

Resiliency with Cloud SQL


Google Cloud's fully managed relational database service





Contents

Resiliency with Cloud SQL	01
Factors Impacting Database Availability	02
Planned Downtime	03
Unplanned Downtime	07
Cloud SQL single-zone instance	08
Backup and Recovery	09
Cloud SQL High Availability	13
Cloud SQL Replication	16
Cloud SQL Disaster Recovery	18
Application considerations for an HA deployment	20



Resiliency with Cloud SQL

Cloud SQL is Google Cloud's fully managed relational database service for MySQL, PostgreSQL, and SQL Server. It provides full compatibility with the source database engines while reducing operations costs by automating database provisioning, storage capacity management, and other time-consuming tasks. Cloud SQL has built-in features to ensure business continuity with reliable and secure services, backed by a 24/7 SRE team providing a 99.95% SLA for the service.

Cloud SQL acts as the database backend for critical business applications deployed by enterprises across the globe. Many of these applications need to be available 24x7, and a key component of these applications is the database storing and managing the data used by the application. Ensuring the availability of these applications and the underlying database is important as downtimes result in loss of revenue, unhappy customers, and potential damage to an organization's reputation.

This in-depth guide discusses the availability features of Cloud SQL and how the service handles planned maintenance events and unplanned outages, including the advantages and controls available to help you minimize application downtime.

We will cover the availability of single-zone Cloud SQL instances, the advantages of a regional HA Cloud SQL deployment, and present an architecture to address disaster recovery requirements.

Each architecture builds on top of the previous deployment method while improving the availability of the Cloud SQL service. In other words, the Cloud SQL HA architecture provides the capabilities of a single-zone Cloud SQL instance deployment, plus additional availability characteristics. We will also briefly discuss the requirements of the applications to manage various outage events.

Note: There may be some differences in the capabilities of the individual database engines in Cloud SQL. Please check the [Cloud SQL documentation](#) for the specific database engine for details.

Factors impacting database availability

All IT systems, including databases, can be subject to unforeseen failures. Sometimes, they need to be taken out of service for planned events like software upgrades, security patches, and hardware and firmware updates. A highly available deployment architecture should ensure data protection that provides zero to low RPO (Recovery Point Objective) and low RTO (Recovery Time Objective) to ensure a fast return to service in case of any type of outage.

Cloud SQL deployments can be architected to meet the various availability requirements of applications. Here are the recommended steps you should follow to build a reliable system:

1. Define the reliability goals for the deployment in terms of service level objectives (SLOs).
2. Design for scale and high availability, factoring in redundancy of components, replication for disaster recovery, multi-region architecture, and handling load.
3. Build and leverage observability into the infrastructure and applications.



There are two types of events that impact availability: **planned downtime** due to maintenance or other activities and **unplanned downtime** due to various outage scenarios.

A planned outage or maintenance downtime is a proactive set of scheduled tasks carried out to improve the database availability, performance, or the security of your Cloud SQL instance or its underlying operating system. This is a scheduled event, and therefore, the impact on availability can be mitigated and controlled.

An unplanned outage, or unplanned downtime, is not scheduled and results from a component or system failure, software bugs, or even human error. The database deployment architecture needs to be designed to handle these outages.



Planned downtime

The planned events that can impact the availability of a Cloud SQL instance fall into two categories:

1. Configuration updates that you initiate to change a Cloud SQL instance's machine type database flags, zone configuration changes etc
2. Maintenance events scheduled by Cloud SQL to patch or update software.

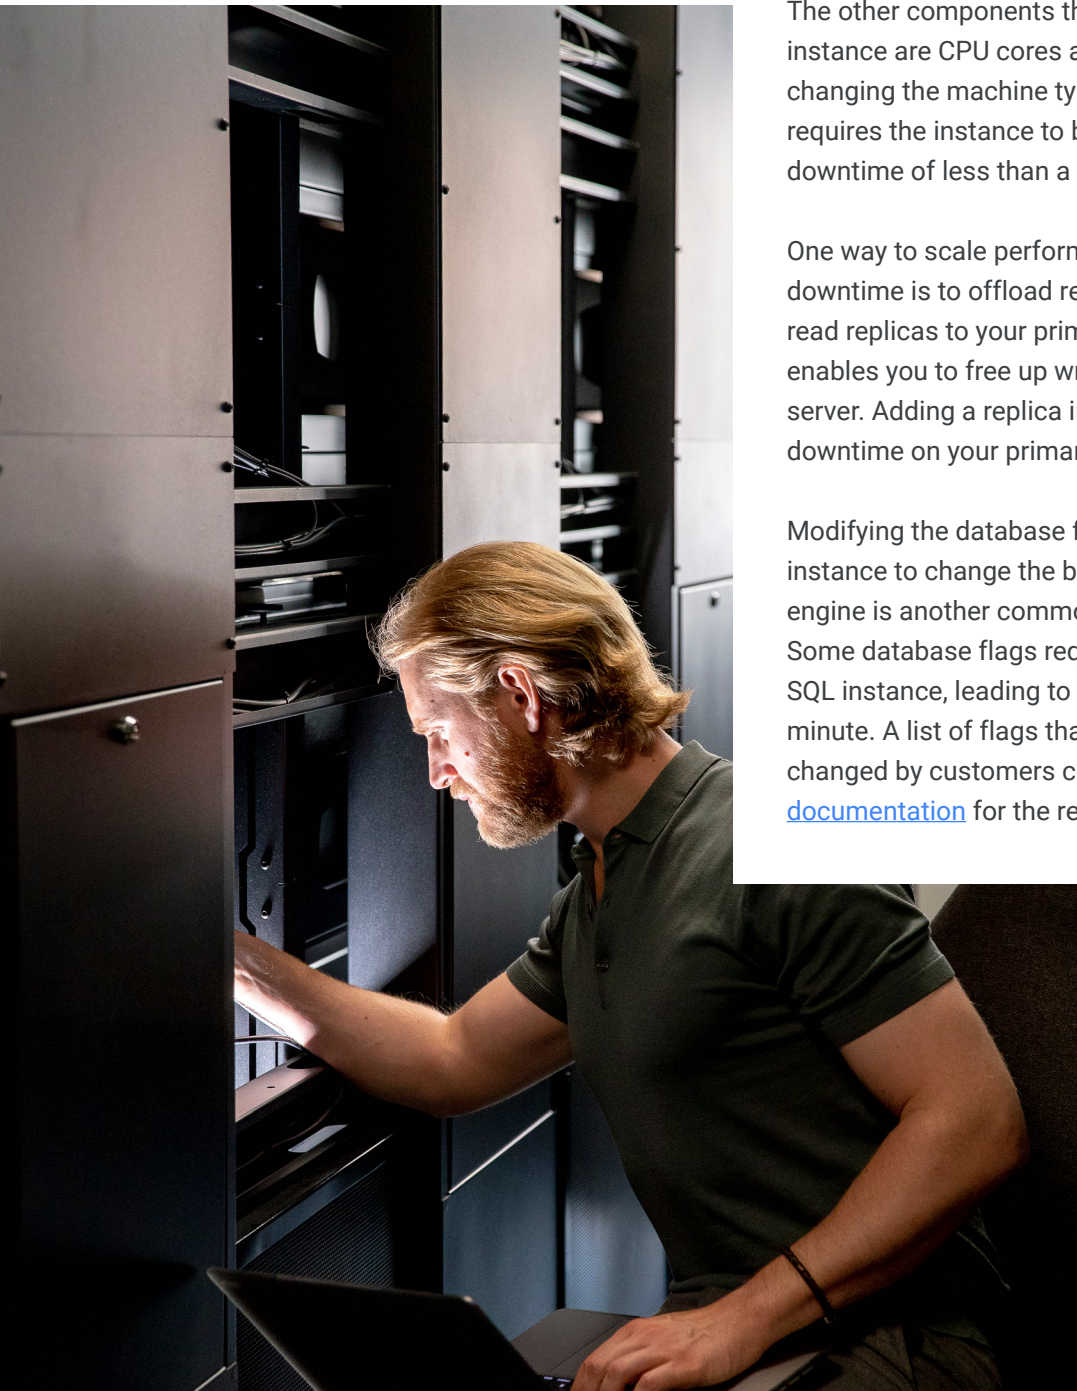
Configuration updates

Scaling an instance up or down in response to changing workload profiles and setting database flags for the respective database engines are types of planned events that you control and schedule, usually during a low period of user activity

