



Simba Google BigQuery ODBC Data Connector

Installation and Configuration Guide

Version 3.3

July 2026

Contents

| | |
|--|-----------|
| Copyright | 5 |
| About This Guide | 6 |
| Purpose | 6 |
| Audience | 6 |
| Knowledge Prerequisites | 6 |
| Document Conventions | 6 |
| About the Simba Google BigQuery ODBC Data Connector | 7 |
| Platform and Data source version support | 8 |
| Windows Connector | 9 |
| Windows System Requirements | 9 |
| Installing the Connector in Windows | 9 |
| Creating a Data Source Name in Windows | 10 |
| Setting Connector-Wide Configuration Options in Windows | 11 |
| Configuring Authentication in Windows | 12 |
| Configuring a Proxy Connection in Windows | 19 |
| Configuring Advanced Options in Windows | 20 |
| Configuring the High-Throughput API in Windows | 22 |
| Configuring Logging Options in Windows | 23 |
| Verifying the Connector Version Number in Windows | 25 |
| macOS Connector | 27 |
| macOS System Requirements | 27 |
| Installing the Connector in macOS | 27 |

- Verifying the Connector Version Number in macOS28
- Linux Connector29**
 - Linux System Requirements29
 - Installing the Connector Using the RPM File29
 - Installing the Connector Using the Tarball Package30
 - Verifying the Connector Version Number in Linux31
- Configuring the ODBC Driver Manager in Non-Windows Machines32**
 - Specifying ODBC Driver Managers in Non-Windows Machines32
 - Specifying the Locations of the Connector Configuration Files33
- Configuring ODBC Connections in Non-Windows Machine35**
 - Creating a Data Source Name in a Non-Windows Machine35
 - Setting Connector-Wide Configuration Options in a Non-Windows Machine37
 - Configuring a DSN-less Connection in a Non-Windows Machine38
 - Configuring Authentication in a Non-Windows Machine39
 - Configuring the High-Throughput API in a Non-Windows Machine48
 - Configuring Logging Options in a Non-Windows Machine49
 - Testing the Connection in Non-Windows Machine50
- Using a Connection String52**
 - DSN Connection String Example52
 - DSN-less Connection String Examples52
- Features55**
 - Data Types55
 - Nested and Repeated Records58

| | |
|---|-----------|
| Arrays | 59 |
| Security and Authentication | 59 |
| Catalog and Schema Support | 60 |
| Large Result Set Support | 60 |
| High-Throughput API | 61 |
| Write-Back | 61 |
| Positional Parameters | 62 |
| ODBC Escapes | 62 |
| Service Endpoints | 62 |
| Google Cloud Dedicated (GCD) | 63 |
| Connector Configuration Properties | 67 |
| Configuration Options Appearing in the User Interface | 67 |
| Configuration Options Having Only Key Names | 83 |
| Third-Party Trademarks | 93 |

Copyright

This document was released in July 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

Purpose

The *Simba Google BigQuery ODBC Data Connector Installation and Configuration Guide* explains how to install and configure the Simba Google BigQuery ODBC Data Connector. The guide also provides details related to features of the connector.

Audience

The guide is intended for end users of the Simba Google BigQuery ODBC Data Connector, as well as administrators and developers integrating the connector.

Knowledge Prerequisites

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

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


Document Conventions

Italics is 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 pencil icon indicates a short note appended to a paragraph.

 **Important:** A text box with an exclamation mark indicates an important comment related to the preceding paragraph.

About the Simba Google BigQuery ODBC Data Connector

The Simba Google BigQuery ODBC Data Connector enables Business Intelligence (BI), analytics, and reporting on data that has been uploaded to Google Storage. The connector complies with the ODBC 3.80 data standard and adds important functionality such as Unicode, as well as 32- and 64-bit support for high-performance computing environments on all platforms.

ODBC 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 ODBC connector, which connects an application to the database. For more information about ODBC, see: <https://insightsoftware.com/blog/what-is-odbc/>. For complete information about the ODBC specification, see the *ODBC API Reference* from the Microsoft documentation: <https://docs.microsoft.com/en-us/sql/odbc/reference/syntax/odbc-api-reference>.

The Simba Google BigQuery ODBC Data Connector is available for Microsoft® Windows®, Linux, and macOS platforms.

The *Installation and Configuration Guide* is suitable for users who are looking to access BigQuery data from their desktop environment. Application developers might also find the information helpful. Refer to your application for details on connecting via ODBC.

Platform and Data source version support

The Simba Google BigQuery ODBC Data Connector supports Windows, macOS, and Linux operating systems. For detailed information on supported operating systems and data source versions, please refer to the connector's release notes.

Windows Connector

This section provides an overview of the Connector in the Windows platform, outlining the required system specifications and the steps for installing and configuring the connector in Windows environments.

Windows System Requirements

Install the connector on client machines where the application is installed. Before installing the connector, make sure that you have the following:

- Administrator rights on your machine.
- 100 MB of available disk space

Before the connector can be used, the Visual C++ Redistributable for Visual Studio 2022 with the same bitness as the connector must also be installed. If you obtained the connector from the Simba website, then your installation of the connector automatically includes this dependency. Otherwise, you must install the redistributable manually. You can download the installation packages for the redistributable at <https://learn.microsoft.com/en-us/cpp/windows/latest-supported-vc-redist?view=msvc-170>.

Installing the Connector in Windows

If you did not obtain this connector from the Simba website, you might need to follow a different installation procedure. For more information, see the *Simba OEM ODBC Connector Installation Guide*.

On 64-bit Windows operating systems, you can execute both 32-bit and 64-bit applications. However, 64-bit applications must use 64-bit connectors, and 32-bit applications must use 32-bit connectors. Make sure that you use a connector whose bitness matches the bitness of the client application:

- `SimbaODBCDriverforGoogleBigQuery32.msi` for 32-bit applications
- `SimbaODBCDriverforGoogleBigQuery64.msi` for 64-bit applications

You can install both versions of the connector on the same machine.

To install the Simba Google BigQuery ODBC Data Connector in Windows:

1. Depending on the bitness of your client application, double-click to run `SimbaODBCDriverforGoogleBigQuery32.msi` or `SimbaODBCDriverforGoogleBigQuery64.msi`.
2. Click **Next**.
3. Select the check box to accept the terms of the License Agreement if you agree, and then click **Next**.
4. To change the installation location, click **Change**, then browse to the desired folder, and then click **OK**. To accept the installation location, click **Next**.
5. Click **Install**.
6. When the installation completes, click **Finish**.

7. If you received a license file through email, then copy the license file into the `\lib` subfolder of the installation folder you selected above. You must have Administrator privileges when changing the contents of this folder.

Creating a Data Source Name in Windows

Typically, after installing the Simba Google BigQuery ODBC Data Connector, you need to create a Data Source Name (DSN).

Alternatively, for information about DSN-less connections, see [Using a Connection String](#).

To create a Data Source Name in Windows:


1. From the Start menu, go to **ODBC Data Sources**.

Note: Make sure to select the ODBC Data Source Administrator that has the same bitness as the client application that you are using to connect to BigQuery.

2. In the ODBC Data Source Administrator, click the **Drivers** tab, and then scroll down as needed to confirm that the Simba Google BigQuery ODBC Data Connector appears in the alphabetical list of ODBC Drivers that are installed on your system.
3. Choose one:
 - To create a DSN that only the user currently logged into Windows can use, click the **User DSN** tab.
 - Or, to create a DSN that all users who log into Windows can use, click the **System DSN** tab.

Note: It is recommended that you create a System DSN instead of a User DSN. Some applications load the data using a different user account, and might not be able to detect User DSNs that are created under another user account.

4. Click **Add**.
5. In the Create New Data Source dialog box, select **Simba Google BigQuery ODBC Data Connector** and then click **Finish**. The Simba Google BigQuery ODBC Data Connector DSN Setup dialog box opens.
6. In the **Data Source Name** field, type a name for your DSN.
7. Optionally, in the **Description** field, type relevant details about the DSN.
8. From the **Encrypt Sensitive Data** drop down list, select one:
 - To encrypt and decrypt the sensitive data by all the users on the system, select **For All Users**.
 - Or, to encrypt and decrypt the sensitive data only by the system user who encrypted them, select **For Current User Only**.
9. Configure authentication using the options in the Authentication area. For more information, see [Configuring Authentication in Windows](#).


10. To allow the connector to access Google Drive so that it can support federated tables that combine BigQuery data with data from Google Drive, select the **Request Google Drive Scope Access** check box.
 11. Choose one:
 - To verify the server using the trusted CA certificates from a specific `.pem` file, specify the full path to the file in the **Trusted Certificates** field and leave the **Use System Trust Store** check box cleared.
 - Or, to use the trusted CA certificates `.pem` file that is installed with the connector, leave the default value in the **Trusted Certificates** field, and leave the **Use System Trust Store** check box cleared.
 - Or, to use the Windows Trust Store, select the **Use System Trust Store** check box and leave the **Trusted Certificates** field cleared.
 12. From the **Minimum TLS** drop-down list, select the minimum version of TLS to use when connecting to your data store.
 13. In the **Catalog (Project)** drop-down list, select the name of your BigQuery project. This project is the default project that the Simba Google BigQuery ODBC Data Connector queries against, and also the project that is billed for queries that are run using the DSN.
-  **Note:** If you are not signed in to your Google account, then you are prompted to sign in.
14. Optionally, in the **Dataset** drop-down list, select the name of the dataset the connector will query by default. For more information, see [Dataset](#).
 15. To configure a connection through a proxy server, click **Proxy Options**. For more information, see [Configuring a Proxy Connection in Windows](#).
 16. To configure logging behavior for the connector, click **Logging Options**. For more information, see [Configuring Logging Options in Windows](#).
 17. To configure advanced connector options, click **Advanced Options**. For more information, see [Configuring Advanced Options in Windows](#).
 18. To test the connection, click **Test**. Review the results as needed, and then click **OK**.
 19. To save your settings and close the Simba Google BigQuery ODBC Data Connector DSN Setup dialog box, click **OK**.
 20. To close the ODBC Data Source Administrator, click **OK**.

Setting Connector-Wide Configuration Options in Windows

When you specify connection settings in a DSN or connection string, those settings apply only when you connect to BigQuery using that particular DSN or string. As an alternative, you can specify certain settings that apply to every connection that uses the Simba Google BigQuery ODBC Data Connector by configuring them in the Windows registry. For information about the Windows registry, see the Microsoft Windows documentation.

Important: Editing the Windows Registry incorrectly can potentially cause serious, system-wide problems that may require re-installing Windows to correct.

To set connector-wide configuration options in Windows:

1. Choose one:
 - If you are using Windows 7 or earlier, click **Start** , then type **regedit** in the **Search** field, and then click **regedit.exe** in the search results.
 - Or, if you are using Windows 8 or later, on the Start screen, type **regedit**, and then click the **regedit** search result.
2. Navigate to the appropriate registry key for the bitness of your connector and your machine:
 - If you are using the 32-bit connector on a 64-bit machine, then browse to the following registry key:
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Simba\Simba BigQuery ODBC Driver\Driver
 - Otherwise, browse to the following registry key:
HKEY_LOCAL_MACHINE\SOFTWARE\Simba\Simba BigQuery ODBC Driver\Driver
3. For each configuration option that you want to configure as a connector-wide setting, create a value by doing the following:
 - a. Right-click the **Driver** subkey and then select **New > String Value**.
 - b. Type the key name of the configuration option and then press **Enter**.
To confirm the key names for each configuration option, see [Connector Configuration Properties](#).
 - c. Right-click the new value and then click **Modify**.
 - d. In the Edit String dialog box, in the **Value Data** field, type the value for the configuration option.
4. Close the Registry Editor.
5. Restart your ODBC application to make sure that the new settings take effect.

The `MaxThreads` property must be configured in the `simba.googlebigqueryodbc.ini` file, or as a Windows Registry key value in Windows. For more information see [MaxThreads](#).

Configuring Authentication in Windows

The Simba Google BigQuery ODBC Data 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 an External Account](#)
- [Using PrivateServiceConnect and GoogleCloudDedicated](#)
- [Using Application Default Credentials](#)

Using a Google User Account

You can configure the connector to authenticate the connection with a Google user account. This authentication method uses the OAuth 2.0 access and refresh tokens associated with the user account as the credentials.

The access token is transmitted with every API call that the connector makes, and it is required for accessing BigQuery data stores. However, the access token expires after a certain amount of time and must be renewed using the refresh token. If the refresh token is stored in your connection information, the connector automatically uses it to renew access tokens when they expire.



Note:

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>.

At minimum, you need to provide the OAuth 2.0 refresh token associated with your account. The connector retrieves and uses an access token based on your specified refresh token.

- If you do not have your refresh token, see [Retrieving a Refresh Token](#).
- If you already have your refresh token, see [Providing a Refresh Token](#).
- If you want to provide a `.json` key file that contains your credentials instead of providing your refresh token directly in your connection information, see [Providing a Key File](#).

Retrieving a Refresh Token

When you authenticate your connection this way, the authentication process provides a temporary confirmation code that you can exchange for an access token and a refresh token.

To configure user account authentication by retrieving a refresh token in Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
2. From the **OAuth Mechanism** drop-down list, select **User Authentication**.
3. Click **Sign In**.
4. In the browser that opens, type your credentials for accessing your BigQuery data and sign in to your account.
5. The connector waits for two minutes (120 seconds) for the sign in to complete. While waiting for the sign in to complete, the configuration dialog will be unresponsive.

6. Click **Accept**, when you are prompted to allow Magnitude Simba Data Connector for BigQuery Client Tools to access your data in Google. The browser will display a message indicating that the dialog successfully retrieved the refresh token.
7. The connector automatically populates the field with your refresh token in the dialog with the retrieved token.

Providing a Refresh Token

If you already have your refresh token, then you can provide the token in your connection information without going through the retrieval process described above.

To configure user account authentication by providing a refresh token in Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
2. From the **OAuth Mechanism** drop-down list, select **User Authentication**.
3. In the **Refresh Token** field, type the refresh token associated with your user account.

Providing a Key File

As an alternative to providing your refresh token directly in your connection information, you can save the token in a `.json` key file and then specify the path to the file in your connection information.

The file must define a JSON object of type `authorized_user` containing the refresh token, client ID, and client secret associated with your user account. For example, the `.json` key file must be written in the following format:

```
{  
  "type": "authorized_user",  
  "client_id": "[YourClientID]",  
  "client_secret": "[YourClientSecret]",  
  "refresh_token": "[YourRefreshToken]"  
}
```

To configure user account authentication by providing a key file in Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
2. From the **OAuth Mechanism** drop-down list, select **Service Authentication**.



Note: Although this is a form of user authentication, the key file must be provided using the service authentication options.

3. In the **Email** field, type your user account email ID.
4. In the **Key File Path** field, type the full path to the `.json` key file.

Using a Google Service Account

You can configure the connector to authenticate the connection with a Google service account. When you authenticate your connection this way, the connector handles authentication on behalf of the service account, so that an individual user account is not directly involved and no user input is required.

To authenticate your connection this way, you must provide a Google service account email address and the full path to a private key file for the service account. You can generate and download the private key file when you set up the service account.



Note:

- For more information about OAuth 2.0 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>.
- For information about obtaining service account keys, see "Creating and Managing Service Account Keys" in the Google Cloud Identity & Access Management documentation: <https://cloud.google.com/iam/docs/creating-managing-service-account-keys>.

To configure service account authentication in Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
2. From the **OAuth Mechanism** drop-down list, select **Service Authentication**.
3. In the **Email** field, type your service account email ID.
4. In the **Key File Path** field, type the full path to the `.p12` or `.json` key file that is used to authenticate the service account ID.

Using an External Account

You can configure the connector to authenticate the connection with an external account (workforce identity federation), with limited support, using Azure AD and Okta identity providers. External account authentication lets the connector exchange a token issued by an external identity provider for a Google Cloud access token, so you do not need to create or manage Google service account keys.

The connector supports two ways of supplying the external subject token:

- File-sourced credentials – The subject token is read from a local file that you keep up to date.
- URL-sourced credentials – The subject token is retrieved at connection time from a local metadata endpoint. Specifications for this credential type are detailed at:

<https://google.aip.dev/auth/4117#determining-the-subject-token-in-microsoft-azure-and-url-sourced-credentials>

To authenticate your connection using external account authorization, you must have workforce identity federation set up with Azure AD or Okta and provide a configuration file in the connection string. The credential configuration file can be generated from the Google API Console (<https://console.developers.google.com/>) or through GCloud CLI.

