



Informatica® PowerExchange for Google
Cloud Spanner

10.5.9

User Guide for PowerCenter

© Copyright Informatica LLC 2018, 2025

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC.

Informatica, the Informatica logo, PowerCenter, and PowerExchange are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners.

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation is subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License.

Portions of this software and/or documentation are subject to copyright held by third parties. Required third party notices are included with the product.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, report them to us at infa_documentation@informatica.com.

Informatica products are warranted according to the terms and conditions of the agreements under which they are provided. INFORMATICA PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2025-09-03

Table of Contents

Preface	5
Informatica Resources.	5
Informatica Network.	5
Informatica Knowledge Base.	5
Informatica Documentation.	5
Informatica Product Availability Matrixes.	6
Informatica Velocity.	6
Informatica Marketplace.	6
Informatica Global Customer Support.	6
 Chapter 1: Introduction to PowerExchange for Google Cloud Spanner.....	 7
PowerExchange for Google Cloud Spanner Overview.	7
Introduction to Google Cloud Spanner.	7
Administration of Google Cloud Spanner.	8
 Chapter 2: PowerExchange for Google Cloud Spanner Configuration.....	 10
PowerExchange for Google Cloud Spanner Configuration Overview.	10
Registering the PowerExchange for Google Cloud Spanner Plug-in.	10
Registering the PowerExchange for Google Cloud Spanner Plug-in from Informatica Administrator.	11
Registering the PowerExchange for Google Cloud Spanner Plug-in from the Command Line Program.	11
 Chapter 3: Google Cloud Spanner Sources and Targets.....	 12
Google Cloud Spanner Sources and Targets Overview.	12
Import Google Cloud Spanner Source and Target Definitions.	12
 Chapter 4: Google Cloud Spanner Mappings.....	 16
Google Cloud Spanner Mappings Overview.	16
Source Filter.	16
Configuring a Source Filter.	17
 Chapter 5: Google Cloud Spanner Sessions.....	 19
Google Cloud Spanner Sessions Overview.	19
Google Cloud Spanner Connections.	19
PowerExchange for Google Cloud Spanner Connections.	20
Configuring a Google Cloud Spanner Connection.	20
Filter Override.	21
Configure Google Cloud Spanner Source Session Properties.	22
Configure Google Cloud Spanner Target Session Properties.	22

Rules and Guidelines for Google Cloud Spanner Targets.	23
Partitioning.	24
Applying Filter Override Conditions for Partitions.	24
Configure the Buffer Memory.	26
Parameterization.	27
Configure the Java Heap Memory.	27
Appendix A: Google Cloud Spanner Data Type Reference.	28
Data Type Reference Overview.	28
Google Cloud Spanner and Transformation Data Types.	28
Index.	30

Preface

Use the *Informatica® PowerExchange® for Google Cloud Spanner User Guide for PowerCenter* to learn how to read from and write to Google Cloud Spanner by using PowerCenter Client. Learn to create a Google Cloud Spanner connection, develop mappings, and run sessions in an Informatica domain.

Informatica Resources

Informatica Network

Informatica Network hosts Informatica Global Customer Support, the Informatica Knowledge Base, and other product resources. To access Informatica Network, visit <https://network.informatica.com>.

As a member, you can:

- Access all of your Informatica resources in one place.
- Search the Knowledge Base for product resources, including documentation, FAQs, and best practices.
- View product availability information.
- Review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to search Informatica Network for product resources such as documentation, how-to articles, best practices, and PAMs.

To access the Knowledge Base, visit <https://kb.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Documentation

To get the latest documentation for your product, browse the Informatica Knowledge Base at https://kb.informatica.com/_layouts/ProductDocumentation/Page/ProductDocumentSearch.aspx.

If you have questions, comments, or ideas about this documentation, contact the Informatica Documentation team through email at infa_documentation@informatica.com.

Informatica Product Availability Matrixes

Product Availability Matrixes (PAMs) indicate the versions of operating systems, databases, and other types of data sources and targets that a product release supports. If you are an Informatica Network member, you can access PAMs at

<https://network.informatica.com/community/informatica-network/product-availability-matrices>.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services. Developed from the real-world experience of hundreds of data management projects, Informatica Velocity represents the collective knowledge of our consultants who have worked with organizations from around the world to plan, develop, deploy, and maintain successful data management solutions.

If you are an Informatica Network member, you can access Informatica Velocity resources at <http://velocity.informatica.com>.

If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that augment, extend, or enhance your Informatica implementations. By leveraging any of the hundreds of solutions from Informatica developers and partners, you can improve your productivity and speed up time to implementation on your projects. You can access Informatica Marketplace at <https://marketplace.informatica.com>.

Informatica Global Customer Support

You can contact a Global Support Center by telephone or through Online Support on Informatica Network.

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:

<http://www.informatica.com/us/services-and-training/support-services/global-support-centers>.

If you are an Informatica Network member, you can use Online Support at <http://network.informatica.com>.

CHAPTER 1

Introduction to PowerExchange for Google Cloud Spanner

This chapter includes the following topics:

- [PowerExchange for Google Cloud Spanner Overview, 7](#)
- [Introduction to Google Cloud Spanner, 7](#)
- [Administration of Google Cloud Spanner, 8](#)

PowerExchange for Google Cloud Spanner Overview

You can use PowerExchange for Google Cloud Spanner for connectivity between PowerCenter and Google Cloud Spanner.

You can use Google Cloud Spanner objects as sources and targets in mappings. When you use Google Cloud Spanner objects in mappings, you must configure properties specific to Google Cloud Spanner.

When you run a session, the PowerCenter Integration Service uses the JAVA client libraries of the Google Cloud Spanner APIs to integrate with Google Cloud Spanner.

Example

An enterprise application uses a relational database to store the customer details on a global scale such as customerID, customerName, accountNum, phoneNum, address, and age. You can move the data from the relational database to Google Cloud Spanner to achieve scalability and high availability of the customer details.

Google Cloud Spanner is a scalable, multi-version, enterprise-grade, globally distributed, and synchronously-replicated database that provides high availability.

Use PowerExchange for Google Cloud Spanner to write data to a Google Cloud Spanner target.