A Cloud SQL instance has **three main components** that may need to be scaled: storage, CPU, and RAM.

Cloud SQL storage can be set to increase dynamically, which means there is no downtime when storage is added to the instance. If the available storage falls below a threshold size, Cloud SQL automatically adds additional storage capacity of up to 64 TB to your instance. Having a database that grows as needed minimizes application downtime by reducing the risk of running out of database space. You can take the guesswork out of capacity sizing without incurring any downtime or performing database maintenance.

While storage size can be increased automatically without downtime, it cannot be decreased. The storage increases are permanent for the life of the instance. There are methods to reduce the size of the database (thereby reducing storage), which could incur some downtime. An example might include migrating to a new instance using our [Database Migration Service \(DMS\)](#) with less storage allocated to the target instance.



The other components that make up the Cloud SQL instance are CPU cores and RAM. Scaling vertically, or changing the machine type to increase CPU and RAM, requires the instance to be restarted, which entails a downtime of less than a minute.

One way to scale performance without incurring downtime is to offload reads to the replicas by adding read replicas to your primary instance. This method enables you to free up write capacity on the primary server. Adding a replica instance does not incur any downtime on your primary instance.

Modifying the database flags of your Cloud SQL instance to change the behavior of the database engine is another common configuration update. Some database flags require you to restart your Cloud SQL instance, leading to a downtime of less than a minute. A list of flags that are allowed to be set or changed by customers can be found in the [Cloud SQL documentation](#) for the respective database engine.





Maintenance events

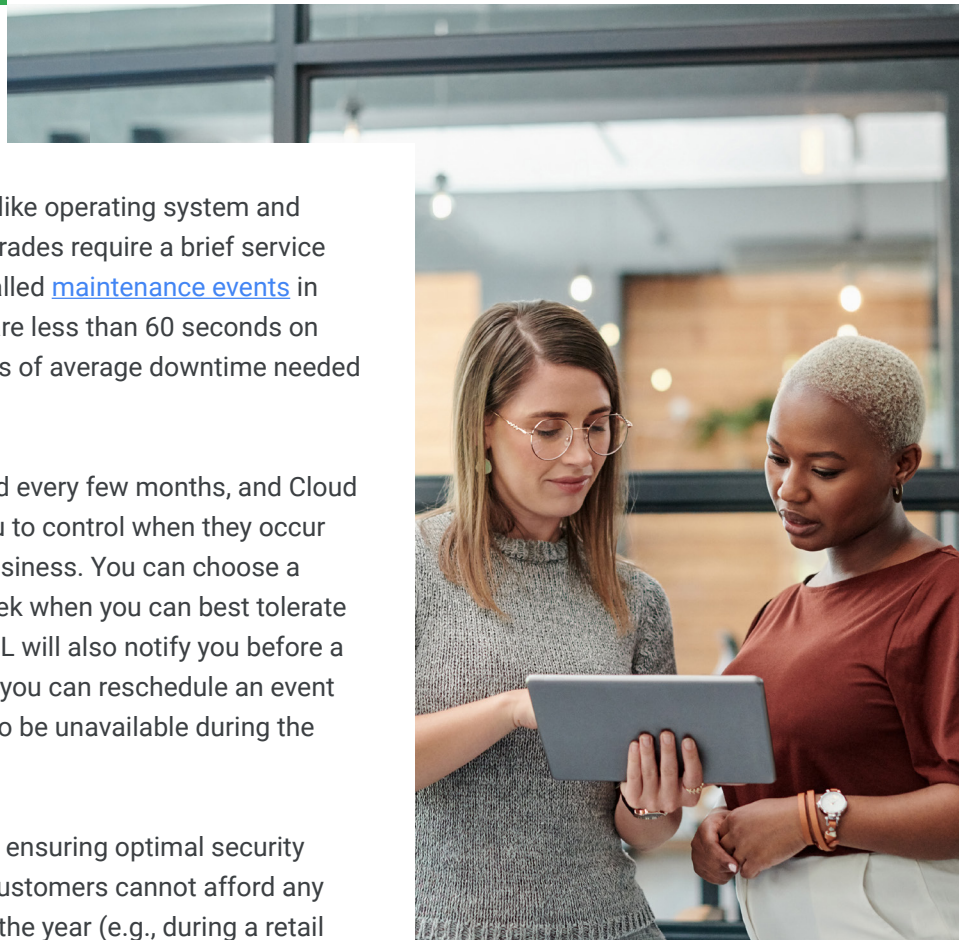
The second essential category of planned maintenance is software maintenance. Given Cloud SQL is a managed service, it automatically updates instances from time to time to ensure that the underlying hardware, operating system, and database engine are reliable, performant, secure, and up to date. We perform many of these updates while the Cloud SQL instance is up and running.

However, certain system updates like operating system and database engine patches and upgrades require a brief service interruption. These updates are called [maintenance events](#) in Cloud SQL. Service interruptions are less than 60 seconds on average, with less than 30 seconds of average downtime needed for PostgreSQL.

Maintenance events are scheduled every few months, and Cloud SQL provides settings to allow you to control when they occur to minimize the impact on your business. You can choose a 1-hour window on a day of the week when you can best tolerate maintenance downtime. Cloud SQL will also notify you before a scheduled maintenance event, so you can reschedule an event if you cannot afford the instance to be unavailable during the planned maintenance schedule.

While these updates are critical to ensuring optimal security and the best functionality, some customers cannot afford any downtime during certain times of the year (e.g., during a retail sale event or a school registration or intake period).

To help, Cloud SQL has a unique feature that enables you to set a deny maintenance period during these times. Setting up deny maintenance periods, which come in blocks of up to 90 days, prevents Cloud SQL from performing automatic maintenance during a deny period. When you configure a deny maintenance period on your primary instance, maintenance for all replicas associated with the primary instance is also denied.



Note: In very rare cases, Cloud SQL might need to schedule maintenance outside of the maintenance settings to patch severe stability issues or time-sensitive vulnerabilities. These updates roll out rapidly, and Cloud SQL counts them as downtime against the SLA.

Maintenance

Preferred window

Sunday, 12:00 AM – 1:00 AM EDT

Order of update 

Later

Notifications

On



Upcoming

This instance is scheduled for routine maintenance on Sunday, 8/15/21 at 12:00 AM EDT. [RESCHEDULE](#)

Deny maintenance period

11/1/21, 12:00 AM - 1/15/22, 11:00 PM EDT. Repeats yearly.