The Credential Configuration File

The credential configuration file is a JSON with a set of shared top-level fields and a `credential_source`. The contents of `credential_source` determine how the connector obtains the subject token (from a file or from a URL).

Shared Fields

The following top-level fields apply to all external account configurations:

| Field | Required | Description |
|--|----------|--|
| <code>type</code> | Yes | Identifies the credential type. This must be <code>external_account</code> or <code>external_account_authorized_user</code> . |
| <code>audience</code> | Yes | The STS audience URI that identifies the identity provider pool. |
| <code>subject_token_type</code> | No | The type of token supplied by the credential source. The default is <code>urn:ietf:params:oauth:token-type:id_token</code> . |
| <code>token_url</code> | No | The endpoint to which a request is sent to exchange the ID token for a Google Cloud Platform access token. The default is <code>https://sts.googleapis.com/v1/token</code> . |
| <code>service_account_impersonation_url</code> | No | The service account impersonation endpoint. If set, the connector exchanges the STS access token for a short-lived service account access token by calling this endpoint. |
| <code>service_account_impersonation</code> | No | Additional service account impersonation options. Only <code>token_lifetime_seconds</code> is currently supported, which sets the requested access token lifetime. |
| <code>workforce_pool_user_project</code> | No | The Google Cloud Platform project to which the token exchange request is executed. |
| <code>credential_source</code> | Yes | Defines how the connector retrieves the subject token. See the file-sourced and URL-sourced sections below. |

Note: Because `subject_token_type` and `token_url` have defaults, you can omit them from the configuration file.

File-Sourced Credentials

Use file-sourced credentials when the subject token is written to a local file by an external process. The connector reads the token from that file at connection time. An example credential file for workforce identity federation with Azure AD or Okta looks like this:

```

    "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": "[PATH_TO_OIDC_CREDENTIALS]"
    }
  
```

The `credential_source` object supports the following fields for file-sourced credentials:

| Field | Required | Description |
|---------------------------------|----------|--|
| file | Yes | The full path to a plaintext file that contains the OIDC subject token. |
| format.type | No | The format of the token file, either text or json. The default is text. |
| format.subject_token_field_name | No | Required when format.type is json. The JSON field name that holds the subject token. |

When the file is in plaintext format, its contents must be the subject token only; any additional content causes the exchange to fail.

Important: Subject tokens expire and cannot be reused to obtain new access tokens. An external process must keep the token file updated with a valid, unexpired subject token.

For additional setup examples, see [Configure Workforce Identity Federation with Azure AD and Sign In Users](https://cloud.google.com/iam/docs/workforce-sign-in-azure-ad) and [Configure Workforce Identity Federation with Okta and Sign In Users](https://cloud.google.com/iam/docs/workforce-sign-in-okta) in the Google Cloud documentation: <https://cloud.google.com/iam/docs/workforce-sign-in-azure-ad> and <https://cloud.google.com/iam/docs/workforce-sign-in-okta>

URL-Sourced Credentials

Use URL-sourced credentials when the subject token is served by a local metadata endpoint, such as the Azure Instance Metadata Service or an OIDC token endpoint. Instead of reading a file, the connector sends an HTTP GET request to the configured URL at connection time and uses the response as the subject token.

The `credential_source` object supports the following fields for URL-sourced credentials:

| Field | Required | Description |
|-------|----------|---|
| url | Yes | The local metadata endpoint that returns the external credential. For Azure, this is the IMDS URL used to retrieve the Azure AD access token. |

| Field | Required | Description |
|---------------------------------|----------|--|
| headers | No | A set of key/value pairs that the connector appends to the GET request headers. |
| format.type | No | The format of the URL response, either text or json. The default is text. |
| format.subject_token_field_name | No | Required when format.type is json. The JSON field name that holds the subject token. |

When format is omitted, the connector treats the response body as a plaintext token. When format.type is json, the connector parses the response as JSON and reads the token from the field named in format.subject_token_field_name.

Azure configurations are a type of URL-sourced credential. For Azure, set format.subject_token_field_name to access_token to retrieve the access token from the Azure Instance Metadata Service.

Configuring External Account Authentication on Windows

To configure external account authentication on Windows using the Key File Path option:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN. Select the DSN, and then click Configure.
2. From the OAuth Mechanism drop-down list, select External Account Authentication.
3. In the Key File Path field, type the full path of the credential configuration file.

Note: For more information about external accounts, see Workforce Identity Federation in the Google Cloud documentation:
<https://cloud.google.com/iam/docs/workforce-identity-federation>

When the connector is configured to use External Account Authentication (OAuthMechanism=4), connection properties are evaluated in the following order of precedence:

4. KeyFile
5. KeyFilePath (or KeyFilePath_Enc if KeyFile is not set)
6. BYOID_properties

Note: BYOID_properties are deprecated. They currently remain supported for backward compatibility, and let you supply external account settings individually for environments that cannot pass a JSON configuration file. New configurations are recommended to use a credential configuration file through KeyFile or KeyFilePath instead.

Using PrivateServiceConnect and GoogleCloudDedicated

To configure Private Service Connect and GoogleCloudDedicated behaviour in a DSN on Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
2. Click **Advanced Options**.

3. Navigate to the **Private Service Connect and Google Cloud Dedicated Options** and select **Enable PSC and GCD Configuration** check box.
4. In the **Private Service name** field, type the name of your private service.
5. In the **Universe Domain** field, type the name of the universe domain.
6. To close the **Advanced Options** dialog box, click **OK**.
7. To save your settings, click **OK**.



Note: Enabling GCD, PSC sets `EnableGCD=1` for all DSNs, but `PrivateService` and `UniverseDomain` apply only to the current DSN.

Using Application Default Credentials

You can configure the connector to authenticate the connection with Application Default Credentials. When you authenticate your connection this way, the connector parses the JSON keyfile specified by the `GOOGLE_APPLICATION_CREDENTIALS` environment variable. No other connection properties are required to use the Application Default Credentials authentication flow.

To configure Application Default Credentials authentication in Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
2. From the **OAuth Mechanism** drop-down list, select **Application Default Credentials Authentication**.



Note: For more information about Application Default Credentials, refer to "How Application Default Credentials works" and "Set up Application Default Credentials" in the Google Cloud documentation: <https://cloud.google.com/docs/authentication/application-default-credentials> and <https://cloud.google.com/docs/authentication/provide-credentials-adc>, respectively.

Metadata Server Authentication for Application Default Credentials

When Application Default Credentials (ADC) are used for authentication and no key file is present, the ODBC driver falls back to retrieving an access token for the default service account from the Google internal metadata server. This fallback mechanism is available only when the driver is running on a Google Cloud Engine virtual machine.



Note: For more details about authenticating through the internal metadata server, see [Application Default Credentials overview](#) and [Restricting service accounts](#).

Configuring a Proxy Connection in Windows

If you are connecting to the data source through a proxy server, you must provide connection information for the proxy server.

To configure a proxy server connection in Windows:

1. To access proxy server options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Proxy Options**.
2. Select the **Use Proxy Server** check box.
3. In the **Proxy Host** field, type the host name or IP address of the proxy server.
4. In the **Proxy Port** field, type the number of the TCP port that the proxy server uses to listen for client connections.
5. To save your settings and close the Proxy Options dialog box, click **OK**.

Note: For proxies, the following URLs are allowed:

- `crl.pki.goog`
- `crls.pki.goog`
- `ocsp.pki.goog`

Configuring Advanced Options in Windows

You can configure advanced options to modify the behavior of the connector.

To configure advanced options in Windows:

1. To access advanced options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Advanced Options**.
2. From the **Language Dialect** drop-down list, select the SQL syntax to use when executing queries:
 - To use standard SQL syntax, select **Standard SQL**.
 - Or, to use the legacy BigQuery SQL syntax, select **Legacy SQL**.
3. To allow query results that are larger than the BigQuery maximum result sizes when using Legacy SQL, select the **Allow Large Result Sets** check box.

Note:

- This option is only available if the Language Dialect drop-down list is set to Legacy SQL. Large result sets are always enabled if the Language Dialect drop-down list is set to Standard SQL.
- For more information about BigQuery maximum result sizes, see "Quotas and Limits" in the Google BigQuery documentation:
https://cloud.google.com/bigquery/quotas#query_jobs.

4. To specify the dataset that stores temporary tables for large result sets, do one of the following:
 - To use the default dataset with the ID `_bqodbc_temp_tables`, select the **Use Default _bqodbc_temp_tables Large Results Dataset** check box.

- Or, to specify a different dataset, clear the **Use Default _bqodbc_temp_tables Large Results Dataset** check box and, in the **Dataset Name For Large Result Sets** field, type the ID of the BigQuery dataset that you want to use.

Note:

- If the dataset does not exist and the data store specifies the US region, the connector creates the dataset.
- These options are only available if the connector is configured to use large result sets.

5. In the **Default temp table expiration time (ms)** field, type the length of time, in milliseconds, for which the temporary tables exist for large results datasets. The minimum value is 3600000.

Note: This option is only available if the connector is configured to use large result sets.

6. In the **Rows Per Block** field, type the maximum number of rows to fetch for each data request.
7. In the **Enable Timestamp Picos** field, type the level of precision for timestamp values, up to picoseconds.
8. To use the High-Throughput API to handle large result sets more efficiently, select the **Allow High-Throughput API for Large Results queries** check box. For more information, including prerequisites and detailed instructions, see [Configuring the High-Throughput API in Windows](#).
9. To enable or disable the Google Cloud Dedicated (GCD), select the **EnableGCD** checkbox. For more information, see [Enable GCD](#).
 - a. In the **Private Service Name** field, type the name of the private service from which it automatically constructs the corresponding URLs. For more information, see [PrivateService](#).
 - b. In the **Universe Domain** field, type the Universe Domain that applies to all services. For more information, see [UniverseDomain](#).
10. To use a customer-managed encryption key (CMEK) when executing queries, in the **Path To CMEK** field, type the resource ID of the CMEK. For more information, see "Protecting Data with Cloud KMS Keys" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/customer-managed-encryption>

Important:

- Do not specify a CMEK unless you are certain that it is the correct value to use.
- If you execute an INSERT statement with an incorrect CMEK, the connector returns an error or corrupts the table. The connector uses the specified CMEK for all queries.
- The `KMSKeyName` property is used only when the connector is configured to use the destination tables.

11. In the **Default String Column Length** field, type the maximum number of characters that can be contained in STRING columns.

12. To create a session ID from the first executed query on a connection, select the **Enable Session** check box.
 - Optionally, in the **Session Location** field, type the desired location.
13. To return data as SQL_WVARCHAR data instead of SQL_VARCHAR data, select the **Use SQL_WVARCHAR Instead Of SQL_VARCHAR** check box.

Note: This option applies only to result set columns that the connector would normally return as SQL_VARCHAR columns. It does not convert all columns into SQL_WVARCHAR.
14. To access public projects and use them as catalogs for the connection, in the **Additional Projects** field, type a comma-separated list of project names.
15. To pass properties down to the server when inserting a job, in the **Query Properties** field, type a comma-separated list of key value pairs.
16. To save your settings and close the Advanced Options dialog box, click **OK**.

Configuring the High-Throughput API in Windows

You can configure the connector to use the High-Throughput API to handle large result sets more efficiently. For more information about the High-Throughput API, see [High-Throughput API](#).

To configure the High-Throughput API in Windows:

1. Make sure that your Google BigQuery project has the Storage API enabled. For more information about the Storage API, see "BigQuery Storage API Overview" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/reference/storage/>.
2. To access the DSN that you want to configure the High-Throughput API for, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**.
3. Choose one:
 - To verify the server using the trusted CA certificates from a specific .pem file, specify the full path to the file in the **Trusted Certificates** field and leave the **Use System Trust Store** check box cleared.
 - Or, to use the trusted CA certificates .pem file that is installed with the connector, leave the default value in the **Trusted Certificates** field, and leave the **Use System Trust Store** check box cleared.

Important: Do not select the Use System Trust Store check box. The High-Throughput API is not compatible with the Windows Trust Store.

4. Click **Advanced Options**.
5. If you are using Legacy SQL (the **Language Dialect** drop-down list is set to **Legacy SQL**), then make sure that the **Allow Large Result Sets** check box is selected.
6. To specify the dataset that stores temporary tables for large result sets and result sets returned by the High-Throughput API, do one of the following:

- To use the default dataset with the ID `_bqodbc_temp_tables`, select the **Use Default _bqodbc_temp_tables Large Results Dataset** check box.
- Or, to specify a different dataset, clear the **Use Default _bqodbc_temp_tables Large Results Dataset** check box and, in the **Dataset Name For Large Result Sets** field, type the ID of the BigQuery dataset that you want to use.



Note: If the dataset does not exist and the data store specifies the US region, the connector creates the dataset. These options are only available if the connector is configured to use large result sets.

7. Optionally, if you want to use large results and query the data with the HTAPI, select the **Allow High-Throughput API for Large Results queries** check box.
8. In the **Activation Threshold for High-Throughput API** field, specify the minimum total pages required to activate reading through the High-Throughput API.
9. To save your settings and close the Advanced Options dialog box, click **OK**.

The connector now uses the BigQuery High-Throughput API instead of the REST API for requests where:

- the number of pages in the results exceeds the Activation Threshold for High-Throughput API value.
- and, if the request uses large result dataset, the Allow High-Throughput API for Large Results Dataset property is set.

Configuring Logging Options in Windows

To help troubleshoot issues, you can enable logging. In addition to functionality provided in the Simba Google BigQuery ODBC Data Connector, the ODBC Data Source Administrator provides tracing functionality.



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

Configuring Connector-wide Logging Options

The settings for logging apply to every connection that uses the Simba Google BigQuery ODBC Data Connector, so make sure to disable the feature after you are done using it. To configure logging for the current connection, see [Configuring Logging for the Current Connection](#).

To enable connector-wide logging in Windows:

1. To access logging options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
2. From the **Log Level** drop-down list, select the logging level corresponding to the amount of information that you want to include in log files:

| Logging Level | Description |
|---------------|--|
| OFF | Disables all logging. |
| FATAL | Logs severe error events that lead the connector to abort. |
| ERROR | Logs error events that might allow the connector to continue running. |
| WARNING | Logs events that might result in an error if action is not taken. |
| INFO | Logs general information that describes the progress of the connector. |
| DEBUG | Logs detailed information that is useful for debugging the connector. |
| TRACE | Logs all connector activity. |

- In the **Log Path** field, specify the full path to the folder where you want to save log files. You can type the path into the field, or click **Browse** and then browse to select the folder.
- In the **Max Number Files** field, type the maximum number of log files to keep.

Note: After the maximum number of log files is reached, each time an additional file is created, the connector deletes the oldest log file.

- In the **Max File Size** field, type the maximum size of each log file in megabytes (MB).

Note: After the maximum file size is reached, the connector creates a new file and continues logging.

- Click **OK**.
- Restart your ODBC application to make sure that the new settings take effect.

The Simba Google BigQuery ODBC Data Connector produces the following log files at the location you specify in the Log Path field:

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

To disable connector logging in Windows:

- Open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
- From the **Log Level** drop-down list, select **LOG_OFF**.
- Click **OK**.
- Restart your ODBC application to make sure that the new settings take effect.

Configuring Logging for the Current Connection

You can configure logging for the current connection by setting the logging configuration properties in the DSN or in a connection string. For information about the logging configuration properties, see

Configuring Logging Options in Windows. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over connector-wide settings.

Note: If the `LogLevel` configuration property is passed in via the connection string or DSN, the rest of the logging configurations are read from the connection string or DSN and not from the existing connector-wide logging configuration.

To configure logging properties in the DSN, you must modify the Windows registry. For information about the Windows registry, see the Microsoft Windows documentation.

Important: Editing the Windows Registry incorrectly can potentially cause serious, system-wide problems that may require re-installing Windows to correct.