Introduction to Google Cloud Spanner

Google Cloud Spanner is a fully managed relational database service that the Google Cloud Platform provides. Google Cloud Spanner is ideal for relational, structured, and semi-structured data that requires high availability, strong consistency, and transactional read and write operations. You can use Google Cloud Spanner for general purpose transactions using SQL (ANSI 2011 with extensions).

Administration of Google Cloud Spanner

Before you use PowerExchange for Google Cloud Spanner, you must complete the following prerequisite tasks:

- Ensure that you have a Google service account to access Google Cloud Spanner.
- Ensure that you have enabled billing for the project.
- Ensure that you have enabled the Cloud Spanner API for your service account. PowerExchange for Google Cloud Spanner uses the Google Cloud Spanner API to integrate with Google Cloud Spanner.
- Ensure that you have the `client_email`, `project_id`, and `private_key` values for the service account and instance id of the Google Cloud Spanner instance. You will need to enter these details when you create a Google Cloud Spanner connection in PowerCenter.
- Ensure that you have selected the regional or multi-region configuration and specified the number of nodes that you want to use for the Google Cloud Spanner instance.

Note: When you increase the number of nodes, the cost charged for a billing account increases. Before you specify the number of nodes, see the Google Cloud Spanner documentation to estimate the cost of using Google Cloud Spanner. To evaluate Google's pricing details for Google Cloud Spanner, click the following URL:

<https://cloud.google.com/spanner/pricing>

- Ensure that you have the database name and table name in the Google Cloud Spanner instance. You will need to enter these details when you run a session in PowerCenter. For more information about creating a Google Cloud Spanner instance, database, and table, click the following URL:

<https://cloud.google.com/spanner/docs/quickstart-console>

- When you read data from or write data to a Google Cloud Spanner table, you must have the following permissions:

- `spanner.databases.beginOrRollbackReadWriteTransaction`
- `spanner.databases.beginPartitionedDmlTransaction`
- `spanner.databases.beginReadOnlyTransaction`
- `spanner.databases.getDdl`
- `spanner.databases.read`
- `spanner.databases.select`
- `spanner.databases.updateDdl`
- `spanner.databases.write`
- `spanner.sessions.create`
- `spanner.sessions.delete`
- `spanner.sessions.get`
- `resourcemanager.projects.get`
- `spanner.databases.list`
- `spanner.instances.get`
- `spanner.instances.list`

- When you only read data from a Google Cloud Spanner table, you must have the following permissions:
 - `spanner.databases.beginReadOnlyTransaction`

- `spanner.databases.getDdl`
- `spanner.databases.read`
- `spanner.databases.select`
- `spanner.sessions.create`
- `spanner.sessions.delete`
- `spanner.sessions.get`
- `resourcemanager.projects.get`
- `spanner.databases.list`
- `spanner.instances.get`
- `spanner.instances.list`

CHAPTER 2

PowerExchange for Google Cloud Spanner Configuration

This chapter includes the following topics:

- [PowerExchange for Google Cloud Spanner Configuration Overview, 10](#)
- [Registering the PowerExchange for Google Cloud Spanner Plug-in, 10](#)

PowerExchange for Google Cloud Spanner Configuration Overview

PowerExchange for Google Cloud Spanner installs with Informatica services.

If you upgrade from a previous version, you must register the PowerExchange for Google Cloud Spanner plug-in with the PowerCenter repository.

Registering the PowerExchange for Google Cloud Spanner Plug-in

To register the plug-in, the repository must be running in exclusive mode. Use the Informatica Administrator or the `pmrep RegisterPlugin` command line program to register the plug-in. If you do not have the correct privileges to register the plug-in, contact the user who manages the PowerCenter Repository Service.

The plug-in file is an `.xml` file that defines the functionality of the adapter. When you install the server component, the installer copies the plug-in file to the following directory:

```
<Informatica installation directory>/server/bin/Plugin
```

The name of the plug-in file for PowerExchange for Google Cloud Spanner is `spannerGoogle_Plugin.xml`.

Registering the PowerExchange for Google Cloud Spanner Plug-in from Informatica Administrator

Register a repository plug-in to add its functionality to the repository.

1. Run the PowerCenter Repository Service in exclusive mode.
2. In the **Navigator**, select the PowerCenter Repository Service to which you want to add the plug-in.
3. In the **Contents** panel, click the **Plug-ins** view.
4. In the **Actions** menu of the **Domain** tab, select **Register Plug-in**.
5. On the **Register Plug-in** page, click the **Browse** button to locate the plug-in file.
6. Enter your user name and password.
7. Click **OK**.

The PowerCenter Repository Service registers the plug-in with the repository. The results of the registration operation appear in the activity log.

8. Run the PowerCenter Repository Service in normal mode.

Registering the PowerExchange for Google Cloud Spanner Plug-in from the Command Line Program

You can use the `pmrep RegisterPlugin` command to register the plug-in from the command line program.

1. Run the PowerCenter Repository Service in exclusive mode.
2. Run the `pmrep Connect` command to connect to the Repository Service using a user account with the Administrator Repository privilege.

The `RegisterPlugin` command uses the following syntax:

```
pmrep connect -r <repository name> -d <domain_name> -n <domain user name> -x  
<domain_password>
```

3. Find `spannerPlugin.xml` in the following directory:
`<Informatica installation directory>\server\bin\Plugin`
4. Run the `pmrep RegisterPlugin` command to update the repository.

The `RegisterPlugin` command uses the following syntax:

```
pmrep registerplugin -i <Informatica installation directory>\server\bin\Plugin  
\spannerGoogle_Plugin.xml -e -N
```

CHAPTER 3

Google Cloud Spanner Sources and Targets

This chapter includes the following topics:

- [Google Cloud Spanner Sources and Targets Overview, 12](#)
- [Import Google Cloud Spanner Source and Target Definitions, 12](#)

Google Cloud Spanner Sources and Targets Overview

You can create a mapping with a Google Cloud Spanner source to extract data from Google Cloud Spanner. You can create a mapping with any source and a Google Cloud Spanner target to load data to Google Cloud Spanner.

When the PowerCenter Integration Service extracts data from the source or loads data to the target, it converts the data based on the data types associated with the source or the target.

