



Failure scenarios and resiliency with Spanner



Table of Contents

Chapter 1

Introduction 3

Chapter 2

Failure modes 4

Chapter 3

Recovery with operation intervention 5

Chapter 4

Recovery outside GCP 9

Chapter 1

Introduction

The Digital Operational Resilience Act ([DORA](#)) and [PRA Outsourcing and Third Party Risk](#) both provide clear guidance that a regulated financial service provider needs to manage risk. Google Cloud has provided 3 scenarios which can be used by financial services customers as guidance for the use and risk management of the [Spanner](#) database service.

The aim of this paper is provide details that enables a regulated financial service provider to know how to handle each scenarios which are provided as follows:

1. Automatically provide intervention for Spanner failure
2. Operator intervention without causing material outage
3. Total failure or stressed exit of Spanner

The next set of sections provides more information on each aspect.

Overview

This document aims to explore the different failure scenarios of Spanner categorized into three levels of severity, with scenario when operating outside of GCP, being the most critical. The three failure scenarios are examined specifically in the context of Spanner, a Google distributed SQL database management and storage service providing features such as global transactions, strongly consistent reads, and automatic multi-site replication and failover.

Spanner is
globally-distributed
relational database
service that
combines the best
of both worlds:
massive scalability
with strong data
consistency

Chapter 2

Failure Modes

This document will discuss the three levels of severity of failure scenarios in Google Cloud, with scenario C being the most critical:

- **A - Scenarios automatically handled by Spanner without operator intervention.**
 - These encompass issues ranging from individual virtual machine (VM) malfunctions to outages affecting a single region.
 - These failures can arise due to hardware or software issues, configuration errors, or localized power or network disruptions.
 - These failures do not require customer intervention, do not impact Spanner availability, and have no impact on business as usual (BAU) operations of a customer.
- **B - Scenarios which require operator intervention but can safely continue running on Spanner.**
 - They occur when a single Multi-Region Spanner configuration is degraded. This means that some Spanner replicas in the configuration are unavailable or performing poorly.
 - These failures do not cause global degradation of Spanner or GCP services, but might require human intervention and may cause disruption to business as usual (BAU) operations of a customer.
- **C - Scenarios which require operating outside of GCP.**
 - They are caused by a global outage or unavailability of Google Cloud Platform (GCP) services.
 - This level requires human intervention and disrupts business as usual (BAU) operations.
NOTE: This is the same approach for a stressed exit for a financial regulated entity.



This document focuses on scenarios which will require intervention (Scenarios B & C) and may cause some form of service degradation. Spanner's inherent service resilience drives scenario A.

Chapter 3

Recovery with Operational Intervention

This section provides a detailed discussion of scenarios which require operator intervention but can safely continue running on Spanner. Such scenarios occur when a quorum of Spanner replicas are unavailable or severely degraded, impacting normal business operations. In this scenario, other Spanner instance configurations may still be available and performing normally. Additionally, Google Cloud Platform (GCP) as a whole is still available, even though a single Spanner instance configuration is down.

Spanner and Failure Domains

Spanner, a Google distributed SQL database management and storage service providing features such as global transactions, strongly consistent reads, and automatic multi-site replication and failover. Internally, Spanner is divided into 'sections' called failure domains.

Customers can deploy their primary and disaster recovery (DR) Spanner instances in independent failure domains. This means that the primary and DR instances are located in different failure domains, which are isolated from each other. This isolation helps protect the DR instance from a failure in the primary failure domain.

Note : This is the recommended configuration for critical systems.



Each failure domain has its own set of resources and Paxos quorum, which is a group of replicas that must agree on all changes to the data. Independent failure domains have built-in mechanisms that protect them from each other, so that a failure in one failure domain is extremely unlikely to affect the others.

Each failure domain corresponds to a specific Spanner instance configuration. For example, the us-east4 and asia-south1 failure domains are independent. Similarly, the us-east4 and nam3 failure domains are independent, even though nam3 hosts a replica from us-east4. Multiple failure domains can overlap geographies (such as nam3 and nam6) or not (such as nam3 and nam8).

Each failure domain is hosted by a separate set of resources, which are updated and released on a staggered schedule, both within and across failure domains. Each failure domain is backed by its own set of Paxos groups, and each failure domain is operated by its own control plane.

Failure domains are designed to be isolated from each other, so that a failure in one failure domain is extremely unlikely to affect the others. If a failure domain fails, other failure domains will remain available with very high probability. Failover to an independent failure domain is a reliable way to resume normal operations.



Impact of Failures

Failures can be categorized into two types: data plane failures and control plane failures. Control planes provide APIs that allow administrators to create, read, update, delete, and list (CRUDL) cloud resources; these are generally administrative functions. The data plane is from a customer point of view the core of the service, providing them the ability to store, query and retrieve data from the service.

Data plane failures generally have a greater impact than control plane failures. Data plane failures can take a variety of forms, such as all operations failing, only new client connections failing, or operations succeeding but with increased latency.

Control plane failures generally have a lower impact than data plane failures because many control plane operations are not as business-critical, such as database creation and schema changes. However, some control plane operations, such as node scaling, can be as impactful as data plane failures.