To add logging configurations to a DSN in Windows:


1. On the Start screen, type **regedit**, and then click the **regedit** search result.
2. Navigate to the appropriate registry key for the bitness of your connector and your machine:
 - 32-bit System DSNs: **HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\ODBC\ODBC.INI***[DSN Name]*
 - 64-bit System DSNs: **HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI***[DSN Name]*
 - 32-bit and 64-bit User DSNs: **HKEY_CURRENT_USER\SOFTWARE\ODBC\ODBC.INI***[DSN Name]*
3. For each configuration option that you want to configure for the current connection, create a value by doing the following:
 - a. If the key name value does not already exist, create it. Right-click the *[DSN Name]* and then select **New > String Value**, type the key name of the configuration option, and then press **Enter**.
 - b. Right-click the key name and then click **Modify**.
To confirm the key names for each configuration option, see [Connector Configuration Options](#).
 - c. In the Edit String dialog box, in the **Value Data** field, type the value for the configuration option.
4. Close the Registry Editor.
5. Restart your ODBC application to make sure that the new settings take effect.

Verifying the Connector Version Number in Windows

If you need to verify the version of the Simba Google BigQuery ODBC Data Connector that is installed on your Windows machine, you can find the version number in the ODBC Data Source Administrator.

To verify the connector version number in Windows:

1. From the Start menu, go to **ODBC Data Sources**.

 **Note:** Make sure to select the ODBC Data Source Administrator that has the same bitness as the client application that you are using to connect to BigQuery.

2. Click the **Drivers** tab and then find the Simba Google BigQuery ODBC Data Connector in the list of ODBC Connectors that are installed on your system. The version number is displayed in the **Version** column.

macOS Connector

This section provides an overview of the Connector in the mac OS platform, outlining the required system specifications and the steps for installing and configuring the connector in mac OS environments.

macOS System Requirements

Install the connector on client machines where the application is installed. Each client machine that you install the connector on must meet the following minimum system requirements:

- 150MB of available disk space
- One of the following ODBC driver managers installed:
 - iODBC 3.52.9 or later
 - unixODBC 2.3.6 or later


Installing the Connector in macOS

If you did not obtain this connector from the Simba website, you might need to follow a different installation procedure. For more information, see the *Simba OEM ODBC Connectors Installation Guide*.

The Simba Google BigQuery ODBC Data Connector is available for macOS as a .dmg file named `SimbaODBCDriverforGoogleBigQuery.dmg`. The connector supports both 32- and 64-bit client applications.

To install the Simba Google BigQuery ODBC Data Connector in macOS:

1. Double-click **SimbaODBC DriverforGoogleBigQuery.dmg** to mount the disk image.
2. Double-click **SimbaODBCDriverforGoogleBigQuery.pkg** to run the installer.
3. In the installer, click **Continue**.
4. On the Software License Agreement screen, click **Continue**, and when the prompt appears, click **Agree** if you agree to the terms of the License Agreement.
5. Optionally, to change the installation location, click **Change Install Location**, then select the desired location, and then click **Continue**.

 **Note:** By default, the connector files are installed in the `/Library/simba/googlebigqueryodbc` directory.

6. To accept the installation location and begin the installation, click **Install**.
7. When the installation completes, click **Close**.
8. If you received a license file through email, then copy the license file into the `/lib` subfolder in the connector installation directory. You must have root privileges when changing the contents of this folder.

For example, if you installed the connector to the default location, you would copy the license file into the `/Library/simba/googlebigqueryodbc/lib` folder.

Next, configure the environment variables on your machine to make sure that the ODBC Driver manager can work with the connector. For more information, see [Configuring the ODBC Driver Manager in Non-Windows Machines](#)

Next, configure the environment variables on your machine to make sure that the ODBC Driver manager can work with the connector. For more information, [Configuring the ODBC Driver Manager in Non-Windows Machines](#).

Verifying the Connector Version Number in macOS

If you need to verify the version of the Simba Google BigQuery ODBC Data Connector that is installed on your macOS machine, you can query the version number through the Terminal.

To verify the connector version number in macOS:

- At the Terminal, run the command:

```
pkgutil --info com.simba.googlebigqueryodbc
```

The command returns information about the Simba Google BigQuery ODBC Data Connector that is installed on your machine, including the version number.

Linux Connector

This section provides an overview of the Connector in the Linux platform, outlining the required system specifications and the steps for installing and configuring the connector in Linux environments.

The Linux connector is available as an RPM file and as a tarball package.

Linux System Requirements

Install the connector on client machines where the application is installed. Each client machine that you install the connector on must meet the following minimum system requirements:

- 150MB of available disk space
- One of the following ODBC driver managers installed:
 - iODBC 3.52.9 or later
 - unixODBC 2.2.14 or later
- glibc 2.17 or later

To install the connector, you must have root access on the machine.

Installing the Connector Using the RPM File

If you did not obtain this connector from the Simba website, you might need to follow a different installation procedure. For more information, see the *Simba OEM ODBC Connectors Installation Guide*.

On 64-bit editions of Linux, you can execute both 32- and 64-bit applications. However, 64-bit applications must use 64-bit connectors, and 32-bit applications must use 32-bit connectors. Make sure that you use a connector whose bitness matches the bitness of the client application:

- `simbagooglegbigquery-[Version]-[Release].i686.rpm` for the 32-bit connector
- `simbagooglegbigquery-[Version]-[Release].x86_64.rpm` for the 64-bit connector

The placeholders in the file names are defined as follows:

- `[Version]` is the version number of the connector.
- `[Release]` is the release number for this version of the connector.

To install the Simba Google BigQuery ODBC Data Connector using the RPM File:

1. Log in as the root user.
2. Navigate to the folder containing the RPM package for the connector.
3. Depending on the Linux distribution that you are using, run one of the following commands from the command line, where `[RPMFileName]` is the file name of the RPM package:

- Or, if you are using SUSE Linux Enterprise Server, run the following command:

```
zypper install [RPMFileName]
```

- Or, if you are using Linux ARM Server, navigate to the `Linux_ARM` folder containing the RPM and run the following command:

```
Linux_ARM/zypper install [RPMFileName]
```

4. If you received a license file through email, then copy the license file into the `/opt/simba/googlebigqueryodbc/lib/32` or `/opt/simba/googlebigqueryodbc/lib/64` folder, depending on the version of the connector that you installed.

Next, configure the environment variables on your machine to make sure that the ODBC Driver manager can work with the connector. For more information, see [Configuring the ODBC Driver Manager in Non-Windows Machines](#)

Installing the Connector Using the Tarball Package

If you did not obtain this connector from the Simba website, you might need to follow a different installation procedure. For more information, see the *Simba ODBC Connectors Installation Guide*.

The Simba Google BigQuery ODBC Data Connector is available as a tarball package named `SimbaODBCDriverforGoogleBigQuery_[Version].[Release]-Linux.tar.gz`, where `[Version]` is the version number of the connector and `[Release]` is the release number for this version of the connector. The package contains both the 32-bit and 64-bit versions of the connector.

On 64-bit editions of Linux, you can execute both 32- and 64-bit applications. However, 64-bit applications must use 64-bit connectors, and 32-bit applications must use 32-bit connectors. Make sure that you use a connector whose bitness matches the bitness of the client application. You can install both versions of the connector on the same machine.

To install the connector using the tarball package:

1. Create the directory of where you want to install the connector, if it does not already exist.
2. Extract the main tarball file to a convenient temporary location.
3. Extract the inner tarball file, corresponding to the version of the connector you want to install (32-bit or 64-bit), to a convenient temporary location.
4. Navigate to the subfolder of the extracted inner tarball file, as done in the previous step, and then copy all the files and folders to the installation directory.
5. Navigate to the temporary location of the extracted main tarball file done in step 2, and then copy the `Tools` folder to the installation directory.
6. Navigate to the temporary location of the extracted main tarball file in step 2, and then copy the `GoogleBigQueryODBC.did` file to the `/lib` subfolder in the installation directory.

7. The Simba ODBC Driver for Google BigQuery path is now `[INSTALLDIR]/lib/[FileName].so`. You can make required changes in the `.ini` files to correct path of connector.
8. Next, configure the environment variables on your machine to make sure that the ODBC Driver manager can work with the connector. For more information, see [Configuring the ODBC Driver Manager in Non-Windows Machines](#)

Verifying the Connector Version Number in Linux

If you need to verify the version of the Simba Google BigQuery ODBC Data Connector that is installed on your Linux machine, you can query the version number through the command-line interface if the connector was installed using an RPM file. Alternatively, you can search the connector's binary file for version number information.

To verify the connector version number in Linux using the command-line interface:

- Depending on your package manager, at the command prompt, run one of the following commands:
 - `yum list 'Simba*' | grep SimbaODBCDriverforGoogleBigQuery`
 - `rpm -qa | grep SimbaODBCDriverforGoogleBigQuery`

The command returns information about the Simba Google BigQuery ODBC Data Connector that is installed on your machine, including the version number.

To verify the connector version number in Linux using the binary file:

1. Navigate to the `/lib` subfolder in your connector installation directory. By default, the path to this directory is: `/opt/simba/googlebigqueryodbc/lib`.
2. Open the connector's `.so` binary file in a text editor, and search for the text `$driver_version_sb$`: The connector's version number is listed after this text.

Configuring the ODBC Driver Manager in Non-Windows Machines

To make sure that the ODBC Driver manager on your machine is configured to work with the Simba Google BigQuery ODBC Data Connector, do the following:

- Set the library path environment variable to make sure that your machine uses the correct ODBC Driver manager. For more information, see [Specifying ODBC Driver Managers in Non-Windows Machines](#).
- If the connector configuration files are not stored in the default locations expected by the ODBC driver manager, then set environment variables to make sure that the Driver manager locates and uses those files. For more information, see [Specifying the Locations of the Connector Configuration Files](#).

After configuring the ODBC Driver manager, you can configure a connection and access your data store through the connector.

Specifying ODBC Driver Managers in Non-Windows Machines

You need to make sure that your machine uses the correct ODBC Driver manager to load the connector. To do this, set the library path environment variable.

macOS

If you are using a macOS machine, then set the `DYLD_LIBRARY_PATH` environment variable to include the paths to the ODBC driver manager libraries. For example, if the libraries are installed in `/usr/local/lib`, then run the following command to set `DYLD_LIBRARY_PATH` for the current user session:

```
export DYLD_LIBRARY_PATH=$DYLD_LIBRARY_PATH:/usr/local/lib
```

For information about setting an environment variable permanently, refer to the macOS shell documentation.

Linux

If you are using a Linux machine, then set the `LD_LIBRARY_PATH` environment variable to include the paths to the ODBC driver manager libraries. For example, if the libraries are installed in `/usr/local/lib`, then run the following command to set `LD_LIBRARY_PATH` for the current user session:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib
```

For information about setting an environment variable permanently, refer to the Linux shell documentation.

Specifying the Locations of the Connector Configuration Files

By default, ODBC Driver managers are configured to use hidden versions of the `odbc.ini` and `odbcinst.ini` configuration files (named `.odbc.ini` and `.odbcinst.ini`) located in the home directory, as well as the `simba.googlebigqueryodbc.ini` file in the `lib` subfolder of the connector installation directory. If you store these configuration files elsewhere, then you must set the environment variables described below so that the driver manager can locate the files.

If you are using iODBC, do the following:

- Set `ODBCINI` to the full path and file name of the `odbc.ini` file.
- Set `ODBCINSTINI` to the full path and file name of the `odbcinst.ini` file.
- Set `SIMBAGOOGLEBIGQUERYODBCINI` to the full path and file name of the `simba.googlebigqueryodbc.ini` file.

If you are using unixODBC, do the following:

- Set `ODBCINI` to the full path and file name of the `odbc.ini` file.
- Set `ODBCSYSINI` to the full path of the directory that contains the `odbcinst.ini` file.
- Set `SIMBAGOOGLEBIGQUERYODBCINI` to the full path and file name of the `simba.googlebigqueryodbc.ini` file.

For example, if your `odbc.ini` and `odbcinst.ini` files are located in `/usr/local/odbc` and your `simba.googlebigqueryodbc.ini` file is located in `/etc`, then set the environment variables as follows:

For iODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCINSTINI=/usr/local/odbc/odbcinst.ini
export SIMBAGOOGLEBIGQUERYODBCINI=/etc/simba.googlebigqueryodbc.ini
```

For unixODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCSYSINI=/usr/local/odbc
export SIMBAGOOGLEBIGQUERYODBCINI=/etc/simba.googlebigqueryodbc.ini
```

To locate the `simba.googlebigqueryodbc.ini` file, the connector uses the following search order:

1. If the `SIMBAGOOGLEBIGQUERYODBCINI` environment variable is defined, then the connector searches for the file specified by the environment variable.
2. The connector searches the directory that contains the connector library files for a file named `simba.googlebigqueryodbc.ini`.

3. The connector searches the current working directory of the application for a file named `simba.googlebigqueryodbc.ini`.
4. The connector searches the home directory for a hidden file named `simba.googlebigqueryodbc.ini` (prefixed with a period).
5. The connector searches the `/etc` directory for a file named `simba.googlebigqueryodbc.ini`.

Configuring ODBC Connections in Non-Windows Machine

The following sections describe how to configure ODBC connections when using the Simba Google BigQuery ODBC Data Connector on non-Windows platforms:

- [Creating a Data Source Name in a Non-Windows Machine](#)
- [Setting Connector-Wide Configuration Options in a Non-Windows Machine](#)
- [Configuring a DSN-less Connection in a Non-Windows Machine](#)
- [Configuring Authentication in a Non-Windows Machine](#)
- [Configuring the High-Throughput API in a Non-Windows Machine](#)
- [Configuring Logging Options in a Non-Windows Machine](#)
- [Testing the Connection in Non-Windows Machine](#)

Creating a Data Source Name in a Non-Windows Machine

When connecting to your data store using a DSN, you only need to configure the `odbc.ini` file. Set the properties in the `odbc.ini` file to create a DSN that specifies the connection information for your data store. For information about configuring a DSN-less connection instead, see [Configuring a DSN-less Connection in a Non-Windows Machine](#).

If your machine is already configured to use an existing `odbc.ini` file, then update that file by adding the settings described below. Otherwise, copy the `odbc.ini` file from the `Setup` subfolder in the connector installation directory to the home directory, and then update the file as described below.

To create a Data Source Name on a non-Windows machine:

1. In a text editor, open the `odbc.ini` configuration file.

Note: If you are using a hidden copy of the `odbc.ini` file, you can remove the period (.) from the start of the file name to make the file visible while you are editing it.

2. In the `[ODBC Data Sources]` section, add a new entry by typing a name for the DSN, an equal sign (=), and then the name of the connector.

For example, on a macOS machine:

```
[ODBC Data Sources]
```

```
Sample DSN=Simba Google BigQuery ODBC Data Connector
```

As another example, for a 32-bit connector on a Linux machine:

```
[ODBC Data Sources]
```

Sample DSN=Simba Google BigQuery ODBC Data Connector 32-bit

3. Create a section that has the same name as your DSN, and then specify configuration options as key-value pairs in the section:
 - a. Set the `Driver` property to the full path of the connector library file that matches the bitness of the application.

For example, on a macOS machine:

```
Driver=/Library/simba/googlebigqueryodbc/lib/libgooglebigqueryodbc_sbu.dylib
```

As another example, for a 32-bit connector on a Linux machine:

```
Driver=/opt/simba/googlebigqueryodbc/lib/32/libgooglebigqueryodbc_sb32.so
```

- b. Set the `Catalog` property to the name of your BigQuery project. This project is the default project that the Simba Google BigQuery ODBC Data Connector queries against, and also the project that is billed for queries that are run using this DSN.

For example:

```
Catalog=testdata
```

- c. Configure authentication using a Google user account or a Google service account. For more information, see [Configuring Authentication in a Non-Windows Machine](#).
 - d. Optionally, to use trusted CA certificates from a `.pem` file other than the default file installed with the connector, set the `TrustedCerts` property to the full path of the file.
 - e. Optionally, to allow the connector to access Google Drive so that it can support federated tables that combine BigQuery data with data from Google Drive, set the `RequestGoogleDriveScope` property to `1`.
 - f. Optionally, set additional key-value pairs as needed to specify other optional connection settings. For detailed information about all the configuration options supported by the Simba Google BigQuery ODBC Data Connector, see [Connector Configuration Properties](#).
4. Save the `odbc.ini` configuration file.



Note: If you are storing this file in its default location in the home directory, then prefix the file name with a period (.) so that the file becomes hidden. If you are storing this file in another location, then save it as a non-hidden file (without the prefix), and make sure that the `ODBCINI` environment variable specifies the location. For more information, see [Specifying the Locations of the Connector Configuration Files](#).

For example, the following is an `odbc.ini` configuration file for macOS containing a DSN that connects to Google BigQuery using a refresh token obtained from a user account:

```
[ODBC Data Sources]
```

```
Sample DSN=Simba Google BigQuery ODBC Data Connector
```

```
[Sample DSN]
```

```
Driver=/Library/simba/googlebigqueryodbc/lib/libgooglebigqueryodbc_sbu.dylib
```

```
Catalog=testdata
```

```
OAuthMechanism=1
```

```
RefreshToken=CH01pcNn/qFcYwUlJpkF_yyufYrqj4O4g7cdXvGgs-zT6
```

As another example, the following is an `odbc.ini` configuration file for a 32-bit connector on a Linux machine, containing a DSN that connects to Google BigQuery using a refresh token obtained from a user account:

```
[ODBC Data Sources]
```

```
Sample DSN=Simba Google BigQuery ODBC Data Connector 32-bit
```

```
[Sample DSN]
```

```
Driver=/opt/simba/googlebigqueryodbc/lib/32/libgooglebigqueryodbc_sb32.so
```

```
Catalog=testdata
```

```
OAuthMechanism=1
```

```
RefreshToken=CH01pcNn/qFcYwUlJpkF_yyufYrqj4O4g7cdXvGgs-zT6
```

You can now use the DSN in an application to connect to the data store.

Setting Connector-Wide Configuration Options in a Non-Windows Machine

When you specify connection settings in a DSN or connection string, those settings apply only when you connect to BigQuery using that particular DSN or string. As an alternative, you can specify certain settings that apply to every connection that uses the Simba Google BigQuery ODBC Data Connector by configuring them in the `simba.googlebigqueryodbc.ini` file.

**Note:**

Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over connector-wide settings.

To set connector-wide configuration options on a non-Windows machine:

1. In a text editor, open the `simba.googlebigqueryodbc.ini` configuration file.
2. In the `[Driver]` section, specify configuration properties as key-value pairs. Start a new line for each key-value pair.

For example, to specify a proxy server and to log information from the `libcurl` library, type the following:

```
Proxy=http://proxy.mycompany.ca:1066
```

```
CURLVerboseMode=true
```

For detailed information about the configuration options supported by the connector at the connector-wide level, see below.

3. Save the `simba.googlebigqueryodbc.ini` configuration file.
4. Restart your ODBC application to make sure that the new settings take effect.

The `MaxThreads` property must be configured in the `simba.googlebigqueryodbc.ini` file, or as a Windows Registry key value in Windows. For more information see [MaxThreads](#).

Configuring a DSN-less Connection in a Non-Windows Machine

To connect to your data store through a DSN-less connection, you need to define the connector in the `odbcinst.ini` file and then provide a DSN-less connection string in your application.

If your machine is already configured to use an existing `odbcinst.ini` file, then update that file by adding the settings described below. Otherwise, copy the `odbcinst.ini` file from the `Setup` subfolder in the connector installation directory to the home directory, and then update the file as described below.

To define a connector on a non-Windows machine:

1. In a text editor, open the `odbcinst.ini` configuration file.

Note: If you are using a hidden copy of the `odbcinst.ini` file, you can remove the period (.) from the start of the file name to make the file visible while you are editing it.

2. In the `[ODBC Drivers]` section, add a new entry by typing a name for the connector, an equal sign (=), and then `Installed`.

For example:

```
[ODBC Drivers]
```

```
Simba Google BigQuery ODBC Data Connector=Installed
```

3. Create a section that has the same name as the connector (as specified in the previous step), and then specify the following configuration options as key-value pairs in the section:
 - a. Set the `Driver` property to the full path of the connector library file that matches the bitness of the application.

For example, on a macOS machine:

```
Driver=/Library/simba/googlebigqueryodbc/lib/libgooglebigqueryodbc_sbu.dylib
```

As another example, for a 32-bit connector on a Linux machine:

```
Driver=/opt/simba/googlebigqueryodbc/lib/32/libgooglebigqueryodbc_sb32.so
```

- b. Optionally, set the `Description` property to a description of the connector.

For example:

```
Description=Simba Google BigQuery ODBC Data Connector
```

4. Save the `odbcinst.ini` configuration file.

Note: If you are storing this file in its default location in the home directory, then prefix the file name with a period (.) so that the file becomes hidden. If you are storing this file in another location, then save it as a non-hidden file (without the prefix), and make sure that the `ODBCINSTINI` or `ODBCSYSINI` environment variable specifies the location. For more information, see [Specifying the Locations of the Connector Configuration Files](#).

For example, the following is an `odbcinst.ini` configuration file for macOS:

```
[ODBC Drivers]
```

```
Simba Google BigQuery ODBC Data Connector=Installed
```

```
[Simba Google BigQuery ODBC Data Connector]
```

```
Description=Simba Google BigQuery ODBC Data Connector
```

```
Driver=/Library/simba/googlebigqueryodbc/lib/libgooglebigqueryodbc_sbu.dylib
```

As another example, the following is an `odbcinst.ini` configuration file for both the 32- and 64-bit connectors in Linux:

```
[ODBC Drivers]
```

```
Simba Google BigQuery ODBC Data Connector 32-bit=Installed
```

```
Simba Google BigQuery ODBC Data Connector 64-bit=Installed
```

```
[Simba Google BigQuery ODBC Data Connector 32-bit]
```

```
Description=Simba Google BigQuery ODBC Data Connector (32-bit)
```

```
Driver=/opt/simba/googlebigqueryodbc/lib/32/libgooglebigqueryodbc_sb32.so
```

```
[Simba Google BigQuery ODBC Data Connector 64-bit]
```

```
Description=Simba Google BigQuery ODBC Data Connector (64-bit)
```

```
Driver=/opt/simba/googlebigqueryodbc/lib/64/libgooglebigqueryodbc_sb64.so
```

You can now connect to your data store by providing your application with a connection string where the `Driver` property is set to the connector name specified in the `odbcinst.ini` file, and all the other necessary connection properties are also set. For more information, see "DSN-less Connection String Examples" in [Using a Connection String](#)

Configuring Authentication in a Non-Windows Machine

The Simba Google BigQuery ODBC Data 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 an External Account](#)
- [Using Application Default Credentials](#)

Using a Google User Account

You can configure the connector to authenticate the connection with a Google user account. This authentication method uses the OAuth 2.0 access and refresh tokens associated with the user account as the credentials.

The access token is transmitted with every API call that the connector makes, and it is required for accessing BigQuery data stores. However, the access token expires after a certain amount of time and must be renewed using the refresh token. If the refresh token is stored in the DSN, the connector automatically uses it to renew access tokens when they expire.

Note:
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>.

At minimum, you need to provide the OAuth 2.0 refresh token associated with your account. The connector retrieves and uses an access token based on your specified refresh token.

- If you do not have your refresh token, see [Retrieving a Refresh Token](#).
- If you have your refresh token, see [Providing a Refresh Token](#).
- If you want to provide a `.json` key file that contains your credentials instead of providing your refresh token directly in your connection information, see [Providing a Key File](#).

Retrieving a Refresh Token

When you authenticate your connection this way, the authentication process provides a temporary authorization code that you can exchange for an access token and a refresh token.

You can retrieve a refresh token by providing your own credentials, or by using a script that uses Simba-provided credentials.

Note:
If you use your credentials to generate a refresh token, you cannot use it in conjunction with the Simba-provided credentials. Conversely, if you use a refresh token generated with the Simba-provided credentials, it cannot be used in conjunction with your user credentials.

To configure authentication by retrieving a refresh token using Simba-provided credentials on a non-Windows machine:

1. In the `[INSTALL_DIR]/Tools` directory, right-click `get_refresh_token.sh` and select **Edit**.
2. From the internal commented section, copy the entire authentication generator URL.
3. In a web browser, navigate to the URL you copied in the previous step.

4. Click **Allow**. The browser redirects you to a page with an authentication token.
5. Copy the authentication token.
6. Using a command line interface, run `get_refresh_token.sh` with your copied authentication token added as the argument to the script. The script generates a refresh token.

Now that you have a refresh token, see [Providing a Refresh Token](#).

To configure user account authentication by retrieving a refresh token on a non-Windows machine:

1. Obtain a refresh token based on your user account:
 - a. In a web browser, navigate to the Google OAuth 2.0 Playground: <https://developers.google.com/oauthplayground/>.
 - b. In the side panel, expand **BigQuery API v2** and select the appropriate scope for the permissions that you need.
- Note:**
For information about the permissions associated with each scope, see "OAuth 2.0 Scopes for Google APIs" in the Google Identity Platform documentation: <https://developers.google.com/identity/protocols/googlescopes>.
- c. Click **Authorize APIs**.
 - d. Sign in to your user account.
 - e. When you are prompted to allow Google OAuth 2.0 Playground to view and manage your data in Google BigQuery, click **Allow**.

The authentication process returns you to the Google OAuth 2.0 Playground, and automatically populates the Authorization Code field with an authorization code.

- f. Click **Exchange Authorization Code for Tokens**.

The Refresh Token and Access Token fields are populated with the appropriate token values.

2. In your DSN or connection string, set the `OAuthMechanism` property to 1.
3. Set the `RefreshToken` property to the refresh token that you obtained from Google.
4. Set the `ClientId` property to your BigQuery client ID.
5. Set the `ClientSecret` property to the corresponding client secret.

Providing a Refresh Token

If you already have your refresh token, then you can provide the token in your connection information without going through the retrieval process described above.

To configure user account authentication by providing a refresh token on a non-Windows machine:

1. Set the `OAuthMechanism` property to 1.
2. Set the `RefreshToken` property to the refresh token associated with your user account.

Providing a Key File

As an alternative to providing your refresh token directly in your connection information, you can save the token in a `.json` key file and then specify the path to the file in your connection information.

The file must define a JSON object of type `authorized_user` containing the refresh token, client ID, and client secret associated with your user account. For example, the `.json` key file must be written in the following format:

```
{
  "type": "authorized_user",
  "client_id": "[YourClientID]",
  "client_secret": "[YourClientSecret]",
  "refresh_token": "[YourRefreshToken]"
}
```

To configure user account authentication by providing a key file on a non-Windows machine:

1. Set the `OAuthMechanism` property to 0.

Note: Although this is a form of user authentication, the connector must be configured to use the service authentication mechanism (`OAuthMechanism=0`) in order to detect and use the key file.

2. Set the `Email` property to your user account email ID.
3. Set the `KeyFilePath` or `KeyFile` property to the full path to the `.json` key file. Alternatively, set the `KeyFile` property to the plain text JSON object.

Using a Google Service Account

You can configure the connector to authenticate the connection with a Google service account. When you authenticate your connection this way, the connector handles authentication on behalf of the service account, so that an individual user account is not directly involved and no user input is required.

To authenticate your connection this way, you must provide a Google service account email address and the full path to a private key file for the service account. You can generate and download the private key file when you set up the service account.

Note:

- For more information about OAuth 2.0 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>.
- For information about obtaining service account keys, see "Creating and Managing Service Account Keys" in the Google Cloud Identity & Access Management documentation: <https://cloud.google.com/iam/docs/creating-managing-service-account-keys>.

To configure service account authentication on a non-Windows machine:

1. Set the `OAuthMechanism` property to 0.
2. Set the `Email` property to your service account email ID.
3. Set the `KeyFilePath` or `KeyFile` property to the full path to the `.p12` or `.json` key file that is used to authenticate the service account ID. Alternatively, set the `KeyFile` property to the plain text JSON object.

Using an External Account

You can configure the connector to authenticate the connection with an external account (workforce identity federation), with limited support, using Azure AD and Okta identity providers. External account authentication lets the connector exchange a token issued by an external identity provider for a Google Cloud access token, so you do not need to create or manage Google service account keys.

The connector supports two ways of supplying the external subject token:

- File-sourced credentials – The subject token is read from a local file that you keep up to date.
- URL-sourced credentials – The subject token is retrieved at connection time from a local metadata endpoint. Specifications for this credential type are detailed at:

<https://google.aip.dev/auth/4117#determining-the-subject-token-in-microsoft-azure-and-url-sourced-credentials>

To authenticate your connection using external account authorization, you must have workforce identity federation set up with Azure AD or Okta and provide a configuration file in the connection string. The credential configuration file can be generated from the Google API Console (<https://console.developers.google.com/>) or through GCloud CLI.

The Credential Configuration File

The credential configuration file is a JSON with a set of shared top-level fields and a `credential_source`. The contents of `credential_source` determine how the connector obtains the subject token (from a file or from a URL).

Shared Fields

The following top-level fields apply to all external account configurations:

| Field | Required | Description |
|---------------------------------|----------|---|
| <code>type</code> | Yes | Identifies the credential type. This must be <code>external_account</code> or <code>external_account_authorized_user</code> . |
| <code>audience</code> | Yes | The STS audience URI that identifies the identity provider pool. |
| <code>subject_token_type</code> | No | The type of token supplied by the credential source. The default is <code>urn:ietf:params:oauth:token-type:id_token</code> . |
| <code>token_url</code> | No | The endpoint to which a request is |

| Field | Required | Description |
|-----------------------------------|----------|--|
| | | sent to exchange the ID token for a Google Cloud Platform access token. The default is https://sts.googleapis.com/v1/token . |
| service_account_impersonation_url | No | The service account impersonation endpoint. If set, the connector exchanges the STS access token for a short-lived service account access token by calling this endpoint. |
| service_account_impersonation | No | Additional service account impersonation options. Only token_lifetime_seconds is currently supported, which sets the requested access token lifetime. |
| workforce_pool_user_project | No | The Google Cloud Platform project to which the token exchange request is executed. |
| credential_source | Yes | Defines how the connector retrieves the subject token. See the file-sourced and URL-sourced sections below. |

Note: Because subject_token_type and token_url have defaults, you can omit them from the configuration file.

File-Sourced Credentials

Use file-sourced credentials when the subject token is written to a local file by an external process. The connector reads the token from that file at connection time. An example credential file for workforce identity federation with Azure AD or Okta looks like this:

```

"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": "[PATH_TO_OIDC_CREDENTIALS]"
}

```

The credential_source object supports the following fields for file-sourced credentials:

| Field | Required | Description |
|---------------------------------|----------|--|
| file | Yes | The full path to a plaintext file that contains the OIDC subject token. |
| format.type | No | The format of the token file, either text or json. The default is text. |
| format.subject_token_field_name | No | Required when format.type is json. The JSON field name that holds the subject token. |

When the file is in plaintext format, its contents must be the subject token only; any additional content causes the exchange to fail.

Important: Subject tokens expire and cannot be reused to obtain new access tokens. An external process must keep the token file updated with a valid, unexpired subject token.

For additional setup examples, see [Configure Workforce Identity Federation with Azure AD and Sign In Users](https://cloud.google.com/iam/docs/workforce-sign-in-azure-ad) and [Configure Workforce Identity Federation with Okta and Sign In Users](https://cloud.google.com/iam/docs/workforce-sign-in-okta) in the Google Cloud documentation: <https://cloud.google.com/iam/docs/workforce-sign-in-azure-ad> and <https://cloud.google.com/iam/docs/workforce-sign-in-okta>

URL-Sourced Credentials

Use URL-sourced credentials when the subject token is served by a local metadata endpoint, such as the Azure Instance Metadata Service or an OIDC token endpoint. Instead of reading a file, the connector sends an HTTP GET request to the configured URL at connection time and uses the response as the subject token.

The credential_source object supports the following fields for URL-sourced credentials:

| Field | Required | Description |
|---------------------------------|----------|---|
| url | Yes | The local metadata endpoint that returns the external credential. For Azure, this is the IMDS URL used to retrieve the Azure AD access token. |
| headers | No | A set of key/value pairs that the connector appends to the GET request headers. |
| format.type | No | The format of the URL response, either text or json. The default is text. |
| format.subject_token_field_name | No | Required when format.type is json. The JSON field name that holds the subject token. |

When format is omitted, the connector treats the response body as a plaintext token. When format.type is json, the connector parses the response as JSON and reads the token from the field named in format.subject_token_field_name.

Azure configurations are a type of URL-sourced credential. For Azure, set format.subject_token_field_name to access_token to retrieve the access token from the Azure Instance Metadata Service.

Configure Workforce Identity Federation Authentication on a Non-Windows Machine Using the KeyFilePath Property

1. Set the OAuthMechanism property to 4.
2. Set the KeyFilePath property to the full path of the external account configuration file. For example:

```
Driver=Simba ODBC Driver for Google
BigQuery;...;OAuthMechanism=4;KeyFilePath=/tmp/path/to/oidc_configuration_file.json
```

Configure Workforce Identity Federation Authentication on a Non-Windows Machine Using the KeyFile Property

1. Set the `OAuthMechanism` property to 4.
2. Set the `KeyFile` property to either the full path of the external account configuration file, or a raw JSON object containing the configuration file contents. For example:

Using a file path:

```
Driver=Simba ODBC Driver for Google
BigQuery;...;OAuthMechanism=4;KeyFile=/tmp/path/to/oidc_configuration_file.json
```

Using a raw JSON object:

```
Driver=Simba ODBC Driver for Google
BigQuery;...;OAuthMechanism=4;KeyFile={"audience":
"//iam.googleapis.com/locations/[LOCATION]/workforcePools/[POOL_
ID]/providers/[PROVIDER_ID]","workforce_pool_user_project": "[PROJECT_
ID]","credential_source": {"file": "[PATH_TO_SUBJECT_TOKEN_FILE]"}}}}
```

Configure Workforce Identity Federation Authentication on a Non-Windows Machine Using BYOID_ Properties

1. Set the `OAuthMechanism` property to 4.
2. Set the appropriate `BYOID_` property to reconstruct the configuration file internally. The following table lists the available properties and their default values:

| Property Name | SQL Default Value |
|-------------------------------------|--|
| <code>BYOID_AudienceUri</code> | None |
| <code>BYOID_CredentialSource</code> | None |
| <code>BYOID_PoolUserProject</code> | None |
| <code>BYOID_SubjectTokenType</code> | <code>urn:ietf:params:oauth:token-type:id_token</code> |
| <code>BYOID_TokenUri</code> | <code>https://sts.googleapis.com/v1/token</code> |

Each of the above properties are strings that represent a field present in the configuration file, with the exception of `BYOID_CredentialSource`, which is a string with a JSON object. For example:

```
Driver=Simba ODBC Driver for Google BigQuery;...;BYOID_
AudienceUri="//iam.googleapis.com/locations/global/workforcePools/testpool-1/providers/okta-provider-
1;BYOID_PoolUserProject=898755234664;BYOID_CredentialSource={"file": "Tests/OAuth/okta_subject_
token.txt"}}
```

The above example shows the minimum properties required in the connection string. The `subject_token_type` and `token_url` properties are absent, to demonstrate that the connector will use the default values. If the corresponding `BYOID_` property (`BYOID_SubjectTokenType` or `BYOID_TokenUri`) is set, the connector will use these properties.

For more information about these properties, see [Configuration Options Having Only Key Names](#).

Note:

- For more information about using external accounts, see "Workforce Identity Federation" in the Google Cloud documentation: <https://cloud.google.com/iam/docs/workforce-identity-federation>.
- When the connector is configured to use External Account Authentication (`OAuthMechanism=4`), connection properties are considered in the following precedence:
 1. `KeyFile`
 2. `KeyFilePath` (or `KeyFilePath_Enc` if the key file is not set)
 3. `BYOID_` properties
- It is recommended to set the corresponding `BYOID_` property in the configuration file. These properties are intended to act as an option for customers who cannot specify `.json` key files to pass to the `KeyFile` property.
- `BYOID_` properties are deprecated. They currently remain supported for backward compatibility, and let you supply external account settings individually for environments that cannot pass a JSON configuration file. New configurations are recommended to use a credential configuration file through `KeyFile` or `KeyFilePath` instead.

Using Application Default Credentials

You can configure the connector to authenticate the connection with Application Default Credentials. When you authenticate your connection this way, the connector parses the JSON keyfile specified by the `GOOGLE_APPLICATION_CREDENTIALS` environment variable. No other connection properties are required to use the Application Default Credentials authentication flow.

To configure Application Default Credentials authentication on a non-Windows machine:

- Set the `OAuthMechanism` property to 3.

Note: For more information about Application Default Credentials, please refer to "How Application Default Credentials works" and "Set up Application Default Credentials" in the Google Cloud documentation: <https://cloud.google.com/docs/authentication/application-default-credentials> and <https://cloud.google.com/docs/authentication/provide-credentials-adc>, respectively.

Metadata Server Authentication for Application Default Credentials

When Application Default Credentials (ADC) are used for authentication and no key file is present, the ODBC driver falls back to retrieving an access token for the default service account from the Google internal metadata server. This fallback mechanism is available only when the driver is running on a Google Cloud Engine virtual machine.

Note: For more details about authenticating through the internal metadata server, see [Application Default Credentials overview](#) and [Restricting service accounts](#).

Configuring the High-Throughput API in a Non-Windows Machine

You can configure the connector to use the High-Throughput API to handle large result sets more efficiently. For more information about the High-Throughput API, see [High-Throughput API](#).

To configure the High-Throughput API on a non-Windows machine:

1. Make sure that your Google BigQuery project has the Storage API enabled. For more information about the Storage API, see "BigQuery Storage API Overview" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/reference/storage/>.

Note: in Linux, if the application running the connector has been built with an RPATH, then you must either include the connector directory in that RPATH, or adjust your `LD_LIBRARY_PATH` to include the connector directory.

2. In the `odbc.ini` file, if you are using Legacy SQL (the `SQLDialect` property is set to 0), then set the `AllowLargeResults` property to 1.
3. To specify the dataset that stores temporary tables for large result sets and result sets returned by the High-Throughput API, do one of the following:
 - To use the default dataset with the `ID_bqodbc_temp_tables`, set the `UseDefaultLargeResultsDataset` property to 1.
 - Or, to specify a different dataset, set the `UseDefaultLargeResultsDataset` property to 0 and set the `LargeResultsDataSetId` property to the ID of the BigQuery dataset that you want to use.

Note: If the dataset does not exist and the data store specifies the US region, the connector creates the dataset.

4. Set the `EnableHTAPI` property to 1.
5. Set the `HTAPI_MinResultsSize` property to the minimum number of table rows required to activate the High-Throughput API.
6. Set the `HTAPI_MinActivationRatio` property to the minimum ratio of total rows to rows in the first page required to activate reading through the High-Throughput API.

Note: If this value is set to 0, then the connector uses the High-Throughput API for all query results that meet the minimum results size specified by the `HTAPI_MinResultsSize` property.

The connector uses the BigQuery High-Throughput API instead of the REST API for requests where both:

- the number of table rows in your query results exceeds the `HTAPI_MinResultsSize` value;
- and the number of pages in the results exceeds the `HTAPI_MinActivationRatio` value.

Configuring Logging Options in a Non-Windows Machine

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.

You can set the connection properties described below in a connection string, in a DSN (in the `odbc.ini` file), or as a connector-wide setting (in the `simba.googlebigqueryodbc.ini` file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over connector-wide settings.

To enable logging on a non-Windows machine:

1. To specify the level of information to include in log files, set the `LogLevel` property to one of the following numbers:

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

2. Set the `LogPath` key to the full path to the folder where you want to save log files.
3. Set the `LogFileCount` key to the maximum number of log files to keep.

Note: After the maximum number of log files is reached, each time an additional file is created, the connector deletes the oldest log file.

4. Set the `LogFileSize` key to the maximum size of each log file in bytes.



Note: After the maximum file size is reached, the connector creates a new file and continues logging.

5. Save the `simba.googlebigqueryodbc.ini` configuration file.
6. Restart your ODBC application to make sure that the new settings take effect.

The Simba Google BigQuery ODBC Data Connector produces the following log files at the location you specify using the `LogPath` key:

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

To disable logging on a non-Windows machine:

1. Set the `LogLevel` key to 0.
2. Save the `simba.googlebigqueryodbc.ini` configuration file.
3. Restart your ODBC application to make sure that the new settings take effect.

Testing the Connection in Non-Windows Machine

To test the connection, you can use an ODBC-enabled client application. For a basic connection test, you can also use the test utilities that are packaged with your driver manager installation. For example, the iODBC driver manager includes simple utilities called `iodbctest` and `iodbctestw`. Similarly, the unixODBC driver manager includes simple utilities called `isql` and `iusql`.

Using the iODBC Driver Manager

You can use the `iodbctest` and `iodbctestw` utilities to establish a test connection with your connector. Use `iodbctest` to test how your connector works with an ANSI application, or use `iodbctestw` to test how your connector works with a Unicode application.



Note: There are 32-bit and 64-bit installations of the iODBC driver manager available. If you have only one or the other installed, then the appropriate version of `iodbctest` (or `iodbctestw`) is available. However, if you have both 32- and 64-bit versions installed, then you need to make sure that you are running the version from the correct installation directory.

For more information about using the iODBC driver manager, see <http://www.iodbc.org>.

To test your connection using the iODBC driver manager:

1. Run `iodbctest` or `iodbctestw`.
2. Optionally, if you do not remember the DSN, then type a question mark (?) to see a list of available DSNs.

3. Type the connection string for connecting to your data store, and then press ENTER. For more information, see [Using a Connection String](#).

If the connection is successful, then the `SQL>` prompt appears.

Using the unixODBC Driver Manager

You can use the `isql` and `iusql` utilities to establish a test connection with your connector and your DSN. `isql` and `iusql` can only be used to test connections that use a DSN. Use `isql` to test how your connector works with an ANSI application, or use `iusql` to test how your connector works with a Unicode application.

Note: There are 32-bit and 64-bit installations of the unixODBC driver manager available. If you have only one or the other installed, then the appropriate version of `isql` (or `iusql`) is available. However, if you have both 32- and 64-bit versions installed, then you need to make sure that you are running the version from the correct installation directory.

For more information about using the unixODBC driver manager, see <http://www.unixodbc.org>.

To test your connection using the unixODBC driver manager:

- Run `isql` or `iusql` by using the corresponding syntax:

```
▪ isql [DataSourceName]
▪ iusql [DataSourceName]
```

`[DataSourceName]` is the DSN that you are using for the connection.

If the connection is successful, then the `SQL>` prompt appears.

Note: For information about the available options, run `isql` or `iusql` without providing a DSN.

Using a Connection String

For some applications, you might need to use a connection string to connect to your data source. For detailed information about how to use a connection string in an ODBC application, refer to the documentation for the application that you are using.

The connection strings in the following sections are examples showing the minimum set of connection attributes that you must specify to successfully connect to the data source. Depending on the configuration of the data source and the type of connection you are working with, you might need to specify additional connection attributes. For detailed information about all the attributes that you can use in the connection string, see [Connector Configuration Properties](#).

DSN Connection String Example

The following is an example of a connection string for a connection that uses a DSN:

```
DSN=[DataSourceName]
```

[DataSourceName] is the DSN that you are using for the connection.

You can set additional configuration options by appending key-value pairs to the connection string. Configuration options that are passed in using a connection string take precedence over configuration options that are set in the DSN.

DSN-less Connection String Examples

Some applications provide support for connecting to a data source using a connector without a DSN. To connect to a data source without using a DSN, use a connection string instead.

The placeholders in the examples are defined as follows, in alphabetical order:

- *[ClientID]* is the Client ID to use when authenticating the connection to BigQuery.
- *[ClientSecret]* is the Client Secret to use when authenticating the connection to BigQuery.
- *[PortNumber]* is the number of the TCP port that the proxy server uses to listen for client connections.
- *[Project]* is the BigQuery project containing the data that you want to use.
- *[Server]* is the IP address or host name of the proxy server to which you are connecting.
- *[ServiceAccount]* is your service account email ID.
- *[ServiceKeyPath]* is the full path to a `.p12` or `.json` key file for service account authentication.
- *[Token]* is the refresh token that you obtain from Google for authorizing access to BigQuery.
- *[UserAccount]* is your user account email ID.

- `[UserKeyPath]` is the full path to a `.json` key file containing your refresh token, client ID, and client secret. For information about the required format of the `.json` file, see [Key File Path](#).

Connecting to Google BigQuery using a User Account

The following is the format of a DSN-less connection string for a user account connection to Google BigQuery:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=1;RefreshToken=[Token];Catalog=[Project];ClientId=[ClientID];ClientSecret=[ClientSecret];
```

For example:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=1;RefreshToken=CH01pcNn/qFcYwUIJpkF_yyufYrqj4O4g7cdXvGgs-  
zT6;Catalog=testdata;ClientId=977385342095.apps.googleusercontent.com;ClientSecret=wbER7576  
mc_1YOIIdGk7jEE;
```

As an alternative to providing your refresh token directly in the string, you can save your credentials in a `.json` key file and then provide the full path to that file in your string. In this case, the connection string must be written in the following format:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=0;Email=[UserAccount];KeyFilePath=[UserKeyPath];Catalog=[Project];
```

For example:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=0;Email=simba@gmail.com;  
KeyFilePath=C:\SecureFiles\UserKeyFile.json;Catalog=testdata;
```

Connecting to Google BigQuery using a Service Account

The following is the format of a DSN-less connection string for a service account connection to Google BigQuery:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=0;Email=[ServiceAccount];KeyFilePath=[ServiceKeyPath];Catalog=[Project];
```

For example:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=0;Email=application-service-  
account@iam.gserviceaccount.com;KeyFilePath=C:\SecureFiles\ServiceKeyFile.p12;Catalog=testdat  
a;
```

Connecting to Google BigQuery through a Proxy Server

The following is the format of a DSN-less connection string for connecting to Google BigQuery with a user account through a proxy server:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=1;RefreshToken=[Token];Catalog=[Project];  
ProxyHost=[Server];ProxyPort=[PortNumber];
```

For example:

```
Driver=Simba Google BigQuery ODBC Data Connector;  
OAuthMechanism=1;RefreshToken=CH01pcNn/qFcYwUIJpkF_yyufYrqj4O4g7cdXvGgs-zT6;  
Catalog=testdata;ProxyHost=192.168.222.160;  
ProxyPort=8000;
```

Features

For more information on the features of the Simba Google BigQuery ODBC Data Connector, see the following:

- [Data Types](#)
- [Nested and Repeated Records](#)
- [Arrays](#)
- [Security and Authentication](#)
- [Catalog and Schema Support](#)
- [Large Result Set Support](#)
- [High-Throughput API](#)
- [Write-Back](#)
- [Positional Parameters](#)
- [ODBC Escapes](#)
- [Service Endpoints](#)
- [Google Cloud Dedicated \(GCD\)](#)

Data Types

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

- [Data type mappings: BigQuery to SQL](#)
- [Data type mappings: SQL to BigQuery](#)

The following table lists the supported data type mappings from BigQuery to SQL.

