

Streaming Integration and SQL-Based Processing for Kafka

The Striim® platform provides an enterprise-grade, non-intrusive streaming integration and SQL-based stream processing solution for moving change data from a wide variety of sources to Kafka. Striim further offers real-time data movement from Kafka to Cloud, Big Data, and other enterprise data targets. While streaming, Striim enables users to query, analyze, and visualize data in Kafka via interactive dashboards.

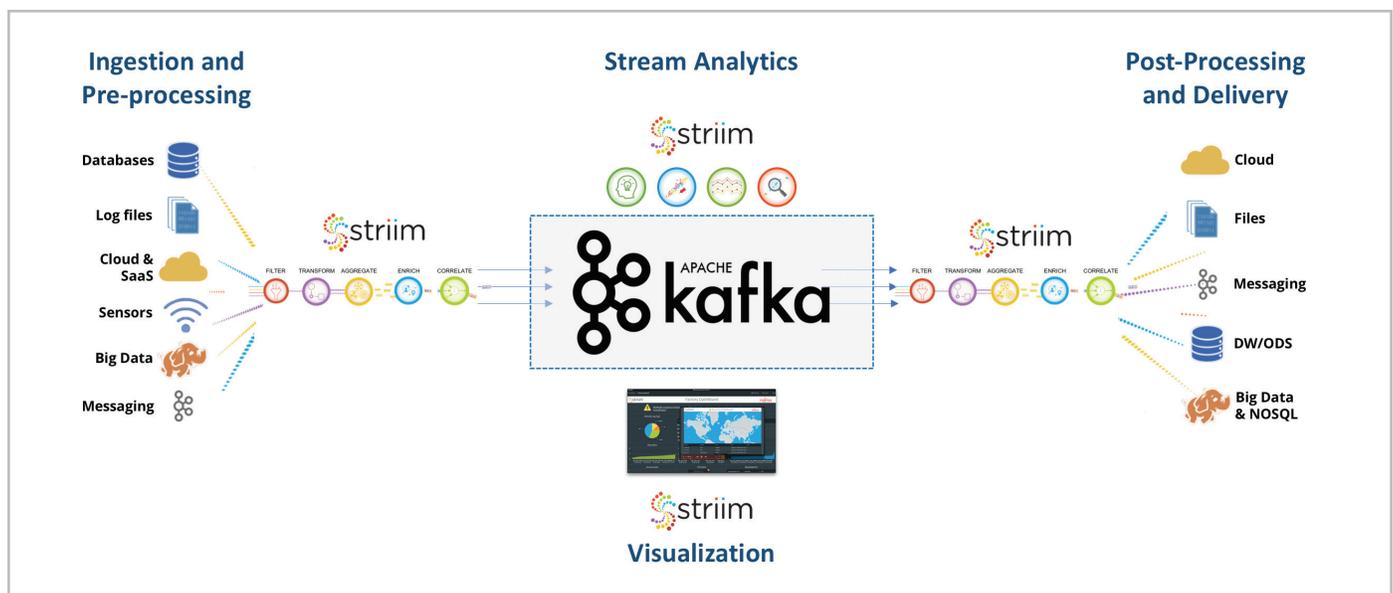
Real-time Data Ingestion with CDC

The Striim platform ingests real-time, streaming data from a variety of sources out-of-the-box, including databases, files, message queues, and devices — on-premises or in the cloud. For enterprise databases such as Oracle, SQL Server, MySQL, HPE NonStop, and MariaDB, the platform offers non-intrusive change data capture (CDC) for efficient and real-time data integration. Striim supports major data formats, including JSON, XML, AVRO, delimited, binary, free text, and change records.

With a drag-and-drop UI and wizards, Striim simplifies creating data flows from popular sources to Kafka. The data can be delivered “as-is,” or go through a series of transformations and enrichments to create exactly the data structure and output needed.

BENEFITS

- Ingest real-time data into Kafka with low impact
- Process data without extensive coding
- Easily deliver Kafka data to the enterprise
- Get immediate insights and alerts
- Visualize data in Kafka with the same product



Striim offers a comprehensive streaming integration solution for Kafka, with built-in analytics and visualization.

SQL-based Stream Processing

In addition to offering an end-to-end platform with enterprise-quality security, scalability, and fault tolerance, Striim makes in-stream data processing for Kafka fast and easy. Through SQL-based continuous queries, Kafka users can process filter, aggregate, transform, join, and enrich multiple streams of real-time data to meet the needs of different Kafka consumers. All without needing APIs or developers. Striim also offers flexible windows that turn unbounded, infinite data streams into continuously changing bounded sets of data. The platform maintains exactly-once-processing capabilities even when working with jumping or sliding time windows.

Delivery to Target On-premises or Cloud

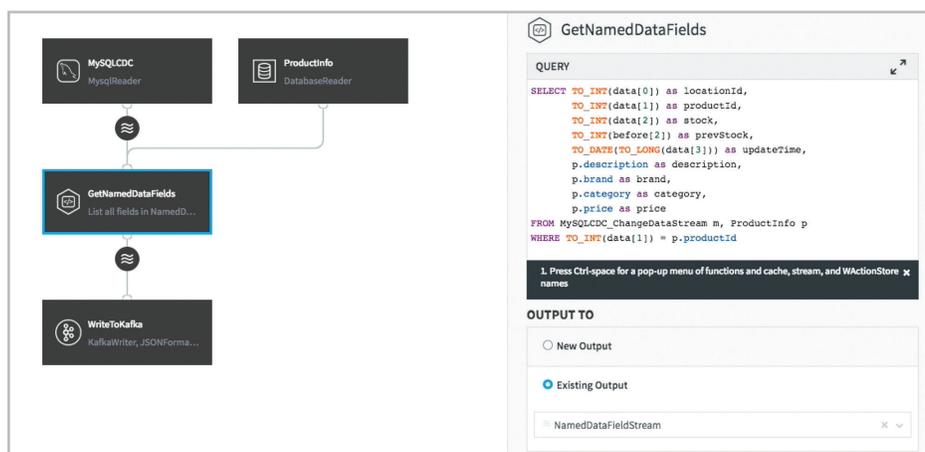
The Striim platform can continuously apply pre-processed, streaming data from Kafka to a broad range of targets, including databases, files, message queues, Hadoop environments, and cloud data stores such as Azure Blob Storage, Azure SQL Database, Amazon Redshift, and Google BigQuery. The data format is configurable and can be applied to the raw Kafka data, or to the results of data processing and analytics.

Analytics and Visualization for Data in Kafka

Striim enables Kafka users to maximize the value from their streaming data by analyzing it in-flight to capture time-sensitive information. The platform performs various analytical processing—including pattern matching, correlation, outlier detection, predictive analytics—in-memory, as the data is flowing through. It pushes the results to live, interactive dashboards to provide real-time visibility and enable data exploration. As with data processing, Striim performs stream analytics via SQL-based continuous queries.

WHY STRIIM?

- Built-in security, scalability, and reliability
- Real-time data integration from a wide variety of data sources
- Non-intrusive CDC from databases with event guarantees
- In-flight enrichment via built-in cache
- Exactly-once processing with time windows and parallelism
- Quick to deploy via SQL-like queries and wizards-based UI



A drag-and-drop UI with wizards and SQL-based language simplify creating data flows from popular sources to Kafka.

For more information, or to schedule a demo, please contact the Striim team at info@striim.com.