Import Google Cloud Spanner Source and Target Definitions

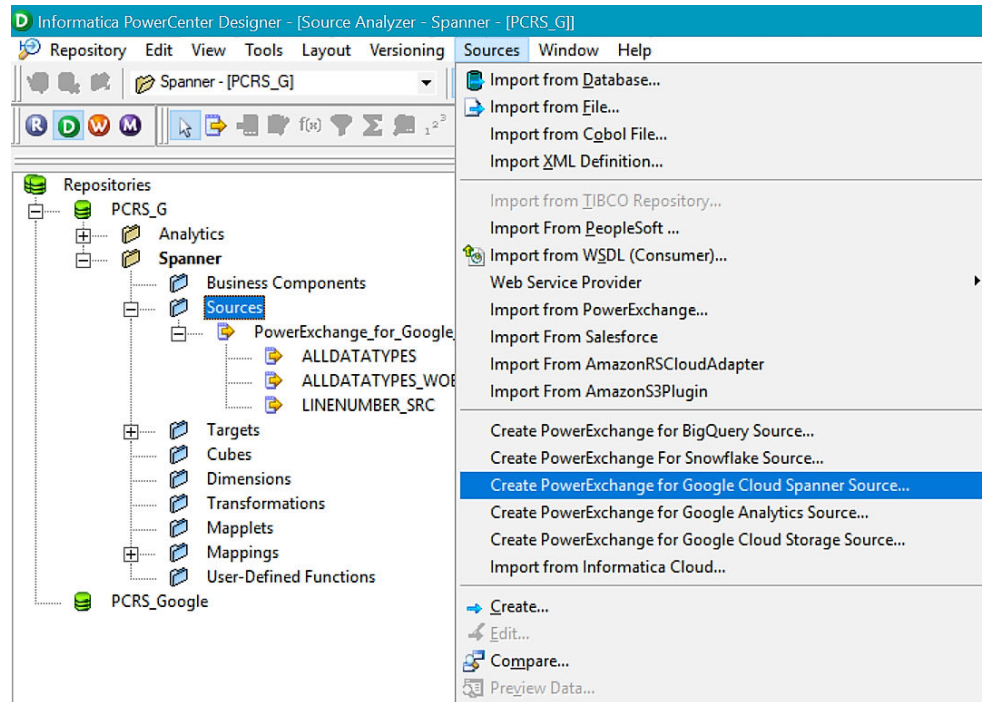
Use the **Create PowerExchange for Google Cloud Spanner Source** or **Create PowerExchange for Google Cloud Spanner Target** wizard to import Google Cloud Spanner source and target definitions into the PowerCenter repository.

You must import Google Cloud Spanner source and target objects before you create a mapping.

1. Start PowerCenter Designer, and connect to a PowerCenter repository.

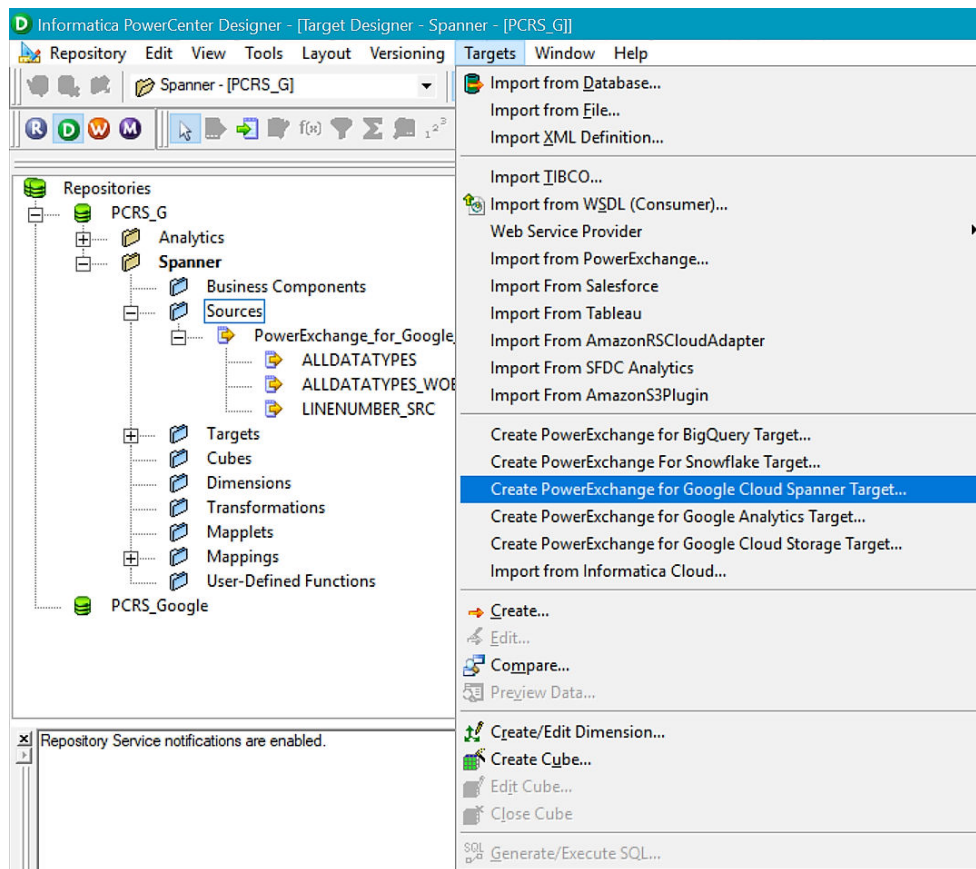
2. Select **Source Analyzer** or **Target Designer**.

- In the Source Analyzer, click **Sources > Create PowerExchange for Google Cloud Spanner Source** as shown in the following image:



The **Select Connection** for wizard appears.

- In the Target Analyzer, click **Targets > Create PowerExchange for Google Cloud Spanner Target** as shown in the following image:



The **Select Connection** for wizard appears.

3. Click **Import from a new connection**.
The **Connection details** dialog box appears.
4. Configure the following connection parameters:

Connection Parameter	Description
Service Account ID	Specifies the client_email value present in the JSON file that you download after you create a service account.
Service Account Key	Specifies the private_key value present in the JSON file that you download after you create a service account.
Project ID	Specifies the project_id value present in the JSON file that you download after you create a service account. If you have created multiple projects with the same service account, enter the ID of the project that contains the dataset that you want to connect to.
Instance ID	Name of the instance that you created in Google Cloud Spanner.

