This option specifies whether the connector executes queries using standard SQL syntax or the legacy BigQuery SQL syntax.

- `SQL`: The connector uses standard SQL.
- `BIG_QUERY`: The connector uses legacy SQL.

Default Value	Data Type	Required
SQL	Enumerated	No

QueryProperties

This option enables you to pass properties through to the server when inserting a job. Properties set in this manner are used for all queries in a connection. The `QueryProperties` list must be in the

following form:

key1=value1, key2=value2, ..., keyN=valueN



Note: To configure the connector to use a default project for datasets, set the `dataset_project_id` property in `QueryProperties` of the connection string to the desired project.

Default Value	Data Type	Required
<code>QueryProperties</code>	None	No

RequestGoogleDriveScope

This option specifies whether the connector requests access to Google Drive. Allowing the connector to access Google Drive enables support for federated tables that combine BigQuery data with data from Google Drive.

- 0: The connector does not request access to Google Drive.
- 1: The connector requests access to Google Drive.

Default Value	Data Type	Required
0	Integer	No

ServiceAccountImpersonationEmail

This option specifies the email address of the target principal (service account) that will perform the impersonation.

Default Value	Data Type	Required
None	String	No

ServiceAccountImpersonationScopes

This option specifies the comma-separated list of scopes.

Default Value	Data Type	Required
None	String	No

ServiceAccountImpersonationTokenLifetime

This option specifies an integer that represents the duration in seconds that an access token is valid.

Default Value	Data Type	Required
3600	Integer	No



Note: The properties `ServiceAccountImpersonationEmail`, `ServiceAccountImpersonationLifetime` and `ServiceAccountImpersonationScopes` only apply when the `OAuthMechanism` value is 0 (service account) or 1 (user account) or 2 (token).

SSLTrustStore

The full path of the Java TrustStore containing the server certificate for one-way SSL authentication.

If the trust store requires a password, provide it using the property `SSLTrustStorePwd`. See [SSLTrustStorePwd](#).

Default Value	Data Type	Required
None	String	No

SSLTrustStorePwd

The password for accessing the Java TrustStore that you specified using the property [SSLTrustStore](#).

Default Value	Data Type	Required
None	String	Yes, if using a TrustStore.

StringColumnLength

The maximum length for string type columns.

Default Value	Data Type	Required
65535	Long	No

Timeout

The length of time, in seconds, for which the connector retries a failed API call before timing out. The specified value must be a positive integer.

Default Value	Data Type	Required
300	Integer	No

TimestampFallback



Important: This connection property is deprecated. Only enable this option if you need to temporarily support connections that used connector version 1.1.0 or earlier, while transitioning your applications to connect using connector version 1.1.1 or later.

This property specifies whether the connector sets the default timezone of the JVM to UTC.

- 1: The connector sets the default timezone of the JVM to UTC. This behavior is non-optimal, but consistent with the behavior from connector versions 1.1.0 and earlier.



Note:

Before enabling `TimestampFallback`, be aware of the following:

- This property applies to all connections that use the Simba Google BigQuery JDBC Connector.
- Enabling this property causes the connector to change JVM settings, affecting other processes that are running on the same JVM.
- Once you have made a connection with this property enabled, all subsequent connections must also have this property enabled. Otherwise, the connector returns an error.

- 0: The connector does not change the timezone settings on the JVM. This is the preferred connector behavior.

Default Value	Data Type	Required
0	Integer	No

Universe Domain

The connector now works with Google's Google Cloud Dedicated (GCD), which helps keep your data within specific regions. You can set it up to use the right Google services (like BigQuery or Secure Token Service) by choosing either:

- `PrivateServiceConnectUris`: If you want to set specific addresses for each service.
- `Universe Domain`: If you want to set the domain for all services URIs.

To configure your Google Cloud Dedicated (GCD) settings in the driver, please refer [Google Cloud Dedicated \(GCD\)](#).

Default Value	Data Type	Required
None	String	No

UnsupportedHTAPIFallback

When the connector uses fetch workflows not supported on the High-Throughput API, this option specifies whether the connector falls back to the REST API or returns an error.

- 1: The connector falls back to the REST API.
- 0: The connector returns an error.

For information on the BigQuery Storage API, see [BigQuery Storage API](#).

Default Value	Data Type	Required
1	Integer	No

useQueryCache

This option specifies whether the connector uses cached query results.

- 1: The connector uses cached query results.
- 0: The connector does not use cached query results.

Default Value	Data Type	Required
1	Integer	No

Third-Party Trademarks

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Google BigQuery, Google, and BigQuery are trademarks or registered trademarks of Google, Inc. or its subsidiaries in Canada, the United States and/or other countries.

All other trademarks are trademarks of their respective owners.