

Streaming Integration for Google BigQuery

The Striim® Platform continuously loads real-time data to BigQuery analytics data warehouse from across the enterprise with minimal impact on data sources. By using Striim for real-time data integration, BigQuery customers have access to timely, pre-processed data from on-premises or cloud data sources. Striim ingests, transforms, and delivers real-time data from a wide variety of sources, including relational databases, data warehouses, log files, sensors, messaging systems, Hadoop, and NoSQL solutions.

Using Striim to build reliable and secure streaming data pipelines with in-flight data processing, Google Cloud customers can easily use BigQuery modern analytics data warehouse for time-sensitive operational decision making that drives business transformation.

Real-Time Data Integration

By enabling streaming data pipelines, Striim offers a secure, reliable, and scalable service for real-time ingestion, stream processing, and delivery of structured, semi-structured, and unstructured data from a variety of on-premises or cloud sources into Google BigQuery.

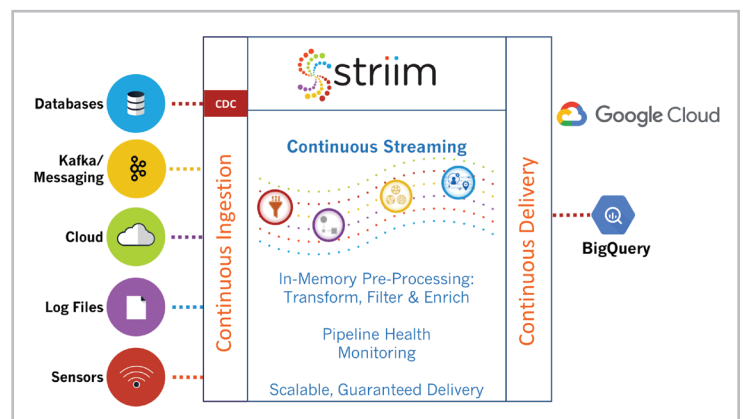
For enterprise databases and data warehouses such as Oracle, Oracle Exadata, SQL Server, MySQL, MariaDB, HPE NonStop, Amazon RDS, and MongoDB, the platform offers non-intrusive change data capture (CDC) to minimize the impact on source systems. Striim also supports widely used enterprise data formats, including JSON, XML, AVRO, delimited, binary, free text, and change records. Other cloud data sources, such as Amazon S3, Salesforce.com, are also among supported data sources.

In-Flight Data Processing

Striim transforms data in flight using SQL-based continuous queries and UI-based operators to filter, aggregate, transform, mask, encrypt, and enrich streams of real-time data in memory before delivering to BigQuery in a consumable format. With in-flight data processing, Striim reduces latency, supports operational workloads, and accelerates time to insight.

BENEFITS

- Ingest real-time data without modifying or slowing down source systems
- Continuously collect a diverse range of data sources, including databases, NoSQL systems, logs, messaging, Hadoop, and sensors
- Easily offload analytical workloads to the Google Cloud by moving data in real time and in the desired format
- Simplify data architecture with in-flight data processing
- Migrate from existing data warehouses or transactional systems without downtime or data loss



Real-Time Operational Intelligence

Striim enables BigQuery customers to reap the benefits of cloud-based analytics by allowing them to use timely, rich, and granular data from a wide range of sources. With access to up-to-date data in BigQuery, users can build time-sensitive analytics solutions such as detecting and predicting fraud and security threats instantaneously, enabling location-based marketing that provide significant operational value to the business. By using granular change data from transactional databases, companies can build accurate machine learning models in BigQuery to optimize business operations.

Simplified Data Architecture

Unlike traditional batch-based ETL solutions, Striim continuously ingests granular and larger data sets for richer analytics. It does so without impacting source systems and enables sub-second latency by processing the data in-memory, while it is streaming. Striim's end-to-end streaming data integration delivers a simplified and scalable data architecture which:

- Minimizes ETL workloads by performing transformations while data is in motion
- Supports compliance with privacy-related regulations via in-flight data masking and encryption
- Optimizes data storage by filtering out unnecessary data
- Enables end-to-end recoverability and full resiliency without needing to manage multiple and varied components, and network hops across these components

Phased Migration Without Downtime

Striim's real-time data synchronization capabilities enable data migration from operational data warehouses and databases to BigQuery without requiring downtime. To prevent data loss, Striim continuously monitors data movement, processing, and provides delivery validation that all data has been successfully moved to the target. Customers can run their legacy data warehouses in parallel with their BigQuery infrastructure for a phased migration to Google Cloud and perform extensive testing to minimize risks.

WHY STRIIM?

- Continuous and real-time data integration from a wide variety of enterprise data sources
- Supports compliance with privacy-related regulations via in-flight encryption, data masking, and filtering
- Designed for high-volume, high-velocity data flows
- Event guarantees with non-intrusive Change Data Capture from databases
- Enterprise-grade with built-in security, scalability, and reliability features
- Quick and easy to deploy via SQL-based queries and wizards-based UI