

Market Insight Report Reprint

Spanning the globe is what Google Cloud Spanner is supposed to do

November 5 2020

James Curtis

Google's Spanner distributed database not only provides SQL transactions, but also guarantees those transactions, regardless of region or location. Released in early 2017, Google Cloud continues to innovate on Spanner because it powers several Google Cloud applications along with a multitude of services.

451 Research



This report, licensed to Google, developed and as provided by 451 Research, LLC, was published as part of our syndicated Market Insight subscription service. It shall be owned in its entirety by 451 Research, LLC. This report is solely intended for use by the recipient and may not be reproduced or reposted, in whole or in part, by the recipient without express permission from 451 Research.

Introduction

Google Cloud Spanner is Google's distributed database that can be deployed within a single region or across multiple regions, including different regions in different countries. A decade in development and released in early 2017, Spanner is Google's go-to database service that provides transactional guarantees in a distributed environment. Despite Spanner being a fairly new cloud database service (generally available in May 2017), Google has internally used it since 2012, and is deeply committed to Spanner's long-term success – because not only does it power some of the world's largest enterprise applications, it also powers several Google applications, from Google Search to Google Maps to Gmail.

THE 451 TAKE

The Spanner paper was published in 2012 after Google had internally developed Spanner and was using it to power internal applications. The Spanner database was released as a GCP service in early 2017. The Google Spanner paper also inspired a handful of other databases of similar ilk. However, what separates Spanner from other distributed SQL databases is how it manages time globally, as well as its greater than, or equal to, five-nines availability SLAs for customers. Built atop Google's massive network infrastructure, Spanner includes a series of atomic clocks and GPS antennas that sit on each node and are coordinated by Google's TrueTime technology to manage time. As such, Spanner can guarantee transactional consistency globally. The idea of a distributed relational database is starting to catch on, even as many organizations may not currently need a global footprint for their applications. Regardless, Google Spanner has proven its influence by inspiring a relatively new database category. And while global may not be everyone's requirement, it's likely that in the near future 'global' will be the new normal.

Context

Google has a long history of data management and database expertise going back to the early 2000s, when the company published a paper on GFS (Google File Storage) in 2003. Google then released a series of papers on data management and processing frameworks that later spawned a variety of technology products, inside and outside of Google. In 2003, it released a paper on MapReduce, which led to Apache Hadoop, and spawned other Hadoop-related open source products. In 2006, a paper was published on Bigtable, a distributed system to manage data, which led to Google Bigtable as well as Apache HBase, in addition to other NoSQL databases.

In 2010, Google wrote a paper on Dremel, which led to Apache Drill, and later inspired the development of Google BigQuery. In 2012, Google published a paper on Spanner that led to the Google Spanner database and inspired other distributed SQL databases. The company also wrote papers on Milwheel, Dataflow and Tensorflow, leading to open source projects such as Tensorflow and Bean, as well as inspiring and contributing to a host of Google cloud products like Pub/Sub, Dataflow, Bigtable, ML and Spanner.

Spanner is part of Google's portfolio of relational databases that also includes Cloud SQL for managed MySQL, PostgreSQL and SQL Server, and Google Bare Metal Service, a managed hardware offering for specialized Oracle workloads. Spanner, however, is considered a cloud-native, fully managed database service that delivers five-nines of availability, according to Google.

From a use case perspective, Google reports Spanner being used across a spectrum of industries, such as gaming, retail, technology, healthcare and financial services. The company also reports enterprises using Spanner in production, many of which have presented at Google conferences. These include Uber (transportation), Square (financial services), ShareChat (social media) and LL Bean (retail), to name a few.

Technology

Spanner is a relational, globally distributed database. It's a fully managed, cloud-native database service that runs on Google's Cloud Platform (or Google Cloud), which means that database administration, among other processes, is greatly reduced. The service includes built-in functionality such as automated failure recovery, sharding and replication. Maintenance updates are also handled automatically without any maintenance downtime. The company promotes Spanner as having zero RPO (recovery point objective) and RTO (recovery time objective) times.

Spanner functions as a distributed database that scales horizontally, including across multiple regions, while also maintaining transactional ACID guarantees (up to 99.999% SLA, according to Google). Spanner's distributed architecture decouples the storage layer from the compute layer. Google Colossus is the storage layer, which is based on Google's early distributed file system, GFS. Colossus not only provides data durability, but allows any Spanner node to read data using an ANSI-compliant SQL dialect. Moreover, the compute nodes hold a metadata copy of the data that tracks where each shard resides. As such, when shards are split or combined, Spanner does not need to move data between nodes to maintain its performance.