Edit maintenance preferences



Edit notification preferences

Above is a screenshot of what the maintenance settings could look like for a customer environment.

It's also important to consider the application's behavior during maintenance windows. Although the downtime during maintenance is very low, applications should be built to handle temporary errors. You should leverage techniques like proxies or connection pooling to minimize the application impact of dropped connections to the database. Also, applications should have error handling and retry logic with exponential backoff built in to handle any query failures or connection drops during maintenance.

For more details, we recommend reading [this article](#) about how Cloud SQL maintenance works.



Unplanned downtime

While downtime related to maintenance events can be controlled, unplanned outages can occur due to a variety of reasons ranging from hardware failures and software bugs to human error. More severe types of outages could be the loss of an entire region or data center due to a natural disaster. It's vital to minimize the downtime associated with these outages and recover with zero or minimal data loss.

In the following sections, we will describe how the features of Cloud SQL help to address various types of unplanned outages and reduce the associated downtime.

The table below lists the high availability Cloud SQL solution for various outage types for unplanned downtime, which we will describe in more detail later.

Architecture	Outage	Cloud SQL Solution	Recovery Point (RPO)	Recovery Window (RTO)
Single Instance	Transient Instance Failure	Reboot	Zero	Minutes
	Rows deleted/ table dropped etc	Backup & recovery + PITR	Seconds/Mins based on last available backup + PITR	Minutes to hours
High Availability	Instance Failure	HA Deployment across Zones	Zero	Minute
	Zone Failure	HA Deployment across Zones	Zero	Minutes
Disaster Recovery	Region Failure	Cloud SQL cross region replicas	Seconds	Minutes

*PITR = Point-in-time recovery



Cloud SQL single-zone instance

The following sections discuss the availability characteristics of a single-zone Cloud SQL instance and the features that help to address recovery. Later, we'll show you how to build on top of that, with additional features, to provide a better availability profile.

A single-zone Cloud SQL instance is likely used for development or testing purposes or for running applications that do not have strong RTO and RPO requirements. Live migration provides enhanced availability and database backups with PITR provide data protection. Please note that Live Migration is not triggerable by the customer and is a GCP-level availability enhancement.



Live Migration (managed by Google Cloud Operations)

In a typical data center environment, any updates to the underlying physical infrastructure like swapping out a defective machine, replacing an old or failing disk, performing BIOS updates, or similar have the potential to bring down the instance undergoing hardware maintenance.

However, Google Cloud performs hardware updates without interruption to a user's application. For example, when updating a database server, Google Cloud uses live migration—an advanced technology that reliably migrates a virtual machine (VM) from the original host to a new one while the VM stays running. There may be a short brownout period during the live migration, and the application should be coded to handle errors and retry.

Live Migration is done on a best-effort basis. If hardware fails completely, or otherwise prevents live migration, the VM automatically crashes and restarts.



If an instance has issues or potentially faces data loss due to corruption or human error, backups help to restore lost data to your Cloud SQL instance. The backup and recovery concepts described below apply to both a single-zone Cloud SQL instance and the HA configuration.

Cloud SQL backups are incremental in nature. The first backup is a full backup with subsequent backups containing only the changes since the last backup. The destination for the backups is Google Cloud Storage, with the option to choose either a multi-region location (backup copies stored in two separate regions) or a custom region location. If you select a custom location, it's helpful to adhere to data residency regulations.

Backup and recovery

In the event of human error or data corruption, Cloud SQL backups protect your data from loss or damage. While backups are a foundational element of an availability strategy, it is important to test recovery using the backups that have been taken.

Create a Backup

Backups allow you to restore instances to recover lost data or undo errors.

To help you save money, backups work incrementally. Each backup stores only the changes to your data since the previous backup. [Learn more](#)

i You'll be billed \$0.08/GB per month for data storage.

Describe this backup (optional)

FirstBackup

You can make a note here to help you identify this backup.

11 / 140

i You can only store backups in locations not restricted by your organization policy "Resource Location Restriction". [Learn more](#)

Choose where to store your backups

Backups are stored in the closest multi-region location to you by default. Only customize if needed.

☒ Multi-region (default)

☐ Region

Location *

us - Data centers in the Un... ▼

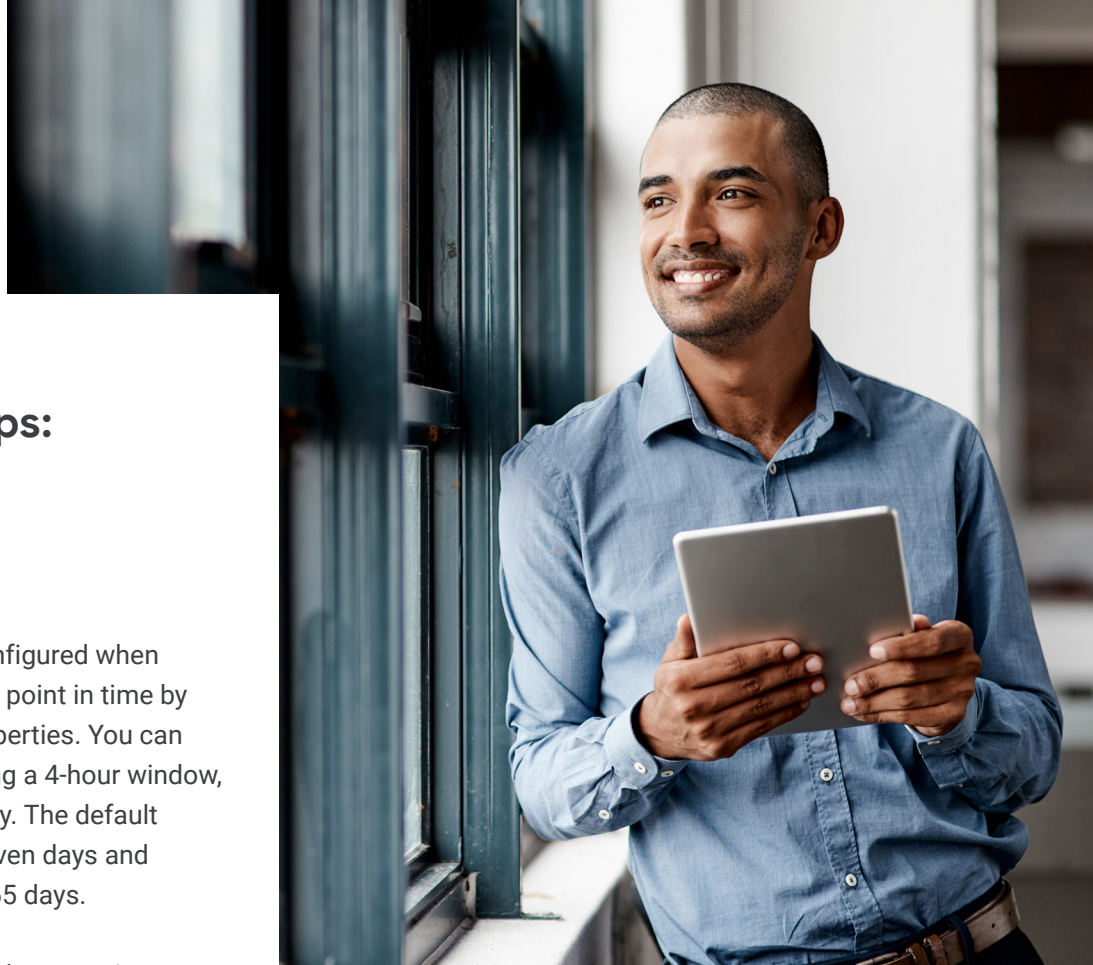
Cloud SQL offers two types of backups:

1. Automated backups
2. On-demand backups

Automated backups are either configured when an instance is created or at a later point in time by editing an instance's instance properties. You can configure backups to initiate during a 4-hour window, which runs automatically every day. The default retention period for backups is seven days and configurable for retention up to 365 days.