5. Click **Test** to test the connection.

6. Click **Save Connection** or **Finish** to add the connection.

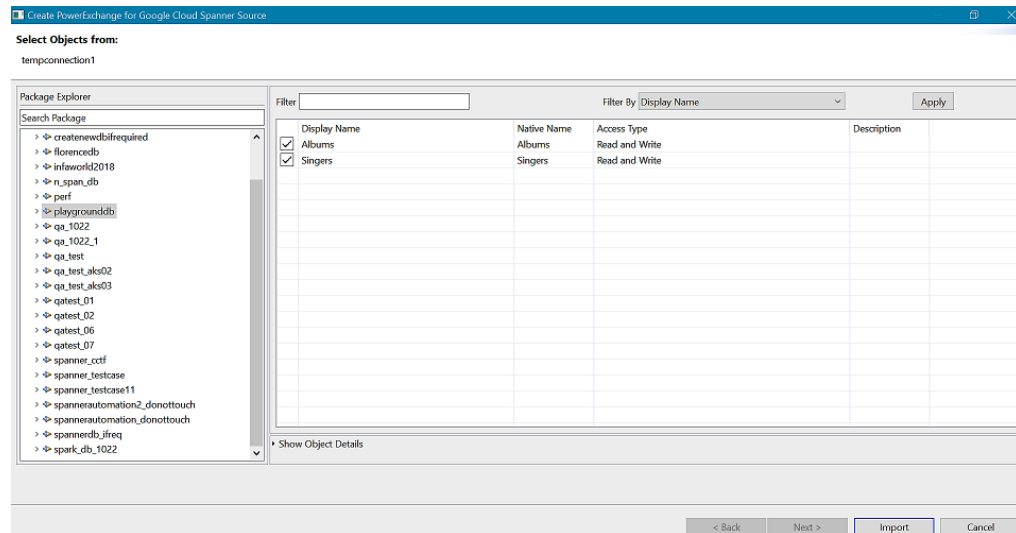
Note: The connection you create is temporary and you cannot edit the connection after you create it.

7. Click **Next**.

The **Select Objects from** tab appears.

8. Select the database in **Package Explorer**.

A list of table appears as shown in the following image:



9. Select the table that you want to import, and then click **Import**. You can import multiple tables from a Google Cloud Spanner database.

To view the table metadata, select the table, and double-click the table name.

CHAPTER 4

Google Cloud Spanner Mappings

This chapter includes the following topics:

- [Google Cloud Spanner Mappings Overview, 16](#)
- [Source Filter, 16](#)
- [Configuring a Source Filter, 17](#)

Google Cloud Spanner Mappings Overview

After you import a Google Cloud Spanner source or target definition into the PowerCenter repository, you can create a mapping to extract data from a Google Cloud Spanner source or load data to a Google Cloud Spanner target.

You can extract data from one or more Google Cloud Spanner sources, and load data to one or more Google Cloud Spanner targets.

You can enter a filter condition to reduce the number of source rows the PowerCenter Integration Service returns from Google Cloud Spanner sources. You can enter a single filter condition or a series of conditions.

Note: You cannot preview data of a Google Cloud Spanner source or target.

Source Filter

You can enter a filter condition to reduce the number of source rows the PowerCenter Integration Service returns from Google Cloud Spanner sources. You can enter a single filter condition or a series of conditions.

Use the source filter in the Application Source Qualifier to retrieve rows from an entity that meet a condition.

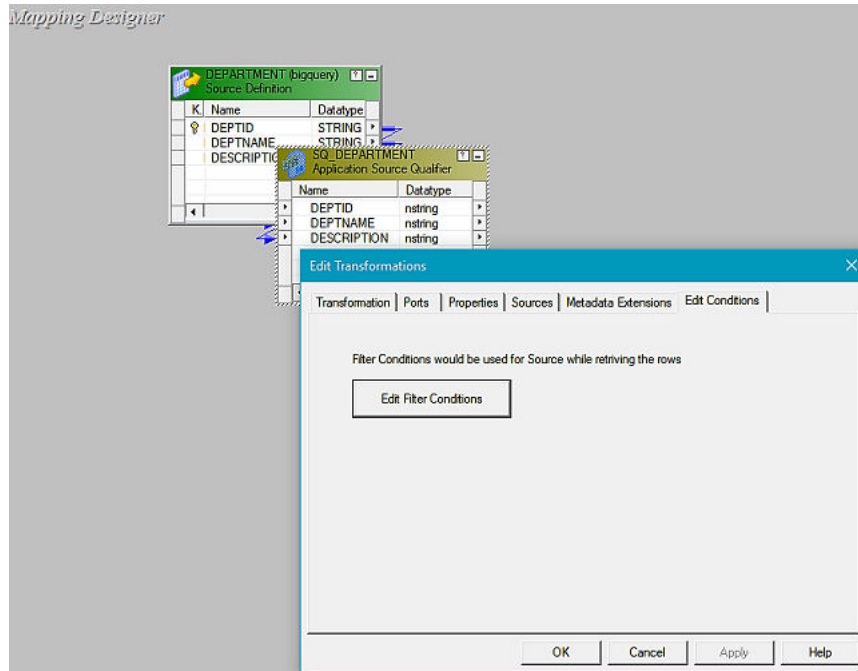
You can provide a source filter to improve the performance when you read data from Google Cloud Spanner.

Configuring a Source Filter

Configure a source filter in the Application Source Qualifier.

1. In the **Mapping Designer**, double-click the Application Source Qualifier.

The **Edit Transformations** dialog box appears as shown in the following image:



2. Click the **Edit Conditions** tab.

3. Click **Edit Filter Conditions**.

The **Add Filter Condition** dialog box appears.

4. Select **Platform Expression** or **Native Expression**.

- If you configure a platform expression, select the filter field and operator that you want to specify in the condition, enter a value for the condition, and click **Add Condition**.

