

DATABASES

DR for cloud: Architecting Microsoft SQL Server with GCP

Christoph Bussler

Solutions Architect, Google Cloud

August 21, 2019



Database disaster recovery (DR) planning is an important component of a bigger DR plan, and for enterprises using Microsoft SQL Server on Compute Engine, it often involves critical data. When you're architecting a disaster recovery solution with Microsoft SQL Server running on [Google Cloud Platform](#) (GCP), you have some decisions to make to build an effective, comprehensive plan.

Microsoft SQL Server includes a variety of disaster recovery strategies and features, such as Always On availability groups or Failover Cluster Instances. And Google Cloud is designed from the start for resilience and availability. There are several types of data centers available within GCP where you can map SQL Server's availability features based on your specific requirements: zones and regions. Zones are autonomous data centers co-located within a GCP region. These regions are available in different geographies such as North America or APAC.

However, there is no single disaster recovery strategy to map Microsoft SQL Server DR features to Google Cloud's data center topology that satisfies every possible combination

of disaster recovery requirements. As a database architect, you have to design a custom disaster recovery strategy based on your specific use cases and requirements.

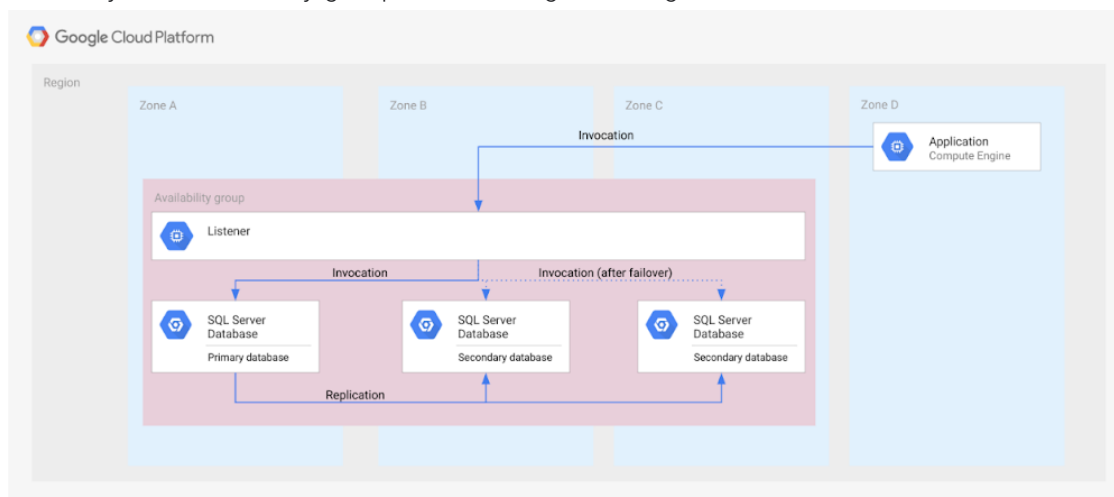
Our new [Disaster Recovery for Microsoft SQL Server](#) solution provides information on Microsoft's SQL Server disaster recovery strategies, and shows how you can map them to zones and regions in GCP based on your business's particular criteria and requirements. One example is deploying an availability group within a region across three zones (shown in the diagram below).

For successful DR planning, you should have a clear conceptual model and established terminology in place. In this solution, you'll find a base set of concepts and terms in context of Google Cloud DR. This includes defining terms like primary database, secondary database, failover, switchover, and fallback.

You'll also find details on recovery point objective, recovery time objective and single point of failure domain, since those are key drivers for developing a specific disaster recovery solution.

Building a DR solution with Microsoft SQL Server in GCP regions

To get started with implementing the availability features of Microsoft SQL Server in the context of Google Cloud, take a look at this diagram, which shows the implementation of an Always On availability group in a GCP region, using several zones:



In the new solution, you'll see other availability features, like log shipping, along with how they map to GCP. In addition, features in Microsoft SQL Server that are not deemed availability features—like server replication and backup file shipping—can actually be used for disaster recovery, so those are included as well.

Disaster recovery features of Microsoft SQL Server do not have to be used in isolation and can be combined for more complex and demanding use cases. For example, you can set up availability groups in two regions with log shipping as the transfer mechanism between the regions.

[Disaster Recovery for Microsoft SQL Server](#) also describes the disaster recovery process itself, how to test and verify a defined disaster recovery solution, and outlines a basic approach, step-by-step.

Learn more about [SQL Server on GCP](#) and [check out all of our solutions](#).

POSTED IN: [DATABASES](#) [GOOGLE CLOUD PLATFORM](#)

RELATED ARTICLES

[No database behind with Cloud SQL
SQL Server](#)

[Cloud Spanner amps up SLA, adds CSV
support, and sharpens monitoring details](#)