On the other hand, on-demand backups run in an ad-hoc manner. They are particularly useful before any risky operation on a database or if you need a backup before the automated backup schedule kicks in. On-demand backups do not have a retention window and persist until you delete them or the instance is deleted.

Backups, automated or on-demand, are a snapshot copy of the database as it was when the backup was run. Customers can restore from backups and get a consistent state as of the time of the backup. There's additional capability to fast-forward the database closer to the current state using point-in-time recovery (PITR). Customers can view a list of backups and the details of a backup for an instance from the cloud console.



Backups, recovery, and high availability

Backups

Automated backups and point in time recovery help protect your data from loss at a minimal cost.

☒ Automate backups

3:00 PM – 7:00 PM

Choose the best window of time for your data to be automatically backed up. May continue outside window until complete. Hours shown in your local time zone (UTC-7).

[Location options](#)

☒ Enable point-in-time recovery

Allows you to recover data from a specific point in time, down to a fraction of a second, via write-ahead log archiving.

Availability

Choice affects cost. You can change this option at any time by editing your instance.

☐ Single zone

In case of outage, no failover. Not recommended for production instances.

☒ High availability (regional)

Automatic failover to another zone within your selected region. Recommended for production instances. [Learn more](#)

Google Cloud

You can restore on-demand or automated backups to the same instance where you took the backup, another instance in the same region, or even a different region—as long as the backups are accessible.

The ability to restore the backup to a different region provides basic disaster recovery capabilities in this single instance scenario. If a primary region fails, you can restore the most recent backup to a different region and the instance will be immediately available for running applications.

During an outage of the region, the backup ID can be determined using the following command:

```
gcloud sql backups list --instance -
```

Use the following command to restore the instance:

```
gcloud sql backups restore BACK_ID \
--restore-instance=TARGET_INSTANCE_
NAME_ \ --backup-instance=SOURCE_
INSTANCE_NAME
```



PITR helps prevent or minimize data loss by recovering an instance to the latest point in time or a specific point in time. For example, if an error like accidental data deletion causes data loss, you can recover a database to its state before the error occurred. The PITR capability leverages the transaction logs of the database engine; automated backups need to be enabled for PITR to be possible. You can configure transaction logs to be retained for one to seven days.

A point-in-time recovery always creates a new instance that inherits the settings of the source instance. It allows both the source instance and new instance to exist while enabling deleted data to be extracted and transferred to the source instance. PITR also provides the ability to create a copy of the instance as it exists currently, making it easy to create instance copies for test and development purposes.





In most cases, backup and recovery deliver an RTO, the time it takes to restore the backup to a new instance, of minutes. The RPO, the time when data loss occurs, depends on whether PITR is enabled on the instance.

If PITR is not enabled, the RPO will be the difference in time between the last backup and the data loss incident. By comparison, the RPO with PITR enabled can be as low as 0 since you can target recovery to the point-in-time just before the data loss incident.

Below are some of the scenarios where you can leverage Backup and Recovery to recover from a failure:

1. **Loss of Instance:** If there is a complete loss of instance due to catastrophic hardware failure, a new instance can be created and the data will be hydrated from the backups.
2. **Data Corruption:** If data corruption or manual data manipulation leads to inconsistencies in the state of the data and database, a new instance can be created. The data must be hydrated from backups and transaction logs (PITR) before the data corruption event occurred.

For more details, we recommend reading [this article](#) about Cloud SQL backup and recovery concepts.

Cloud SQL High Availability

The Cloud SQL High Availability configuration provides next-level availability for Cloud SQL instances. It builds on top of our foundational availability features described above.

Many applications need higher availability than what is provided by a single-zone instance. Examples of this might include business-critical applications where the availability of the database, with no loss of data and very low recovery time in case of any failures, is essential for an organization to function.

Database clustering needs to be set up for the Cloud SQL instance in this case. We call this a **High Availability (HA) configuration**, which can be set up when you create an instance or at a later point in time.

An HA configuration provides redundancy, minimizing downtime when a zone or instance becomes unavailable. HA configurations have two instances (a primary and a standby) located in a primary and secondary zone within a configured region. An HA instance is called a **regional instance**. Only the instance in the primary or secondary zone is active and accepts read and write connections at a time.



Google Cloud

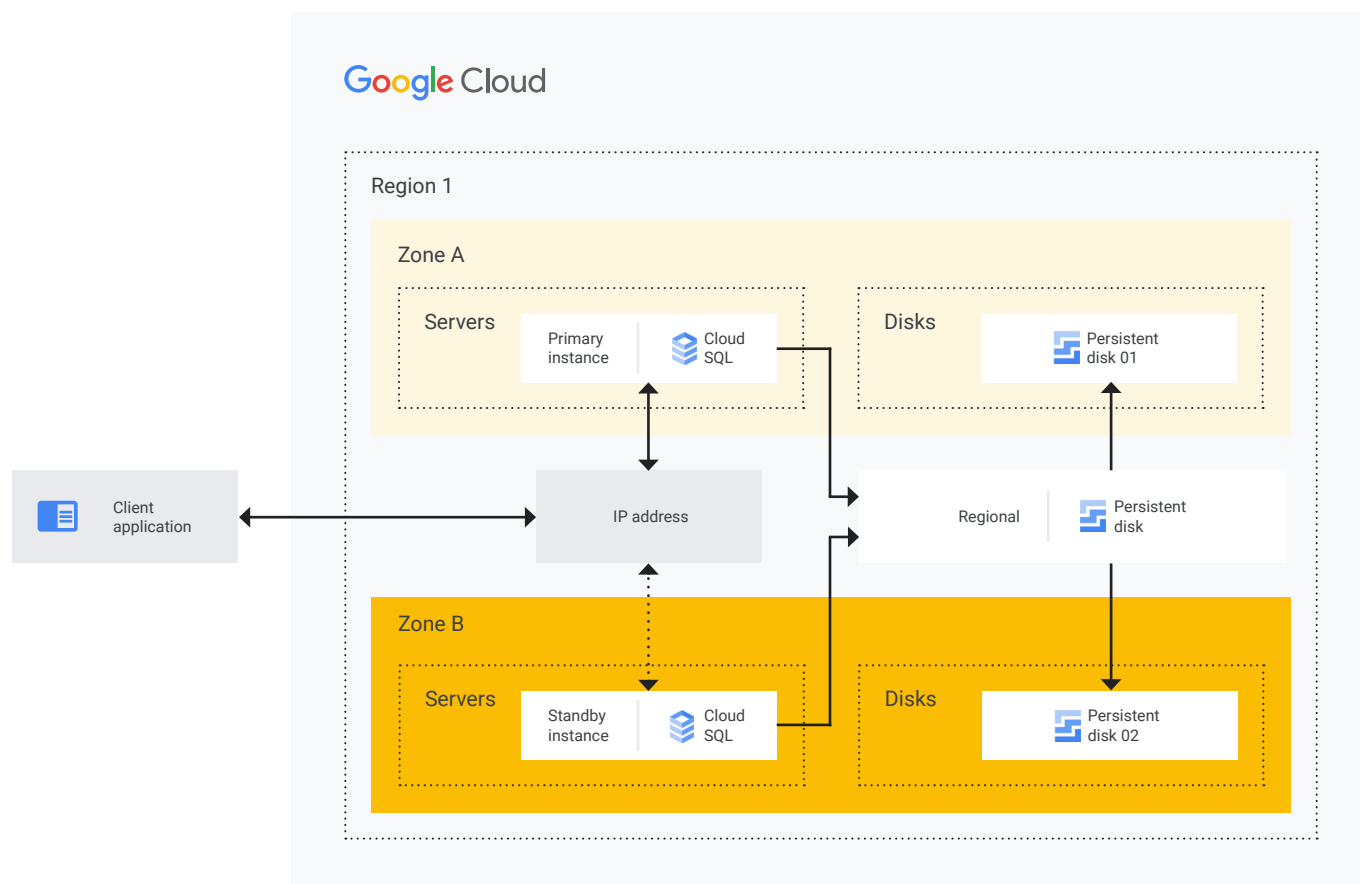
HA instances use a concept called Regional Persistent Disk (RePD). All writes made to the primary instance are replicated synchronously to disks in both zones before a transaction is reported as committed, providing data redundancy. If the primary fails and a standby instance needs to take over as the new primary, you will not lose any data.