The following image shows a platform filter expression configured for a Google Cloud Spanner source:

The screenshot shows a 'Filter' dialog box with a title bar. Below the title bar, it says 'Specify filter conditions for the object:' followed by 'ALLDATATYPES_WOBYTE'. There are two radio buttons: 'Platform Expression' (selected) and 'Native Expression'. Below these is a section titled 'Add Filter Condition' with three input fields: 'Select Field' (a dropdown menu), 'Select Op' (a dropdown menu), and 'Value' (a text input field). To the right of these fields is a green plus icon. Below this section is a table with four columns: 'Field', 'Operator', 'Value', and 'Remove'. The first row contains the values 'ALLDATATYPES_WOBYTE.INT', '>', '5', and a red 'X' icon. There are two empty rows below it. At the bottom right of the dialog are 'Finish' and 'Cancel' buttons.

Field	Operator	Value	Remove
ALLDATATYPES_WOBYTE.INT	>	5	✖

- If you configure a native expression, specify a filter expression in the following format:
<ColumnName><Operator><Value>

You can use AND, OR, or nested conditions in the filter expression. The expression that you enter becomes the WHERE clause in the query used to retrieve records from the source.

The following image shows a native filter expression configured for a Google Cloud Spanner source:

The screenshot shows a 'Filter' dialog box with a title bar. Below the title bar, it says 'Specify filter conditions for the object:' followed by 'ALLDATATYPES_WOBYTE'. There are two radio buttons: 'Platform Expression' and 'Native Expression' (selected). Below these is a section titled 'Enter your query here:' with a large text input field. The text 'INT>50' is entered in the field. At the bottom right of the dialog are 'Finish' and 'Cancel' buttons.

5. Click **Finish** to add the filter condition.
6. Click **OK**.

CHAPTER 5

Google Cloud Spanner Sessions

This chapter includes the following topics:

- [Google Cloud Spanner Sessions Overview, 19](#)
- [Google Cloud Spanner Connections, 19](#)
- [Filter Override, 21](#)
- [Configure Google Cloud Spanner Source Session Properties, 22](#)
- [Configure Google Cloud Spanner Target Session Properties, 22](#)
- [Partitioning, 24](#)
- [Parameterization, 27](#)
- [Configure the Java Heap Memory, 27](#)

Google Cloud Spanner Sessions Overview

After you create mappings, you can create a session to extract or load data.

You must configure a Google Cloud Spanner connection in the **Workflow Manager** to extract data from or load data into a Google Cloud Spanner table. You can define properties in a session to determine how the PowerCenter Integration Service must extract data from or load data into a Google Cloud Spanner table.

Google Cloud Spanner Connections

Create a Google Cloud Spanner connection to read data from a Google Cloud Spanner source and write data to a Google Cloud Spanner target. You must create a connection for each database that you want to connect

to. You can use Google Cloud Spanner connections in mappings. When you create a Google Cloud Spanner connection, you can configure a connection mode based on how you want to read and write the data.

PowerExchange for Google Cloud Spanner Connections

When you configure a Google Cloud Spanner connection, you define the connection attributes that the PowerCenter Integration Service uses to connect to the Google Cloud Spanner.

The following table describes the Google Cloud Spanner connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	String that the PowerCenter Integration Service uses to identify the connection. The ID is not case sensitive. The ID must be 255 characters or fewer and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Google Cloud Spanner .
Project ID	Specifies the project_id value present in the JSON file that you download after you create a service account. If you have created multiple projects with the same service account, enter the ID of the project that contains the bucket that you want to connect to.
Service Account ID	Specifies the client_email value present in the JSON file that you download after you create a service account.
Service Account Key	Specifies the private_key value present in the JSON file that you download after you create a service account.
Instance ID	Name of the instance that you created in Google Cloud Spanner.

Configuring a Google Cloud Spanner Connection

Configure a Google Cloud Spanner connection in the **Workflow Manager** to define the connection attributes that the PowerCenter Integration Service uses to connect to Google Cloud Spanner.

1. In the Workflow Manager, click **Connections > Application**.
The **Application Connection Browser** dialog box appears.
2. Click **New**.
The **Select Subtype** dialog box appears.
3. Select **Google Cloud Spanner** and click **OK**.
The **Application Connection Editor** dialog box appears.
4. Enter a name for the Google Cloud Spanner connection.

5. Enter the Google Cloud Spanner connection attributes.
6. Click **OK** to create a Google Cloud Spanner connection.

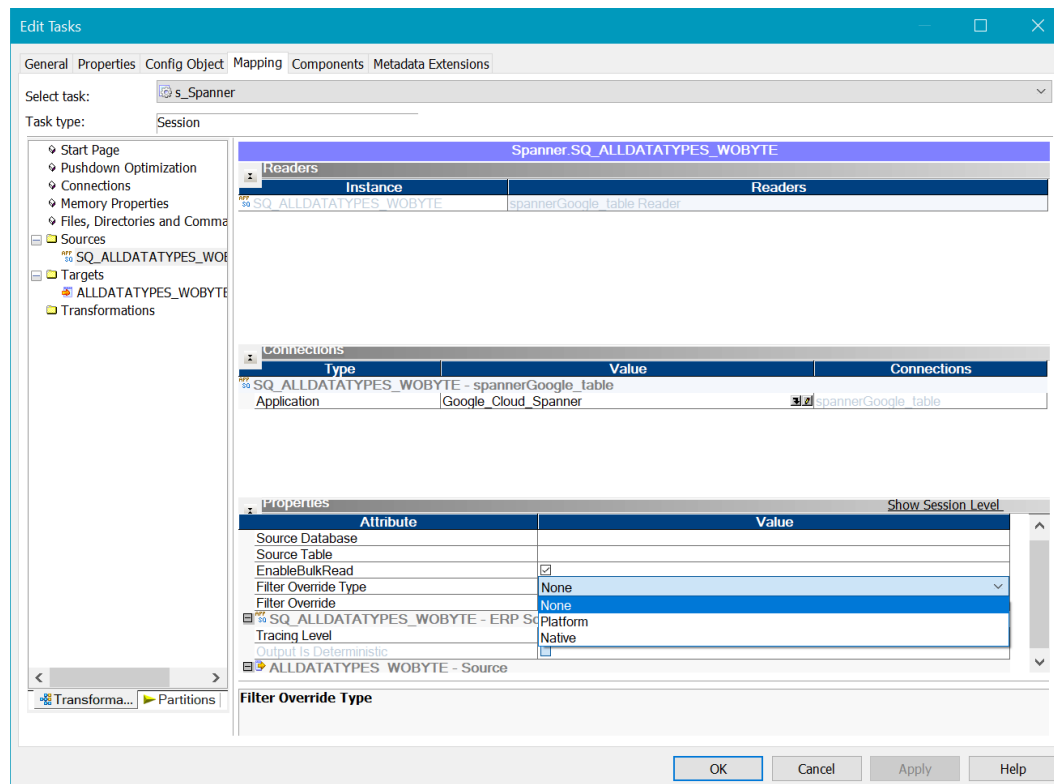
Filter Override

When you read data from a Google Cloud Spanner source, you can specify the filter override type and filter override condition in the Google Cloud Spanner source session properties to override the filter condition you specify in the source qualifier.

