This tutorial walks you through the following steps using the Cloud Spanner client library for PHP:

- Create a Cloud Spanner instance and database.

- Write, read, and execute SQL queries on data in the database.

- Update the database schema.

- Update data using a read-write transaction.

- Add a secondary index to the database.

- Use the index to read and execute SQL queries on data.

- Retrieve data using a read-only transaction.

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 (/spanner/pricing).

1. Complete the steps described in Set up (/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` (/sdk/gcloud/reference/auth/application-default/login) to set up your local development environment with authentication credentials.

2. Follow the steps in Service accounts (/docs/authentication#service_accounts) to set up a service account as your Application Default Credentials. Following those steps, you should obtain both a service account key file (in JSON) and a `GOOGLE_APPLICATION_CREDENTIALS` environment variable that will allow you to authenticate to the Cloud Spanner API.

3. Install the following on your development machine if they are not already installed:

   - PHP (http://php.net/)

   - composer (https://getcomposer.org/)

   - gRPC (/php/grpc)

4. Clone the sample app repository to your local machine:

   Alternatively, you can download the sample (https://github.com/GoogleCloudPlatform/php-docs-samples/archive/master.zip) as a zip file and extract it.

5. Change to the directory that contains the Cloud Spanner sample code:

6. Install dependencies:

   This will install the Cloud Spanner client library for PHP, which you can add to any project by running `composer require google/cloud-spanner`.

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 an *instance configuration*, which determines where your data is stored, and also the number of nodes to use, which determines the amount of serving and storage resources in your instance.

Execute the following command to create a Cloud Spanner instance in the region `us-central1` with 1 node:

Note that this creates an instance with the following characteristics:

- Instance ID `test-instance`

- Display name `Test Instance`

- Instance configuration `regional-us-central1` (Regional configurations store data in one region, while multi-region configurations distribute data across multiple regions. Learn more in Instances (/spanner/docs/instances).)

- Node count of 1 (`node_count` corresponds to the amount of serving and storage resources available to databases in the instance. Learn more in Node count (/spanner/docs/instances#node_count).)

You should see:

The samples repo contains a sample that shows how to use Cloud Spanner with PHP.

Take a look at the functions in `src/create_database.php` and `src/add_column.php`, which show how to create a database and modify a database schema. The data uses the example schema

shown in the Schema and data model

 (/spanner/docs/schema-and-data-model#creating-interleaved-tables) page.

Create a database called `example-db` in the instance called `test-instance` by running the
following at the command line.

You should see:

You have just created a Cloud Spanner database. The following is the code that created the
database.

ps://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/create_database.php)

The code also defines two tables, `Singers` and `Albums`, for a basic music application. These
tables are used throughout this page. Take a look at the example schema
(/spanner/docs/schema-and-data-model#creating-interleaved-tables) if you haven't already.

The next step is to write data to your database.

To do reads and writes, you need to obtain an instance of `Google\Cloud\Spanner\Database`
(https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database).

GitHub (https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/quickstart.php)

You can think of a `Database` as a database connection: all of your interactions with Cloud Spanner must go through a `Database`. Typically you create a `Database` when your application starts up, then you re-use that `Database` to read, write, and execute transactions.

Each client uses resources in Cloud Spanner, so you must call **Database::close** (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?method=close) to clean up the client's resources, including network connections.

Read more in the **Database**
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database)
reference.

You can insert data using Data Manipulation Language (DML) in a read-write transaction.

You use the `executeUpdate()` method to execute a DML statement.

github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/write_data_with_dml.php)

Run the sample using the `write-data-with-dml` command.

You should see:

You can also insert data using [mutations](/spanner/docs/modify-mutation-api) (/spanner/docs/modify-mutation-api).

You write data using the [Database::insertBatch](https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?method=insertBatch)
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?
method=insertBatch)
method. `insertBatch` adds new rows to a table. All inserts in a single batch are applied
atomically.