If your primary instance becomes unresponsive, Cloud SQL automatically switches to serving data from a standby instance. During failover, the new primary continues serving data even when the original primary comes back online to avoid another outage for the application. Applications with retry logic will be able to tolerate the switching times, ensuring business continuity as well as a few minutes of recovery time.

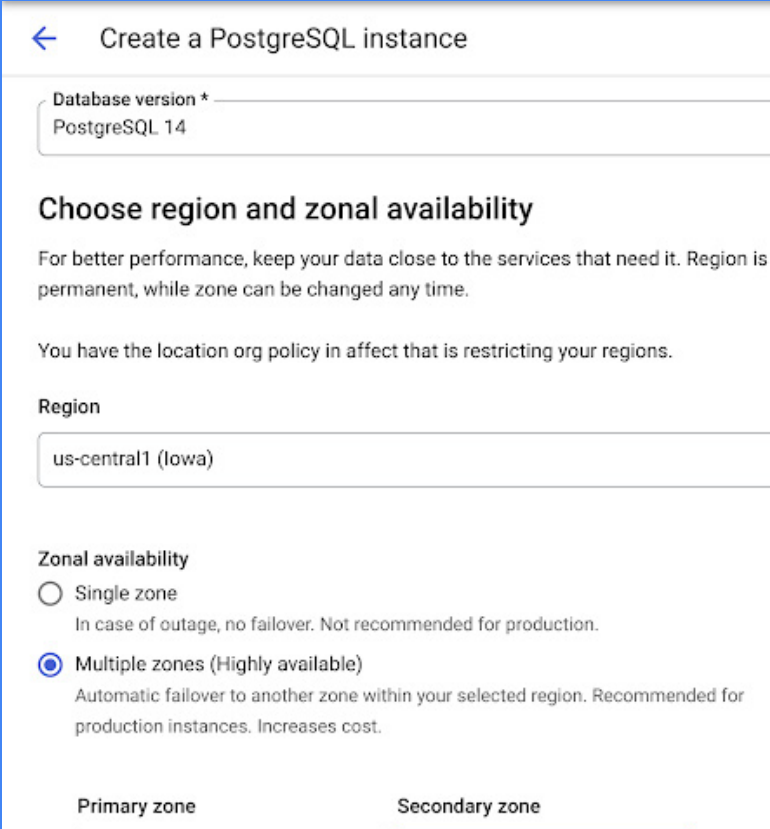
However, if the primary instance needs to be in the zone that had the outage, you can fail back once the original zone and instance is back online by initiating a failback operation.

The Cloud SQL HA configuration ensures there is no data loss if an instance or zone fails, providing zero RPO and roughly a minute of failover to the standby instance, thus minimizing the RTO.

Below is the HA configuration schematic:



Here's how to create a Cloud SQL HA configuration:



← Create a PostgreSQL instance

Database version *
PostgreSQL 14

Choose region and zonal availability

For better performance, keep your data close to the services that need it. Region is permanent, while zone can be changed any time.

You have the location org policy in affect that is restricting your regions.

Region

us-central1 (Iowa)

Zonal availability

☐ Single zone
In case of outage, no failover. Not recommended for production.

☒ Multiple zones (Highly available)
Automatic failover to another zone within your selected region. Recommended for production instances. Increases cost.

Primary zone Secondary zone

Cloud SQL Replication

Replication is the ability to create copies of a Cloud SQL instance or an on-premises database and offload work to the copies. A read replica is a copy of the primary instance that reflects changes to the primary in almost real time. While the main use case for replicas is to offload read workloads from the primary instance, they also improve availability.

