

[Cloud Spanner](https://cloud.google.com/spanner/) (<https://cloud.google.com/spanner/>)

[Documentation](https://cloud.google.com/spanner/docs/) (<https://cloud.google.com/spanner/docs/>) [Guides](#)

Getting started with Cloud Spanner using gcloud

Objectives

This tutorial walks you through the following steps using the `gcloud` (<https://cloud.google.com/sdk>) command-line tool.

- Create a Cloud Spanner instance, database, and schema
- Write data to the database and execute SQL queries that data
- Clean up by deleting the database and instance

For the complete `gcloud spanner` reference, see [gcloud spanner](https://cloud.google.com/sdk/gcloud/reference/spanner/) (<https://cloud.google.com/sdk/gcloud/reference/spanner/>).

Costs

This tutorial uses Cloud Spanner, which is a billable component of the Google Cloud. For information on the cost of using Cloud Spanner, see [Pricing](https://cloud.google.com/spanner/pricing) (<https://cloud.google.com/spanner/pricing>).

Before you begin

1. Complete the steps described in [Set up](https://cloud.google.com/spanner/docs/getting-started/set-up) (<https://cloud.google.com/spanner/docs/getting-started/set-up>), which covers creating and setting a default Google Cloud project, enabling billing, enabling the Cloud Spanner API, and setting up OAuth 2.0 to get authentication credentials to use the Cloud Spanner API. In particular, ensure that you run `gcloud auth application-default login` (<https://cloud.google.com/sdk/gcloud/reference/auth/application-default/login>) to set up your local development environment with authentication credentials.

Set a default project

If you haven't already done so, set the ID of a Google Cloud Platform project as the default project for the gcloud command-line tool:

```
gcloud config set project PROJECT_ID
```

If you don't set the default project, you must pass `--project PROJECT_ID` to each of the commands below as the first argument to `gcloud spanner`. For example:

```
gcloud spanner --project=PROJECT_ID instance-configs list
```

Instances

When you first use Cloud Spanner, you must create an instance, which is an allocation of resources that are used by Cloud Spanner databases. When you create an instance, you choose where your data is stored and how many nodes are used for your data.

Instances and instance configurations

To create an instance, you must select an instance configuration, which is like a blueprint for your instance that defines the geographic placement and replication of your Cloud Spanner data.

List instance configurations

When you create an instance, you specify an *instance configuration*, which defines the geographic placement and replication of your databases in that instance. You can choose a regional configuration, which stores data in one region, or a multi-region configuration, which distributes data across multiple regions. Learn more in [Instances](https://cloud.google.com/spanner/docs/instances) (<https://cloud.google.com/spanner/docs/instances>).

To see the set of instance configurations that are available for your project:

```
gcloud spanner instance-configs list
```

You should see a list of regional and multi-region configurations.

Regional configurations distribute data in a single region, while multi-region configurations distribute data geographically across multiple regions. Read more about these in [Instances](https://cloud.google.com/spanner/docs/getting-started/gcloud/instances) (<https://cloud.google.com/spanner/docs/getting-started/gcloud/instances>).

Create an instance

To create an instance named `test-instance` with the display name `My Instance` using the regional instance configuration `regional-us-central1` with 1 nodes:

```
gcloud spanner instances create test-instance --config=regional-us-central1 \
  --description="My Instance" --nodes=1
```

In the command above, the instance name is set to `test-instance` and `--description` sets the display name of the instance. Both of these values must be unique within a Google Cloud Platform project.

Note: Use the instance ID, not the display name, when referring to an instance in `gcloud spanner` commands.

Set the default instance

You can set the default instance that Cloud Spanner uses when you have not specified an instance in your command. To set the default instance:

```
gcloud config set spanner/instance test-instance
```

Create a database

Create a database named `example-db`.

```
gcloud spanner databases create example-db
```

Create a schema

Use Cloud Spanner's [Data Definition Language](https://cloud.google.com/spanner/docs/data-definition-language)

(<https://cloud.google.com/spanner/docs/data-definition-language>) (DDL) to create, alter, or drop tables, and to create or drop indexes.

Let's create two tables

```
gcloud spanner databases ddl update example-db \  
  --ddl='CREATE TABLE Singers ( SingerId INT64 NOT NULL, FirstName STRING(1024), Las
```

```
gcloud spanner databases ddl update example-db \  
  --ddl='CREATE TABLE Albums ( SingerId INT64 NOT NULL, AlbumId INT64 NOT NULL, Albu
```

Write data

Let's add some sample data to our database

```
gcloud spanner rows insert --database=example-db \  
  --table=Singers \  
  --data=SingerId=1,FirstName=Marc,LastName=Richards  
  
gcloud spanner rows insert --database=example-db \  
  --table=Singers \  
  --data=SingerId=2,FirstName=Catalina,LastName=Smith  
  
gcloud spanner rows insert --database=example-db \  
  --table=Singers \  
  --data=SingerId=3,FirstName=Alice,LastName=Trentor  
  
gcloud spanner rows insert --database=example-db \  
  --table=Albums \  
  --data=SingerId=1,AlbumId=1,AlbumTitle="Total Junk"  
  
gcloud spanner rows insert --database=example-db \  
  --table=Albums \  
  --data=SingerId=2,AlbumId=1,AlbumTitle="Green"  
  
gcloud spanner rows insert --database=example-db \  
  --table=Albums
```

```
--table=Albums \  
--data=^:^^SingerId=2:AlbumId=2:AlbumTitle="Go, Go, Go"
```

By default, comma is used to delimit items in lists. In the last insert command, we specified colon (^:^) as the delimiter so that we could use comma in the album title.

Query data using SQL

Execute a query on the command line:

```
gcloud spanner databases execute-sql example-db \  
--sql='SELECT SingerId, AlbumId, AlbumTitle FROM Albums'
```

For the Cloud Spanner SQL reference, see [Query syntax](https://cloud.google.com/spanner/docs/query-syntax) (<https://cloud.google.com/spanner/docs/query-syntax>).

To see a list of flags you can use with the `execute-sql` command, see [gcloud spanner databases execute-sql](https://cloud.google.com/sdk/gcloud/reference/spanner/databases/execute-sql) (<https://cloud.google.com/sdk/gcloud/reference/spanner/databases/execute-sql>).

Cleanup

To avoid incurring additional charges to your Google Cloud account for the resources used in this tutorial, drop the database and delete the instance that you created.

Drop a database

To delete an existing instance:

```
gcloud spanner databases delete example-db
```

Delete an instance

To delete an existing instance:

```
gcloud spanner instances delete test-instance
```



Note that deleting an instance also drops all of the databases in that instance. Deleting an instance is not reversible.

Except as otherwise noted, the content of this page is licensed under the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/) (https://creativecommons.org/licenses/by/4.0/), and code samples are licensed under the [Apache 2.0 License](https://www.apache.org/licenses/LICENSE-2.0) (https://www.apache.org/licenses/LICENSE-2.0). For details, see our [Site Policies](https://developers.google.com/terms/site-policies) (https://developers.google.com/terms/site-policies). Java is a registered trademark of Oracle and/or its affiliates.

Last updated January 24, 2020.