

[MySQL](#) (/sql/docs/mysql/replication/create-replica) | **PostgreSQL** | [SQL Server](#)

This page describes how to create a read replica for a Cloud SQL instance.

A read replica is a copy of the master that reflects changes to the master instance in almost real time. You create a replica to offload read requests or analytics traffic from the master. You can create multiple read replicas for a single master instance.

Read replicas are read-only. You cannot write to them.

Read replicas do not provide failover capability. To provide failover capability for a PostgreSQL instance, see [Configuration for High Availability](#) (/sql/docs/postgres/configure-ha).

For more information about how replication works, see [Requirements and Tips for Configuring Replication](#) (/sql/docs/postgres/replication/tips#read-replica).

All references to binary logging in this document apply to the master instance. Cloud SQL does not support binary logging on a read replica instance.

By default, the replica is created with the same number of CPUs and the same amount of memory as the master instance. You can increase these values for the replica, but you cannot decrease them.

- Learn about [requirements and best practices for replication](/sql/docs/postgres/replication/tips) (/sql/docs/postgres/replication/tips).

- Learn how to [manage replicas](/sql/docs/postgres/replication/manage-replicas) (/sql/docs/postgres/replication/manage-replicas).