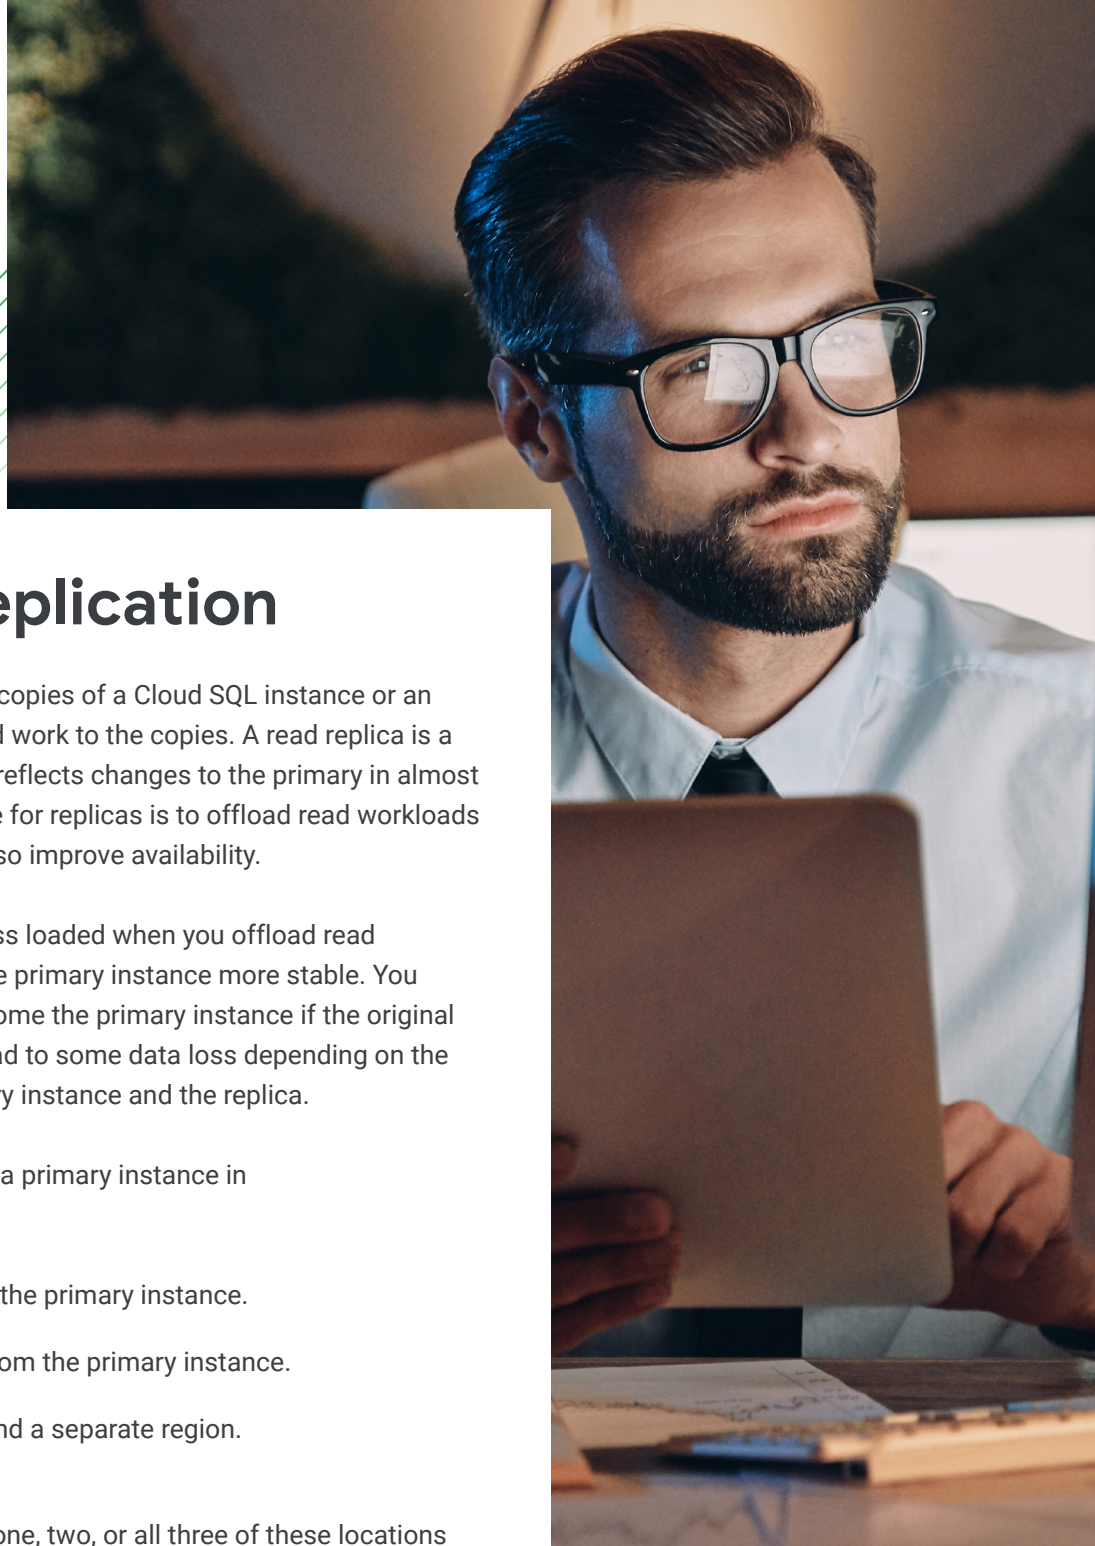
The primary instance becomes less loaded when you offload read workloads to a replica, keeping the primary instance more stable. You can also promote a replica to become the primary instance if the original is corrupted. However, this can lead to some data loss depending on the replication lag between the primary instance and the replica.

There can be multiple replicas for a primary instance in various combinations:

- A replica in the same zone as the primary instance.
- A replica in a separate zone from the primary instance.
- A replica in a separate zone and a separate region.

It's possible to create a replica in one, two, or all three of these locations if needed. We recommend putting your replicas in a different zone. This ensures that replicas will continue to operate even if there is an outage in the zone that contains the primary instance.

Cloud SQL replicates data to replicas asynchronously. As a result, a replica could be behind the primary by seconds or minutes, depending on certain load characteristics and any transmission lag due to the replica's location. If a replica gets promoted as the primary, the RPO can be non-zero.



Here are some scenarios where Cloud SQL read replicas help:

1. **Workload segregation between reads and writes (offloading):** Instead of all transactions being processed by the primary instance, reads can be served by replicas. This ensures the primary instance only processes critical writes or reads as part of transactions, improving the performance of the database and the application.



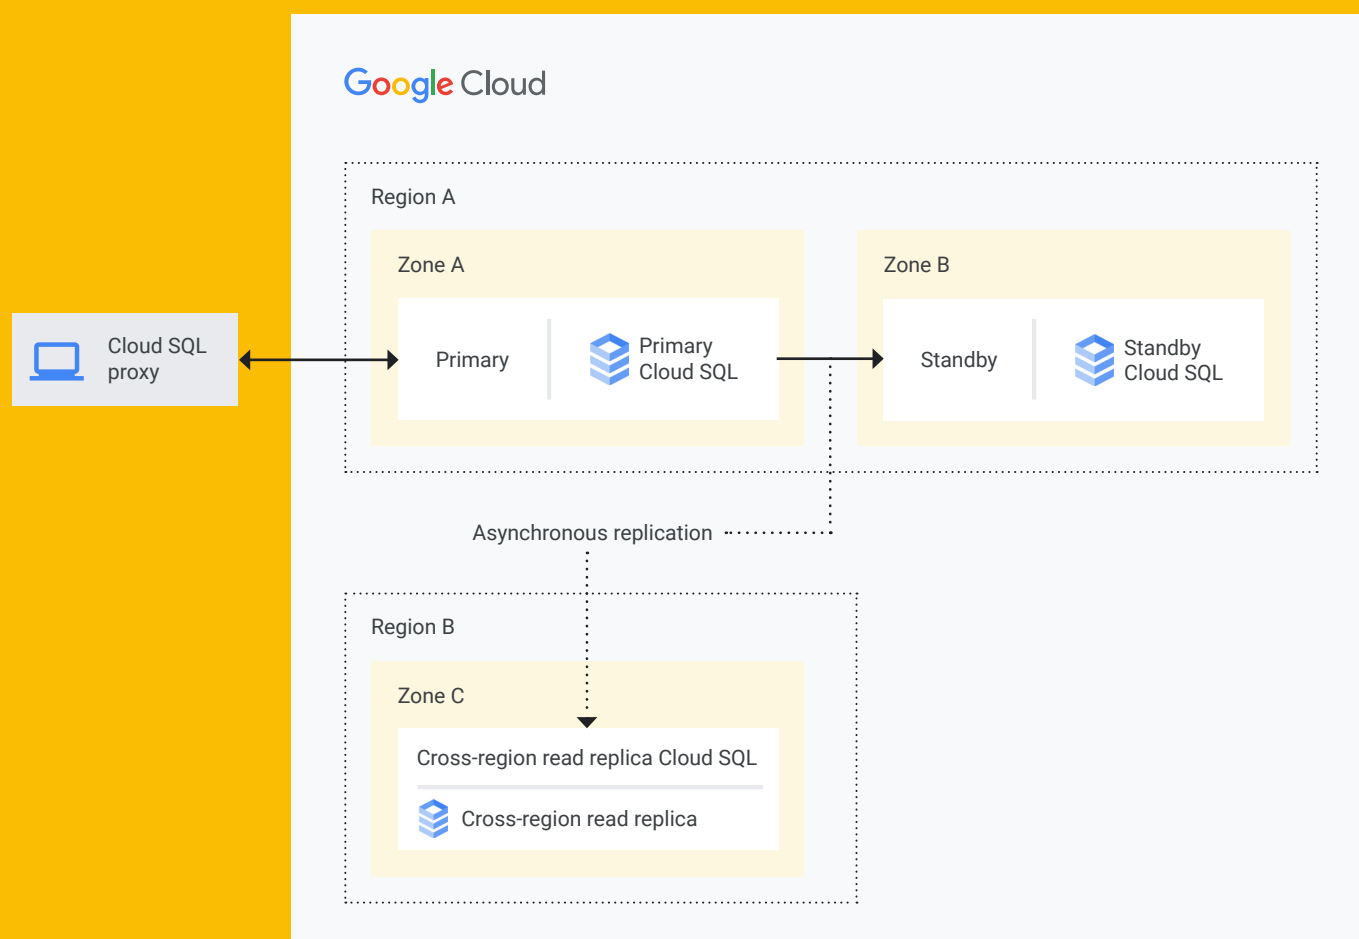
2. **Better RPO and RTO of the database:** Many single instance databases suffer from longer RPO due to the time it takes to restore and recover from backups, which may not be acceptable from a business perspective. Read replicas help keep RPO down to a few seconds to improve availability and continuity for single-instance configurations. Similarly, promoting a read replica to a primary takes a few seconds, which also helps improve RTO. While the HA configuration is a better solution from a RPO and RTO perspective, replication gives users the ability to use replicas for reads which may be more suitable for certain use cases.