This code shows how to write the data using mutations:

[⊃ (https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/insert_data.php)](https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/insert_data.php)

Run the sample using the `insert-data` command.

You should see:

Cloud Spanner supports a native SQL interface for reading data, which you can access on the command line using the `gcloud` command-line tool or programmatically using the Cloud Spanner client library for PHP.

Execute the following SQL statement to read the values of all columns from the `Albums` table:

Note: See SQL syntax (/spanner/docs/query-syntax) for the Cloud Spanner SQL reference.

The result should be:

In addition to executing a SQL statement on the command line, you can issue the same SQL statement programmatically using the Cloud Spanner client library for PHP.

Use Database::execute()
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?
method=execute)
to run the SQL query.

Here's how to issue the query and access the data:

b (https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/query_data.php)

Run the sample using the `query-data` command.

You should see the following result:

Your results will not necessarily be in this order. If you need to ensure the ordering of the result, use an `ORDER BY` clause, as documented in SQL best practices (/spanner/docs/sql-best-practices#use_order_by_to_ensure_the_ordering_of_your_sql_results).

You can include custom values in SQL statements using parameters. Here is an example of using `@lastName` as a parameter in the `WHERE` clause to query records containing a specific value for `LastName`.

com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/query_data_with_parameter.php)

Run the sample using the query-data-with-parameter command.

You should see the following result:

In addition to Cloud Spanner's SQL interface, Cloud Spanner also supports a read interface.

Use `Database::read()`
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?
method=read)
to read rows from the database. Use a `KeySet` object to define a collection of keys and key
ranges to read.

Here's how to read the data:

<u>**ub**</u> (https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/read_data.php)

Run the sample using the `read-data` command.

You should see output similar to:

Assume you need to add a new column called `MarketingBudget` to the `Albums` table. Adding a new column to an existing table requires an update to your database schema. Cloud Spanner supports schema updates to a database while the database continues to serve traffic. Schema updates do not require taking the database offline and they do not lock entire tables or columns; you can continue writing data to the database during the schema update. Read more about supported schema updates and schema change performance in Schema updates (/spanner/docs/schema-updates).

You can add a column on the command line using the `gcloud` command-line tool or programmatically using the Cloud Spanner client library for PHP.

Use the following ALTER_TABLE (/spanner/docs/data-definition-language#alter_table) command to add the new column to the table:

You should see:

Use **Database::updateDdl**
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?
method=updateDdl)
to modify the schema:

(https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/add_column.php)

Run the sample using the `add-column` command.

You should see:

The following code writes data to the new column. It sets `MarketingBudget` to `100000` for the
row keyed by `Albums(1, 1)` and to `500000` for the row keyed by `Albums(2, 2)`.

(https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/update_data.php)

Run the sample using the `update-data` command.

You should see:

You can also execute a SQL query or a read call to fetch the values that you just wrote.

Here's the code to execute the query:

om/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/query_data_with_new_column.php)

To execute this query, run the sample using the `query-data-with-new-column` argument.

You should see:

You can update data using DML in a read-write transaction.

You use the `executeUpdate()` method to execute a DML statement.

/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/write_data_with_dml_transaction.php)

Run the sample using the `write-data-with-dml-transaction` command.

You should see:

Note: You can also underline:update data using mutations
(/spanner/docs/modify-mutation-api#updating_rows_in_a_table).

Suppose you wanted to fetch all rows of `Albums` that have `AlbumTitle` values in a certain range. You could read all values from the `AlbumTitle` column using a SQL statement or a read call, and then discard the rows that don't meet the criteria, but doing this full table scan is expensive, especially for tables with a lot of rows. Instead you can speed up the retrieval of rows when searching by non-primary key columns by creating a secondary index (/spanner/docs/secondary-indexes) on the table.

