

[Cloud SQL](https://cloud.google.com/sql/) (<https://cloud.google.com/sql/>)

[Documentation](https://cloud.google.com/sql/docs/) (<https://cloud.google.com/sql/docs/>)

[MySQL](https://cloud.google.com/sql/docs/mysql/) (<https://cloud.google.com/sql/docs/mysql/>) [Guides](#)

# Connecting to Cloud SQL from Cloud Functions

---

**MySQL** | [PostgreSQL](https://cloud.google.com/sql/docs/postgres/connect-functions) (<https://cloud.google.com/sql/docs/postgres/connect-functions>) | [SQL Server](#)

## Beta

This feature is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](https://cloud.google.com/products/#product-launch-stages) (<https://cloud.google.com/products/#product-launch-stages>).

This page contains information and examples for connecting to a Cloud SQL instance from a service running in Cloud Functions.

Cloud SQL is a fully-managed database service that makes it easy to set up, maintain, manage, and administer your relational PostgreSQL and MySQL databases in the cloud.

[Cloud Functions](https://cloud.google.com/functions/docs/concepts/overview) (<https://cloud.google.com/functions/docs/concepts/overview>) is a lightweight compute solution for developers to create single-purpose, stand-alone functions that respond to Cloud events without the need to manage a server or runtime environment.

## Setting up a Cloud SQL instance

### 1. [Create a Cloud SQL for MySQL instance](https://cloud.google.com/sql/docs/mysql/create-instance)

(<https://cloud.google.com/sql/docs/mysql/create-instance#create-2nd-gen>).

★ **Note:** These instructions require your Cloud SQL instance to have a public IP address. If you want to use a private IP address, see [Configuring Serverless VPC Access](https://cloud.google.com/vpc/docs/configure-serverless-vpc-access) (<https://cloud.google.com/vpc/docs/configure-serverless-vpc-access>) and connect directly using the private IP.

2. Find the ***INSTANCE\_CONNECTION\_NAME*** for the instance on the **Instance details** page. It uses the format ***PROJECT\_ID:REGION:INSTANCE\_ID***, and is used to identify the Cloud

SQL instance you are connecting to.

3. Enable the Cloud SQL Admin API, if you haven't already done so:

**ENABLE THE API** ([HTTPS://CONSOLE.CLOUD.GOOGLE.COM/FLOWS/ENABLEAPI?APIID=SQLADMIN](https://console.cloud.google.com/flows/enableapi?APIID=SQLADMIN))

## Configuring Cloud Functions

Cloud Functions does not require any special configuration beyond making sure the service account used has the correct permissions.

Cloud Functions uses a service account to authorize your connections to Cloud SQL. This service account must have the correct IAM permissions to successfully connect. Unless otherwise configured, the default service account is in the format `service-YOUR_PROJECT_NUMBER@gcf-admin-robot.iam.gserviceaccount.com`.

When connecting resources in two different projects, make sure that both projects have enabled the correct IAM roles and have given the service account the correct permissions.

Ensure that the service account for your service has one of the following [IAM roles](https://cloud.google.com/iam/docs/understanding-roles#cloud-sql-roles) (<https://cloud.google.com/iam/docs/understanding-roles#cloud-sql-roles>):

- `Cloud SQL Client` (preferred)
- `Cloud SQL Editor`
- `Cloud SQL Admin`

Or, you can manually assign the following [IAM permissions](https://cloud.google.com/storage/docs/access-control/using-iam-permissions)

(<https://cloud.google.com/storage/docs/access-control/using-iam-permissions>):

- `cloudsql.instances.connect`
- `cloudsql.instances.get`

For detailed instructions on adding IAM roles to a service account, see [Granting Roles to Service Accounts](https://cloud.google.com/iam/docs/granting-roles-to-service-accounts) (<https://cloud.google.com/iam/docs/granting-roles-to-service-accounts>).

## Connecting to Cloud SQL

Once correctly configured, you can connect your service to your Cloud SQL instance using the Unix domain socket located at `/cloudsql/INSTANCE_CONNECTION_NAME`. These connections are automatically encrypted without any additional configuration.

**Warning:** Linux based operating systems have a max socket path length of 107 characters. If the total length of the path exceeds this length, you will not be able to connect with a socket from Cloud Functions.