Cloud SQL disaster recovery

The Cloud SQL HA configuration provides protection against zonal failures. However, a failure that affects an entire region, typically due to a natural disaster or multiple catastrophic failures, could render both the primary and standby instances unavailable.

To recover from this type of failure, you can set up cross-region read replicas in a different region from where a primary is located.

Below is an example of a Cloud SQL DR schematic:




RTO and RPO for cross-region read replicas are minutes rather than hours. The application tier is also multi-regional, which means an application can continue processing even when the primary region has a failure.

There are a several critical considerations to take into account when using this configuration:

1. Failover to a read replica in the standby region is not automatic when your primary region fails, although it can be automated using a script..
2. Your operations team should determine if the replica needs to be promoted to the new primary and initiate the promotion of the replica.
3. Once the new primary instance is active, you will need to reconfigure the clients to connect to the new primary.
4. You will need to convert to an HA configuration and create a new cross-region replica in a third Google Cloud region to protect a newly promoted primary. We are working on reducing the number of steps to make this stage easier.
5. You can create a new cross-region replica and initiate a failover process to fail back to the original primary region once the outage in the primary region resolves. As the failover to the original primary region is manually triggered and not based on an outage, you can choose a day and time for this maintenance activity to reduce the impact on your customers.

For more details, we recommend reading [this article](#) on Cloud SQL disaster recovery.



Application considerations for an HA deployment

Databases do not operate in isolation. They provide data services to applications and constitute a larger ecosystem of data being moved across systems. Apart from leveraging Cloud SQL's high availability capabilities to create reliable deployments, it's also critical to monitor your database and applications using observability tools.

You should always maximize observability in applications and ensure they can handle load to avoid instability and failure. Monitoring should be implemented early in the development cycle, well before you deploy the system in production. Observability should also go beyond monitoring to include tracing, profiling, and debugging.

Cloud SQL provides [Query Insights](#) that helps you detect, diagnose, and prevent query performance problems in Cloud SQL databases. It is currently available for the PostgreSQL database engine, and we plan to extend support to other engines. At the time of the writing of this paper, Cloud SQL Insights for MySQL was available in preview.

Suppose your applications have direct connections to your database, with each application opening multiple connections simultaneously. In this case, your database could be overwhelmed with a flood of connections.

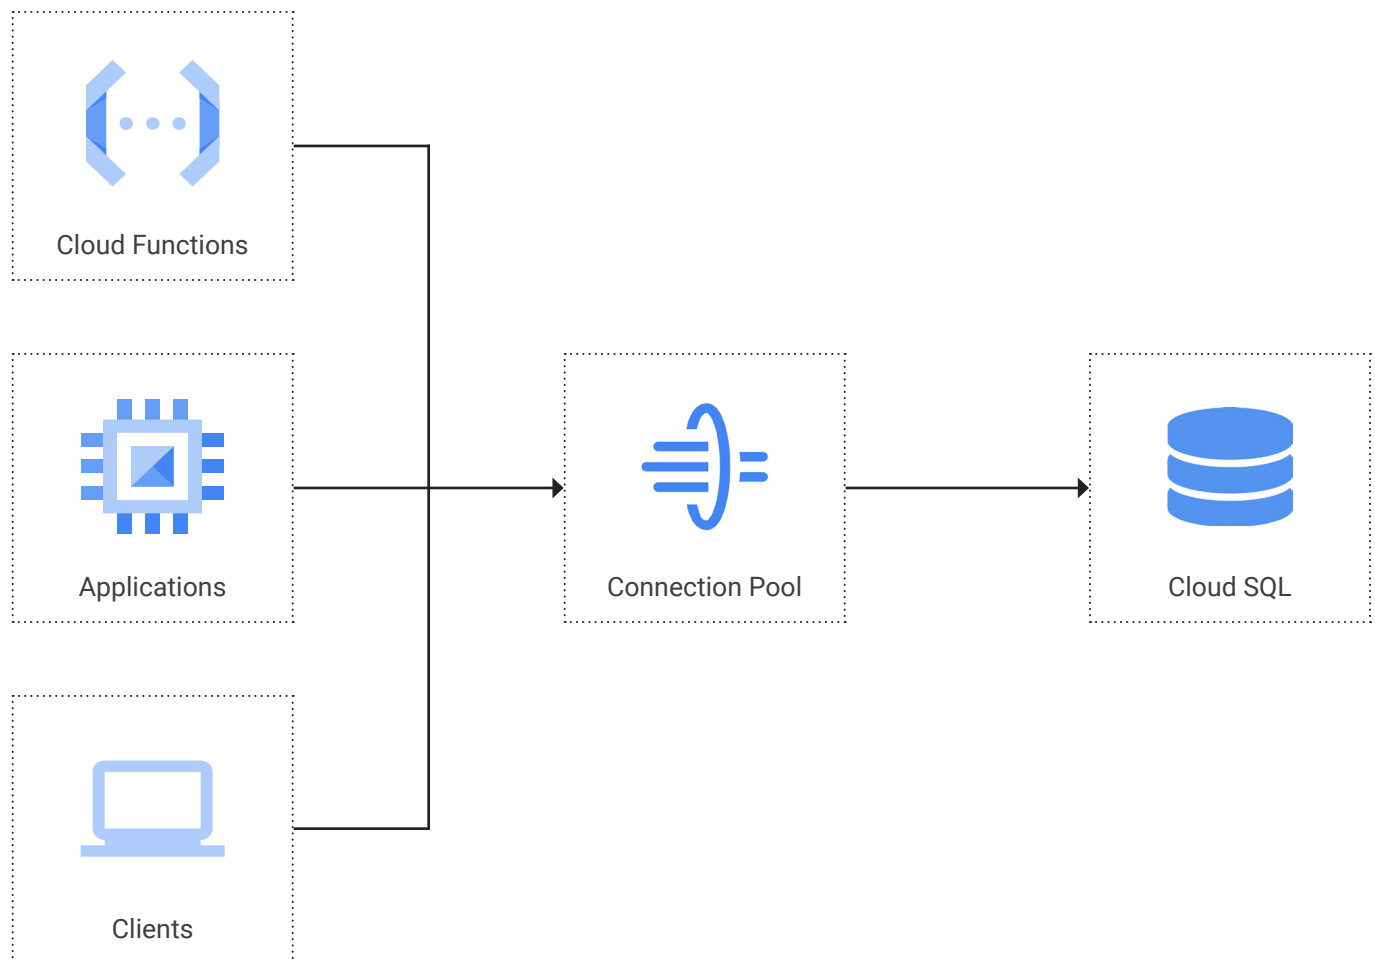
Connections consume memory, CPU resources, operating system processes, and more. Allowing too many connections at once can contribute to out of memory (OOM) events, triggering reboots (unexpected downtime) and impacts Cloud SQL stability. You should connect applications to your database through a connection pooling mechanism to avoid this situation.



