Cloud SQL (https://cloud.google.com/sql/)
Documentation (https://cloud.google.com/sql/docs/)
MySQL (https://cloud.google.com/sql/docs/mysql/) Guides

Overview of backups

MySQL | PostgreSQL (https://cloud.google.com/sql/docs/postgres/backup-recovery/backups) |
 SQL Server (https://cloud.google.com/sql/docs/sqlserver/backup-recovery/backups)

This page describes how backups of your Cloud SQL instance work, and how you can use them
to restore your data to the same or another instance.

For step-by-step directions for scheduling backups or creating an on-demand backup, see
Creating and Managing On-Demand and Automatic Backups
 (https://cloud.google.com/sql/docs/mysql/backup-recovery/backing-up).

What backups provide

Backups provide a way to restore your Cloud SQL instance to recover lost data or recover from
a problem with your instance. You should enable automated backups for any instance that
contains data that you need to protect from loss or damage.

Enabling automated backups, along with binary logging, is also required for some operations,
such as clone and replica creation.

What backups cost

Cloud SQL retains up to 7 automated backups for each instance. For First Generation
instances, the cost of these backups is included in your instance cost; their storage space does
not count against the allotted storage for the instance. For Second Generation instances, the
storage used by backups is charged at a reduced rate. See the pricing page
 (https://cloud.google.com/sql/docs/mysql/pricing#2nd-gen-storage-networking-prices) for more
information.

For First Generation instances using the per use (https://cloud.google.com/sql/pricing#per_use)
billing plan, scheduled backups incur a small incremental charge, because the instance is

activated to take the backup.

> **Note:** Second Generation is replacing First Generation; support for First Generation instances ends January 30, 2020. To upgrade a First Generation instance to Second Generation, see [Upgrading a First Generation Instance to Second Generation](https://cloud.google.com/sql/docs/mysql/upgrade-2nd-gen) (https://cloud.google.com/sql/docs/mysql/upgrade-2nd-gen).

About backup size

Backups for Second Generation instances are incremental; they contain only data that has changed since the previous backup was taken. This means that your oldest backup is a similar size to your database, but the sizes of subsequent backups depend on the rate of change of your data. When the oldest backup is deleted, the size of the next oldest backup increases so that a full backup still exists.

Types of backups

Cloud SQL performs two types of backups:

- [On-demand backups](#on-demand-backups) (#on-demand-backups)
- [Automated backups](#automated-backups) (#automated-backups)

On-demand backups

For Second Generation instances, you can create a backup at any time. This could be useful if you are about to perform a risky operation on your database, or if you need a backup and you do not want to wait for the backup window. You can create on-demand backups for any Second Generation instance, whether the instance has automatic backups enabled or not.

On-demand backups are not automatically deleted the way automated backups are. They persist until you delete them or until their instance is deleted. Because they are not automatically deleted, on-demand backups can have a long-term effect on your billing charges if you do not delete them.

Automated backups

When you enable automated backups, you specify a 4-hour backup window. The backup starts during the backup window. When possible, schedule backups when your instance has the least activity.

Where backups are stored

Backups locations include:

- Default locations (#default-backup-location) that Cloud SQL selects, based on the location of the original instance
- Custom locations (#custom-backup-location) that you choose when you do not want to use the default location

Default backup locations

MySQL Second Generation instances:

By default, Cloud SQL stores backup data in two regions for redundancy. If there are two regions in a continent, the backup data remains on the same continent. Because there is only one region in Australia, backup data from the Sydney region is stored in a location in Asia. For the São Paulo region, backup data is stored in a US-based location.

Note: You cannot restore from any backup, including multi-region backups, if the region where the backup was originally taken is unavailable.

MySQL First Generation instances: Backup data is stored in the continent where the instance resides.

Custom backup locations

MySQL Second Generation:

When you select a custom location for your backup data, Google stores the backup data in that location. For a complete list of valid regional values, see Instance Locations (https://cloud.google.com/sql/docs/mysql/locations#location-r). For a complete list of multi-regional values, see Multi-regional locations (https://cloud.google.com/sql/docs/mysql/locations#location-mr).

See Setting and viewing a custom location for backups
 (https://cloud.google.com/sql/docs/mysql/backup-recovery/backing-up#locationbackups).

Can I export a backup?

No, you cannot export a backup. You can only export instance data. See Exporting data from
Cloud SQL (https://cloud.google.com/sql/docs/mysql/import-export/exporting).

About the special backup user

Cloud SQL creates a special database user, `cloudsqladmin`, for each instance, and generates a
unique instance-specific password for it. Cloud SQL logs in as the `cloudsqladmin` user to
perform automated backups.

How backups affect instance operations

First Generation instances: For First Generation instances, backups are taken by using `FLUSH`
`TABLES WITH READ LOCK` to create a snapshot. This means that writes to the instance are
blocked while the backup is taken. (The instance remains online, and reads are unaffected.)

Typically, backups complete within a few seconds, but if a large amount of data has been
written since the last backup, the backup takes longer to complete.

If there is a pending operation at the time of the backup attempt, Cloud SQL usually makes
several attempts within the window to complete the backup. Operations that block backup are
long-running operations such as import, export, update (for example, an instance metadata
change), and instance restart.

During a long-running operation, such as loading data, you can temporarily disable automated
backups.

Note: Second Generation is replacing First Generation; support for First Generation instances ends January
30, 2020. To upgrade a First Generation instance to Second Generation, see Upgrading a First Generation
Instance to Second Generation (https://cloud.google.com/sql/docs/mysql/upgrade-2nd-gen).

Second Generation instances: For Second Generation instances, the `FLUSH TABLES WITH READ LOCK` flag is not used for backups. This means that writes and other operations are unaffected by backup operations.

What's next

- Schedule automated backups
  (https://cloud.google.com/sql/docs/mysql/backup-recovery/backing-up#schedulebackups).

- Create an on-demand backup
  (https://cloud.google.com/sql/docs/mysql/backup-recovery/backing-up#on-demand).

- Restore from a backup
  (https://cloud.google.com/sql/docs/mysql/backup-recovery/restoring#restorebackups).

- Perform a point-in-time recovery
  (https://cloud.google.com/sql/docs/mysql/backup-recovery/restoring#pitr).