Spanner has extensive external and internal monitoring to ensure the reliability and performance of both control and data plane services.

- External monitoring: Spanner's console metrics include end-to-end latency, CPU utilization, error rates, and top transaction statistics. Users can also build dashboards and automated alerts on top of a wide array of metrics using OpenCensus/Open Telemetry integration.
- Internal monitoring: Spanner's SLA is backed by extensive internal monitoring across a wide array of mechanisms, including customer traffic monitoring, probes, and resource health metrics. Spanner operations are backed by 24/7 on-call rotations managed by geographically distributed teams, with secondary and "all hands on deck" escalation paths.

How to operate in such scenarios

To operate in the context of scenarios where a single Spanner configuration is unavailable but other configurations are available, customers should have:

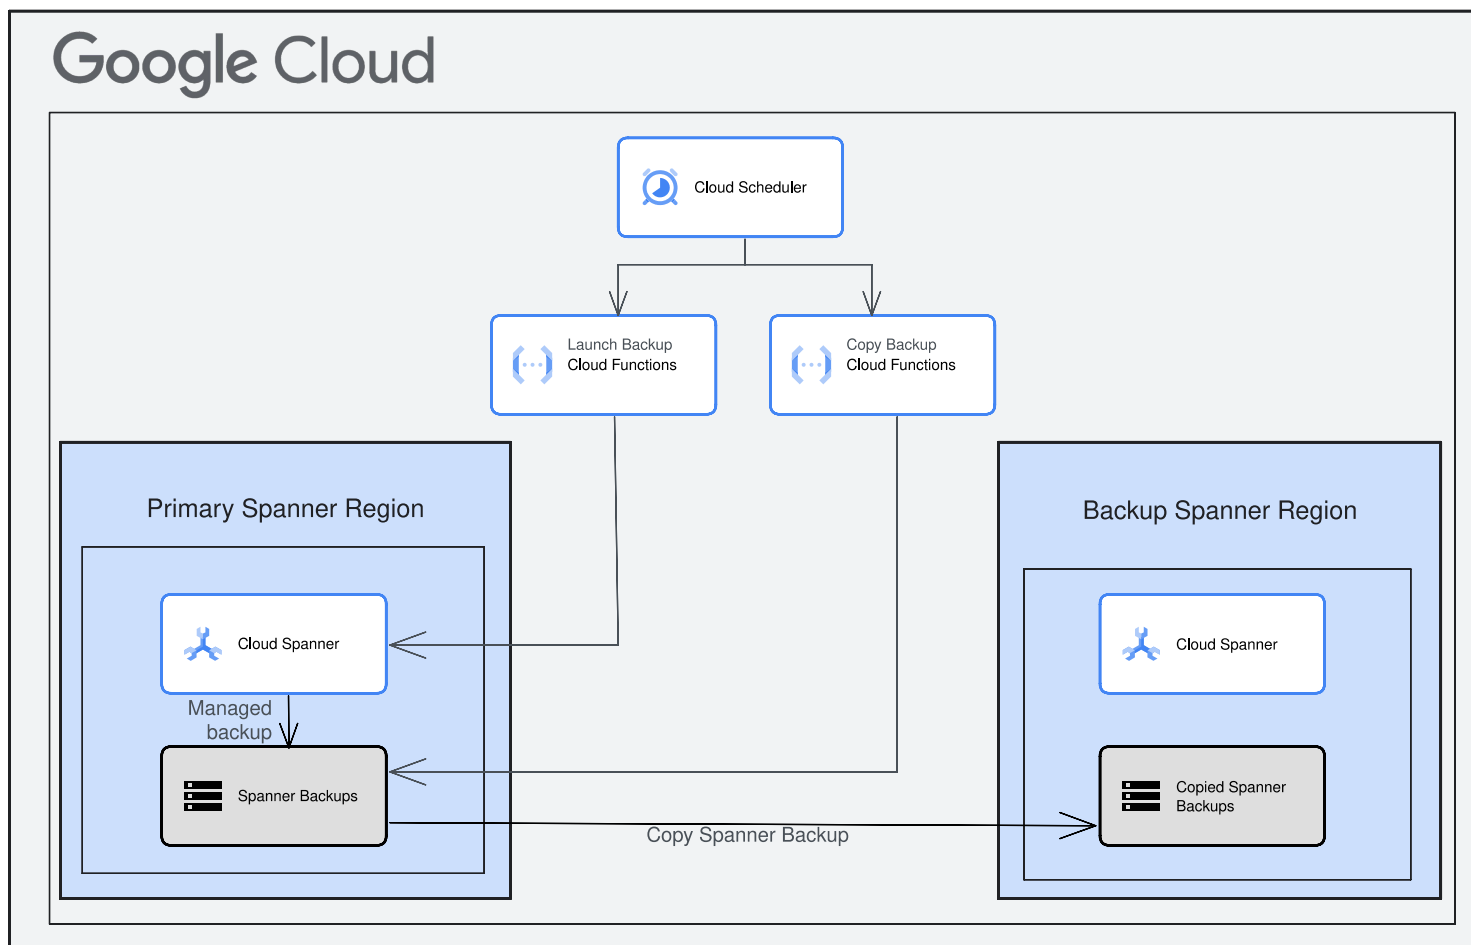
- Have regular scheduled [managed backups](#).
- Have an automatic task to [copy the backups to another instance](#) in a Spanner backup failure domain (in another region)

The existence of these backups in a separate region serve as a safeguard against data loss and ensure that critical data can be restored promptly in the event of extended downtime or outage impacting a single failure domain.

