

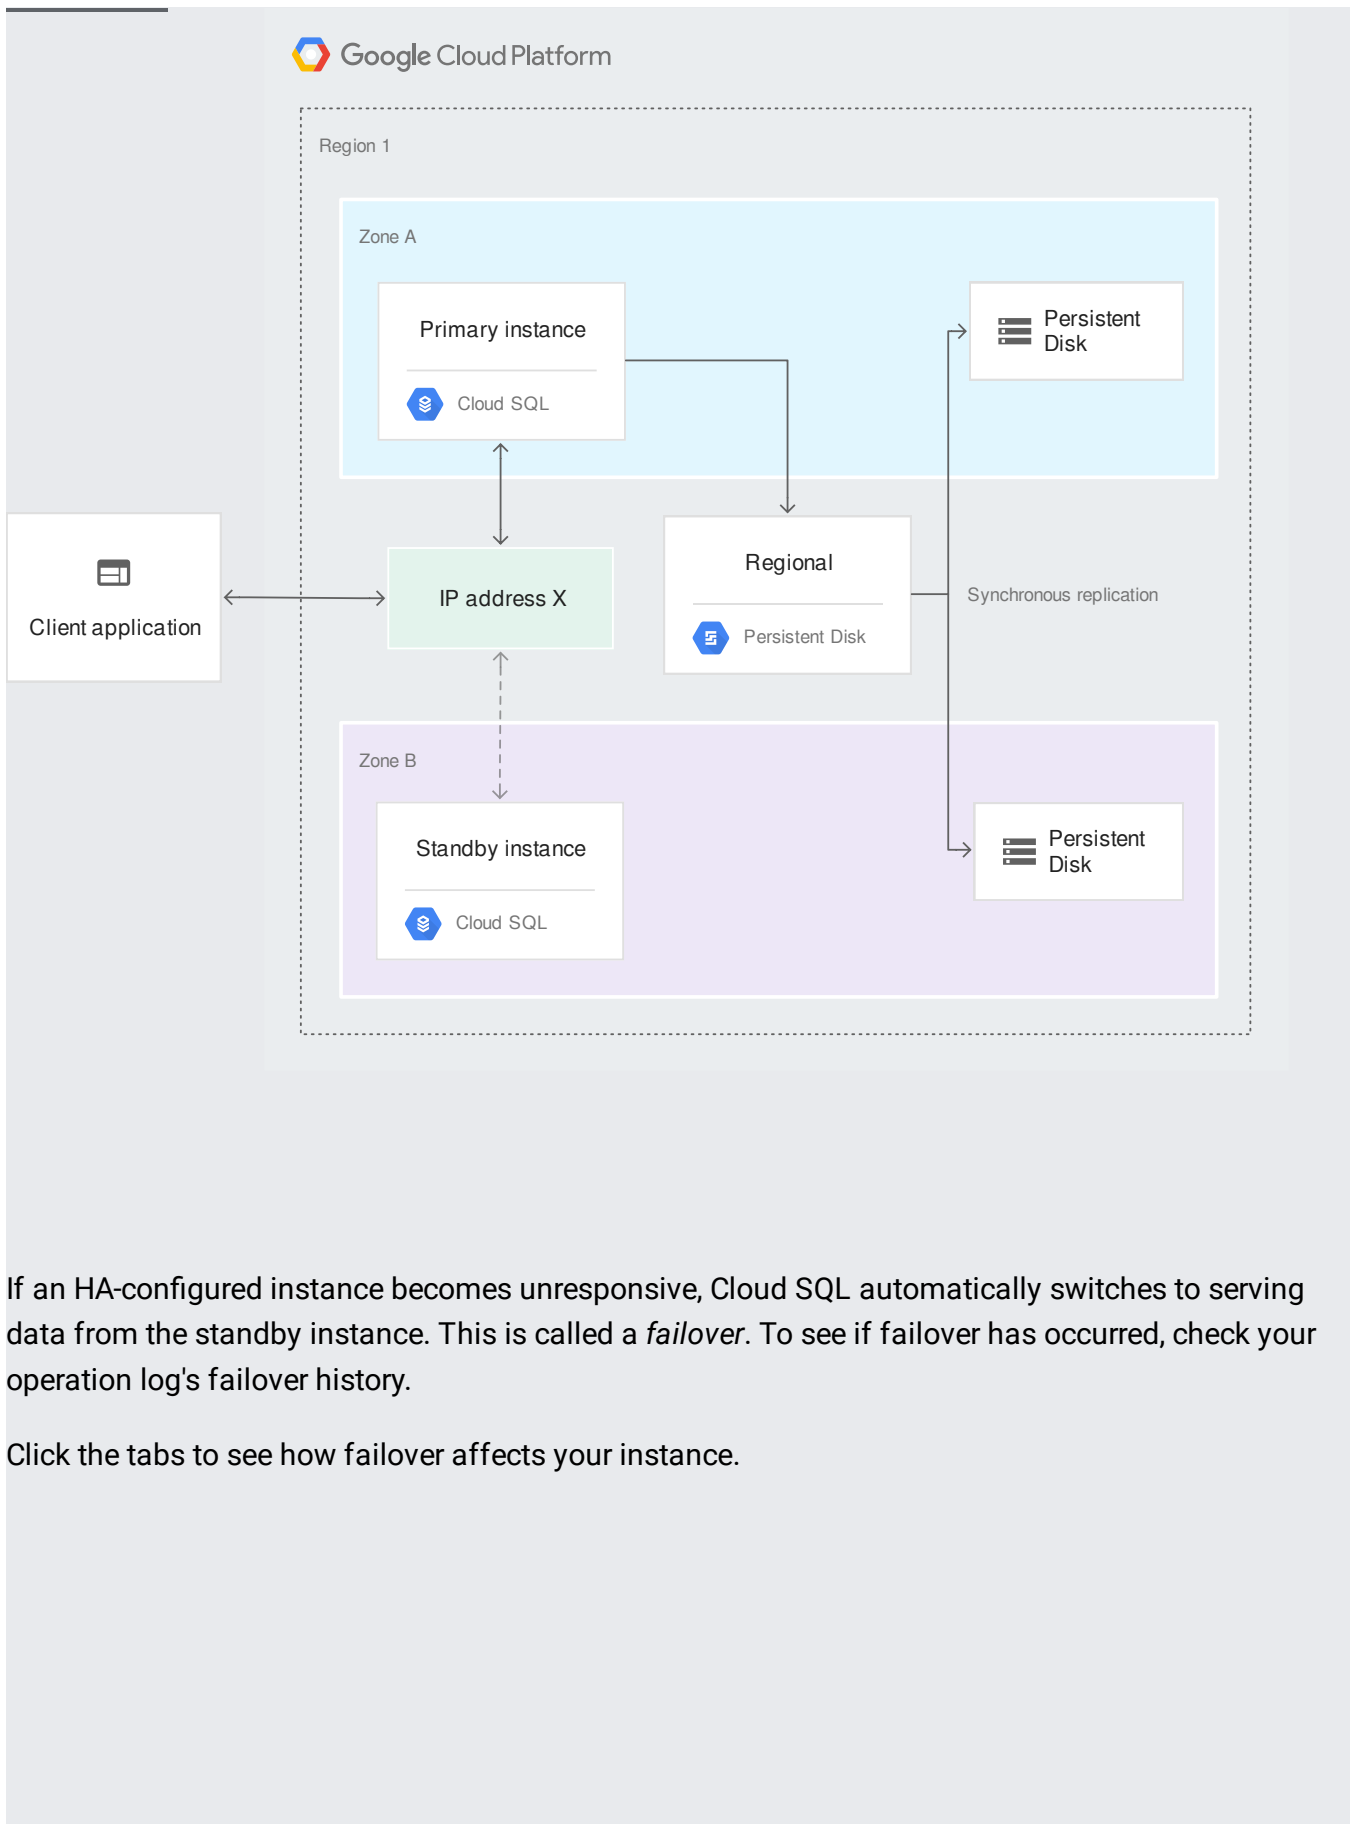
---

[MySQL](/sql/docs/mysql/high-availability/) (/sql/docs/mysql/high-availability) | **PostgreSQL** | [SQL Server](/sql/docs/sqlserver/high-availability/)  
(/sql/docs/sqlserver/high-availability)

This page is an overview of the high availability (HA) configuration for Cloud SQL instances. To configure a new instance for HA, or to enable HA on an existing instance, see [Enabling and disabling high availability on an instance](/sql/docs/postgres/configure-ha/) (/sql/docs/postgres/configure-ha).