**Note:** By default, Cloud Functions **does not** support connecting to the Cloud SQL instance using TCP. Your code should **not** try to access the instance using an IP address (such as `127.0.0.1` or `172.17.0.1`) unless you have configured [Serverless VPC Access](https://cloud.google.com/vpc/docs/configure-serverless-vpc-access) (<https://cloud.google.com/vpc/docs/configure-serverless-vpc-access>).

PYTHON

NODE.JS

[cloud-sql/mysql/sqlalchemy/main.py](https://github.com/GoogleCloudPlatform/python-docs-samples/blob/master/cloud-sql/mysql/sqlalchemy/main.py)  
(<https://github.com/GoogleCloudPlatform/python-docs-samples/blob/master/cloud-sql/mysql/sqlalchemy/main.py>)

UDPLATFORM/PYTHON-DOCS-SAMPLES/BLOB/MASTER/CLOUD-SQL/MYSQL/SQLALCHEMY/MAIN.PY)

```
# The SQLAlchemy engine will help manage interactions, including automatically
# managing a pool of connections to your database
db = sqlalchemy.create_engine(
    # Equivalent URL:
    # mysql+pymysql://<db_user>:<db_pass>@/<db_name>?unix_socket=/cloudsql/<cloud_
    sqlalchemy.engine.url.URL(
        drivename='mysql+pymysql',
        username=db_user,
        password=db_pass,
        database=db_name,
        query={
            'unix_socket': '/cloudsql/{}'.format(cloud_sql_connection_name)
        }
    ),
    # ... Specify additional properties here.
    # ...
)
```

To see this snippet in the context of a web application, view [the source code on GitHub](https://github.com/GoogleCloudPlatform/python-docs-samples/blob/master/cloud-sql/mysql/sqlalchemy/) (<https://github.com/GoogleCloudPlatform/python-docs-samples/blob/master/cloud-sql/mysql/sqlalchemy/>)

## Best Practices & Other Information

You can use the [Cloud SQL proxy](https://cloud.google.com/sql/docs/mysql/sql-proxy) (<https://cloud.google.com/sql/docs/mysql/sql-proxy>) when testing your application locally. See the [quickstart for using the proxy for local testing](https://cloud.google.com/sql/docs/mysql/quickstart-proxy-test) (<https://cloud.google.com/sql/docs/mysql/quickstart-proxy-test>) for detailed instructions.

### Connection Pools

Connections to underlying databases may be dropped, either by the database server itself, or by the infrastructure underlying Cloud Functions. To mitigate this, we recommend that you use a client library that supports connection pools that automatically reconnect broken client connections.

Additionally, we recommend using a globally scoped connection pool as this improves the likelihood that your function reuses the same connection for subsequent invocations of the function, and closes the connection naturally when the instance is evicted (auto-scaled down).

For more detailed examples on how to use connection pools, see [Managing database connections](https://cloud.google.com/sql/docs/mysql/manage-connections) (<https://cloud.google.com/sql/docs/mysql/manage-connections>).

### Connection Limits

Cloud SQL imposes a maximum limit on concurrent connections, and these limits may vary depending on the database engine chosen (see [Cloud SQL Quotas and Limits](https://cloud.google.com/sql/docs/quotas#fixed-limits) (<https://cloud.google.com/sql/docs/quotas#fixed-limits>)).

When using a connection pool, it is important to set the maximum connections to 1. This may seem counter-intuitive, however, Cloud Functions limits concurrent executions to 1 per instance. This means you never have a situation where two requests are being processed by a single function instance at the same time. This means in most situations only a single database connection is needed.

Where possible, you should take care to only initialize a connection pool for functions that need to use it. If a deployed function initializes a connection pool it doesn't need, it could create unused connections that count towards your quota. For more details on dealing with global variables in Cloud Functions, see [Tips & Tricks](#)

([https://cloud.google.com/functions/docs/bestpractices/tips#use\\_global\\_variables\\_to\\_reuse\\_objects\\_in\\_future\\_invocations](https://cloud.google.com/functions/docs/bestpractices/tips#use_global_variables_to_reuse_objects_in_future_invocations))

For more detailed examples on how to limit the number of connections, see [Managing database connections](#) (<https://cloud.google.com/sql/docs/mysql/manage-connections#count>).

---

*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 December 5, 2019.*