

INDUSTRY

Gaming

TECHNOLOGIES

- Google Cloud Spanner
- Google Compute Engine
- Google Compute Firewall
- Cloud Identity & Access Management
- Stackdriver, Apache NiFi
- Apache JMeter, Terraform.

BUSINESS NEED

A leading video game design studio needed to modernize its backend infrastructure stack and game database, hosted in on-premise Microsoft SQL Server Enterprise Edition, to allow for more distributed computing, easy scalability, better performance, and optimized costs.

SOLUTION

Pythian had the ability to quickly understand the unique complexities and requirements of the client's source and target databases. A POC was developed to evaluate Cloud Spanner's feasibility and performance under load, the ease of migration, the impact to existing architecture and design, and to identify any risks or issues.

RESULTS

By modernizing their backend infrastructure stack on the cloud, the client was in a position to reap the benefits of distributed computing across different geographical regions, along with flexible scaling as they continue to add new games to their portfolio.

VIDEO GAME STUDIO LEVERAGES CLOUD SPANNER TO IMPROVE DATABASE PERFORMANCE, WITH HELP FROM PYTHIAN

BUSINESS NEED

A leading video game design studio needed to modernize its backend infrastructure stack and game database, hosted in on-premise Microsoft SQL Server Enterprise Edition, to allow for more distributed computing, easy scalability, better performance, and optimized costs. The client's existing SQL Server databases, combined with an intermediate data access layer in C++ for database communications and object marshalling, weren't able to scale up and down quickly enough—leading to frustrating bottlenecks as players were forced to queue to log in and wait for matches.

Because on-premise game databases are often extremely complex and can be a challenge to scale and manage, even when well optimized, the client wanted to evaluate the feasibility of Google Cloud Spanner by partnering on a proof of concept (POC) with Pythian, an expert in both on-premise and cloud databases. Cloud Spanner is a horizontally scalable, globally consistent, relational database service that allows for horizontal scaling across rows, regions, and continents. It's often favored as a gaming database because it combines the benefits of a relational database with non-relational horizontal scalability.

SOLUTION

With expertise in multiple database technologies, along with 20-plus years of on-prem, hybrid, and cloud data management experience, Pythian had the ability to quickly understand the unique complexities and requirements of the client's source and target databases. A POC was developed to evaluate Cloud Spanner's feasibility and performance under load, the ease of migration, the impact to existing architecture and design, and to identify any risks or issues. Evaluation of the cloud solution began with a two-node configuration, with Pythian gradually adding up to four nodes to achieve

the best performance at the lowest cost. Next, a migration approach was established involving a sample of 65 database tables moved from SQL Server to Cloud Spanner. These tables were directly related to the database's most commonly executed, yet slowest-running queries. This was followed by an analysis of the client's SQL Server environment, to identify which items needed technical workarounds or refactoring to provide similar functionality on Cloud Spanner.

To extract data from the client's SQL Server, Apache NiFi was deployed on a Google Cloud VM. This was accompanied by a helper script, to streamline implementation by creating a pipeline for SQL Server objects, with the data staged in AVRO format on Google Cloud. Deployment of Google infrastructure—including Cloud Spanner, Google Compute Engine for NIFI, Google Compute Engine for SQL, and Google Compute Firewall for Resources – was performed through Terraform, which allows the deployment of infrastructure as code (IaC). Cloud Identity and Access Management (IAM) was recommended for security purposes, as it allows admins to grant permission to users or groups without modifying Cloud Spanner instances or individual database permissions.

RESULT

The POC successfully demonstrated that migrating to Cloud Spanner would provide better query performance than the client's SQL Server production environment. In addition, it showed Cloud Spanner's impressive scalability, since it allows DBAs to easily tweak the number of nodes using Google Cloud Console (each Cloud Spanner node provides up to 2 TB of storage). Furthermore, Pythian demonstrated to the client that it can easily export data from its SQL Server environment and import it into Cloud Spanner; therefore, opening up further possibilities including disaster recovery and ingestion for analytics through other services such as Google BigQuery.

The POC included daily standups alongside the client to help them understand and learn more about Cloud Spanner and the migration process. Pythian also conducted a session each Friday to show progress, share the latest code, and demonstrate other aspects of the

project. Through these means, Pythian educated the client on how to optimize the performance of Cloud Spanner instances, including the following best practices:

- Always design a schema that prevents hotspots and other performance issues as much as possible
- Always place critical compute resources within the same region as your Cloud Spanner Instance
- Always provision enough Cloud Spanner nodes to keep high priority total CPU utilization under 65%

Performance testing of the POC was conducted for feasibility, performance and scalability using Apache JMeter, while database monitoring of various metrics such as CPU utilization and latency was conducted using Stackdriver.

Following delivery of the POC, it was confirmed that Cloud Spanner's scalability of both performance and storage would make it an ideal choice should the client migrate its game database to the cloud. Cloud Spanner's combination of a relational database structure with horizontal scaling has proven an excellent fit for replacing on-premise game state and authentication databases, while the ability to scale easily by adding nodes provides flexibility as player bases grow and games become more complex.

By modernizing their backend infrastructure stack on the cloud, the client would be in a position to reap the benefits of distributed computing across different geographical regions, along with flexible scaling as they continue to add new games to their portfolio.

ABOUT PYTHIAN

Pythian excels at helping businesses around the world use data and the cloud to transform how they compete and win in the data economy. From cloud automation to machine learning, Pythian leads the industry with proven innovative technologies and deep data expertise. For more than 20 years Pythian has built its reputation by delivering solutions to the toughest data challenges faster and better than anyone else.

WORLDWIDE OFFICES

Ottawa, Canada
New York City, USA
London, England
Sydney, Australia
San Francisco, USA
Hyderabad, India

© The Pythian Group Inc., 2020