From a user perspective, when a database instance is created, whether within a single region or across multiple regions in different geographies, the data (called shards or splits) is distributed synchronously to multiple replica nodes (read-write, read-only, witness) in different zones that leverage a Paxos-based replication scheme. This synchronous replication strategy ensures that each replica (read-write or read-only) has the latest data that can service incoming reads or writes.

The backbone of Spanner, however, is Google's Cloud network infrastructure that spans some 24 regions and 73 zones globally, consisting of Americas, Asia-Pacific, Europe, Middle East and Africa. Leveraging this global network gets to the heart of how Spanner manages global time. Google refers to its global time management technology as TrueTime, which consists of integrated and connected GPS antennas and atomic oscillators (clocks) that reside on each node in the network. The job of TrueTime is to ensure that time is consistent throughout the network, and to make corrections in case of any 'clock drifts.' TrueTime also commits a timestamp to every transaction when it occurs. Additional functionality in Spanner manages the ordering of the transactions such that users can perform reads across nodes, even when nodes may span different cloud regions.

Competition

Given Google Cloud Spanner's focus on relational workloads and transactional guarantees, Spanner is likely to be compared to a host of traditional relational vendors, many that provide scaling along with distributed deployment capabilities. Vendors include Oracle, IBM, Amazon, Microsoft and MariaDB, as well as a few open source products such as MySQL and PostgreSQL, along with the vendors that support these databases.

Scaling is offered for many of these databases, but it can vary between vendors. Oracle, for instance, has its RAC (Real Application Clusters) technology, but it also provides Oracle Sharding. IBM points to its pureScale technology and Db2 on Cloud for scaling. Microsoft provides scaling for Azure SQL Database using Elastic Database tools, but scaling is also available for Azure Database for PostgreSQL, using Hyperscale, which leverages technology from the Citus Data acquisition. Amazon Aurora is available on AWS.

MariaDB is another that provides a distributed deployment option for SkySQL called MariaDB Xpand, based on the company's acquisition of Clustrix. Percona, which provides support for MySQL, offers the Percona Distribution for MySQL, which includes tooling and products to run MySQL in a distributed architecture.

However, Google Spanner can also be categorized among a group of databases that not only provide SQL support but also are highly horizontally scalable. This group of databases also favors data consistency within their architectures to support ACID transactions. Vendors include CockroachDB, Yugabyte, NuoDB, VoltDB, PingCAP and MemSQL.

A collection of NoSQL vendors is also worth mentioning, with some providing (varying) support for ACID transactions in distributed environments. Open source offerings include Apache Cassandra and Redis, as well as DataStax and Redis Labs, which commercially leverage Cassandra and Redis, respectively. However, other Spanner competitors might include Microsoft's Azure Cosmos, Amazon DynamoDB, MongoDB, Couchbase, InterSystems, Aerospike and FairCom, to name a few.

SWOT Analysis

STRENGTHS

A key element of Spanner is its ability to manage time globally using Google's TrueTime technology, which leverages node-based atomic clocks and GPS antennas that sit on each node that make up Google's massive global cloud network infrastructure.

OPPORTUNITIES

Running transactional (OLTP) workloads with required guarantees are a staple of many organizations. However, it is often difficult to predict future workload and system demand. Existing legacy systems often struggle to adapt to changing demands. While positioned as a globally distributed database, Spanner provides the ability to scale as part of future needs within an organization.

WEAKNESSES

Although Spanner is a managed cloud service that incorporates cloud-native technologies, it is not serverless, however, because customers are still required to provision their own compute instances when scaling.

THREATS

The idea of a distributed relational database can have great appeal for many organizations for a variety of reasons, whether it's providing transactional guarantees within a region or across multiple regions. But differentiation, between the numerous database products/services can get a bit fuzzy as vendors incorporate contrasting architectural approaches when handling scaling, availability, ACID guarantees, failover, recovery and so forth. This can potentially challenge organizations as they make purchasing decisions.

CONTACTS

The Americas +1 877 863 1306 market.intelligence@spglobal.com

Europe, Middle East & Africa +44 20 7176 1234 market.intelligence@spglobal.com

Asia-Pacific +852 2533 3565 market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers. (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its Web sites, www.standardandpoors.com (free of charge) and www. ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.