STRIIM BLOG

Online Enterprise Database Migration to Google Cloud

Christoph Bussler, Google Cloud and Ferhat Hatay, Edward Bell, Striim
August 17, 2020 · 5 minute read

Migrate to cloud

Migrating existing workloads to the cloud is an formidable step in the journey of digital transformation for enterprises. Moving an enterprise application from on premises to run in the cloud, or modernizing with the best use of cloud-native technologies, is only part of the challenge. A major part of this task is to move the existing enterprise databases while business continuously operate at full speed.

Pause never

How the data is extracted and loaded into the new cloud environment plays a big role in keeping the business critical systems performant. Particularly for enterprise databases supporting mission-critical applications, avoiding downtime is a must-have requirement during migrations to minimize both the risk and operational disruption.

For business critical applications, the acceptable downtime precipitously approaches zero. All the while, moving large amounts of data, and essential testing of the business critical applications can take days, weeks, or even months.

Keep running your business

The best practice in enterprise database migration, to minimize and even altogether eliminate the downtime, is to use online database migration that keeps the application running.

In the online migration, changes from the enterprise source database are captured non-intrusively as real-time data streams using Change Data Capture (CDC) technology. This capability is available for most major databases, including Oracle, Microsoft SQL Server, HPE NonStop, MySQL, PostgreSQL, MongoDB, and Amazon RDS, but has to be harnessed in the correct way.

In online database migration, first, you initially load the source database to the cloud. Then, any changes in the source database that have happened since you were executing the initial load are applied to the target cloud database continuously from the real-time data stream. The source and target databases will remain up to date until you are ready to completely cut over. You will also have the option to fallback to the source all along, further minimizing risks.

Integrate continuously

Online database migration also provides essential data integration services for the new application development in the cloud. The change delivery can be kept running while you develop and test the new cloud applications. You may even choose to keep the target and source databases in sync indefinitely typically for continuous database replication in hybrid or multi-cloud use cases.

Keep fresh

Once the real-time streaming data pipelines to the cloud are set up, businesses can easily build new applications, and seamlessly adopt new cloud services to get the most operational value from the cloud environment. Real-time streaming is a crucial element in all such data movement use cases, and it can be widely applied to hybrid or multi-cloud architectures, operational machine learning, analytics offloading, large scale cloud analytics, or any other scenario where having up-to-the-second data is essential to the business.

Change Data Capture

Striim, in strategic partnership with Google Cloud, offers online database migrations and real-time hybrid cloud data integration to Google Cloud through non-intrusive Change Data Capture (CDC) technologies. Striim enables real-time continuous data integration from on-premises and other cloud data sources to BigQuery, Cloud Spanner, Cloud SQL for PostgreSQL, for MySQL, and for SQL Server, as well as Cloud Pub/Sub and Cloud Storage as well as other databases running in the Google Cloud.

Replicate to Google Cloud

In addition to data migration, data replication is an important use case as well. In contrast to data migration, data replication continuously replicates data from a source system to a target system "forever" without the intent to shut down the source system.

An example target system in the context of data replication is BigQuery. It is the data analytics platform of choice in Google Cloud. Striim supports continuous data streaming (replication) from an on-premises database to BigQuery in Google Cloud in case the data has to remain on-premises and cannot be migrated. Striim bridges the two worlds and makes Google Cloud data analytics accessible by supporting the hybrid environment.

Transform in flight

Data migration and continuous streaming in many cases transports the data unmodified from the source to the target systems. However, many use cases require data to be transformed to match the target systems, or to enrich and combine data from different sources in order to complement and complete the target data set for increased value and expressiveness in a simple and robust architecture. This method is frequently referred to as Extract Transform Load, or ETL.

Striim provides a very flexible and powerful in-flight transformation and augmentation functionality in order to support use cases that go beyond simple one-time data migration.

More to migrate? Keep replicate!

Enterprises in general have several data migration and online streaming use cases at the same time. Often data migration takes place for some source databases, while data replication is ongoing for others.

A single Striim installation can support several use cases at the same time, reducing the need for management and operational supervision. The Striim platform supports high-volume, high velocity data with built-in validation, security, high-availability, reliability, and scalability as well as backup-driven disaster recovery addressing enterprise requirements and operational excellence.

The following architecture shows an example where migration and online streaming is implemented at the same time. On the left, the database in the Cloud is migrated to the Cloud SQL database on the right. After a successful migration the source database is going to be removed. In addition, the two source databases on the left in an on-premises data center are continuously streamed (replicated) to BigQuery for analytics and Cloud Spanner for in-Cloud processing.

Keep going

In addition, Striim as the data migration technology is implemented in a high-availability configuration. The three servers on Compute Engine form a cluster, and each of the servers is executing in a different zone, making the cluster highly available and protecting the migration and online streaming from zone failures or zone outages.

Accelerate Cloud adoption

As organizations modernize their data infrastructure, integrating mission-critical databases is essential to ensure information is accessible, valuable, and actionable. Striim and Google Cloud's partnership supports Google customers with a smooth data movement and continuous integration solutions, accelerating Google Cloud adoption and driving business growth.

Learn more

To learn more about the enterprise cloud data integration questions, feel free to reach out to Striim and check out these references:?

Google Cloud Solution Architecture: Architecting database migration and replication using Striim

Blog: Zero downtime database migration and replication to and from Cloud Spanner

Tutorial: Migrating from MySQL to BigQuery for Real-Time Data Analytics

Striim Google Virtual Hands-On Lab: Online Database Migration to Google Cloud using Striim

Self-paced Hands-on Lab: Online Data Migration to Cloud Spanner using Striim

Further Reading


Striim Expands Partnership with Google Cloud to Offer Database Migration Service

Striim Announces Real-Time Data Migration to Google Cloud Spanner


How to Migrate Oracle Database to Google Cloud SQL for PostgreSQL with Streaming Data Integration