| BigQuery Data Type | SQL Data Type |
|--------------------|---------------|
| ARRAY | SQL_VARCHAR |

| BigQuery Data Type | SQL Data Type |
|--------------------|--|
| BIGNUMERIC | SQL_NUMERIC |
| | SQL_DECIMAL |
| | <p>Note:</p> <p>The connector sends SQL_DECIMAL data to BigQuery as BIGNUMERIC data, because BigQuery does not support a DECIMAL data type.</p> <p>The connector always returns BIGNUMERIC data as SQL_NUMERIC data, and sends SQL_NUMERIC data to BigQuery as BIGNUMERIC data.</p> |
| BOOL | SQL_BIT |
| BOOLEAN | SQL_BIT |
| BYTES | SQL_VARBINARY |
| DATE | SQL_DATE |
| DATETIME | SQL_TYPE_TIMESTAMP |
| | <p>Note:</p> <p>For ODBC versions prior to ODBC 3, the connector uses SQL_TIMESTAMP.</p> |
| FLOAT64 | SQL_DOUBLE |
| GEOGRAPHY | SQL_VARCHAR or SQL_WVARCHAR. |
| | <p>Note:</p> <ul style="list-style-type: none"> For information about whether GEOGRAPHY data is returned as SQL_VARCHAR or SQL_WVARCHAR, see Use SQL_WVARCHAR instead of SQL_VARCHAR. Geography data can only be inserted using specific functions. For more information, see: https://cloud.google.com/bigquery/docs/reference/standard-sql/geography_functions. |
| INTEGER | SQL_BIGINT |
| INTERVAL | SQL_VARCHAR |
| INT64 | SQL_BIGINT |
| JSON | SQL_VARCHAR |

| BigQuery Data Type | SQL Data Type |
|--------------------|---|
| NUMERIC | SQL_NUMERIC |
| | SQL_DECIMAL |
| | <p>Note:</p> <p>The connector sends SQL_DECIMAL data to BigQuery as NUMERIC data, because BigQuery does not support a DECIMAL data type.</p> <p>The connector always returns NUMERIC data as SQL_NUMERIC data, and sends SQL_NUMERIC data to BigQuery as NUMERIC data.</p> |
| STRING | SQL_VARCHAR or SQL_WVARCHAR |
| | <p>Note:</p> <p>For information about whether STRING data is returned as SQL_VARCHAR or SQL_WVARCHAR, see Use SQL_WVARCHAR instead of SQL_VARCHAR.</p> |
| STRUCT | SQL_VARCHAR |
| TIME | SQL_TIME |
| TIMESTAMP | SQL_TYPE_TIMESTAMP |
| | <p>Note:</p> <p>For ODBC versions prior to ODBC 3, the connector uses SQL_TIMESTAMP.</p> |

The following table lists the supported data type mappings from SQL to BigQuery.

| SQL Data Type | BigQuery Data Type |
|-------------------|--------------------|
| SQL_BIGINT | INT64 |
| SQL_BIT | BOOL |
| SQL_CHAR | STRING |
| SQL_DATE | DATE |
| SQL_DECIMAL | NUMERIC |
| SQL_DOUBLE | FLOAT64 |
| SQL_INTEGER | INT64 |
| SQL_LONGVARBINARY | BYTES |
| SQL_LONGVARCHAR | STRING |
| SQL_NUMERIC | NUMERIC |
| SQL_SMALLINT | INT64 |

| SQL Data Type | BigQuery Data Type |
|--------------------|--------------------|
| SQL_TIME | TIME |
| SQL_TIMESTAMP | TIMESTAMP |
| SQL_TINYINT | INT64 |
| SQL_TYPE_DATE | DATE |
| SQL_TYPE_TIME | TIME |
| SQL_TYPE_TIMESTAMP | TIMESTAMP |
| SQL_VARBINARY | BYTES |
| SQL_VARCHAR | STRING |
| SQL_WLONGVARCHAR | STRING |
| SQL_WVARCHAR | STRING |

Nested and Repeated Records

The Simba Google BigQuery ODBC Data Connector partially supports nested and repeated records.

In Standard SQL, the record is returned as a BigQuery string, and converted to SQL_VARCHAR. The Standard SQL syntax represents the sub-components of record data as nested sub-types. In the example below, `city` and `years` belong to the base record type of `address`, and `name` does not.

```
{
  "v": { // "column" object
    "f": [
      { // "address" value
        "v": [ // "address" array
          { // "city" value
            "v": "Vancouver"
          },
          { // "years" value
            "v": "5"
          }
        ]
      },
      { // "name" value
        "v": "Google"
      }
    ]
  }
}
```

```
        }
    ]
}
}
```

In Legacy SQL, the record is flattened by BigQuery before it is returned by the connector. The query returns the data as the schema of a flat table. All STRUCT members are their own columns, and all ARRAY members are expanded into as many rows as there are array elements.

Arrays

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

For example, the SQL statement `SELECT [1,2,3]` returns the following JSON:

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

Security and Authentication

To protect data from unauthorized access, BigQuery data stores require all connections to be authenticated using the OAuth 2.0 protocol and encrypted using TLS 1.2 with one-way authentication. The Simba Google BigQuery ODBC Data Connector protects your data by providing support for these authentication protocols and further obscuring data from unwanted access by fetching it in a non-text format. The data is compressed using zlib and encrypted using TLS.

The connector provides mechanisms that allow you to complete an OAuth 2.0 authentication flow using a Google user account or a Google service account. The connector retrieves a token based on the account credentials specified in your DSN or connection string, and then uses the token to authenticate

the connection to BigQuery. For detailed configuration instructions, see [Configuring Authentication in Windows](#) or [Configuring Authentication in a Non-Windows Machine](#).

Additionally, the connector automatically encrypts all connections with TLS. TLS encryption protects data and credentials when they are transferred over the network, and provides stronger security than authentication alone. By default, the connector uses the trusted CA certificates file that is included during installation, but you can configure the connector to use a different file by setting the `TrustedCertificates` option (the `TrustedCerts` property). In Windows machines, you can configure the connector to use the system trust store by setting the `Use System Trust Store` option (the `UseSystemTrustStore` property). For detailed configuration instructions, see [Creating a Data Source Name in Windows](#) or [Creating a Data Source Name in a Non-Windows Machine](#).

Catalog and Schema Support

The Simba Google BigQuery ODBC Data Connector supports both catalogs and schemas to make it easy for the connector to work with various ODBC applications. Projects are mapped to catalogs, and table datasets are mapped to schemas. For more information, see [Catalog \(Project\)](#).

Large Result Set Support

The Simba Google BigQuery ODBC Data Connector supports the `AllowLargeResults` option in BigQuery job configurations, enabling result sets greater than the BigQuery maximum result sizes. To store large query results, the connector creates temporary tables in BigQuery under the dataset ID specified using the `Dataset Name For Large Result Sets` connector configuration property. These temporary tables exist for a limited time, specified using the `Temporary Table Expiration Time` connector configuration property, before they are deleted.



Note:

For more information about BigQuery maximum result sizes, see "Quotas and Limits" in the Google BigQuery documentation: https://cloud.google.com/bigquery/quotas#query_jobs.

Large result sets are always supported if Standard SQL is used. If Legacy SQL is used, large result sets are only supported if the `Allow Large Result Sets` check box is selected or the `AllowLargeResults` option is set to 1.

For more information about large result sets and the limitations of this feature, see the following sections in the BigQuery documentation:

- "Queries" in *Quota Policy*: <https://developers.google.com/bigquery/quota-policy>.
- "Returning large query results" in *Query Data*: <https://developers.google.com/bigquery/querying-data>.



Note:

To enable the High-Throughput API for large result sets, select the `Enable High-Throughput API` check box or set the `EnableHTAPI` property to 1. For more information, see [High-Throughput API](#).

High-Throughput API

The High-Throughput API is a new feature of the Simba Google BigQuery ODBC Data Connector. This API enables the connector to leverage the BigQuery Storage API, allowing higher data throughput than the standard REST API. With this API, the connector can handle large result sets more efficiently. For more information about the Storage API, see "BigQuery Storage API Overview" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/docs/reference/storage/>.

The connector checks the number of rows in an incoming result set table and extrapolates the number of pages needed to retrieve all the results. If the number of table rows and the ratio of rows to pages both exceed the defined thresholds, the connector uses the BigQuery Storage API. Should the connector encounters any issues with initializing the storage API for retrieval, it falls back to using the standard REST API, unless this is a permissions issue.

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:

- [AllowHtapiForLargeResults](#)
- [Activation Threshold for High-Throughput API](#)

Requirements

Before you can use the High-Throughput API, you must make certain that your system meets the following requirements:

- The BigQuery project that 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>.

- in Windows, the connector must not be configured to use the trust store, that is, the Use System Trust Store check box must not be selected or the `UseSystemTrustStore` property must not be set to 1.
- To enable the High-Throughput API for use with large result sets, the Enable High-Throughput API check box must be selected or the `EnableHTAPI` property must be set to 1.

The INTERVAL data type is not supported on the Read API. When retrieving data from a column of the INTERVAL type, the connector returns an error. To enable the connector to fallback to the REST API, set `UnsupportedHTAPIFallback` to 1.

Write-Back

The Simba Google BigQuery ODBC Data Connector supports Data Manipulation Language (DML) statements such as INSERT, MERGE, 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.

The Simba Google BigQuery ODBC Data 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=?
```

ODBC Escapes

The Simba Google BigQuery ODBC Data Connector supports a subset of the ODBC escape syntaxes. For a complete list of the escapes that the connector supports, call SQLGetInfo from the connector.

For more information about ODBC escapes, see "ODBC Escape Sequences" in the Programmer's Reference: [https://msdn.microsoft.com/en-us/library/ms711838\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/ms711838(v=vs.85).aspx).

For information about known issues that occur for specific ODBC escape use cases, see the "Known Issues" section in the *Simba Google BigQuery ODBC Data Connector Release Notes*.

Service Endpoints

The Simba Google BigQuery ODBC Data Connector uses the following API service endpoints:

| API Service | Service Endpoint |
|--|--|
| Default server name | www.googleapis.com |
| Default API URL | https://bigquery.googleapis.com/bigquery/v2 |
| Default access token request URL (User authentication) | https://oauth2.googleapis.com/ token |
| Default access token request URL (Service authentication) | https://accounts.google.com/ o/oauth2/auth? |
| Default OAuth scope | https://www.googleapis.com/ auth/cloud-platform |
| Google Drive scope | https://www.googleapis.com/ auth/drive |
| Google Client X509 certification URL | https://www.googleapis.com/ robot/v1/metadata/x509/ |
| Storage API URL | https://bigquerystorage.googleapis.com/ |

In addition, if you are behind a proxy, you must make sure that the proxy allows communication to `curl.pki.google`, `crls.pki.google`, and `ocsp.pki.google`.

Google Cloud Dedicated (GCD)

Google Cloud Dedicated (GCD) is Google’s solution to data sovereignty. This section assumes that the reader has already interacted with Google’s sales and support teams and is familiar with the process of setting up GCD for their services.

For example, a company might set up a BigQuery service with Google that uses a different web address than the standard one. Instead of the usual BigQuery URL, it might look like `https://bigquery.not-google-apis.google`. In this case, `not-google-apis.google` is called the **UniverseDomain**, and `bigquery` is the Service (or **PrivateService**, if it is customized, like `bigquery-example`). You can set up the connector to connect to these service endpoints.

If the connector finds a mismatch between the domain in your credentials and the domain you are trying to access like a difference between the `UniverseDomain` setting and the `"universe_domain"` in your credentials file, it will show a 'Domain Mismatch' error. This indicates that there is a conflict in the domains the connection is configured to query. That is because credentials cannot be used across different universe domains.

Note: Note: Google Cloud Dedicated does not support all services, and some supported services may not have full implementations yet.

As of July 2025, the following OAuth2 services are known to be unsupported:

- User Account authentication (OAuthMechanism values 1 and 2) are not supported. Any attempt to connect using this method will fail.

Using PrivateService and UniverseDomain to configure Google Cloud Dedicated (GCD)

The **PrivateService** and **UniverseDomain** properties can be configured in an ODBC connection string or a DSN.

As an example (Presuming `EnableGCD`` is ``1`` in the product INI.):

```
DSN=BigQueryOdbcDsn; OAuthMechanism=4; KeyFile=/path/to/config_file.json;
UniverseDomain=test-domain.google; PrivateService=test-service;
```

The example connection string uses External Account authentication to access Google services and then uses that authorization to make HTTP requests. The configuration file includes a `token_url` and may also have a `"universe_domain"`. The connector reads both the config file and the connection string, compares the values, and figures out which URL to use for Google’s Secure Token Service. Based on this, it will use the following URLs when connecting to Google services:

| Service | Original URI | New URI |
|----------------------|---|--|
| Secure Token Service | <code>https://sts.googleapis.com</code> | <code>https://sts-test-service.p.test-domain.google</code> |

| Service | Original URI | New URI |
|----------------------------------|--|---|
| Identity Access Management (IAM) | https://iamcredentials.googleapis.com | https://iamcredentials-test-service.p.test-domain.goog |
| BigQuery | https://bigquery.googleapis.com | https://bigquery-test-service.p.test-domain.goog |
| BigQuery Storage | https://bigquerystorage.googleapis.com:443 | https://bigquerystorage-test-service.p.test-domain.goog:443 |

Using PrivateServiceConnectUris to configure Google Cloud Dedicated (GCD)

Using **PrivateServiceConnectUris** to set up Google Cloud Dedicated is allowed in the initial release, but no longer be supported in future updates. This option is available only for backward compatibility.

PrivateServiceConnectUris is the older method used to connect to Google’s Private Service Connect. It has been updated to let users directly specify the URIs for their Google Cloud Dedicated service endpoints. This property takes a comma-separated list of key=value pairs, where each key is a name and each value is a URI. The connector reads this list and uses those URIs instead of the default ones. Unlike the newer method, these URIs are not based on the PrivateService and UniverseDomain values.

When using GCD through the `PrivateServiceConnectUris` property, the scheme and authority of the URI must be provided for the below URIs:

| URI | Default Value |
|----------|--|
| OAUTH2 | https://oauth2.googleapis.com |
| STS | https://sts.googleapis.com |
| IAM | https://iamcredentials.googleapis.com |
| BIGQUERY | https://bigquery.googleapis.com |
| READ_API | https://bigquerystorage.googleapis.com:443 |

If a URI is left out of the list, the connector treats it as the default when making the connection. So, a connection string like the one below will result in the following possible URIs:

```
DSN=BigQueryOdbcDsn; OAuthMechanism=4; KeyFile=/path/to/config_file.json; PrivateServiceConnectUris=STS=https://sts-test-service.p.test-domain.goog, BIGQUERY=https://bigquery-test-service.p.test-domain.goog;
```

| URI | Value |
|--------|---|
| OAUTH2 | https://oauth2-test-domain.p.test-domain.goog |

| URI | Value |
|----------|--|
| STS | https://sts-test-domain.p.test-domain.goog |
| IAM | https://iamcredentials-test-domain.p.test-domain.goog |
| BIGQUERY | https://bigquery-test-domain.p.test-domain.goog |
| READ_API | https://bigquerystorage-test-domain.p.test-domain.goog:443 |

To avoid domain validation errors, it is recommended to provide values for all required URIs.