Adding a secondary index to an existing table requires a schema update. Like other schema updates, Cloud Spanner supports adding an index while the database continues to serve traffic. Cloud Spanner automatically backfills the index with your existing data. Backfills might take a few minutes to complete, but you don't need to take the database offline or avoid writing to the indexed table during this process. For more details, see index backfilling (/spanner/docs/secondary-indexes#adding_an_index).

After you add a secondary index, Cloud Spanner automatically uses it for SQL queries that are likely to run faster with the index. If you use the read interface, you must specify the index that you want to use.

You can add an index on the command line using the `gcloud` command line tool or programmatically using the Cloud Spanner client library for PHP.

Use the following **CREATE INDEX** (/spanner/docs/data-definition-language#create_index) command to add an index to the database:

You should see:

Use **Database::updateDdl**
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?
method=updateDdl)
to add an index:

(https://github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/create_index.php)

Run the sample using the `create-index` command.

Adding an index can take a few minutes. After the index is added, you should see:

For SQL queries, Cloud Spanner automatically uses an appropriate index. In the read interface, you must specify the index in your request.

To use the index in the read interface, use the **Database::read**
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?
method=read)
method.

github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/read_data_with_index.php)

Run the sample using the `read-data-with-index` command.

You should see:

You might have noticed that the read example above did not include reading the `MarketingBudget` column. This is because Cloud Spanner's read interface does not support the ability to join an index with a data table to look up values that are not stored in the index.

Create an alternate definition of `AlbumsByAlbumTitle` that stores a copy of `MarketingBudget` in the index.

Adding an index can take a few minutes. After the index is added, you should see:

Use `Database::updateDdl`
 (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?method=updateDdl)
to add an index with a `STORING` clause:

'github.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/create_storing_index.php)

Run the sample using the `create-storing-index` command.

You should see:

Now you can execute a read that fetches all `AlbumId`, `AlbumTitle`, and `MarketingBudget` columns from the `AlbumsByAlbumTitle2` index:

om/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/read_data_with_storing_index.php)

Run the sample using the `read-data-with-storing-index` command.

You should see output similar to:

Suppose you want to execute more than one read at the same timestamp. Read-only transactions (/spanner/docs/transactions#read-only_transactions) observe a consistent prefix of the transaction commit history, so your application always gets consistent data. Use a Snapshot (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/snapshot) object for executing read-only transactions. Use the `Database::snapshot` (https://googleapis.github.io/google-cloud-php/#/docs/google-cloud/master/spanner/database?method=snapshot) method to get a `Snapshot` object.

The following shows how to run a query and perform a read in the same read-only transaction:

ithub.com/GoogleCloudPlatform/php-docs-samples/blob/master/spanner/src/read_only_transaction.php)

Run the sample using the `read-only-transaction` command.

You should see output similar to:

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.

If you delete an instance, all databases within it are automatically deleted. This step shows how to delete a database without deleting an instance (you would still incur charges for the

instance).

1. Go to the **Cloud Spanner Instances** page in the Google Cloud Console.

   [Go to the Instances page](https://console.cloud.google.com/spanner/instances) (https://console.cloud.google.com/spanner/instances)

2. Click the instance.

3. Click the database that you want to delete.

4. In the **Database details** page, click **Delete**.

5. Confirm that you want to delete the database and click **Delete**.

Deleting an instance automatically drops all databases created in that instance.

1. Go to the **Cloud Spanner Instances** page in the Google Cloud Console.

   [Go to the Instances page](https://console.cloud.google.com/spanner/instances) (https://console.cloud.google.com/spanner/instances)

2. Click your instance.

3. Click **Delete**.

4. Confirm that you want to delete the instance and click **Delete**.

- <u>Access Cloud Spanner in a virtual machine instance</u>
(/spanner/docs/configure-virtual-machine-instance): create a virtual machine instance with access
to your Cloud Spanner database.

- Learn about authorization and authentication credentials in <u>Getting started with</u>
<u>authentication</u> (/docs/authentication/getting-started).

- Learn more <u>Cloud Spanner concepts</u> (/spanner/docs/concepts).