Scheduling of managed backups can currently be achieved using a [Cloud Scheduler and Cloud Functions task](#).

Automatic copying of backups to another region can also be achieved using Cloud Scheduler and a Cloud Functions task to monitor the existence of backups in the main region and copy them to a backup region whenever new backup sets are created.





In the case of a failure scenario where the main region becomes unavailable, these backups in the backup region can be used to restore and populate a Spanner database in the backup region.

Recovery outside GCP

This section provides a detailed discussion of scenarios which require operating outside of GCP to resolve. Such scenarios occur when for whatever reason, customers can no longer use any Spanner configuration on GCP. The customer needs to ensure their data is available outside of Google Cloud and is in a recognized format. Such scenarios require the customer to move their workloads and data outside Google Cloud.

Some of these scenarios may allow the customer to perform a managed exit, where the Google Cloud services remain available but the customer has decided to move to another provider, which would allow a graceful shutdown and export of data. Other scenarios would be a stressed exit, where Google Cloud services are shut down abruptly and the customer's data in Google Cloud is no longer available.

Managed/planned exit

- Given advance notification of requirement to exit
- Time to choose exit plan, and dates/times to exit
- Application can be shut down cleanly
- Data on GCP remains accessible
- Expected time from notification/decision to completion of exit plan: Months to Year

Stressed/unplanned exit

- No advance notification
- Application stopped suddenly/uncleanly
- Data in GCP no longer accessible
- Expected time from notification/decision to completion of exit plan: Days to Weeks

In both cases, manual intervention will be required, and business as usual (BAU) operations will be disrupted for an extended period.

Exporting data in a managed exit scenario

In a managed exit scenario, the customer still has access to the Spanner database and Google Cloud Services, and wishes to take a copy of their data so that they can move it elsewhere.

Spanner team maintains a [Spanner Export Dataflow template](#) which can make a snapshot copy of the entire database into a set of [Apache Avro](#) files (which is an open file format) on Google Cloud Storage. From there, users can copy their data to any external storage service, for portability.

In the case of a managed exit, the customer can retrieve the data stored in spanner by:

- Stopping the application workloads, or putting the application into a read-only mode thus preventing any further writes to the Spanner database
- Launch a Spanner Export task to create a copy of their data in a Google Cloud Storage bucket.
- Copy the data files generated by this export job from the Google Cloud Storage bucket to an external storage service.


Exporting data in a stressed exit scenario

Customers may want to maintain a near real time copy of all Spanner data outside of Google Cloud. This could be used at any time to rebuild the data in another provider/service without any dependencies on Google Cloud Platform.

Note: This scenario is understood to be a very long tail risk event, where customers are forced to provision services on another platform due to Google Cloud being unavailable.

The [Spanner Export Dataflow template](#) can be performed as a scheduled task to create snapshot copies of the database in ([Apache Avro](#)) format on Google Cloud Storage. For the incremental changes to the database between each snapshot, Spanner's native Change Data Capture system - [Spanner Change Streams](#) - can be used with another Google-provided streaming Dataflow template which continuously takes the [change stream records and stores them as Apache Avro](#) files in Google Cloud Storage.

With these two Google-provided and supported templates, the customer will have both snapshot and incremental database updates in an open file format ([Apache Avro](#)) on Google Cloud Storage. The customer can then set up a task to automatically replicate these files as they are generated to an external storage service, so that in the scenario that they are disconnected from Google Cloud Platform, they will have a near real-time copy of their data on this external storage service.



Summary

This document provides a deep dive into the various failure scenarios for Spanner and how to mitigate them for regulated industries. Spanner has a rigorous approach to ensuring service resiliency, with built-in mechanisms that automatically handle the vast majority of failures. However, there are scenarios where operator intervention may be required, and this document details those as well as the most severe failure scenario, a total failure of Spanner, requiring operation outside of GCP. It also discusses how to recover from such failures, including a managed exit and a stressed exit.



References

- Cloud Spanner documentation: This is the official documentation for Cloud Spanner, and it contains a comprehensive overview of the service, including its features, architecture, and how to use it.
<https://cloud.google.com/spanner/docs>
- Spanner: Google's Globally-Distributed Database: This research paper provides a detailed technical overview of the design and implementation of Spanner.
<https://research.google.com/archive/spanner-osdi2012.pdf>
- Databases overview | Cloud Spanner - Google Cloud: This page provides an overview of Cloud Spanner databases, including their structure and how they are created and managed.
<https://cloud.google.com/spanner>
- Backup and Disaster Recovery solutions with Google Cloud: <https://cloud.google.com/solutions/backup-dr>