

This white paper shows you how to implement an inexpensive disaster recovery solution for Microsoft SQL Server on Google Cloud Platform. This paper is for you if you're:

- A system administrator.
- A database administrator.
- Looking to protect a production SQL Server instance.
- Don't have or don't want to use a separate data center or other colocation site.

In this white paper, you'll learn to:

- Copy your SQL Server databases to Google Cloud Storage.
- Spin up a SQL Server virtual machine (VM) on Google Compute Engine to test your backups.
- Use some free scripts to restore your backups to the most recent point in time.
- Test your backups to help make sure they're not corrupted.

With conventional database transaction-log shipping, every time the primary SQL Server takes a backup of a log, another SQL Server notices the newly created backup file and restores it. This way, you always have a ready-to-go standby SQL Server. But this approach can become expensive. Using a cloud-based strategy, you can still take log backups at a regular interval and then sync those files to the cloud. When disaster strikes, or if you want only to test your backups, you spin up a SQL Server in the cloud, restore those backups, and then test your databases to check for corruption.

Syncing the backup files up into the cloud can be the most challenging step in the process. Here are some issues to keep in mind:

- These files might not be small, especially for your full backups.
- The files might not be stored on your SQL Server. For example, you might be writing your backups to a file share or UNC path.
- Your Internet bandwidth might be limited.

- You need to rely on command-line tools for the syncing.

To read the full white paper, click the button:

[Download the PDF \(/solutions/sql-server-disaster-recovery-plan-compute-engine.pdf\)](/solutions/sql-server-disaster-recovery-plan-compute-engine.pdf)