```
DSN=BigQueryOdbcDsn; OAuthMechanism=4; KeyFile=/path/to/config_file.json;
PrivateServiceConnectUris=OAUTH2=https://oauth2-test-domain.p.test-
domain.goog, STS=https://sts-test-domain.p.test-
domain.goog, IAM=https://iamcredentials-test-domain.p.test-
domain.goog, BIGQUERY=https://bigquery-test-domain.p.test-domain.goog, READ_
API=https://bigquerystorage-test-domain.p.test-domain.goog:443
```

URI Precedence and Resolution

It is important to understand which URI takes priority when connecting to the connector, since URIs can be defined in multiple places. The three main sources are:

- **PrivateServiceConnectUris** connection property.
- **UniverseDomain** and/or **PrivateService** connection properties when **EnableGCD=1**.
- The key file or configuration file for authentication flows (for example: **KeyFile** connection property).

The order listed above shows the general priority the connector follows when choosing which URI to use. If **PrivateServiceConnectUris** is set, it takes priority over the **UniverseDomain** and **PrivateService** settings. In that case, the connector will ignore the later.

For the supported Service Account and External Account authentications, the following URI resolution process is applied:

1. Get the URI from the file (if there is a "universe_domain" in the file, try to apply that universe domain to the URI in the file).
2. Make sure the universe domain from the config file does not conflict with the one from the connection settings.
3. If the universe domain is the default `googleapis.com`, use the URI from the config file for authentication. Otherwise, use the one from the connection settings.

 **Note:**

- If the connection settings include a universe domain, the same domain must also be in the credentials file. If they do not match, the connection will fail with a DomainMismatch error.
- If the connection string does not specify a universe domain, the universe domain of the connection is considered to be `googleapis.com`.
- If the credentials file does not have a `"universe_domain"` field, the universe domain of the credentials file is considered to be `googleapis.com`.

Connector Configuration Properties

Connector Configuration Options lists the configuration options available in the Simba Google BigQuery ODBC Data Connector alphabetically by field or button label. Options having only key names, that is, not appearing in the user interface of the connector, are listed alphabetically by key name.

When creating or configuring a connection from a Windows machine, the fields and buttons described below are available in the following dialog boxes:

- Simba Google BigQuery ODBC Data Connector DSN Setup
- Advanced Options
- Logging Options

When using a connection string or configuring a connection from a non-Windows machine, use the key names provided below.

Configuration Options Appearing in the User Interface

The following configuration options are accessible via the Windows user interface for the Simba Google BigQuery ODBC Data Connector, or via the key name when using a connection string or configuring a connection from a Linux or macOS computer:

- [Activation Threshold for High-Throughput API](#)
- [Additional Projects](#)
- [AllowHtapiForLargeResults](#)
- [Allow Large Result Sets](#)
- [Catalog \(Project\)](#)
- [Dataset Name For Large Result Sets](#)
- [Dataset](#)
- [Default String Column Length](#)
- [Email](#)
- [Enable GCD](#)
- [Enable High-Throughput API](#)
- [Enable Session](#)
- [OAuth Mechanism](#)
- [Path to CMEK](#)
- [PrivateService](#)
- [Proxy Host](#)
- [Proxy Password](#)
- [Proxy Port](#)
- [Proxy Username](#)
- [Query Properties](#)
- [Ratio of Results to Rows Per Block](#)
- [Refresh Token](#)
- [Request Google Drive Scope Access](#)
- [Rows Fetched Per Block](#)

- Enable Timestamp Picos
- HTAPI_SessionRetryTimeout
- JobCreationMode
- Key File Path
- Language Dialect
- Log Level
- Log Path
- Max File Size
- Max Number Files
- Minimum TLS
- SessionLocation
- Default temp table expiration time (ms)
- Trusted Certificates
- UniverseDomain
- Use Default_bqodbc_temp_tables Large Results Dataset
- Use SQL_WVARCHAR instead of SQL_VARCHAR
- Use System Trust Store

Activation Threshold for High-Throughput API

The minimum total pages (ratio of total rows to rows in the first page) required to activate reading through the High-Throughput API.

Important: The `HTAPI_MinActivationRatio` property has been deprecated and the `HTAPI_ActivationThreshold` property is now the primary alias. The primary alias has precedence and is used when both are present.

| Key Name | Default Value | Required |
|---------------------------|---------------|----------|
| HTAPI_ActivationThreshold | 3 | No |

Additional Projects

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

| Key Name | Default Value | Required |
|--------------------|---------------|----------|
| AdditionalProjects | None | No |

Allow Large Result Sets

This option specifies the connector's response to query results larger than the BigQuery maximum result sizes when using Legacy SQL.

- Enabled (1): The connector allows query results that are larger than the BigQuery maximum result sizes.
- Disabled (0): The connector returns an error when query results are larger than the BigQuery maximum result sizes.

Note:

- If this option is enabled, you must also specify a dataset for storing the temporary tables. For more information see [Dataset Name For Large Result Sets](#).
- For more information about BigQuery maximum result sizes, see "Quotas and Limits" in the Google BigQuery documentation: https://cloud.google.com/bigquery/quotas#query_jobs.

For additional information on SQL see [Language Dialect](#).

| Key Name | Default Value | Required |
|-------------------|---------------|----------|
| AllowLargeResults | Clear (0) | No |

AllowHtapiForLargeResults

This option specifies whether the connector uses the BigQuery High-Throughput API for large result sets. For detailed information about the High-Throughput API, see [High-Throughput API](#).

- Enabled (1): The connector uses the High-Throughput API for large result sets that exceed the activation threshold for High-Throughput API.
- Disabled (0): The connector does not use the High-Throughput API for large result sets.

Important:

- Pricing for the High-Throughput API is different than pricing for the standard REST API. For more information, see "BigQuery Storage API Pricing" in the Google BigQuery documentation: <https://cloud.google.com/bigquery/pricing#storage-api>.
- The `EnableHTAPI` property has been deprecated and the `AllowHtapiForLargeResults` property is now the primary alias. The primary alias has precedence and is used when both are present.

| Key Name | Default Value | Required |
|---------------------------|---------------|----------|
| AllowHtapiForLargeResults | Clear (0) | No |

Catalog (Project)

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

the DSN.

Simba ODBC Connector for Google BigQuery supports multiple catalogs, the equivalent of Google BigQuery projects.

For queries, tables in the projection list must be fully qualified, in the format of `catalog.schema.table`. If the catalog is not specified, the connector will assume the project specified by the **Catalog** connection option.

For catalog functions, to retrieve information from the desired catalog, the ODBC **SQLSetConnectAttr** method must be called with **SQL_ATTR_CURRENT_CATALOG** set to the desired catalog.

Note: To use the High-Throughput API, the specified project 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.

| Key Name | Default Value | Required |
|----------|---------------|----------|
| Catalog | None | Yes |

Dataset Name For Large Result Sets

The ID of the BigQuery dataset that you want to use to store temporary tables for large result sets. Only specify a value for this option if you want to enable support for large result sets.

- Note:**
- If the `UseDefaultLargeResultsDataset` property is enabled, it overrides this value.
 - If this dataset does not exist and the data store specifies the US region, the connector creates the dataset.
 - If you want to use large result sets with Legacy SQL, you must also enable the **Allow Large Result Sets** option. See [Allow Large Result Sets](#).

| Key Name | Default Value | Required |
|-----------------------|--|----------|
| LargeResultsDataSetId | None (but see Use Default _ bqodbc_temp_tables Large Results Dataset) | No |

Dataset

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. Additionally, it treats the default

dataset as part of the project that is being billed. For information about specifying the project to bill, see [Catalog \(Project\)](#).

| Key Name | Default Value | Required |
|----------------|---------------|----------|
| DefaultDataset | None | No |

Default String Column Length

The maximum number of characters that can be contained in STRING columns.

| Key Name | Default Value | Required |
|---------------------------|---------------|----------|
| DefaultStringColumnLength | 16384 | No |

Email

When configuring Service Authentication, set this option to the service account email ID.

When configuring User Authentication with a `.json` key file, set this option to your user account email ID.

| Key Name | Default Value | Required |
|----------|---------------|--|
| Email | None | Yes, if OAuth Mechanism is set to Service Authentication (OAuthMechanism=0) with a <code>.p12</code> key file. |

Enable GCD

This option specifies whether the connector supports a driver-level boolean flag, Enable GoogleCloudDedicated (`EnableGCD`), which allows enabling or disabling Google Cloud Dedicated (GCD) functionality.

- Enabled (1): The connector reads the `PrivateService` and `UniverseDomain`.
- Disabled (0): The connector reads only the `PrivateServiceUri`.

Note:

- The `EnableGCD` property can now be set at both the connection level (the connection string or connection properties) and the driver level (driver properties in the product INI file). When both are specified, the connection-level setting takes precedence.
- If `PrivateServiceConnectUri` is specified in the connection properties, the `PrivateService` and `UniverseDomain` properties will be ignored. For more information, see the [PrivateServiceConnectUri](#).

| Key Name | Default Value | Required |
|-----------|---------------|----------|
| EnableGCD | off | No |

When the `EnableGCD` property is set to `off`, it will ignore the `UniverseDomain` and `PrivateService` properties, and will use the default service endpoints for its operations.

| Service | URI Endpoint |
|----------------------------------|--|
| BigQuery | https://bigquery.googleapis.com |
| OAuth2 | https://oauth2.googleapis.com |
| Identity Access Management (IAM) | https://iamcredentials.googleapis.com |
| Secure Token Service (STS) | https://sts.googleapis.com |
| BigQuery Storage | https://bigquerystorage.googleapis.com:443 |

When the `EnableGCD` property is set to `True`, the `UniverseDomain` and `PrivateService` property settings help set up the service endpoints it needs to connect to.

For example:

| PrivateService | UniverseDomain | BigQuery Service URI |
|----------------|------------------|---|
| (empty) | (empty) | https://bigquery.googleapis.com |
| testservice | (empty) | https://bigquery-testservice.p.googleapis.com |
| (empty) | test-domain.goog | https://bigquery.test-domain.goog |
| testservice | test-domain.goog | https://bigquery-testservice.p.test-domain.goog |

These values are also used to set up the endpoint URLs for other services like OAuth2 and IAM. This ensures that when the connector makes a request to one of these services, it uses the correct Google Cloud Dedicated and Private Service Connect settings.

Enable High-Throughput API

This option specifies whether the connector uses the BigQuery High-Throughput API for large result sets. For detailed information about the High-Throughput API, see [High-Throughput API](#).

- Enabled (1): The connector uses the High-Throughput API for large result sets that exceed the activation threshold for High-Throughput API.
- Disabled (0): The connector does not use the High-Throughput API for large result sets.



Important:

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

| Key Name | Default Value | Required |
|-------------|---------------|----------|
| EnableHTAPI | Clear (0) | No |

Enable Session

This property 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.

- Enabled (1): The connector creates a session ID from the first executed query on that connection.
- Disabled (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.
- The `EnableSession` property is required when using transactions.

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

| Key Name | Default Value | Required |
|---------------|---------------|----------|
| EnableSession | Disabled (0) | No |

Enable Timestamp Picos

Simba ODBC Connector for Google BigQuery supports timestamp values with picoseconds precision.



Note: This feature is supported in both HTAPI and Rest API connection workflows.

| Key Name | Default Value | Required |
|----------------------|---------------|----------|
| EnableTimestampPicos | 0 | No |

HTAPI_SessionRetryTimeout

Specifies the duration, in seconds, for which create read session requests that return the gRPC status code UNAVAILABLE will be retried. Setting this connection property to 0 will disable the retry behavior

for UNAVAILABLE status codes returned by create read session requests.

| Key Name | Default Value | Required |
|---------------------------|---------------|----------|
| HTAPI_SessionRetryTimeout | 600 | No |

JobCreationMode

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



Note:

- If JobCreationMode=1, use JOB_CREATION_REQUIRED is required.
- If JobCreationMode=2, use low latency or JOB_CREATION_OPTIONAL is optional.

| Key Name | Default Value | Required |
|-----------------|---------------|----------|
| JobCreationMode | 2 | No |

Key File Path

When configuring Service Authentication, set this option to the full path to the .p12 or .json key file that is used to authenticate the service account email address.

When configuring User Authentication with a .json key file, set this option to the full path to the .json key file containing your OAuth 2.0 credentials. The file must define a JSON object of type `authorized_user` containing the refresh token, client ID, and client secret associated with your user account. For example, the .json key file must be written in the following format:

```
{
  "type": "authorized_user",
  "client_id": "[YourClientID]",
  "client_secret": "[YourClientSecret]",
  "refresh_token": "[YourRefreshToken]"
}
```

| Key Name | Default Value | Required |
|-------------|---------------|---|
| KeyFilePath | None | Yes, if OAuth Mechanism is set to Service Authentication (OAuthMechanism=0) and KeyFile is not set. |

Language Dialect

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

- Standard SQL (1): The connector uses standard SQL.
- Legacy SQL (0): The connector uses legacy SQL.

| Key Name | Default Value | Required |
|------------|------------------|----------|
| SQLDialect | Standard SQL (1) | No |

Log Level

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.
- When logging with connection strings and DSNs, this option only applies to per-connection logs.

Set the property to one of the following values:

- OFF (0): Disable all logging.
- FATAL (1): Logs severe error events that lead the connector to abort.
- ERROR (2): Logs error events that might allow the connector to continue running.
- WARNING (3): Logs events that might result in an error if action is not taken.
- INFO (4): Logs general information that describes the progress of the connector.
- DEBUG (5): Logs detailed information that is useful for debugging the connector.
- TRACE (6): Logs all connector activity.

When logging is enabled, the connector produces the following log files at the location you specify in the Log Path (`LogPath`) property:

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

| Key Name | Default Value | Required |
|----------|---------------|----------|
| LogLevel | OFF (0) | No |

Log Path

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

! **Important:** When logging with connection strings and DSNs, this option only applies to per-connection logs.

| Key Name | Default Value | Required |
|----------|---------------|-----------------------------|
| LogPath | None | Yes, if logging is enabled. |

Max File Size

The maximum size of each log file in bytes. After the maximum file size is reached, the connector creates a new file and continues logging.

If this property is set using the Windows UI, the entered value is converted from megabytes (MB) to bytes before being set.

! **Important:** When logging with connection strings and DSNs, this option only applies to per-connection logs.

| Key Name | Default Value | Required |
|-------------|---------------|----------|
| LogFileSize | 20971520 | No |

Max Number Files

The maximum number of log files to keep. After the maximum number of log files is reached, each time an additional file is created, the connector deletes the oldest log file.

! **Important:** When logging with connection strings and DSNs, this option only applies to per-connection logs.

| Key Name | Default Value | Required |
|--------------|---------------|----------|
| LogFileCount | 50 | No |

Minimum TLS

The minimum version of TLS/SSL that the connector allows the data store to use for encrypting connections. For example, if TLS 1.1 is specified, TLS 1.0 cannot be used to encrypt connections.

- TLS 1.0 (1.0): The connection must use at least TLS 1.0.
- TLS 1.1 (1.1): The connection must use at least TLS 1.1.
- TLS 1.2 (1.2): The connection must use at least TLS 1.2.

| Key Name | Default Value | Required |
|----------|---------------|----------|
| Min_TLS | TLS 1.3 (1.3) | No |

OAuth Mechanism

The OAuth 2.0 authentication mechanism used to authenticate the connector.

- Service Authentication (0): The connector authenticates as a service, through a Google service account.
- User Authentication (1): The connector authenticates as a user, through a Google user account.
- Application Default Credentials Authentication (3): The connector authenticates as a user, through a GOOGLE_APPLICATION_CREDENTIALS environment variable.
- External Account Authentication (4): The connector authenticates using workforce identity federation.

Note:

- When the connector is configured to use External Account Authentication (OAuthMechanism=4), the flow now supports the native SAI that is included in the external account configuration specification and the connection properties are considered in the following precedence:
 1. KeyFile
 2. KeyFilePath (or KeyFilePath_Enc if the key file is not set)
 3. BYOID_properties
- When the connector is configured to use Application Default Credentials (OAuthMechanism=3), the flow now supports impersonated_service_account keyfiles.
- SAI_Email, SAI_Lifetime, and SAI_Scopes properties are applicable when the OAuthMechanism= 0, or 1, or 2.

| Key Name | Default Value | Required |
|----------------|-------------------------|----------|
| OAuthMechanism | User Authentication (1) | No |

Path to CMEK

The resource ID 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:

- The `KMSKeyName` property is used only when the connector is configured to use the destination tables.
- 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.

| Key Name | Default Value | Required |
|-------------------------|---|----------|
| <code>KMSKeyName</code> | None. The connector uses the default encryption key from Google. | No |

PrivateService

This options allows users to specify the name of their private service, and the connector will construct the URLs from them.

| Key Name | Default Value | Required |
|-----------------------------|---------------|----------|
| <code>PrivateService</code> | None | No |

Proxy Host

The host name or IP address of a proxy server that you want to connect through.

| Key Name | Default Value | Required |
|------------------------|---------------|--|
| <code>ProxyHost</code> | None | Yes, if connecting through a proxy server. |

Proxy Password

The password that you use to access the proxy server.

| Key Name | Default Value | Required |
|----------|---------------|--|
| ProxyPwd | None | Yes, if connecting to a proxy server that requires authentication. |

Proxy Port

The number of the port that the proxy server uses to listen for client connections.

| Key Name | Default Value | Required |
|-----------|---------------|--|
| ProxyPort | None | Yes, if connecting through a proxy server. |

Proxy Username

The user name that you use to access the proxy server.

| Key Name | Default Value | Required |
|----------|---------------|--|
| ProxyUid | None | Yes, if connecting to a proxy server that requires authentication. |

Query Properties

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.

This option also supports the `time_zone` connection property in the following form:

`time_zone=America/Los_Angeles`

For detailed information, see "Supported Time Zone Values" in the Google BigQuery documentation: https://cloud.google.com/dataprep/docs/html/Supported-Time-Zone-Values_66194188.



Important:

The `QueryProperties` property corresponds to the BigQuery "connectionProperties" field present in both the `QueryRequest` and `JobConfigurationQuery` objects. Each item in the comma-separated list is translated to a `ConnectionProperty` object. For detailed information, see the Google BigQuery documentation for these objects:

- `ConnectionProperty`:
<https://cloud.google.com/bigquery/docs/reference/rest/v2/ConnectionProperty>
- `QueryRequest`:
<https://cloud.google.com/bigquery/docs/reference/rest/v2/jobs/query#QueryRequest>
- `JobConfigurationQuery`:
<https://cloud.google.com/bigquery/docs/reference/rest/v2/Job#JobConfigurationQuery>

| Key Name | Default Value | Required |
|------------------------------|---------------|----------------------------------|
| <code>QueryProperties</code> | None | Yes, if using a default project. |

Ratio of Results to Rows Per Block

The minimum total pages required to activate reading through the High-Throughput API.



Important:

This property has been deprecated and the `HTAPI_ActivationThreshold` property is now the primary alias. The primary alias has precedence and is used when both are present.

| Key Name | Default Value | Required |
|---------------------------------------|---------------|----------|
| <code>HTAPI_MinActivationRatio</code> | 3 | No |

Refresh Token

The refresh token that you obtain from Google for authorizing access to BigQuery.

When you configure a DSN with the Windows connector, the refresh token is generated automatically after you provide the confirmation code.

When you configure a DSN with the Linux or macOS versions of the connector, you can use the Google OAuth 2.0 Playground to generate the token. For more information, see [Using a Google User Account](#).

| Key Name | Default Value | Required |
|---------------------------|---------------|--|
| <code>RefreshToken</code> | None | Yes, if authenticating through a user account. |

Request Google Drive Scope Access

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.

- Enabled (1): The connector requests access to Google Drive.
- Disabled (0): The connector does not request access to Google Drive.

| Key Name | Default Value | Required |
|-------------------------|---------------|----------|
| RequestGoogleDriveScope | Clear (0) | No |

Rows Fetched Per Block

The maximum number of rows that the connector can fetch for each data request.

| Key Name | Default Value | Required |
|---------------------|---------------|----------|
| RowsFetchedPerBlock | 100000 | No |

SessionLocation

This option specifies the region or location in which the queries executed by the connection will be executed. For more information on locations in BigQuery, consult the Google BigQuery documentation: [BigQuery locations](#).

| Key Name | Default Value | Required |
|-----------------|---------------|----------|
| SessionLocation | None | No |

Default temp table expiration time (ms)

The default table expiration time, in milliseconds, for large results data sets created by the connector.

| Key Name | Default Value | Required |
|-------------------------------------|---------------|----------|
| LargeResultsTempTableExpirationTime | 3600000 | No |

Trusted Certificates

The full path of the .pem file containing trusted CA certificates, for verifying the server.

If this option is not set, then the connector defaults to using the trusted CA certificates .pem file installed by the connector. To use the trusted CA certificates in the .pem file, set the UseSystemTrustStore property to 0 or clear the Use System Trust Store check box in the SSL Options dialog.

| Key Name | Default Value | Required |
|--------------|--|----------|
| TrustedCerts | The cacerts .pem file in the \lib subfolder within the connector's installation directory. The exact file path varies depending on the version of the connector that is installed. For example, the path for the Windows connector is different from the path for the macOS connector. | No |

UniverseDomain

This options allows the user to specify the universe domain that applies to all services.

| Key Name | Default Value | Required |
|----------------|---------------|----------|
| UniverseDomain | None | No |

Use Default _bqodbc_temp_tables Large Results Dataset

This option specifies whether the connector uses a default dataset to store temporary tables for large result sets if no dataset is specified and the data store specifies the US region.

- Enabled (1): The connector uses a default large result dataset of _bqodbc_temp_tables.
- Disabled (0): The connector does not use a default large result dataset. If no dataset is specified, large result support is not enabled.



Note:

- If this value is specified, it overrides the LargeResultsDatasetId property.
- If this option is selected, the data store specifies the US region, and the _bqodbc_temp_tables dataset does not exist, the connector creates this dataset.
- If you want to use large result sets with Legacy SQL, you must also enable the **Allow Large Result Sets** option. See [Allow Large Result Sets](#).

| Key Name | Default Value | Required |
|-------------------------------|---------------|----------|
| UseDefaultLargeResultsDataset | Clear (0) | No |

Use SQL_WVARCHAR instead of SQL_VARCHAR

This option specifies how data types are mapped to SQL.

- Enabled (1): The connector returns data as SQL_WVARCHAR data instead of SQL_VARCHAR data.
- Disabled (0): The connector returns data as SQL_VARCHAR data.



Note:

This option applies only to result set columns that the connector would normally return as SQL_VARCHAR columns. It does not convert all columns into SQL_WVARCHAR.

| Key Name | Default Value | Required |
|-------------|---------------|----------|
| UseWVarChar | Clear (0) | No |

Use System Trust Store

This option specifies whether to use a CA certificate from the system trust store, or from a specified .pem file.

- Enabled (1): The connector verifies the connection using a certificate in the system trust store.
- Disabled (0): The connector verifies the connection using a specified .pem file. For information about specifying a .pem file, see [Trusted Certificates](#).



Note: This option is only available in Windows.

| Key Name | Default Value | Required |
|---------------------|---------------|----------|
| UseSystemTrustStore | | No |

Configuration Options Having Only Key Names

The following configuration options do not appear in the Windows user interface for the Simba Google BigQuery ODBC Data Connector. They are accessible only when you use a connection string or configure a connection in macOS or Linux.

- [BYOID_AudienceUrl](#)
- [BYOID_CredentialSource](#)
- [BYOID_PoolUserProject](#)
- [BYOID_SubjectTokenType](#)
- [BYOID_TokenUrl](#)
- [ClientId](#)
- [ClientSecret](#)

- [FilterTablesOnDefaultDataset](#)
- [IgnoreTransactions](#)
- [Driver](#)
- [FilterTablesOnDefaultDataset](#)
- [KeyFile](#)
- [PrivateServiceConnectUris](#)
- [Service Account Impersonation Email \(SAI_Email\)](#)
- [Service Account Impersonation Auth Scopes \(SAI_Scopes\)](#)
- [Service Account Impersonation Token Lifetime \(SAI_Lifetime\)](#)
- [Timeout](#)
- [UnsupportedHTAPIFallback](#)
- [UseQueryCache](#)

The following connector-wide property must be configured as a Windows Registry key value, or in the `simba.googlebigqueryodbc.ini` file for macOS or Linux.

- [MaxThreads](#)

BYOID_AudienceUrl

When configuring External Account Authentication, set this option to the full path of the full audience URL. This option corresponds to the `audience` property in the configuration file. For example, the URI must be written in the following format:

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

| Key Name | Default Value | Required |
|-------------------|---------------|--|
| BYOID_AudienceUrl | None | Yes, if OAuth Mechanism is set to External Account Authentication (OAuthMechanism=4) and other BYOID_ properties, or KeyFile or KeyFilePath are not set. |

BYOID_CredentialSource

When configuring External Account Authentication, set this option to the full path of the `.json` configuration file that is used to specify the source from which to retrieve a subject token to exchange for an access token. This option corresponds to the `credential_source` property in the configuration file.



Note:

- The key file must be in raw `.json` format.
- Only "file" ODBC credential sources are supported (files from which the subject token can be read).
- 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-lived-credentials#generate_a_configuration_file.

| Key Name | Default Value | Required |
|------------------------|---------------|--|
| BYOID_CredentialSource | None | Yes, if OAuth Mechanism is set to External Account Authentication (OAuthMechanism=4) and other BYOID_ properties, or KeyFile or KeyFilePath are not set. |

BYOID_PoolUserProject

When configuring External Account Authentication, set this option to the full path of the project ID or name to use for external account authentication. This option corresponds to the `workforce_pool_user_project` key in the configuration file.

| Key Name | Default Value | Required |
|-----------------------|---------------|--|
| BYOID_PoolUserProject | None | Yes, if OAuth Mechanism is set to External Account Authentication (OAuthMechanism=4) and other BYOID_ properties, or KeyFile or KeyFilePath are not set. |

BYOID_SubjectTokenType

When configuring External Account Authentication, set this option to the full path of the subject token type. This option corresponds to the `subject_token_type` key in the configuration file.

| Key Name | Default Value | Required |
|------------------------|---|--|
| BYOID_SubjectTokenType | urn:ietf:params:oauth:token-type:id_token | Yes, if OAuth Mechanism is set to External Account Authentication (OAuthMechanism=4) and other BYOID_ properties, or KeyFile or KeyFilePath are not set. |

BYOID_TokenUrl

When configuring External Account Authentication, set this option to the full path of the STS token URL to query when retrieving a short-lived access token. This option corresponds to the `token_url` key in the configuration file.

| Key Name | Default Value | Required |
|----------------|-------------------------------------|--|
| BYOID_TokenUrl | https://sts.googleapis.com/v1/token | Yes, if OAuth Mechanism is set to External Account Authentication (OAuthMechanism=4) and other BYOID_ properties, or KeyFile or KeyFilePath are not set. |

ClientId

The default value depends on the package that you received. It is recommended to use your own Client ID and Client Secret values. When using User Account, the value of Client ID should be the value that refresh token was created with.

| Key Name | Default Value | Required |
|----------|------------------------|----------|
| ClientId | Branded default Client | No |

ClientSecret

The default value depends on the package that you received. It is recommended to use your own Client ID and Client Secret values. When using User Account, the value of Client Secret should be the value that refresh token was created with.

| Key Name | Default Value | Required |
|--------------|---------------|----------|
| ClientSecret | *** | No |

Driver

In Windows, the name of the installed connector for (Simba Google BigQuery ODBC Data Connector).

On other platforms, the name of the installed connector as specified in `odbcinst.ini`, or the absolute path of the connector shared object file.

| Key Name | Default Value | Required |
|----------|---|----------|
| Driver | Simba BigQuery ODBC Driver when installed in Windows, or the absolute path of the connector shared object file when installed on a non-Windows machine. | Yes |

FilterTablesOnDefaultDataset

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

- **FALSE:** The connector returns all tables in the `SQLTables` call and all columns in the `SQLColumns` call.
- **TRUE:** 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 details, see [Dataset](#).

When this option is set to `TRUE`, the connector behaves as described below for the functions `SQLTables` and `SQLColumns`.

For the function `SQLTables`:

| Catalog | Schema | Table | Table Type | Returned List |
|---------|----------|-----------|------------|--|
| NULL | NULL | NULL or % | NULL or % | All tables that belong to the default dataset under the default catalog |
| % | NULL | NULL or % | NULL or % | All tables that belong to the default dataset under all catalogs |
| NULL | % | NULL or % | NULL or % | All tables that belong to all schemas under the default catalog |
| % | % | NULL or % | NULL or % | All tables that belong to all schemas under all catalogs |
| NULL | <schema> | NULL or % | NULL or % | All tables that belong to the specified schema under the default catalog |

| Catalog | Schema | Table | Table Type | Returned List |
|-----------|----------|-----------|------------|--|
| <catalog> | <schema> | NULL or % | NULL or % | All tables that belong to the specified schema under the specified catalog |

For the function SQLColumns:

| Catalog | Schema | Table | Column | Returned List |
|-----------|----------|-------|--------|--|
| NULL | NULL | NULL | NULL | All columns of all tables that belong to the default dataset under the default catalog |
| <catalog> | NULL | NULL | NULL | All columns of all tables that belong to the default dataset under the specified catalog |
| NULL | % | NULL | NULL | All columns of all tables that belong to all datasets under the default catalog |
| <catalog> | % | NULL | NULL | All columns of all tables that belong to all datasets under the specified catalog |
| NULL | <schema> | NULL | NULL | All columns of all tables that belong to the specified dataset under the default catalog |
| <catalog> | <schema> | NULL | NULL | All columns of all tables that belong to the specified dataset under the specified catalog |

| Key Name | Default Value | Required |
|------------------------------|---------------|----------|
| FilterTablesOnDefaultDataset | FALSE | No |

IgnoreTransactions

This option determines whether the connector ignores attempts to perform transactions.

- 0: Attempts to perform transactions produce a user alert.
- 1: The connector ignores attempts to perform transactions. No alerts are generated for these calls.

| Key Name | Default Value | Required |
|--------------------|---------------|----------|
| IgnoreTransactions | 0 | No |

KeyFile

When configuring Service Authentication, set this option to the full path to the plain text JSON object or .p12 or .json key file that is used to authenticate the service account email address.

| Key Name | Default Value | Required |
|----------|---------------|---|
| KeyFile | None | Yes, if OAuth Mechanism is set to Service Authentication (OAuthMechanism=0) and KeyFilePath is not set. |

MaxThreads

Defines the maximum number of threads that the connector can initialize and run concurrently in a threadpool.

To configure this property as a connector-wide setting for a non-Windows connector, you must use the `simba.googlebigqueryodbc.ini` file.

| Key Name | Default Value | Required |
|------------|---------------|----------|
| MaxThreads | 8 | No |

PrivateServiceConnectUris

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

- SERVER_NAME:** The baseline Google APIs service URI, used to perform authentication retries for proxy issues in Windows only. The default value is: `https://www.googleapis.com`.

For example:

`SERVER_NAME=https://<myprivateserver>.p.googleapis.com`

- ACCOUNTS:** The baseline accounts service URI, used only for interactive authentication (hidden `OAuthMechanism` 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`

- 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
```

- **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.
- When the connector is configured to use Service Authentication (`OAuthMechanism=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. `PrivateServiceConnectUris=..., OAuth2=<YOUR_OAuth2_URL>, ...`
 3. Default
- When the connector is configured to use External Account Authentication (`OAuthMechanism=4`), the connector prioritizes the `STS` URI from either the configuration file specified in the `KeyFile` or `KeyFilePath` property, or, from the `BYOID_TokenUri` property. In order, the precedence is:
 1. `KeyFile/KeyFilePath{ _Enc }`
 2. `BYOID_TokenUri`
 3. `PrivateServiceConnectUris=..., STS=<YOUR_STS_URL>, ...`
 4. 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>.
- The `IMPERSONATION` value is now supported by the `PrivateServiceConnectUris` connection property.

| Key Name | Default Value | Required |
|--------------------------|---------------|----------|
| PrivateServiceConnectUri | None | No |

Service Account Impersonation Email (SAI_Email)

This option specifies an email address for the service account to impersonate.

| Key Name | Default Value | Required |
|----------------------------------|---------------|----------|
| ServiceAccountImpersonationEmail | None | No |

Service Account Impersonation Auth Scopes (SAI_Scopes)

This option specifies the comma-separated list of scopes.

| Key Name | Default Value | Required |
|-----------------------------------|---------------|----------|
| ServiceAccountImpersonationScopes | None | No |

Service Account Impersonation Token Lifetime (SAI_Lifetime)

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

| Key Name | Default Value | Required |
|-------------------------------------|---------------|----------|
| ServiceAccountImpersonationLifetime | 3600 | 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 an integer. A value of 0 indicates no timeout.

| Key Name | Default Value | Required |
|----------|---------------|----------|
| Timeout | 300 | 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 High-Throughput API, see [High-Throughput API](#).

| Key Name | Default Value | Required |
|--------------------------|---------------|----------|
| UnsupportedHTAPIFallback | 1 | No |

UseQueryCache

This option determines whether the connector uses the query cache when retrieving results.

- 1: The connector uses cached query results, if they are available.
- 0: The connector does not use the query cache.

For detailed information about cached query results, see "Using Cached Query Results" in the Google Cloud Platform documentation: <https://cloud.google.com/bigquery/docs/cached-results>.

| Key Name | Default Value | Required |
|---------------|---------------|----------|
| UseQueryCache | 1 | No |

Third-Party Trademarks

Debian is a trademark or registered trademark of Software in the Public Interest, Inc. or its subsidiaries in Canada, United States and/or other countries.

Linux is the registered trademark of Linus Torvalds in Canada, United States and/or other countries.

Mac, macOS, Mac OS, and OS X are trademarks or registered trademarks of Apple, Inc. or its subsidiaries in Canada, United States and/or other countries.

Microsoft, MSDN, Windows, Windows Server, Windows Vista, and the Windows start button are trademarks or registered trademarks of Microsoft Corporation or its subsidiaries in Canada, United States and/or other countries.

Red Hat, Red Hat Enterprise Linux, and CentOS are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in Canada, United States and/or other countries.

SUSE is a trademark or registered trademark of SUSE LLC or its subsidiaries in Canada, United States and/or other countries.

Ubuntu is a trademark or registered trademark of Canonical Ltd. or its subsidiaries in Canada, United States and/or other countries.

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.