You must select the corresponding Native or Platform expression as the **Filter Override Type** in the Google Cloud Spanner source session properties based on the expression used in the source qualifier. Default is **None**.

After you select the filter expression, specify the filter condition in the **Filter Override** field. When you run the session, the PowerCenter overrides the filter condition you specified in the source qualifier and uses the condition you specified in the session properties to filter the Google Cloud Spanner data.

The following image shows the filter override type and the filter override condition fields in the Google Cloud Spanner source session properties:



Configure Google Cloud Spanner Source Session Properties

You can configure the session properties for a Google Cloud Spanner source on the **Mapping** tab. Define the properties for the source instance in the session.

The following table describes the session properties that you can configure for a Google Cloud Spanner source session:

Property	Description
Index Directive	Specify the index directive in an SQL statement to indicate the index name to query a table. To specify a index directive, use the following syntax: <code>@{FORCE_INDEX=TableIndex}</code> For more information about index directive, click the following URL: https://cloud.google.com/spanner/docs/secondary-indexes
Source Database	Optional. Overrides the Google Cloud Spanner database name specified in the Google Cloud Spanner source definition.
Source Table	Overrides the Google Cloud Spanner table name specified in the Google Cloud Spanner source definition.
Batch Size	Minimum number of rows that PowerExchange for Google Cloud Spanner reads in a batch. Default is 0.
EnableBulkRead	Select this property to read data in bulk from large Google Cloud Spanner tables. Note: If you select this property, ensure that the Batch Size source session property is set to the default value.
Filter Override Type	The type of filter expression that you want to override in the source qualifier. You can select Native or Platform filter expression type. Default is None .
Filter Override	The filter condition that overrides the filter condition you specify in the source qualifier. If you select the Native filter override type, specify a filter expression in the following format: <code><tableName.columnName><Operator><Value></code> If you select the Platform filter override type, specify a filter expression in the following format: <code><columnName><Operator><Value></code>

Configure Google Cloud Spanner Target Session Properties

You can configure the session properties for a Google Cloud Spanner target on the **Mapping** tab. You can perform insert, update, upsert, and delete operations on a Google Cloud Spanner target. Define the properties for the target instance in the session.

The following table describes the session properties that you can configure for a Google Cloud Spanner target session:

Property	Description
UpdateMode	<p>Determines the mode that the PowerCenter Integration Service uses to update rows in the Google Cloud Spanner target.</p> <p>If you select an update mode, you must select Update for the Treat Source Rows As session property in the Properties page.</p> <p>You can select one of the following modes:</p> <ul style="list-style-type: none"> - Update As Update. The PowerCenter Integration Service updates all rows flagged for update in the Update Strategy transformation if the entries exist. - Update Else Insert. The PowerCenter Integration Service first updates all rows flagged for update in the Update Strategy transformation if the entries exist in the target. If the entries do not exist, the PowerCenter Integration Service inserts the entries. <p>Default is Update As Update.</p>
Target DataBase	Optional. Overrides the Google Cloud Spanner database name specified in the Google Cloud Spanner target definition.
Create Database If Required	If the database specified in the Target DataBase field is not present in the Google Cloud Spanner instance, PowerExchange for Google Cloud Spanner creates the target database with the name that you specify in the Target DataBase field.
Target Table	Optional. Overrides the Google Cloud Spanner target table name specified in the Google Cloud Spanner target definition.
Create Table If Required	<p>If the table is not present in the target database, PowerExchange for Google Cloud Spanner creates the target table with the name that you specify in the Target Table field.</p> <p>PowerExchange for Google Cloud Spanner creates the table with the same schema as the target table that you specified in the Google Cloud Spanner target definition.</p>
Batch Size	<p>Minimum number of rows in a batch.</p> <p>To improve the performance of PowerExchange for Google Cloud Spanner when writing data to a Google Cloud Spanner target table, you must increase the batch size. Enter a number greater than 0.</p> <p>Default is 100.</p>
Write Truncate	<p>When you enable this property, PowerExchange for Google Cloud Spanner overwrites the existing data in the target table.</p> <p>Note: The Write Truncate option is applicable only when you perform an insert operation on a Google Cloud Spanner target.</p>

Rules and Guidelines for Google Cloud Spanner Targets

Consider the following rules and guidelines for Google Cloud Spanner targets:

- To configure upsert operations to write data to a Google Cloud Spanner target, you must select **Update Else Insert** as the update mode under the target session properties.
- When you configure update or upsert operations to write data to a Google Cloud Spanner target, you must map all the primary key fields of the Google Cloud Spanner target. Otherwise, the session fails. The **Update Columns** field is optional.
- If you specify the target table name in the target session properties and the table does not exist in the database, you must select the **Create Table If Required** option. Otherwise, the session fails.

- If you select the **Create Table If Required** option and specify the target table name in the target session properties, the target table that you specify must have the same schema as the target table that you specified in the Target transformation. Otherwise, the session fails.
- If you select the **Create Database If Required** option and the database is not present in the Google Cloud Spanner instance, you must select the **Create Table If Required** option. Otherwise, the session fails.
- If you select the **Create Database If Required** option and the target table name is different from the table name that you specified in the Target transformation, you must specify the target table name in the target session properties. You can specify a new target database name, if required.
- You cannot create a target at run time by using the **Create Target** option.
- When you write large datasets to a Google Cloud Spanner target, the session runs successfully. However, the PowerCenter Integration Service returns an incorrect number of success rows in the session log.