Google Cloud

All of our Cloud SQL database engines have several open source and third-party proxy and connection pooling solutions that you can use based on the functionality needed. In the future, we plan to offer a managed connection pooling solution for Cloud SQL.

A sample connection pooling schematic is shown below:

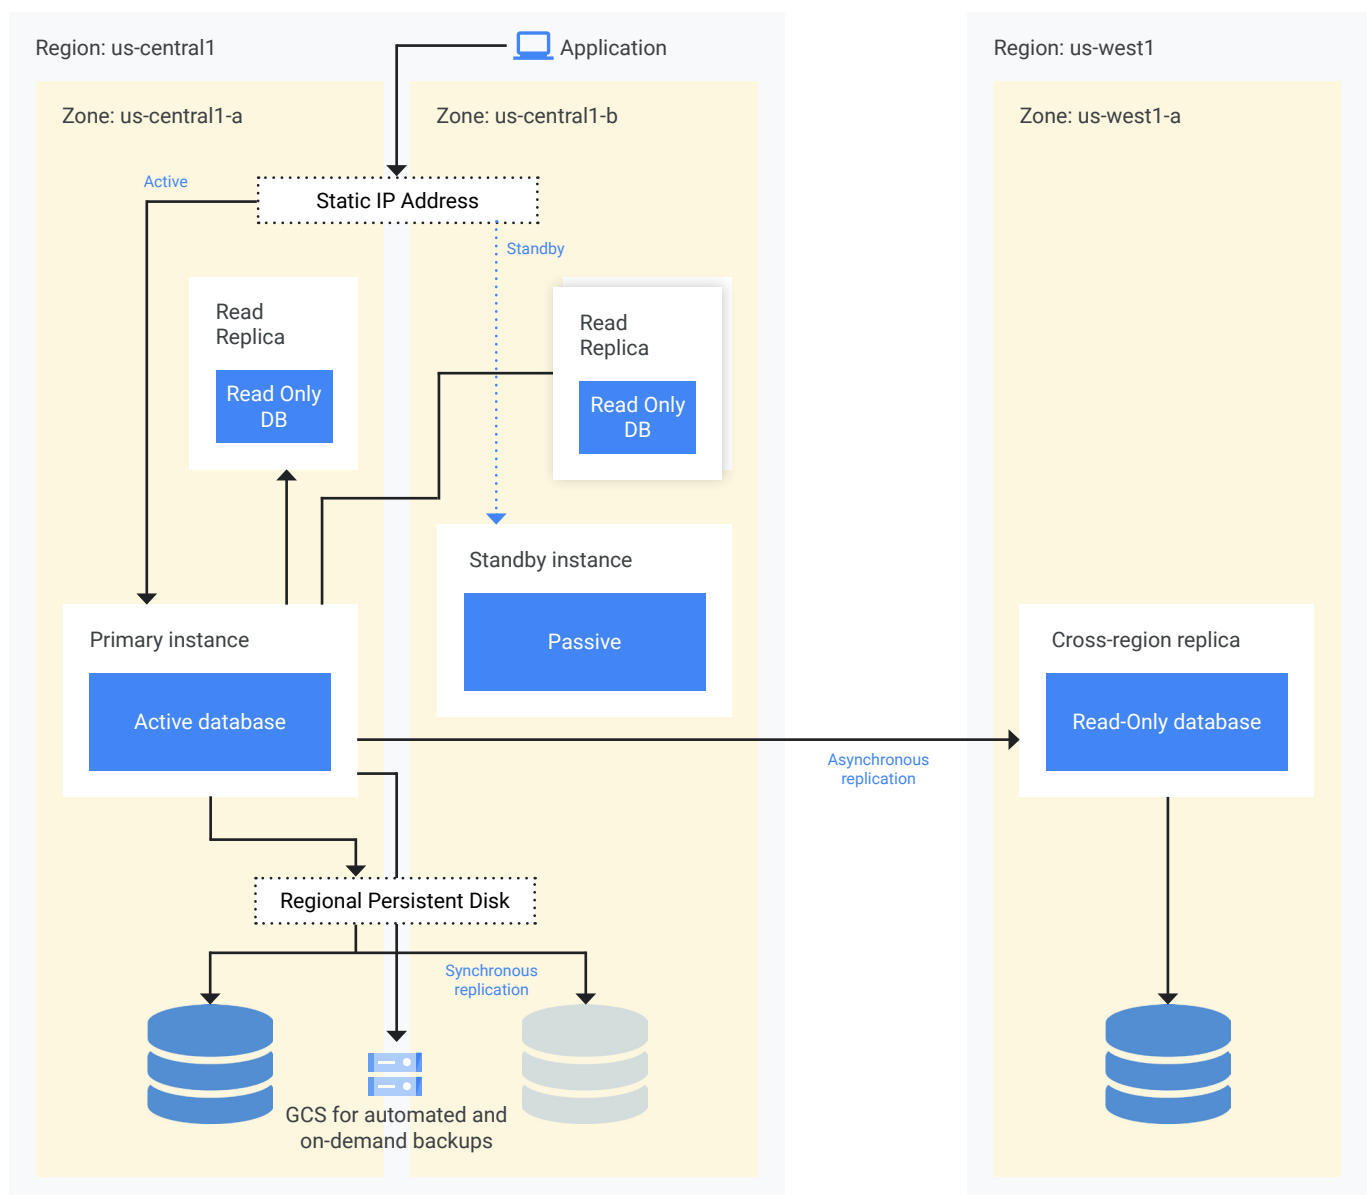


Google Cloud

Rather than failing, you should always design applications to degrade gracefully under overload by serving partial responses or allowing limited functionality. In our experience, it's also important to implement exponential backoff with randomization in the error retry logic of client applications. This prevents applications from sending continuous requests to the database, which can make the outage worse.

For more details, we recommend following [this framework](#) to help you build and deploy reliable systems.

The High Availability and Disaster Recovery architecture discussed throughout this guide is summarized in the following diagram.



You should always test the deployment of your regional database HA configuration and the cross-region replica for disaster recovery. This ensures that the HA solution functions as designed, enforces a standard database HA failover procedure, and establishes data consistency.

After the initial HA testing of your database, we highly recommend testing the availability characteristics of the entire system, including failover of the application and the database together.

Periodic business continuity drills with defined standard operating procedures are also necessary to ensure that the high availability solution works seamlessly in the event of a disaster.



Resources

[Cloud SQL maintenance blog](#)

[Cloud SQL backup and recovery concepts](#)

[Cloud SQL Disaster Recovery](#)

[Google Cloud Architecture framework - Reliability](#)