The HA configuration, sometimes called a *cluster*, provides data redundancy. A Cloud SQL instance configured for HA is also called a *regional instance* and is located in a primary and secondary zone within the configured region. Within a regional instance, the configuration is made up of a *primary instance* (master) and a *standby instance*. Through [synchronous replication](/compute/docs/disks/#repds) (/compute/docs/disks/#repds) to each zone's persistent disk, all writes made to the primary instance are also made to the standby instance. In the event of an instance or zone failure, this configuration reduces downtime, and your data continues to be available to client applications.


Regional PD support for Cloud SQL and the Cloud SQL HA configuration are GA with full [SLA](/sql/sla) (/sql/sla) coverage. An HA-configured instance is charged at double the price of a standalone instance. This includes CPU, RAM, and storage. For more information, see [the pricing page](/sql/docs/postgres/pricing/) (/sql/docs/postgres/pricing).



If an HA-configured instance becomes unresponsive, Cloud SQL automatically switches to serving data from the standby instance. This is called a *failover*. To see if failover has occurred, check your operation log's failover history.

Click the tabs to see how failover affects your instance.





The following process occurs:

- The primary instance or zone fails.

Each second, the primary instance writes to a system database as a heartbeat signal. If multiple heartbeats aren't detected, failover is initiated. This occurs if the primary instance is unresponsive for approximately 60 seconds or the zone containing the primary instance experiences an outage.

- The standby instance now serves data upon reconnection.

Through a shared static IP address with the primary instance, the standby instance now serves data from the secondary zone.

- Read replicas continue to serve data from a healthy zone.

If failover occurs, read replicas do not change zones; they continue to serve data even if they are in a different zone than the primary instance. You can initiate another failover, in this case called a *failback*, to resume serving data from the primary's original zone (if the zone is in a healthy state).

For Cloud SQL to allow a failover, the configuration must meet the following requirements:

- The primary instance must be in a normal operating state (not stopped, undergoing maintenance, or performing a long-running operation).
- The secondary zone and standby instance must both be in a healthy state. When the standby instance is unhealthy and/or replication to the secondary zone is interrupted, failover operations are blocked. After Cloud SQL repairs the standby instance and the secondary zone is available, replication resumes and Cloud SQL allows failover.

If both the primary and standby instances are unhealthy, Cloud SQL does not allow failover.

Configuring an instance for HA does not affect your need for backups or how you create them.

There is no difference in working with non-HA and HA instances, so your application does not need to be configured in any particular way. When failover occurs, any existing connections to the primary instance and read replicas are closed, and it will take approximately 2-3 minutes for connections to be reestablished. Your application reconnects using the same connection string or IP address, so you do not need to update your application after failover.

To see exactly how your applications are affected by failover, you should [manually initiate failover \(/sql/docs/postgres/configure-ha#test\)](/sql/docs/postgres/configure-ha#test).

Maintenance events affect primary instances configured with HA in the same way as any other instance. You can expect primary instances to be down during this time. To minimize impact to your service, you can set a maintenance window to control when downtime occurs.

- [Enabling and disabling high availability on an instance \(/sql/docs/postgres/configure-ha\)](/sql/docs/postgres/configure-ha).
- [Initiate failover \(/sql/docs/postgres/configure-ha#test\)](/sql/docs/postgres/configure-ha#test).
- [Learn more about managing your database connections \(/sql/faq#connections\)](/sql/faq#connections).
- Learn more about [regions and zones \(/sql/docs/instance-locations\)](/sql/docs/instance-locations) in Cloud SQL.