Partitioning

When you read from Google Cloud Spanner, you can configure pass-through partitioning to optimize the session performance at run time.

When you specify pass-through partitioning for a Google Cloud Spanner Source Qualifier transformation, you can specify filter conditions in the Google Cloud Spanner session properties to override the filter condition you specify in the source qualifier. The PowerCenter Integration Service uses the filter condition you specify in the session properties when it filters data from the source.

To configure pass-through partitioning, select the Source Qualifier transformation, and add a partition point from the **Mapping** tab of the session properties. Add the number of partitions you require and select the partition type as **Pass Through** for each of the partitions.

Based on the number of partitions you add, the PowerCenter Integration Service adds those many number of partition fields for the **Filter Override** attribute in the session properties. Specify the filter override condition for each of the partitions. The PowerCenter Integration Service uses the filter conditions you specify to pass data through the appropriate partition.

You can configure pass-through partitioning for the columns of the following data types:

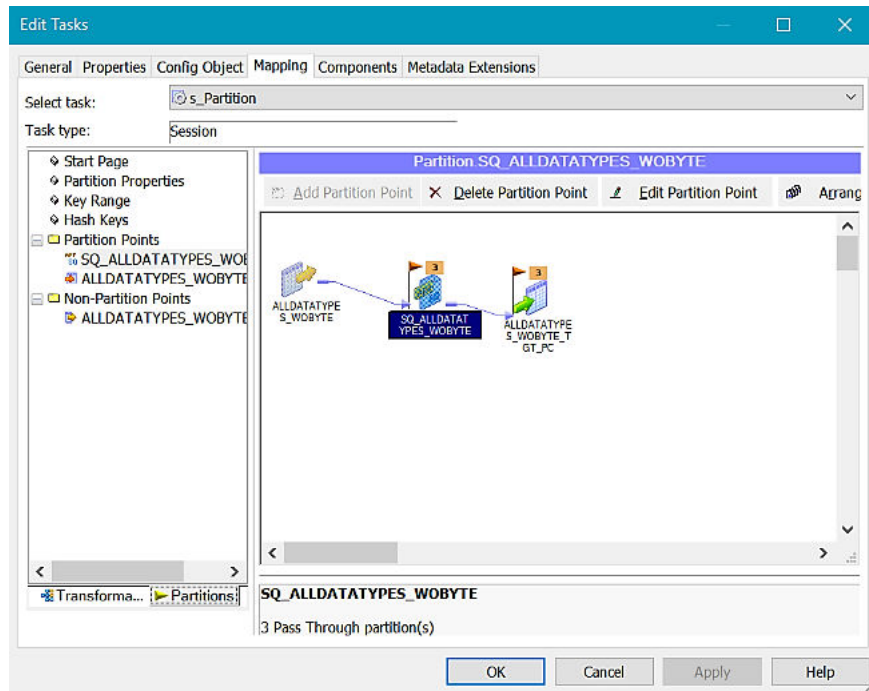
- Integer
- Sting
- Timestamp

Applying Filter Override Conditions for Partitions

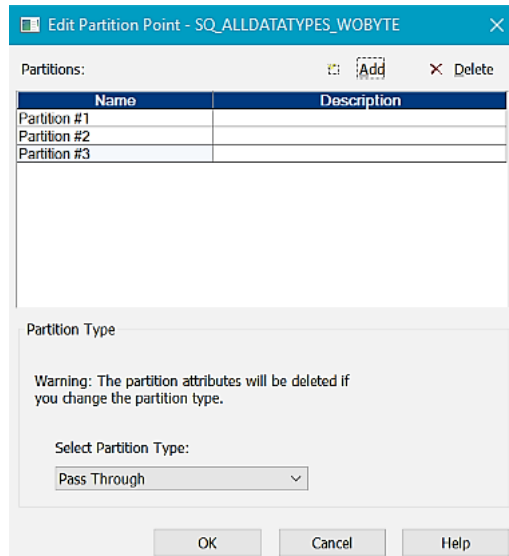
Configure pass-through partitioning for a Google Cloud Spanner session. After you add the number of partitions, you can specify a filter override condition for each of the partitions.

1. In the **Workflow Designer**, open the session properties.
2. In the **Partitions** view, click **Add Partition Point**.

The transformation name appears under the **Partition Points** node.



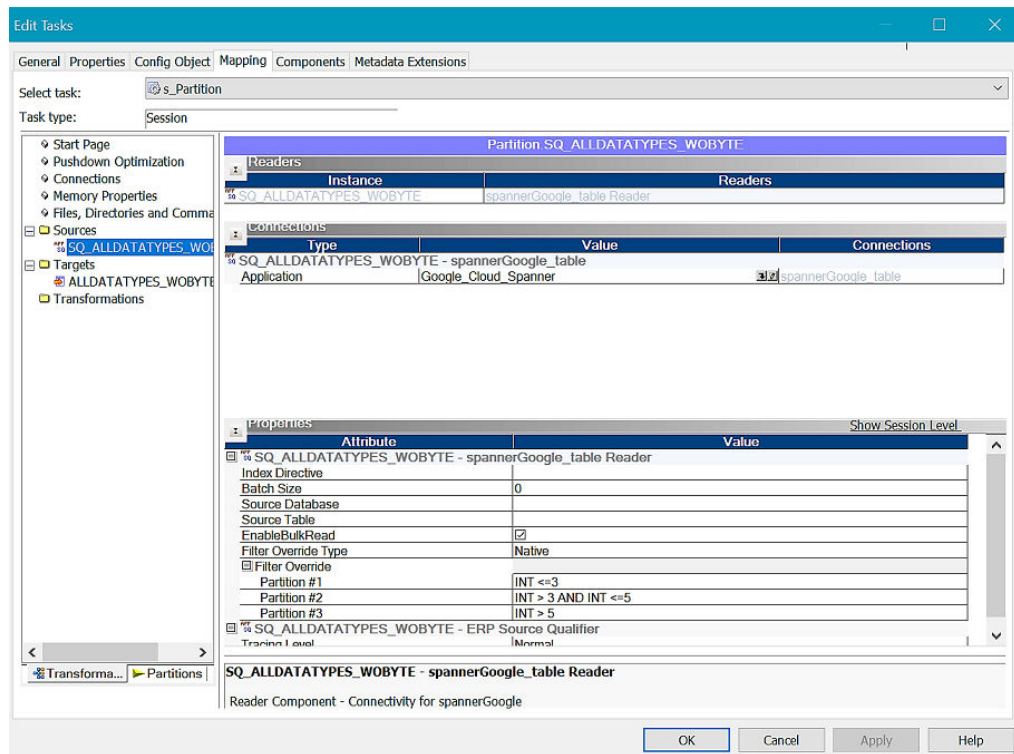
3. Select the Source transformation, and click **Edit Partition Point**.
The **Edit Partition Point** dialog box appears.
4. Click **Add** to add the partitions and enter a description for each partition.
5. Select the partition type as **Pass Through** for each of the partitions.



6. Click **OK**.
7. Click the **Transformations** tab in the session properties.
Based on the number of partitions you create, those many number of partition fields appear for the filter override property.

8. Specify the filter override condition for each partition.

For example, in the following image, you can configure filter override conditions for each of the three partitions configured:



When you run the session, based on the condition you specify for each partition, the PowerCenter Integration Service performs the following tasks:

- Passes all rows that contain INT values less than or equal to 3 to the first partition.
- Passes all rows that contain INT greater than 3 and lesser than or equal to 5 to the second partition.
- Passes all rows that contain INT values greater than or equal to 5 to the third partition.

9. Click **OK**.

Configure the Buffer Memory

When you configure pass-through partitioning for a Google Cloud Spanner session, you must increase the **DTM buffer size** and the **Default buffer block size** session property. Otherwise, the session fails with the following error:

ERROR: The session does not have enough buffer blocks for targets and sources. Increase the DTM buffer size to at least [value of buffer size] bytes. Error creating buffer pool. Total Buffer Pool size is [value of buffer pool size] bytes and Block size is [value of buffer block size] bytes.

To overcome this issue, you must configure the buffer memory settings in the session properties.

1. In the **Mapping Designer**, open the mapping for the session.
2. Click the **Config Object** tab of the session properties.
3. Edit the **Default buffer block size** in the Advanced settings.

Specify the value mentioned in the error message.

The default unit is bytes. Append KB, MB, or GB to the value to specify other units.

4. Click the **Properties** tab.

5. Edit the **DTM buffer size** property.

Specify the value mentioned in the error message.

The default unit is bytes. Append KB, MB, or GB to the value to specify other units.

Parameterization

You can parameterize the Google Cloud Spanner connection and session properties. You must create a mapping parameter in a mapping in the Designer and use the mapping parameters in the source qualifier in a session. Before you run a session, you must define values for the mapping parameters in the parameter file.

You can parameterize the following source session properties:

- Index Directive
- Batch Size
- Source Database
- Source Table

You can parameterize the following target session properties:

- Target DataBase
- Target Table
- Batch Size

Note: Before you run a session, you must not specify the parameters in the session properties and define values for the parameters in the parameter file. Otherwise, the session fails. However, the PowerCenter Integration Service does not return any error rows in the session log

Configure the Java Heap Memory

When you extract data from or load data to a large Google Cloud Spanner table, you must configure the memory for the Java heap size in the node that runs the PowerCenter Integration Service. Otherwise, the session fails.

1. In the Administrator tool, navigate to the PowerCenter Integration Service for which you want to change the Java heap size.
2. Click the **Processes** tab.
3. Edit the **General Properties** section.
4. Specify the maximum heap size in **Java SDK Maximum Memory** limit based on the amount data you want to process.
5. Click **Ok**.
6. Restart the PowerCenter Integration Service.

APPENDIX A

Google Cloud Spanner Data Type Reference

This appendix includes the following topics:

- [Data Type Reference Overview, 28](#)
- [Google Cloud Spanner and Transformation Data Types, 28](#)

Data Type Reference Overview

PowerCenter uses the following data types in Google Cloud Spanner mappings:

- Google Cloud Spanner native data types. Google Cloud Spanner data types appear in Google Cloud Spanner definitions in a mapping.
- Transformation data types. Set of data types that appear in the transformations. They are internal data types based on ANSI SQL-92 generic data types, which the PowerCenter Integration Service uses to move data across platforms. They appear in all transformations in a mapping.

When the PowerCenter Integration Service reads source data, it converts the native data types to the comparable transformation data types before transforming the data. When the PowerCenter Integration Service writes to a target, it converts the transformation data types to the comparable native data types.

Google Cloud Spanner and Transformation Data Types

The following table lists the Google Cloud Spanner data types that PowerCenter supports and the corresponding transformation data types:

Google Cloud Spanner Data Type	Transformation Data Type	Range and Description for the Transformation Data Type
BOOLEAN	Nstring	True or False
BYTE	Binary	1 to 104,857,600 bytes

Google Cloud Spanner Data Type	Transformation Data Type	Range and Description for the Transformation Data Type
DATE	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
FLOAT64	Double	Precision 15
INT64	BigInt	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
STRING	Nstring	1 to 104,857,600 characters
TIMESTAMP	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)

Note: PowerExchange for Google Cloud Spanner does not support the Array data type.

INDEX

C

- configuration
 - overview [10](#)
- configuring source filter [17](#)

F

- filter
 - application source qualifier [16](#)
 - source filter [16](#)
- filter conditions [17](#)

G

- Google Cloud Spanner
 - mappings [16](#)
 - overview [7](#)
 - source [12](#)
 - target [12](#)
- Google Cloud Spanner connections
 - properties [20](#)
 - connection configuration [20](#)
 - overview [20](#)
- Google Cloud Spanner data types
 - mapping with transformation data types [28](#)

I

- Import
 - connection properties [12](#)

- Import (*continued*)
 - source definition [12](#)
 - target definition [12](#)

M

- mapping
 - filter [16](#)
 - source [16](#)
 - target [16](#)

P

- plug-in registration
 - administrator tool [11](#)
- PowerExchange for Google Cloud Spanner
 - administration [8](#)
 - configuration [10](#)
 - overview [7](#)
 - plug-in registration [10](#)

S

- session
 - default buffer block size [26](#)
 - DTM buffer size [26](#)
 - java heap size [27](#)
 - overview [19](#)
 - source properties [22](#)
 - target properties [22](#)